Energy Policy 38 (2010) 4921–4932

Contents lists available at ScienceDirect

Energy Policy

journal homepage: www.elsevier.com/locate/enpol

A symmetric safety valve

Dallas Burtraw *, Karen Palmer, Danny Kahn

Resources for the Future, 1616 P Street NW, Washington, DC 20036, USA

ARTICLE INFO

Article history:
Received 26 February 2009
Accepted 25 March 2010
Available online 8 May 2010

Keywords:
Climate change
Cost management
Cap and trade

ABSTRACT

How to set policy in the presence of uncertainty has been central in debates over climate policy. Concern about costs has motivated the proposal for a cap-and-trade program for carbon dioxide, with a "safety valve" that would mitigate against spikes in the cost of emission reductions by introducing additional emission allowances into the market when marginal costs rise above the specified allowance price level. We find two significant problems, both stemming from the asymmetry of an instrument that mitigates only against a price increase. One is that most important examples of price volatility in cap-and-trade programs have occurred not when prices spiked, but instead when allowance prices collapsed. Second, a single-sided safety valve may have unintended consequences for investment. We illustrate that a symmetric safety valve provides environmental and welfare improvements relative to the conventional one-sided approach.

1. Introduction

Policymakers advance economic efficiency when they set policy goals at levels that equate the marginal costs of additional pollution controls with the marginal benefits of improvements in environmental quality. Increasingly policymakers employ incentive-based approaches, such as tradable allowances or taxes, to achieve these goals in a least cost manner. However, when attempting to set goals, policymakers face a great deal of uncertainty about the costs and benefits to society of achieving a particular goal and, in particular, how those costs and benefits are likely to change over time. The presence of uncertainty affects the choice of policy instruments from an efficiency perspective (Weitzman, 1974; Roberts and Spence, 1976; Pizer, 2002).

The issue of how to set policy in the presence of uncertainty has been particularly salient in climate policy, where meaningful efforts to control emissions could prove much more costly than prior regulatory efforts to limit emissions of air pollution, and where the costs and benefits of controlling emissions of greenhouse gases are highly uncertain. One proposal to neutralize the possibility of unexpected increases in cost in a cap-and-trade program is a "safety valve" that serves as a ceiling on the price of an emission allowances by increasing the provision of emission allowances in the market if and when a price ceiling is achieved (Pizer, 2002; Kopp et al., 2002).[1] This proposal gained practical relevance for a cap on carbon dioxide (CO_2) beginning in 2004 when it was incorporated in the climate policy section of the comprehensive energy policy advanced by the National Commission on Energy Policy and incorporated into draft legislative language by Senator Bingaman (D-NM). The safety valve ceiling on the allowance price was also a feature of a climate cap-and-trade legislative proposal introduced in the House of Representatives by Representatives Udall (D-NM) and Petri (R-WI). A safety valve provision also was incorporated in the Bush Administration's proposed Clear Skies Act that would have imposed national caps on emissions of sulfur dioxide (SO_2), nitrogen oxides (NO_x) and mercury (Hg) from electricity generators. Today it is a central element of climate policy discussions, and quite often a controversial one, in part because of mistrust about the level at which the safety valve price ceiling would be set. Murray et al. (2009) offer an amended version of a safety valve policy by providing a quantitative limit on the number of allowances that could be issued by the safety valve. Related approaches involving quantity-constrained strategic allowance reserves have been incorporated into the recent legislation passed by the U.S. House of Representatives (Waxman, D-CA, and Markey, D-MA, H.R. 2454) and in the Senate (Kerry, D-MA, and Boxer, D-CA, S. 1733).

Most advocates of the safety valve approach focus exclusively on the situation where realized costs of reducing pollution turn out to be higher than expected and thus the original emissions cap would no longer be viewed as efficient. However, in virtually every case when an incentive-based form of regulation like emissions cap and trade has been used, costs have been overestimated rather than underestimated prior to the regulation taking effect (Harrington et al., 2000). The most common explanation is that in many cases baseline emission levels were overestimated and the emissions reductions necessary to achieve a target level wound up being less than anticipated, thus reducing the costs of compliance.

* Corresponding author. Tel.: +1 202 328 5087.
E-mail address: burtraw@rff.org (D. Burtraw).

[1] Whenever we use the term "safety valve" without modification we refer to a "high-side" safety valve, or price ceiling.

0301-4215/$ - see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.enpol.2010.03.068

Unanticipated changes in allowance price have short and long-run characteristics. In the short run, prices may be variable due to phenomenon such as volatility in the allowance market, economic activity, weather, and fuel prices. Prices in previous environmental markets have been surprisingly volatile, which can erode political support even if prices revert to their expected value. In the long run, prices could ascend or descend to levels not anticipated and stay there due perhaps to these same phenomena and technological change. This event might justify revisions to the environmental program such as a change in the emissions target. However, such modifications may be politically challenging and anticipation of such adjustments might affect short run investment behavior. Furthermore, policy developments have taken a different direction, especially in the U.S. where proposals have focused on long-run compliance periods to lessen uncertainty stemming from revisions to policy and to better enable long-run investment planning. The proposal for comprehensive economy-wide climate policy that passed the U.S. House of Representatives (H.R. 2454) and the leading proposal in the U.S. Senate (S. 1733) both involve compliance periods, with emissions caps and allocation scenarios, that reach until 2050. Clearly such long-run goals can and may likely be revisited by policymakers. Nonetheless, the historic experience of growing intervals between successive amendments to the U.S. Clean Air Act suggests it becomes difficult to do so when policies become embodied in regulatory infrastructure and investment planning. This may be especially true for market-based regulation, where banked allowances also have value under the existing program. In this context, proposals to introduce a safety valve price ceiling have been motivated by the desire to replace the need for administrative or legislative adjustments with a decision rule. This paper examines the use of a safety valve and proposes an alternative.

We find two fundamental economic problems with the introduction of a one-sided safety valve price ceiling in a market for emissions allowances. One is that probabilistically the instrument can be expected to be triggered at some point in time, e.g. the allowance price ceiling would be achieved, bringing new emission allowances into the market. If the intent of policymakers initially was to balance costs and benefits in deciding the stringency of an emissions cap, then the introduction of a one-sided safety valve changes the emissions that can be expected and affects the ex ante cost–benefit calculation in a way that is not usually anticipated.

Furthermore, the single-sided safety valve that serves to cap one-side (the high-side) of the allowance price does not provide appropriate insurance against uncertainty. In particular, if the costs of emissions reductions turn out to be "too high" compared to ex ante information, then from an efficiency perspective, emissions would be too low and the safety valve would allow for additional emissions. However, if costs turn out to be "too low", as has been more frequently the case in previous trading programs, then emissions are too high and the one-sided safety valve will not provide remedy.

Indeed, the economic benefits of insuring against the prospect of costs that are lower than expected appear at least as important as the benefits of insuring against costs that are higher than expected, based on experience with cap-and-trade programs to date. A key example is the SO_2 cap under Title IV of the 1990 Clean Air Act Amendments. We calculate that a safety valve price floor protecting legislative intent against the prospect that costs would be substantially lower than expected would have improved economic welfare by $1.5 billion to $8.25 billion per year in each year since 1995. (All values are reported in 2004 dollars.) Emission allowance prices also fell dramatically at the end of the first round of the EU ETS program due to a lack of inter-period banking and non-binding caps (Ellerman and Joskow, 2008), and allowance prices for VOCs have persistently been much lower than expected in the Chicago VOC trading program (Evans and Kruger, 2007). Having a symmetric safety valve in these situations would help prevent price collapse.

Second, even if the one-sided safety valve never does bind, its introduction to a cap-and-trade program affects the expectation of future emissions levels and allowance prices and thereby the expectations about the payoff from various investment strategies. Using a detailed simulation model of the electricity sector under a cap-and-trade policy and accounting for uncertainty in the future price of natural gas, we find that the single-sided safety valve is likely to reduce investment in nonemitting technology and thereby increase expected emissions.[2]

We introduce a revised instrument labeled a "symmetric safety valve" that provides a floor as well as a ceiling on the price of emission allowances.[3] By construction, this design does a better job of insuring against price volatility than does a one-sided safety valve. Moreover it has the potential to address both problems that we identify with a one-sided safety valve. First, it can recover the ex ante expected level of emissions, based on ex ante expected costs and benefits. Second, it can recover the ex ante expected payoff to investments in nonemitting technologies. A symmetric safety valve would preserve all of the virtue while avoiding the unfortunate unintended consequences of a single-sided approach. We show that it can recover expected levels of emissions and produce welfare gains relative to a one-sided instrument. In so doing, it may help repair the political coalitions between environmental managers and economists that have been somewhat fractured by the discussion of safety valves.

