Congressional Budget Office

Testimony

Statement of
Douglas W. Elmendorf
Director

Flexibility in the Timing of Emission Reductions Under a Cap-and-Trade Program

before the
Committee on Ways and Means
U.S. House of Representatives

March 26, 2009

This document is embargoed until it is delivered at 10:00 a.m. (EDT) on Thursday, March 26, 2009. The contents may not be published, transmitted, or otherwise communicated by any print, broadcast, or electronic media before that time.

CONGRESSIONAL BUDGET OFFICE
SECOND AND D STREETS, S.W.
WASHINGTON, D.C. 20515

Chairman Rangel, Ranking Member Camp, and Members of the Committee, thank you for the invitation to discuss ways to reduce the economic cost of a cap-and-trade program for greenhouse-gas emissions. That cost would depend importantly on firms' flexibility in the timing of their emission reductions. Analysts have developed a number of options for increasing timing flexibility, and this testimony reviews the advantages and disadvantages of leading options.

Accumulating evidence about the pace and potential extent of global warming has heightened policymakers' interest in cost-effective ways to achieve substantial reductions in emissions of greenhouse gases. Although the potential damage from climate change is large, the potential cost of avoiding change is large as well. Meaningfully reducing the risk of damage would require that the United States and other nations make fundamental changes in the way that energy is produced and used. Those changes could include replacing carbon dioxide-emitting fossil fuels with appropriate renewable fuels or nuclear power; reducing energy use, perhaps through major gains in energy efficiency; and capturing and storing greenhouse gases on a large scale.

Many analysts agree that the most cost-effective way to spur significant changes in the production and use of energy is to put a price on carbon emissions. By establishing such a price—rather than by dictating specific technologies or changes in behavior—the government would encourage households and firms to reduce emissions in the least costly ways. Either a carbon tax or a cap-and-trade program would effectively put a price on carbon emissions and lead to emission reductions *where* and *how* it was least costly to achieve them.

Allowing flexibility about *when* emissions were reduced would further lower the costs—and would do so without lowering the benefits—because climate change depends not on the amount of greenhouse gases released in a given year but on their buildup in the atmosphere over decades. To successfully capture all of the potential cost savings from shifting emission-cutting efforts from high-cost years to low-cost years, a cap-and-trade program would have allowance prices that did not fluctuate in response to *temporary* factors that affect compliance costs—such as the weather, economic activity, and disruptions in critical fuel markets—but that did respond to new information about more *permanent* factors that affect compliance costs over a period of many years—such as the introduction of a new technology. In practice, differentiating between temporary and permanent factors could be difficult.

My testimony makes the following key points about incorporating flexibility in the timing of emission reductions in a cap-and-trade program:

- First, permitting firms to "bank" allowances—that is, to save allowances for use in the future—has helped lower compliance costs (relative to those under inflexible annual caps) in existing cap-and-trade programs. However, the cost savings from banking are limited by the difficulty of predicting future allowance prices. Specifically, firms need to distinguish between temporary and permanent factors affecting allowance prices.

- Second, permitting firms to borrow future allowances, as well as to bank them, could further lower compliance costs. Existing cap-and-trade programs typically preclude borrowing, in part because of concerns that firms that borrow allowances might be unable to pay them back later.

- Third, permitting firms to purchase allowances from a public "reserve pool"—composed of allowances that were borrowed from future years or that supplemented the initial supply—could partially substitute for allowing borrowing by individual firms. The reserve pool could help reduce costs by giving firms the opportunity to exceed annual caps in years when the cost of complying was temporarily high. Its effectiveness in realizing cost savings would depend on the size of the pool and the threshold price at which firms could purchase the reserve allowances. The cost savings would be limited, as with banking and borrowing, by the difficulty of predicting future allowance prices.

- Fourth, setting a floor and ceiling for the price of allowances would also lower firms' compliance costs, but it would not ensure a particular level of emissions. Adjusting the floor and ceiling prices over time in order to achieve a long-term target for emissions would be complicated by firms' efforts to anticipate those shifts and to bank or borrow allowances in advance of them.