2. Literature review

The literature addressing instrument choice for environmental policy in the presence of uncertainty about the costs and/or benefits of regulation is extensive. Early work by Weitzman (1974) identifies conditions under which price instruments would be preferable to quantity instruments and vice versa. In his model, the more efficient instrument hinges on the relative slopes of the marginal benefit and marginal cost curves. He shows that quantity instruments are preferred to price instruments if the marginal benefits curve is steeper than the marginal cost curve. Roberts and Spence (1976) analyze a combination instrument, which is specified as a licensed target level of emissions, a per unit subsidy for reductions in emissions below the firm's licensed target level and a penalty for emissions above the target where the penalty is weakly greater than the subsidy. Roberts and Spence prove that this hybrid approach yields lower social costs (defined as the sum of pollution damages and clean-up costs) than would result from using either instrument in isolation. Note that if damages are constant, then a pure tax would be optimal, and Robert and Spence find a similar result. Pizer (2002) uses a computable general equilibrium simulation model to analyze the welfare consequences of using different instruments to reduce CO_2 emissions. His work shows that the expected welfare gains from a price approach to climate policy are 5 times higher than expected gains with a pure quantity approach. The CO_2 emissions cap coupled with a safety valve ceiling on the price of CO_2 allowances yields slightly higher net social benefits than a tax policy by itself because, in this case, the climate benefits function

[2] Blythe et al. (2007) also look at the role of an allowance price floor in encouraging investment in green technology.

[3] The symmetric safety valve is analogous to a financial collar, which is a derivative instrument that can be used in the face of uncertain interest payments to fix payments within a certain range.

is slightly convex, and it substantially outperforms the pure quantity approach.

Pizer dismisses the use of a floor on allowance prices, suggesting implementation of such a policy in the form of a commitment to buy-back allowances once the price falls to the level of the floor would have adverse dynamic properties. This adverse effect is discussed in Chapter 14 of Baumol and Oates (1988), who argue that subsidizing emissions reductions in a competitive industry will typically lead to decreased output at the firm level, but increased output at the industry level and can also lead to higher emissions. They also cite Wenders (1975) who argues that subsidizing emission reductions can reduce incentives for the adoption of a new pollution-reducing innovation if firms anticipate that adopting the new technology will reduce subsidy payments. In both these cases, the assumption is that firms have the property rights to emissions and the government is buying them back.

However, the symmetric safety valve need not be implemented as a government buy-back of allowances that were previously distributed for free. For example, if some portion of allowances is being sold in an auction instead of distributed gratis, the low-side safety valve could take the form of a reserve price in the auction (Hepburn et al., 2006). If the willingness to pay for emission allowances for all bidders were to fall below that floor, a given *lot* of emissions allowances would not be sold. If willingness to pay exceeds the reserve price for a portion of the allowances offered in the auction, then those allowances will be sold at the reserve price and the rest of the lot will remain unsold.[4] The academic literature and numerous notorious examples of failed auctions point to a credible and efficient reserve price as an important aspect of auction design (Binmore and Klemperer, 2002; Ausubel and Cramton, 2004; Burtraw and Palmer, 2006). For example, in the recent 700 MHz spectrum auction the FCC set reserve prices that total more than $10 billion, and one commonly finds such a feature in commercial auction platforms such as EBay. A reserve price is incorporated into the CO_2 allowance auction used in the Regional Greenhouse Gas Initiative and in the December 2009 auction of RGGI allowances, the clearing price of allowances with a 2012 vintage fell to the level of the price floor and only roughly 3/4 of the allowances offered in the auction were sold at the reserve price (RGGI, 2009).[5] An alternative to an auction reserve price, if compliance periods extend for multiple years, would be to adjust allocations in later years if low prices prevail during the early years. If allowance banking were allowed, one would expect a reduction to latter year allocations to raise prices in earlier years.

An important insight from the Weitzman (1974) paper, which has typically been ignored in subsequent work, is the role of correlation of benefits and costs in the identification of optimal instruments. Stavins (1996) shows that when benefits and costs are statistically correlated, benefit uncertainty can affect instrument choice and the extent of that effect depends on several parameters. When benefits and costs are positively correlated, a quantity approach to regulation tends to be preferred to a price approach and when the correlation is negative, the tax approach will tend to be preferred. Stavins argues that positive correlation is more likely and that in general correlation of benefits and costs tends to favor emissions caps over emissions taxes. Evans (2007) considers correlation among the costs of control for different pollutants and reductions in those pollutants. He finds that Weitzman's advice regarding the choice of quantity or price instrument does not hold in general, and the efficient choice of instrument for one pollutant will depend on the choice for the other pollutant.

Another strand of the literature looks at the potential for emissions intensity regulation to outperform fixed quantities or prices in the presence of uncertainty. Quirion (2005) finds that with uncertainty about business-as-usual emission levels and about the slope of the marginal cost curve, an absolute cap on emissions produces slightly higher expected welfare than a cap on emissions intensity, but a price instrument yields substantially higher expected welfare than an intensity cap. Pizer (2005) suggests that indexing emissions targets to a measure of economic growth is a good approach for dealing with economic growth and unexpected changes in economic fortunes. Newell and Pizer (2008) analyze the use of indexed regulation for climate policies and identify conditions (related to the first and second moments of the index and the *ex post* optimal quantity level of the emissions cap) under which indexing will improve welfare as compared to both fixed quantities and fixed emissions taxes.

In addition to protecting against unexpected extreme prices, a safety valve may limit price volatility. Fell and Morgenstern (2009) use a dynamic model to compare a tax with cap and trade when it is coupled with a safety valve or price collar. They show a trade-off between expected abatement costs and variance of emissions. There is a long-standing debate about whether commodity price volatility, e.g. the standard deviation of prices in a period, affects investment behavior in general (Sauter and Awerbuch, 2003). For example, Mohn and Osmundsen (2008) find a negative link between oil price volatility and exploration efforts, suggesting that price stability provides a benefit to oil-importing countries. Zhao (2003) also finds investment incentives decrease with cost uncertainty, but in the context of emissions policies tradable allowances may maintain incentives for investment in abatement technology better than emissions fees. Philibert (2008) argues that a safety valve allows for a more ambitious target in the face of uncertainty about costs because it prevents costs in excess of acceptable levels. Further, there is persistent interest in the U.S. policy debate about possible manipulation of the allowance market. Although such manipulation seems improbable given the size of the market, illegal manipulation of large commodity markets has occurred in the past. The symmetric safety valve provides some comfort in this regard because it limits the potential profitability from price manipulation.

A safety valve has obvious relevance with respect to the ability to respond to changes over time, and therefore has some relation to the opportunity to bank emission allowances. The relationship between emissions banking and a safety valve is little explored in the literature. Jacoby and Ellerman (2004) suggest that banking will provide less protection from upside cost shocks than would a safety valve, particularly during the early years of a policy when no bank has yet accumulated for firms to draw on (assuming borrowing from the future is prohibited). However, they also point out that banking can provide greater price support in the case of lower than average cost, because the safety valve proposals usually do not include a price floor on allowances. Fell and Morgenstern (2009) show that restrictions on banking, including the provision of mandated interest rates, can limit the value of the policy. Banking provides a way to capitalize on a short run decline in marginal abatement cost by enabling extra emission reductions in that period that can be banked for use in later periods when costs may be higher. However, if the decline in cost is long-term in nature then the price will fall in every period and banking will provide little price support. Hence, a symmetric safety valve is not envisioned as an alternative to emissions banking, but as a complement.

[4] The reserve price could work in this way with a variety of different auction types including an ascending clock auction or a uniform price auction.

[5] The major legislative proposals in the U.S. (H.R. 2454 and S. 1733) use an auction to distribute a portion of the emissions allowances and specify a reserve price that should be a feature of that auction.

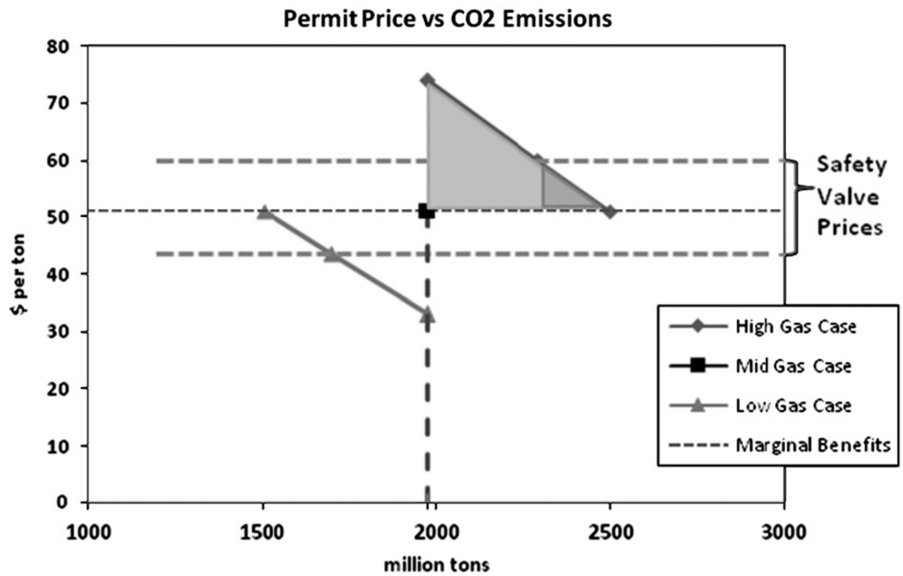

Fig. 1. Illustration of the safety valve.

3. The single-sided safety valve

To illustrate the effects of a safety valve, we examine a CO_2 cap-and-trade policy in the electricity sector. We evaluate the important case of natural gas price uncertainty using RFF's detailed simulation model of the electricity sector. The model divides the nation into 20 regions, 9 of which are assumed to yield electricity prices based on market outcomes, and the rest are assumed to be under cost of service regulation. We assume implementation of the Clean Air Interstate Rule (CAIR) and the Clean Air Mercury Rule (CAMR) policies for NO_x, SO_2, and Hg emissions. We assume a discount rate of 8%, with 2030 as the forecast horizon year.[6]

The central case is built upon EIA (2006) forecasts of natural gas prices (as discussed previously) and other parameters reported in the *Annual Energy Outlook*. We assume gas price is normally distributed. We consider two alternatives that incorporate a 30% increase and decrease in gas prices, respectively. High and low prices are picked to represent prices that are approximately one standard deviation away from the mean based on the authors' assessment of recent forecasts.

The example is illustrated in Fig. 1, where the horizontal axis represents the aggregate emissions of CO_2 from the electricity sector in 2020. We assume the marginal benefit of emissions reductions is known with certainty to be $51.10/ton, and assume the regulator sets an emissions target where the known marginal benefits are equal to the expected but uncertain marginal costs. The central case revolves around the mid-value (expected value) for natural gas price of $6.31/mmBtu under a moderate climate policy (drawn from EIA (2006) projections), the regulator determines an emissions target (E^*) of 1973 million tons. The central point in the figure illustrates this price and quantity target.

The downward sloped line that lies above the expected permit price in Fig. 1 illustrates a realization of a marginal abatement cost schedule if natural gas prices turn out to be higher than expected, at a level we conjecture to be equal to $8.21/mmBtu.[7] Shifting from coal-fired to gas-fired generation is an important way that emission reductions are achieved, so the higher price of natural gas leads the cost of achieving emission reductions to be higher than in the mid case. For the high gas price case the marginal cost of achieving the emission target E^* increases to P^h=$74.1/ton. Since marginal benefits are assumed constant and equal to P^*, given the high cost of emission reductions in the high gas price case the target quantity of emissions is too low and there is a welfare loss equal to the large shaded triangle.

Were there a safety valve in place, say at a level between the mid and high allowance price outcomes, e.g. P^{HSV}=$ 59.7/ton, it would cap the level to which marginal costs could rise by issuing additional emission allowances. This would lead to emissions of 2293 million tons. Compared to the target where realized marginal cost equals expected marginal cost of $51.1, there would still be a welfare loss at P^{HSV} indicated by the smaller cross-hatched triangle, but this welfare cost would be less than from the strict quantity instrument without the safety valve.

4. The symmetric safety valve

If the safety valve were to apply only when marginal costs are higher than expected, the emission target would not respond if natural gas price turns out to be lower than expected. The lower negatively sloped line segment in Fig. 1 illustrates this outcome with an assumed natural gas price in 2020 equal to $4.42/mmBtu. At the original emissions target of E^*, the marginal abatement cost decreases to P^L=$33/ton. The example illustrates that the welfare consequences of the drop in natural gas price can be just as great as when gas price is higher than expected due to the difference between marginal benefits and marginal costs. Since marginal benefits are assumed constant at P^* equal to $51.1/ton, there is a welfare cost analogous to the large shaded triangle in the previous

[6] Further detail on the model can be found in Paul and Burtraw (2002).

[7] There is significant likelihood that fuel prices will change over the course of environmental programs. For example, Title IV of the Clean Air Act Amendments regulating SO_2 passed Congress in 1990; the first phase of the program took effect in 1995, and the second phase took effect in 2000. In 1990 the delivered price of natural gas for electric utilities was about $3.15/kcf, but by 1999 it had fallen to $2.89/kcf (2004$). Similarly, the average price for low sulfur subbituminous coal fell from about $12.81/ton in 1990 to $7.56/ton in 1999, while the consumption of low sulfur coal grew tremendously (EIA, 2005, Tables 6.8 and 7.8). Investment decisions to comply with the legislation crafted by Congress in 1990 continued to take shape more than a decade later.

example, but in this case the loss is due to the fact that from an efficiency perspective the quantity of emissions is too high given the low cost of emission reductions.

A low-side safety valve could correct for the unexpected decline in compliance cost. For example, if there were a safety valve (allowance price floor) at a level that was between the low and mid allowance price outcomes, P^{LSV}=\$43.5/ton, it would cap the extent to which marginal costs could fall by reducing the number of allowances provided to the market. As illustrated, a reduction in emissions below $E^{*}-E^{LSV}=1699$ million tons would be achieved. There would still be a welfare cost compared to the efficient outcome *ex post*, but the cost would be less than under the strict quantity instrument without the low-side safety valve.

There has been little attention given to how a safety valve would function. In the case of a high-side safety valve, advocates have suggested that the regulator could issue additional allowances at the safety valve price level through direct sale, auction or potentially through free allocation. The low-side safety valve could have the same structure. If allowance price falls to the level of the floor, the regulator would reduce the provision of allowances in future periods. An example of this approach is embodied in the Clean Air Interstate Rule (CAIR) where the allowable quantity of future emissions changes the value (ton of emissions per allowance) of future allowances without changing the quantity of the allowances. The promulgation of the CAIR rule in 2005 led to a reduction in the allowable emissions in the future and to an increase in the price of allowances in future and in the current period. As noted previously, an even more direct way to implement the low-side safety valve would be through the use of a reserve price in an auction, if an auction is used to initially distribute a portion of the allowances. When the safety valve policy combines a high-side and a low-side safety valve, we refer to it as a symmetric safety valve.

5. Historical experience

Historically, the failure to have a safety valve on the low side in the event that compliance costs are lower than expected has had larger consequences than the failure of a safety valve on the high-side. The only important example of unexpected outcomes within a cap-and-trade program that may have been remedied by a safety valve on the high-side has been the RECLAIM program in southern California, where prices skyrocketed in 2000 due to unexpected demand for emission allowances reflecting very high marginal cost of compliance, which led to suspension of trading in the program in 2001. In that program, however, emission allowance banking also could have helped remedy the market disruption.

In contrast, the most prominent economic failure of any cap-and-trade program occurred in the SO_2 program under Title IV—a program generally noted for its many successful aspects. The SO_2 program is credited with success in facilitating the reduction in compliance costs compared to prescriptive regulatory approaches (Carlson et al., 2000; Ellerman et al., 2000), demonstrating on a large scale the effectiveness of an economic approach to pollution control (Stavins, 1998; Joskow et al., 1998), and achieving billions of dollars in environmental and public health net benefits (Burtraw et al., 1998; Banzhaf et al., 2004).

However, the expensive failing of the SO_2 program has been its inability to adjust to new information. In 1990, at the adoption of Title IV, Portney (1990), the only economist who ventured an opinion about the benefits and costs of the amendments, concluded that the benefits of Title IV about equaled the cost. By the first year of the program's implementation in 1995, it had become clear that the benefits would be an order of magnitude greater than costs (Burtraw et al., 1998). Unfortunately the program was unable to adapt to this new information until the adoption of CAIR, now scheduled to take effect in 2010, fifteen years after the launch of the program.