- Finally, a "managed-price" approach would allow for substantial cost savings by eliminating short-term volatility in the price of allowances while accommodating longer-term shifts in prices that would be necessary to keep emissions within a multidecade cap set by legislators. In a managed-price arrangement, firms could purchase allowances from the government each year at a price specified by regulators. The policy would be similar to a tax in that respect. Unlike a tax, however, the policy would not require firms to purchase *all* of their allowances from the government: Policymakers could choose to distribute some allowances to firms for free and could allow them to comply by purchasing "offsets," or credits for emission reductions made in sectors not covered by the cap. Regulators would establish a path of rising prices for allowances, with the goal of complying with the cumulative cap that legislators set. That path would be adjusted periodically if new information indicated that future compliance costs were going to be higher or lower than anticipated or if progress in meeting the cumulative cap was less than expected. A key issue would be how frequently to adjust the price path: Increasing the length of time between adjustments would help policymakers distinguish between temporary and permanent cost factors but could also require larger adjustments to keep emissions on track to meet the long-run cap.

Minimizing Costs Over Time in a Cap-and-Trade Program

Under a cap-and-trade program, policymakers would typically set an annual cap on emissions for each of the years covered by the policy and allocate an allowance for each ton of emissions permitted under the cap. The cap would become more stringent over time, resulting in a decline in the number of allowances allocated.

Compared with regulations requiring the use of particular technologies or reductions from particular sources of emissions, the cap-and-trade program would reduce the cost of obtaining the desired level of emissions by allowing firms to reduce emissions where and how it was least costly to do so. After the allowances were initially distributed, firms would be free to trade them. As a result, firms that could reduce emissions most cheaply would profit by selling their excess allowances to firms facing higher emission-cutting costs.

Offering the firms subject to the cap additional flexibility as to *when* cuts in greenhouse gases were made—by requiring that they meet the annual caps only on average—could result in substantial additional cost savings, while producing the same effect on the climate. That opportunity to reduce the cost of a cap-and-trade program without reducing its benefits stems from the long-run nature of climate change. Climate change results from the accumulation of greenhouse gases in the atmosphere over many decades and centuries, yet year-to-year fluctuations in emissions have little effect. By contrast, the economic cost of reducing emissions can vary a lot from year to year—depending on the weather, economic activity, and the prices of fossil fuels. In order to minimize the cost of achieving a cap, firms would want to weigh the cost of reducing emissions today against the cost of reducing them later, and allocate their emission-cutting efforts accordingly. In making that trade-off, firms would discount future costs at a rate equal to the rate of return on alternative investments that they might make in lieu of reducing emissions.[1]

To capture all of the potential cost savings from shifting emission-cutting efforts from high-cost years to low-cost years, the allowance price should not fluctuate in response to temporary factors. If prices did respond to temporary factors, additional costs would be incurred because too many emission reductions would be made when prices were high and too few when prices were low. The variation in the cost of reductions could disrupt production processes and economic activity and make planning difficult for firms and households.[2]

Although prices in an efficient cap-and-trade program would not fluctuate in response to temporary cost factors, they would change in response to new information about lasting or permanent factors that affected compliance costs over many years. For example, if a new emission-reducing technology proved to be more expensive than previously anticipated, then current and future allowance prices would have to

1. For a discussion of why a cost-minimizing tax would increase at that rate, see Gilbert E. Metcalf and others, *Analysis of a Carbon Tax to Reduce U.S. Greenhouse Gas Emissions* (Cambridge, Mass.: Massachusetts Institute of Technology Joint Program on the Science and Policy of Global Change, 2008), pp. 28–29.

2. Metin Celebi and Frank Graves, "CO_2 Price Volatility: Consequences and Cures" (discussion paper, The Brattle Group, January 2009).

be higher in order to keep emissions within the desired cap. In contrast, if a technology proved to be less costly than anticipated, allowance prices could be lower. Unlike the effect of temporary cost factors, the new information about future costs would generally lead to a one-time increase or decrease in both current and future prices.

While new information about future compliance costs would require an adjustment to both current and future prices, those adjustments would not lead to a change in the emissions made over the period that the policy was in effect. In contrast, new information indicating that the damage from climate change might be greater or lesser than previously anticipated would necessitate both a change in the cumulative cap and a shift in the path for prices.