Why did the estimates of benefits and costs change so dramatically? First, the anticipated benefits of emission reductions grew tremendously with new information about the damage to human health from fine particulates associated with emissions of SO_2 and NO_x. Second, and more important to this discussion, the estimates of the costs of emission reductions fell sharply, due in large part to the flexibility in compliance options afforded by the program.

The fact that information can change so dramatically and so quickly leads one to ask: To what extent does policy reflect scientific information about both the benefits and costs of regulation? Scientific and economic information is fundamentally uncertain. How policymaking interprets the data, and how the policy system responds when scientific information evolves, is of vital importance. Typically, once regulators reach a decision, it becomes exceedingly difficult to modify that decision (Center for International Studies, 1998). For instance, the Clean Air Act was amended in 1977, again in 1990, and has not been amended further since. Statutory regulation such as Title IV put regulators' feet into cement. It is very difficult to change statutory direction given new scientific information.

The policy system could benefit from the use of decision rules that automatically incorporate new information. It is understandable that the policy system would be slow to incorporate new information about the benefits of Title IV, because information about benefits is not readily observable outside of the process of scientific research and peer review, which may take years to achieve general acceptance. However, cap-and-trade programs are uniquely designed to generate information about costs, in the form of allowance prices, which instantaneously provide a summary statistic of pollution control costs that is widely accessible.

Before adoption of CAIR, which directly influences compliance with the SO_2 trading program, estimates suggested that the expected SO_2 emissions in 2010 were to be about 9.18 million tons (Banzhaf et al., 2004). In 1990 the EPA estimated that the marginal cost of achieving the emission reduction targets in Phase II around the year 2010 would be \$718–942/ton (2004\$) (ICF Resources Inc., 1990). However, as the program unfolded it quickly became apparent that the marginal costs as reflected in the price of emission allowances were dramatically below expectations. Fig. 2 illustrates that the price of an SO_2 allowance has been well below \$200/ton throughout most years until the CAIR proposal was announced in 2004.

Let us imagine that it was Congress' intent to roughly balance marginal benefits with marginal costs, and that a low-side safety valve had been in place that would reduce the provision of allowances if the price were to fall below \$567/ton, about 33% below the mid-value of the range of expected costs. Banzhaf et al. estimate that an SO_2 allowance price of \$567/ton in 2010 would yield total national annual emissions of 7.1 million tons, about 2.08 million tons less than under Title IV in the baseline (and in the absence of CAIR).

What would have been the value of a low-side safety valve that led to additional emission reductions? Banzhaf et al. use estimates of marginal benefits of \$3968/ton. This is substantially less than those used by the EPA in Regulatory Impact Assessment because Banzhaf et al. use a lower value of statistical life. Using the Banzhaf et al. estimates, the additional annual health benefits from placing a floor on the allowance price would total \$8.25 billion in 2010 (2004\$). However, perhaps one may reason that Congress could not have expected benefits of this magnitude from a safety valve, because it did not expect the health benefits per ton of emissions reduction to be this large. One could say that if Congress acted to equate marginal

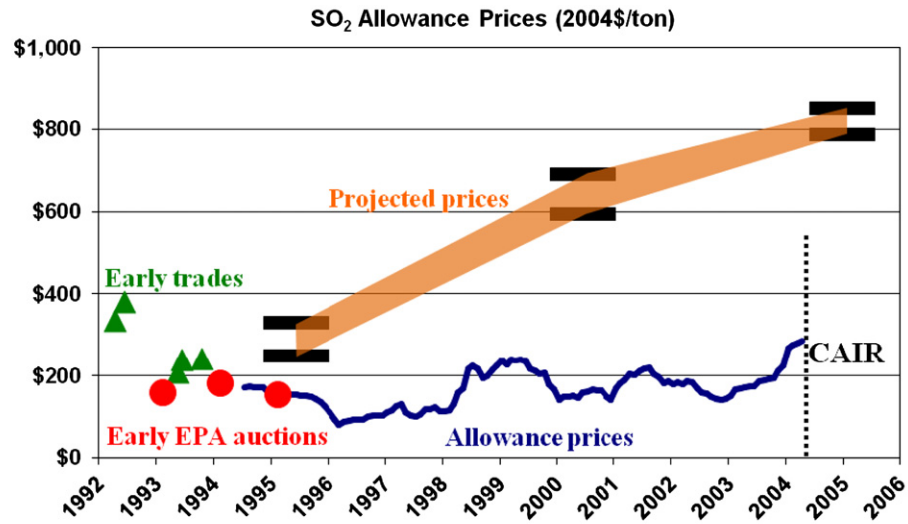

Fig. 2. Variation in SO_2 prices.

benefits and marginal costs then they would value additional emission reductions at an expected value of $718–942/ton. At this value, the anticipated additional health benefits from the safety valve set 33% below expected marginal cost would be between $1.5 billion and $1.95 billion in 2010. Arguably, the legislative intent of Congress was to capture these benefits, but they did not have the policy tools available at the time to anticipate and flexibly adjust to changes in scientific information. The symmetric safety valve provides such a tool.

6. The safety valve affects expectations and investment

The example illustrated in Fig. 1 has a fundamentally *naïve* characterization of behavior because, as illustrated, the regulator makes decisions on the basis of expected values. She does not account for how the safety valve affects expected values. Consequently the imposition of a one-sided safety valve will influence the market equilibrium and affect the decisions of investors, with unintended and potentially negative consequences that could undermine policy goals.[8]

To illustrate this point, we simplify the multi-period problem into an instantaneous present value calculation. Consider the profit function for a single firm that offers nonemitting electricity generation:

$$\pi = qP(Q,P_A) - C(q) \tag{1}$$

where q is the quantity produced by the potential investment, Q is the aggregate quantity of electricity in the market and P_A is the price of allowances. Cost is a function of quantity of the production. We assume these functions are increasing in their arguments.

The firm maximizes profits by choosing quantity (q). Under the assumption that the facility's output is too small to make an impact on the aggregate production and price $(\delta P/\delta q) = 0$ the firm maximizes profits by choosing q such that marginal revenues equal marginal costs:

$$P(Q,P_A) = \frac{\delta C}{\delta q} \tag{2}$$

[8] By analogy, the provision of insurance affects the behavior of investors because the insurance changes the expectations over potential payoffs. Here we find something similar—investors can be expected to respond to the safety valve, which leads to a different market equilibrium.

In general we expect the aggregate quantity and price of allowances to be uncertain, so that $Q = \tilde{Q}$ and $P_A = \tilde{P}_A$ (where the tilde represents uncertain variables), which cause the product price to be uncertain $(P = \tilde{P})$. Assuming the firm is risk neutral, the profit maximization condition would require the firm to equate expected marginal revenue with marginal cost: $E(P) = \delta C/\delta q$. We assume the potential distribution f of allowance prices stretches from a bound of zero at its minimum to infinity: $\tilde{P}_A \sim f(0,\infty)$, which along with the distribution of potential aggregate generation determine the expected electricity price.

The high-side safety valve intentionally alters the distribution of the potential allowance price, so that the price cannot rise above the safety valve level (SV). If we *naïvely* ignore the interaction of the allowance price and the investment decisions of other firms and consider only the role of the safety valve on allowance price (holding the decisions of other firms constant, e.g. $(\delta\tilde{Q}/\delta P_A) = 0$, then the allowance price with the one-sided safety valve has the distribution:

$$\begin{aligned} \tilde{P}_A^{SV} &\sim h(0,SV) \\ &= f \text{ for } P_A \le SV \text{ and} \\ &\quad SV \text{ for } P_A > SV \end{aligned} \tag{3}$$

Letting F and H be the cumulative distribution functions for $\tilde{P}_A^{SV}$ and $\tilde{P}_A$, then $F \le H$ over their entire range, and $\tilde{P}_A^{SV}$ will have an expected value that is strictly less than P_A, $E(P_A^{SV}) < E(P_A)$. The change in the distribution of potential allowance price is illustrated in Fig. 3. The top panel illustrates a probability distribution for allowance price with the dotted curve. Allowance price in the absence of a safety valve is designated by "P". The expected value for the allowance price is designated $E(P)$, shown by a dotted line. The addition of the safety valve censors the potential distribution of allowance prices. If the distribution is otherwise unaffected, as described in Eq. (3), the mean shifts to the left, as indicated by the dashed line $E[PSV]_{naïve}$.