Options for Providing Intertemporal Flexibility

Reducing the potential risk of climate change would entail reducing the emissions of greenhouse gases that could occur over multiple decades. Options for granting flexibility as to when the emission reductions occur fall into two categories: The first category, which includes banking and borrowing allowances or creating a public reserve pool of them, would permit firms to transfer allowances across time. The second category, which includes setting a floor and a ceiling for the price of allowances or using a managed-price approach to specify a path for allowance prices over time, would permit regulators to set allowance prices in a manner that induced a cost-effective time pattern of emissions.

Banking

Banking would provide firms with the ability to take advantage of the cost-saving opportunities created when the cost of reducing emissions to meet an annual cap was unusually low. Firms could trim emissions further and bank allowances in that situation and then use the banked allowances, or sell them at a profit, when trimming emissions was more difficult and compliance costs were high. Because banking would increase the demand for allowances in low-cost periods (when firms wanted to accumulate them) and increase the supply of allowances in high-cost periods (when firms withdrew allowances from the bank), banking would reduce price fluctuations and lower the cost of achieving a cumulative cap.

Experience with existing programs that did *not* include banking illustrates the crucial role it can play in helping to prevent major disruptions in allowance prices. For example, the lack of banking contributed to an extreme spike in allowance prices in the Regional Clean Air Incentives Market (RECLAIM), a program that was designed to address nitrogen oxide, a contributor to ozone pollution in the Los Angeles basin region. In the initial years after RECLAIM went into effect in 1994, light trading occurred and allowance prices were low. But prices changed radically during the summer of 2000: Allowance prices in that year averaged $45,609—more than 10 times higher than in previous years—and spot prices reached as high as $90,000.

Although the cost of meeting the cap in the initial years of the program was low, firms were not allowed to reduce emissions below the cap and build up a bank of allowances. When a heat wave in the summer of 2000 caused a surge in demand for electricity in California and a spike in electricity prices, the demand for allowances increased.[3] But because only a limited supply of allowances existed, prices shot up. As a result, regulators removed energy generators from the cap-and-trade program and returned them to relatively costly command-and-control regulations.

The success of banking in minimizing compliance costs over time depends on firms' ability to discern the difference between temporary cost factors that lead to a short dip in compliance costs and more permanent cost factors that result in a sustained reduction in compliance costs. Firms would find it cost-effective to bank allowances in the former case but not in the latter. But discerning the difference between temporary and permanent cost factors can be difficult (for both firms and policymakers). As a result, banking would allow firms to capture some, but not all, of the cost savings that could be obtained from shifting emission-cutting efforts across time. Recent research indicates that the cost savings from banking in a greenhouse-gas cap-and-trade program would be larger if shocks in compliance costs did not persist over time.[4]

To the extent that firms were successful in differentiating between temporary and permanent cost factors, banking would decrease volatility in allowance prices. Even successful banking, however, would not eliminate price changes. If new information indicated that the cost of complying with caps in the future was going to be higher than previously anticipated—for example, because of information that a new technology was going to cost more than anticipated—then the current price of allowances would increase and firms would appropriately reduce emissions further today.

Yet even with banking in place, volatility in allowance prices in existing cap-and-trade programs appears to be greater than can be explained by changes in expectations about future compliance costs. In the Acid Rain Program, which capped sulfur dioxide emissions beginning in 1995, prices for allowances fluctuated considerably within short periods of time. Prices varied from less than $75 to more than $200 in roughly three years, in spite of the fact that firms were allowed unlimited banking (see Figure 1).[5] Similarly, allowance prices have fluctuated considerably in a cap-and-trade program designed to reduce carbon dioxide emissions in the European Union—even

3. Carol Coy and others, "Stabilization of NOx RTC Prices" (white paper, South Coast Air Quality Management District, January 11, 2001), p. 12.

4. See Harrison Fell, Ian A. MacKenzie, and William A. Pizer, *Prices Versus Bankable Quantities,* Discussion Paper DP 08-32-Rev (Washington, D.C.: Resources for the Future, July 2008). Another key factor is the rate at which firms discount future costs. A lower rate induces more banking and generally lowers aggregate compliance costs over the period of the policy.