A consequence of the change in the allowance price would be a change in the equilibrium in the electricity market, leading to a lower price under the safety valve, $E(P^{SV}) < E(P)$. The individual profit maximizing investor described in Eq. (1) would choose a level of production under the safety valve where

$$E(P^{SV}) = \frac{\delta C}{\delta q^{SV}} < \frac{\delta C}{\delta q} = E(P) \tag{4}$$

leading to a reduction in its investment and output, $q^{SV} < q$.

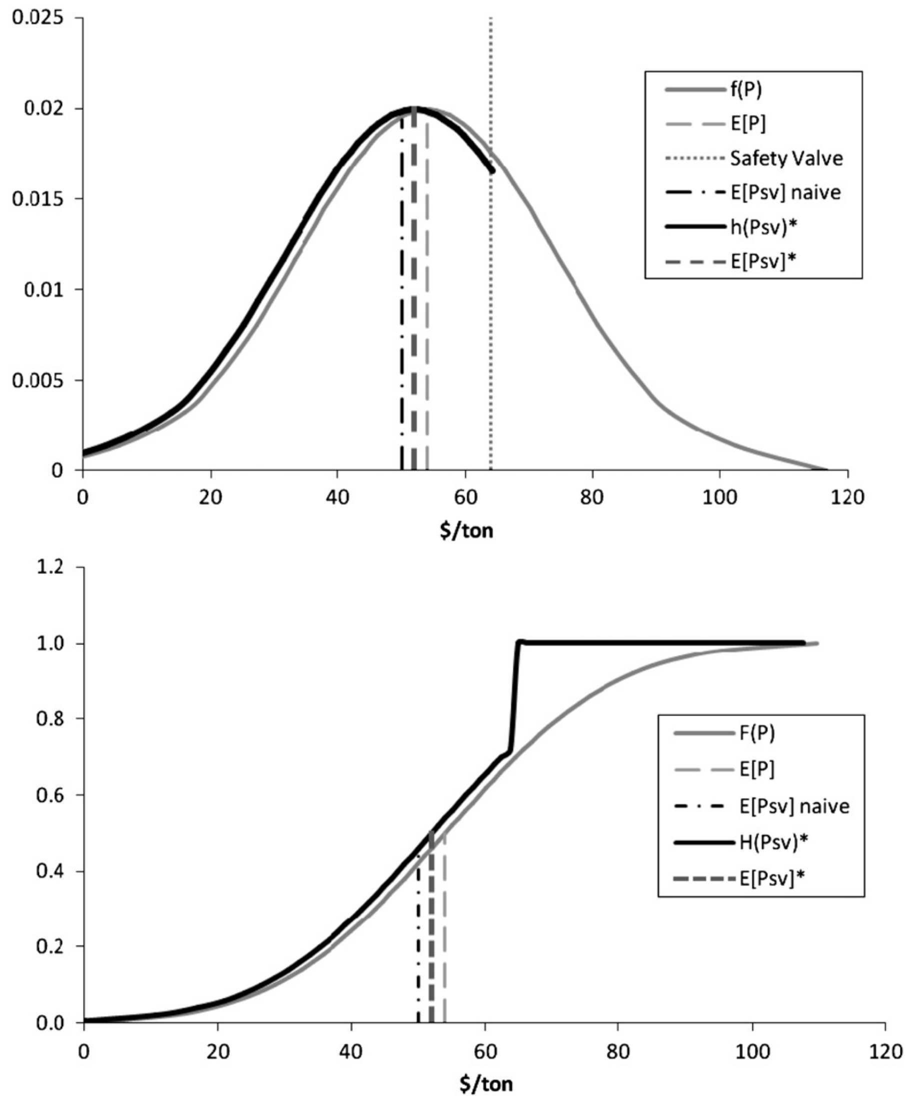

Fig. 3. Illustration of the distribution of allowance prices associated with uncertain gas price outcomes.

The consequence of the high-side safety valve in this example is to reduce investment in the nonemitting facility. One can conjecture that in the aggregate the policy leads to less investment in renewable technology or low-emitting technology that may suffer a price disadvantage when the external social costs of electricity generation are not included in electricity price. The cap-and-trade program serves as a mechanism to internalize into investment decisions the social cost of technology choices and "level the playing field," as many observers have suggested. However, the single-sided safety valve would appear to provide an asymmetric influence that would tilt the playing field away from investments in nonemitting sources.

7. The safety valve equilibrium in a simulation model with perfect foresight

The formulation above assumes that the behavior of other investors or actors in the market do not respond to the change in expectations; however, clearly there would be a response. For instance, one could imagine that a lower allowance price would lead to more fossil generation and a lower electricity price, reinforcing the effect described above. However, the lower allowance price also might increase the emission intensity of generation for any given level of production, which would cause a bounce back in the price of allowances. Whatever the underlying source of uncertainty for allowance price is, it is likely to directly affect the cost and aggregate quantity of production.

To identify the equilibrium outcome we return to results from the simulation modeling that was introduced in Section 3. In doing so, we conjecture *a priori* that the high-side safety valve should lead to less investment in nonemitting and low-emitting sources of generation than in the absence of the safety valve, as well as a lower expected allowance price, a lower electricity price and greater expected emissions.

In this exercise natural gas and coal prices are set at forecast values (low, mid, high) in each year. Changes in the demand for fuel would have an effect on fuel price leading to changes in producer surplus (rents) outside the electricity sector. To distinguish our results from those effects, changes in fuel prices are disabled in the simulation.[9] We freeze the level of electricity

[9] Changes in relative fuel prices, especially between natural gas and coal, could have a significant effect on the cost of mitigation in the model, and fuel price changes naturally would be endogenous to the policy choice. Taking fuel prices as exogenous *ex ante*, obscures the fact that the introduction of a price on CO_2 would put upward pressure on relative fuel prices. If prices were endogenous it would shift up the expected marginal cost of the policy in each of our scenarios.

consumption at the central case levels for each simulation year in order to avoid second best issues in the welfare calculation that are associated with differences between price and marginal cost. For all of the policy scenarios, emission allowances are allocated to emitters on the basis of historic generation (which affects the electricity price and level of electricity demand) and additional allowances are purchased at the safety valve price.

The equilibria that were described graphically in Section 3 are summarized in Table 1. The model is deterministic and actors behave as though they have certain and perfect foresight—e.g. they know the future path of natural gas prices and respond accordingly. The middle column represents the mid case for gas prices. The first and last columns represent the outcome for low and high gas prices, respectively, in the absence of a safety valve. There is little change in CO_2 emissions, but it is interesting to note that low gas prices lead to a modest increase in emissions because there is new gas generation in lieu of new investment in renewables. High gas prices also lead to an increase in emissions, as gas-fired generation falls and there is an increase in coal-fired generation that more than offsets the new investment in renewables. Allowance price ranges widely from a low of $33 under the low gas scenario to a high of $74. Electricity price also ranges widely. Fig. 4 illustrates the change in electricity price relative to the mid case for each simulation year in the model.

The fourth column of Table 1 represents the influence of a single-sided safety valve on the high side. The second column represents the influence of a safety valve on the low side. On either side, the safety valve has a direct effect on CO_2 emissions, as would be expected because it affects the quantity of emissions directly. As a consequence, under different gas price scenarios the variation in other variables such as electricity price and renewable generation is reduced, compared to the absence of the safety valve.

Note also that in either case, the safety valve improves welfare relative to the baseline. Welfare is calculated as the sum of changes in producer and consumer surplus, plus the change in environmental benefits associated with changes in emissions relative to the emission quantity target, valued at their expected cost of $51/ton. The change in welfare for each case is measured relative to the mid gas price case. The greatest improvement comes from adding a safety valve in the high gas price case. Relative to the mid gas price case, which is normalized to a value of zero, the high gas price case leads to a loss of over $23 billion. The high-side safety valve reduces this loss to about $7 billion because it closes the gap between marginal benefits and marginal costs by allowing an increase in emissions. In the low gas case, welfare improves by over $37 billion, due to lower cost of production. In the low-side safety valve case, welfare improves further to $40 billion by reducing emissions below the emission target, thereby taking advantage of the relatively low marginal cost of abatement.

8. Modeling uncertainty

The simulation model is deterministic, meaning that it incorporates certain foresight about potentially uncertain

Table 1
Deterministic model with certain foresight, 2020.

	Low gas price	Low gas price w/low-side safety valve	Mid gas price	High gas price w/ high-side safety valve	High gas price
Gas price ($/mmBtu)		4.42	6.31	8.21	
CO_2 emissions[a] (Mtons)	2009	1699	1973	2293	1999
Welfare[b] (Billion $)	37.66	39.94	0	-6.85	-23.25
Electricity price ($/MWh)	82.1	84.8	93.1	99.1	102.4
Allowance price ($/ton)	33.0	43.5	51.1	59.7	74.1
Renewable generation[c] (BkWh)	313	360	394	568	581

[a] Differences in CO_2 emissions in the first, third and fifth data columns reflect variation in model convergence.
[b] Welfare compared to mid gas price case.
[c] Underestimate because renewables include only changes in wind, landfill gas, geothermal. Biomass is held constant.

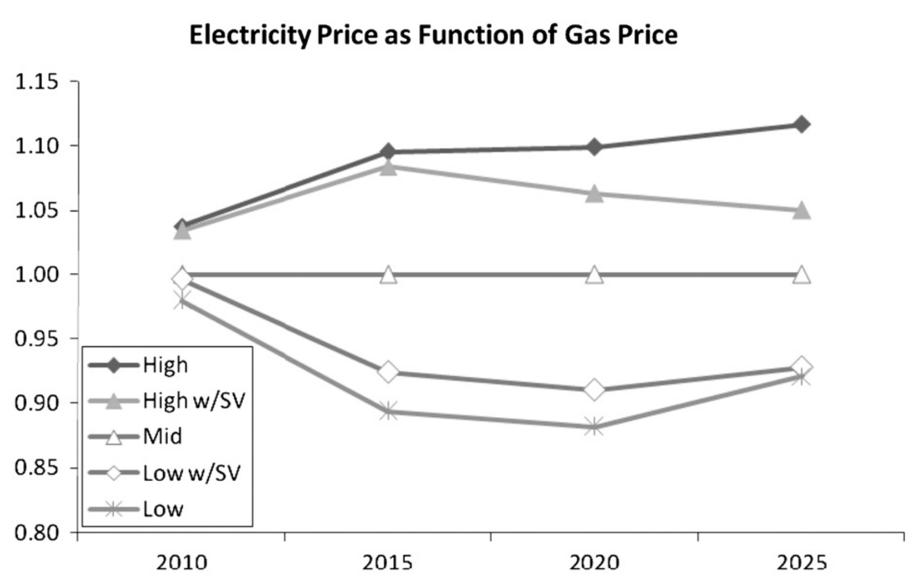

Fig. 4. Variation in electricity prices in the deterministic model with perfect foresight.

Table 2
Delta method approximation of key variables in model with uncertainty, 2020.

Expected value $E[\phi(\tilde{g})]$	**No safety valve**	**High-side safety valve**	**Symmetric safety valve**
Gas price (\$/mmBtu)	6.31	6.31	6.31
CO_2 emissions (Mtons)	1983	2313	2015
Welfare[a] (Billion \$)	13.82	29.62	31.80
Electricity price (\$/MWh)	91.43	88.19	90.78
Allowance price (\$/ton)	55.66	41.88	52.00
Renewable generation[b] (BkWh)	494.5	482.3	528.2

[a] Welfare is the difference relative to the mid gas price case.
[b] Underestimate because renewables include only changes in wind, landfill gas, geothermal. Biomass is held constant.

variables. Investment decisions are made as though each actor knows for certain the future values of every variable, as well as the decisions of every other actor, so there is no uncertainty taken into account in the model solution. However, we can use a collection of model solutions to make a mathematical inference about the outcome of the market equilibrium when investors make decisions taking uncertainty into account.

Using the results from the deterministic model for various realizations of the underlying uncertain parameter, we construct a linearization using the delta approach, which is a variation of a Taylor series expansion. The expected value of a function ϕ of a random variable $\tilde{g}$ with expected value $\overline{g}$ and variance σ_g^2, can be approximated by

$$E\left[\phi(\tilde{g})\right] \cong \phi(\overline{g})+\phi'(\overline{g})E[(\tilde{g}-\overline{g})]+\tfrac{1}{2}\,\phi''(\overline{g})E[(\tilde{g}-\overline{g})^2]=\phi(\overline{g})+\tfrac{1}{2}\phi''(\overline{g})\sigma_g^2 \quad (5)$$

where ϕ' and ϕ'' are first and second derivatives of the function.

The function ϕ can represent a variety of measures including aggregate economic welfare, electricity price, allowance price or the installed nonemitting generation capability. For this experiment, the random variable $\tilde{g}$ is the natural gas price. We consider low, mid and high values of \$4.42/mmBtu, \$6.31/mmBtu and \$8.21/mmBtu in 2020 (2004\$). As described previously, we assume it is common knowledge that these prices are distributed normally with an expected value of \$6.31/mmBtu and a standard deviation of \$1.90/mmBtu, so the mid-value in this experiment is the mean value of the natural gas price and the low and high values are both one standard deviation from the mean.

The results from this experiment are reported in Table 2, for the case of no safety valve, a high-side (only) safety valve, and a symmetric safety valve that includes a safety valve on both the high-side and low side. The high-side safety valve leads to the expectation of greater emissions than in the no safety valve case because with some probability the safety valve will be triggered, thereby placing extra allowances on the market. As a consequence the allowance price and electricity price are lower. All variables except welfare are normalized using the no safety valve case as a numeraire (the value is set equal to one). For welfare, the difference between the no safety valve case and the mid case in the deterministic model is normalized as a numeraire because only changes in welfare have economic relevance. A potentially important unintentional result is that the lower expected allowance price leads to lower expected payoffs to investment in renewable technologies. Consequently we see a decline in renewable generation.[10]

Many observers have criticized the high-side safety valve because it might undermine the environmental targets of the program, and that is the result we obtain. Emissions are higher and investments in clean technology and associated generation levels are lower as a result of the safety valve. The reduction in investments initiates a cascade of consequences, as there is less learning as a result of the decline in investment, so the costs of renewable technologies remain above their levels in the absence of the safety valve.

However, the unintended consequences are fully remedied when the safety valve is characterized as a symmetric instrument. In this case, emissions fall back to virtually the same level as in the absence of a safety valve, and renewable generation increases to above its level in the absence of a safety valve. The results for the high-side and symmetric safety valve are compared visually in Fig. 5. The figure shows that not only do measures of interest to environmental advocates return to their intended levels, but welfare improves even further than in the case with only a high-side safety valve. Also, electricity price and allowance price return to nearly the same level as in the absence of the safety valve.

9. Surprise in a model with certain but imperfect foresight

The delta method could be applied in a different way by assuming a different information structure. In the previous example, the finite differences are calculated using the model with perfect foresight. An alternative would be certain but imperfect foresight; wherein investment decisions made under one set of assumptions could prove imprudent if conditions were to change unexpectedly. For example, if gas prices were to deviate from expectations after investment decisions have been made, then generators could experience large losses and welfare could be negatively affected. Since the safety valve is a policy to mitigate the welfare costs of surprises such as this one, we consider a case where investors' expectations are incorrect, and use these data to calculate finite differences.

The scenario involves a surprise in natural gas prices in 2015. Investors make an investment plan based beginning in the first simulation year in 2010 and based on certain but imperfect foresight about the future path of gas prices. In 2015, investors learn that gas prices are on a different path. Taking existing investments as sunk, investors solve the perfect foresight problem with the new data. We ran simulation scenarios that include a gas price surprise to determine the effect both a one-sided and symmetric safety valve would have on the expected value of several key variables.

Fig. 6 illustrates the path of electricity prices under the surprise in natural gas prices, compared against the expected price path for prices that was illustrated previously in Fig. 4. The surprise in 2015 leads to a precipitous change in electricity prices in the absence of a safety valve, especially when natural gas prices rise unexpectedly.

The surprise in gas prices leads to comparable variations across the different policy scenarios than were obtained in the previous example for most variables. Table 3 illustrates these differences.

[10] Here, only a subset of renewable technologies is allowed to change. If biomass also allowed to change one would see even more of an effect on renewable generation.

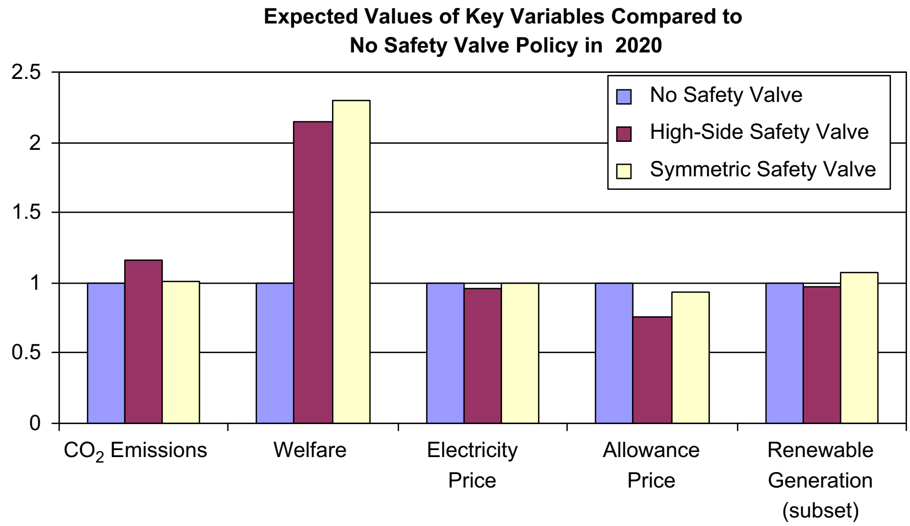

Fig. 5. Delta method approximations of outcomes under uncertainty.

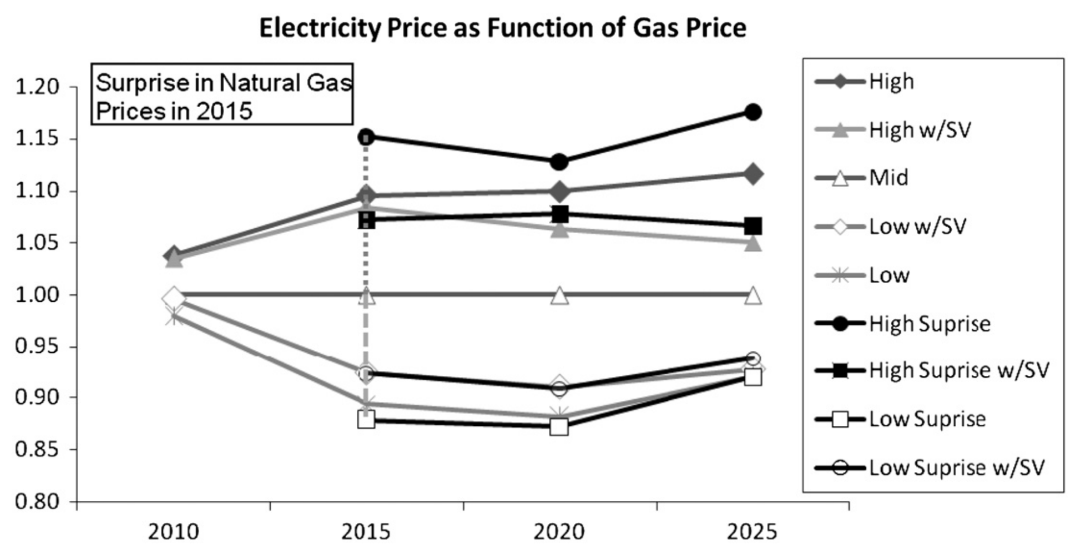

Fig. 6. Variation in electricity prices in the model with certain but imperfect foresight.

Table 3
Model with certain but imperfect foresight and a gas price surprise in 2015; results for 2020.

	Low gas price	Low gas price w/symmetric safety valve	Mid gas price	High gas price w/high-side safety valve	High gas price
Gas price ($/mmBtu)		4.42	6.31	8.21	
CO_2 emissions (Mtons)	1983	1690	1973	2264	1983
Welfare[a] (Billion $)	35.25	38.01	0	−13.85	−21.99
Electricity price ($/MWh)	81.25	84.67	93.14	100.40	105.10
Allowance price ($/ton)	32.14	43.28	51.14	59.42	68.42
Renewable generation[b] (BkWh)	296	328	394	473	472

[a] Welfare is the difference relative to the mid gas price case.
[b] Underestimate because renewables include only changes in wind, landfill gas, geothermal. Biomass is held constant.

One outcome that is interesting is the increase in renewable generation in the high gas price case with a high-side safety valve. The reason is that although dedicated biomass does not change in the model, co-fired biomass is allowed to change. The high gas price leads to more coal-fired generation, and with that comes a greater amount of co-fired biomass.

We apply the delta method to this set of results to replicate the experiment of a first-order approximation to behavior in a model with uncertainty. Table 4 reports these results, and they are illustrated visually in Fig. 7. Again, the variables of interest return to their approximate levels in the absence of the safety valve. The effect on renewable generation is greater than in the previous

Table 4
Delta method approximation of key variables in model with certain but imperfect foresight; results for 2020.

Expected value E[$\varphi(\tilde{g})$]	**No safety valve**	**High-side safety valve**	**Symmetric safety valve**
Gas price ($/mmBtu)	6.31	6.31	6.31
CO_2 emissions (Mtons)	1992	2256	1982
Welfare[a] (Billion $)	12.67	19.97	22.59
Electricity price ($/MWh)	93.18	88.84	92.04
Allowance price ($/ton)	49.47	41.15	51.57
Renewable generation[b] (BkWh)	374.6	377.1	407.4

[a] Welfare is the difference relative to the mid gas price case.
[b] Underestimate because renewables include only changes in wind, landfill gas, geothermal. Biomass is held constant.

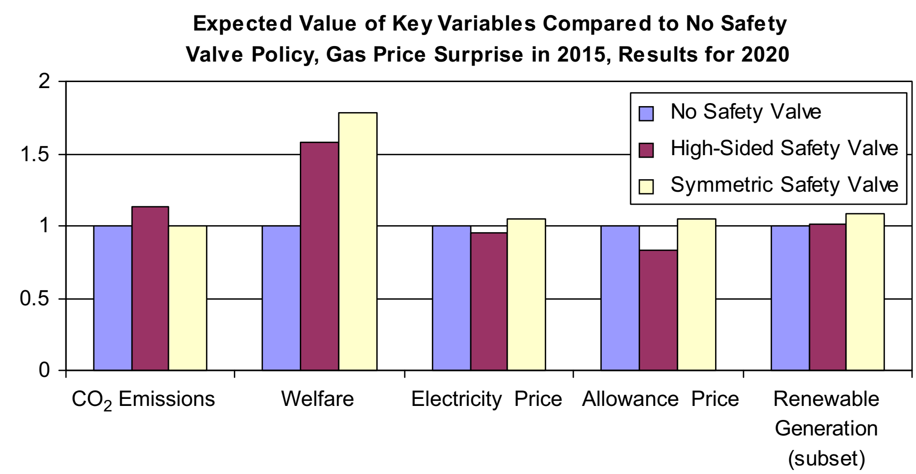

Fig. 7. Delta method approximations of outcomes in a model with certain but imperfect foresight.

example. Also, the welfare contribution of a symmetric safety valve is greater relative to the high-side safety valve.

10. Conclusion

Significant attention has been directed to price stabilization measures in emission allowance trading programs. In particular, attention has focused on the introduction of a single-sided safety valve that would mitigate potential price spikes by introducing additional emission allowances into the market when costs rise above the specified "safety valve" level. If abatement costs rise unexpectedly and the one-sided safety valve is triggered, it will increase emissions, but if costs fall unexpectedly there is no change in emissions. Consequently, one can expect the consequence of a one-sided safety valve to be erosion in the environmental stringency of an emissions cap.

The experience of most cap-and-trade programs to date indicates that the most important examples of price volatility have occurred not when allowance prices rose but when allowance prices fell below their expected values. For example, in the case of SO_2 emission trading, the inability of the trading program to adjust to the fall in allowance prices led to welfare losses of between $1.5 and $8 billion dollars per year.

A second reason to be interested in price stabilization when prices fall below expectations is the influence that a low price has on investment. In the absence of a safety valve, investors will take risks given expectations over a distribution of potential payoffs for their investment. A high-side safety valve that prevents spikes in allowance prices will have the unintended consequence of lowering the expected allowance price, and as a consequence the overall expected return on an investment in nonemitting technology. Indeed, the effects on investment have produced growing support for a price floor in the EU Emissions Trading System.[11]

A symmetric safety valve is a price stabilization policy that works in the case of unanticipated spikes or drops in allowance price. In the case when allowance price falls below the safety valve floor, the safety valve would contract the number of allowances issued in the market. The reduction in the quantity of allowances can be implemented in a variety of ways, but the simplest way may be through a change in the portion of emission allowances that is initially distributed through auction. In fact, good design suggests that an auction should have a reservation price, which is a floor below which the allowances will not be sold. Such a price floor serves directly to implement the low-side safety valve.

We use a linear approximation representing a Taylor series expansion around the mid case to model uncertain natural gas prices in a detailed electricity market model. We show that a high-side safety valve can be expected to increase emissions and decrease investment in nonemitting technologies, relative to the absence of a safety valve. However, the symmetric safety valve returns the expected value for these and other key parameters to the vicinity of their levels in the absence of a safety valve. In addition, although a high-side safety valve improves welfare, a symmetric safety valve improves welfare even further. In summary, we find a symmetric safety valve can improve the performance of allowance trading programs, improve welfare, and may help overcome political objections from environmental advocates who have opposed the use of a safety valve.

[11] "Support builds for EU carbon price floor," *Point Carbon*, February 5, 2009.

Further areas remain to be developed in this analysis. One has to do with a method for determining the breadth of a safety valve around expected marginal cost (Fell et al., 2010). When marginal benefits are constant, as in the examples we use, then the most efficient safety valve would be one exactly equal to the value of marginal benefits. In other words, the efficient policy is a tax. However, when marginal benefits are not flat but vary over a range then intuition suggests the efficient safety valve would vary from the expected level of marginal benefits.

Acknowledgement

This research is funded by EPA Agreement Number RD 83099001 and by the Mistra-funded Climate Policy Research Programme (Clipore). The authors are grateful to Erica Myers for comments and Anthony Paul for technical assistance.

Appendix A. Supplementary material

Supplementary data associated with this article can be found in the online version at doi:10.1016/j.enpol.2010.03.068.

References

Ausubel, L., Cramton, P., 2004. Vickrey auctions with reserve pricing. Economic Theory 23 (3), 493–505.
Baumol, W.J., Oates, W.E., 1988. The Theory of Environmental Policy second ed Cambridge University Press.
Banzhaf, H.S., Burtraw, D., Palmer, K., 2004. Efficient emission fees in the US electricity sector. Resource and Energy Economics 26, 317–341.
Binmore, K., Klemperer, P., 2002. The biggest auction ever: the sale of the British 3G telecom licenses. The Economic Journal 112, C74–C76.
Blythe, W., Ming, Y., Bradley, R., 2007. Climate Policy Uncertainty and Investment Risk. International Energy Agency, Paris.
Burtraw, D., Palmer, K., 2006. Summary of the workshop to support implementing the minimum 25 percent public benefit allocation in the regional greenhouse gas initiative. Resources for the Future Discussion Paper 06-45.
Burtraw, D., Krupnick, A.J., Mansur, E., Austin, D., Farrell, D., 1998. The costs and benefits of reducing air pollutants related to acid rain. Contemporary Economic Policy 16 (October), 379–400.
Carlson, C., Burtraw, D., Cropper, M., Palmer, K., 2000. SO_2 control by electric utilities: what are the gains from trade? Journal of Political Economy 108, 1292–1326.
Center for International Studies, 1998. Environmental policymaking: a workshop on scientific credibility, risk and regulation. Massachusetts Institute of Technology, Cambridge, September 24–25. Summary prepared by Brian Zuckerman and Sanford Weiner.
Ellerman, A.D., Joskow, P.L.I., 2008. The European Union's Emissions Trading System in Perspective. Pew Center on Global Climate Change, May.
Ellerman, A.D., et al., 2000. Markets for Clean Air. Cambridge University Press.
Energy Information Administration, 2005. Annual Energy Review 2005. Report No. DOE/EIA-0384 (2005). Posted: July 27, 2006 ⟨http://www.eia.doe.gov/emeu/aer/txt/ptb0608.html⟩.
Energy Information Administration, 2006. Annual Energy Outlook 2006 with Projections to 2030. Report #: DOE/EIA-0383(2006) (February).
Evans, D., 2007. Integrated environmental regulation with multiple pollutants and uncertain joint abatement: theory and an application to electric utilities. Ph.D. Dissertation, Department of Economics, University of Maryland.
Evans, D., Kruger, J., 2007. Where are the sky's limits? Lessons from Chicago's cap-and-trade program. Environment 49 (2), 18–32.
Fell, H., Morgenstern, R.D., 2009. Alternative approaches to cost containment in a cap-and-trade system. Resources for the Future Discussion Paper 09-14, Washington, DC.
Fell, H., Burtraw, D., Morgenstern, R., Palmer, K., Preonas, L., 2010. Soft and hard price collars in a cap-and-trade system: a comparative analysis. Resources for the Future Discussion Paper, Washington, DC.
Harrington, W., Morgenstern, R.D., Nelson, P., 2000. On the accuracy of regulatory cost estimates. Journal of Policy Analysis and Management 19 (2), 297–317.
Hepburn, C., Grubb, M., Neuhoff, K., Matthes, F., Tse, M., 2006. Auctioning of EU ETS Phase II allowances: how and why? Climate Policy 6 (1), 135–158.
ICF Resources Inc., 1990. Comparison of the economic impacts of the acid rain provisions of the Senate Bill (S.1630) and the House Bill (S.1630[sic]). Draft report prepared for the U.S. Environmental Protection Agency, July.
Jacoby, H.D., Ellerman, A.D., 2004. The safety valve and climate policy. Energy Policy 32 (4), 481–491.
Joskow, P.L., Schmalensee, R., Bailey, E.M., 1998. The market for sulfur dioxide emissions. American Economic Review 88 (4), 669–685.
Kopp, R., Morgenstern, R.D., Pizer B., Toman, M., 2002. A proposal for credible early action in U.S. climate policy. Resources for the Future (available at ⟨www.weathervan.rff.org/features/feature060.html⟩ accessed 8/14/02).
Mohn, K., Osmundsen, P., 2008. Exploration economics in a regulated petroleum province: the case of the norwegian continental shelf. Energy Economics 30, 303–320.
Murray, B.C., Newell, R.G., Pizer, W.A., 2009. Balancing cost and emissions certainty: an allowance reserve for cap-and-trade. Review of Environmental Economics and Policy 3 (1), 84–103.
Newell, R.G., Pizer, W.A., 2008. Indexed regulation. Journal of Environmental Economics and Management 56 (3), 221–233.
Paul, A., Burtraw, D., 2002. The RFF Haiku Electricity Market Model. RFF Report (May 22).
Philibert, C., 2008. Price caps and price floors in climate policy: a quantitative assessment, International Energy Agency Information Paper (December).
Pizer, W.A., 2002. Combining price and quantity controls to mitigate global climate change. Journal of Public Economics 85 (3), 409–434.
Pizer, W.A., 2005. The case for intensity targets. Climate Policy 5 (4), 455–462.
Portney, P.R., 1990. Economics and the clean air act. Journal of Economic Perspectives 4 (4), 173–181.
Quirion, P., 2005. Does uncertainty justify intensity emission caps? Resource and Energy Economics 27, 343–353.
Regional Greenhouse Gas Initiative, 2009. RGGI states complete sixth successful CO_2 auction, news release. Available at ⟨www.rggi.org⟩ (accessed 12/10/09).
Roberts, M.J., Spence, M., 1976. Effluent charges and licenses under uncertainty. Journal of Public Economics 5 (3-4), 193–208.
Sauter, R., Awerbuch, S., 2003 (draft). Oil price volatility and economic activity: a survey and literature review. IEA Research Paper. International Energy Agency, Paris.
Stavins, R.N., 1996. Correlated uncertainty and policy instrument choice. Journal of Environmental Economics and Management 30 (2), 218–232.
Stavins, R.N., 1998. What can we learn from the grand policy experiment? Lessons from SO_2 allowance trading. Journal of Economic Perspectives 12 (3), 69–88.
Weitzman, M.L., 1974. Prices vs. quantities. Review of Economic Studies 41 (4), 477–491.
Wenders, J.T., 1975. Methods of pollution control and the rate of change in pollution abatement technology. Water Resources Research II (June), 343–346.
Zhao, J., 2003. Irreversible abatement investment under cost uncertainties: tradable emissions permits and emissions charges. Journal of Public Economics 87, 2765–2789.