5. Prices after 2003 were affected by an anticipated change in the stringency of future caps. Banking enabled allowance prices to adjust in anticipation of that tightening, but that adjustment could have ultimately increased, not decreased, compliance costs because the increase in allowance prices far exceeded estimates of the cost of meeting the tighter cap.

Figure 1.

Prices for Sulfur Dioxide Allowances in the Acid Rain Program

(1994 dollars per ton)

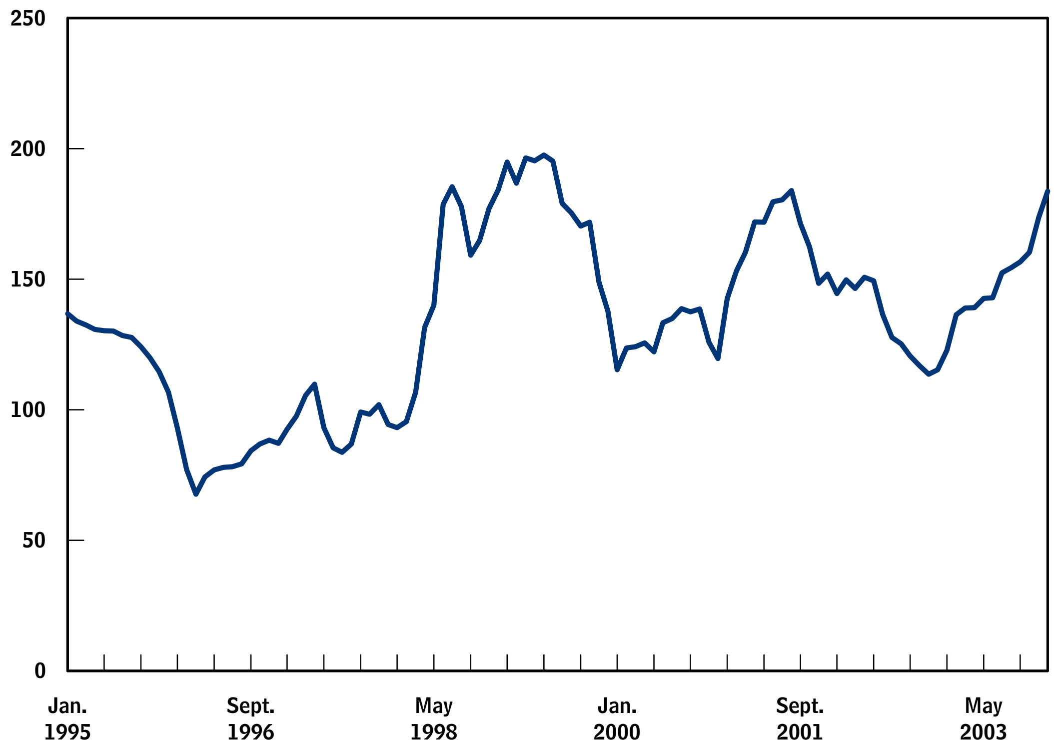

Source: Congressional Budget Office based on data from Resources for the Future.

though firms can bank allowances both within the phase of the program covering 2008 to 2012 and into the next phase, and they can borrow allowances from one year ahead. For example, prices for what are termed "vintage 2009" allowances, which may be used to comply with emission requirements in either 2008 or 2009, varied considerably throughout 2008. They started the year at less than $20, rose to over $28, dropped to roughly $16 by the end of 2008, and continued to fall in the beginning of 2009 (see Figure 2). Those prices were closely correlated with oil prices.

Borrowing

Firms would have a motivation to borrow future allowances whenever they thought that current prices for allowances were high relative to future prices. That circumstance could be the result of a temporary spike in current prices or the expectation that a new technology would fundamentally lower compliance costs in the future. Some proposals for cap-and-trade programs would allow firms to meet a limited amount of their current requirement (typically no more than 15 percent) to reduce emissions with allowances borrowed from the future (typically no more than five years out). Like banking, the extent to which borrowing would lower firms' compliance

Figure 2.

Allowance Prices in the European Union's Cap-and-Trade Program

(Dollars per ton)

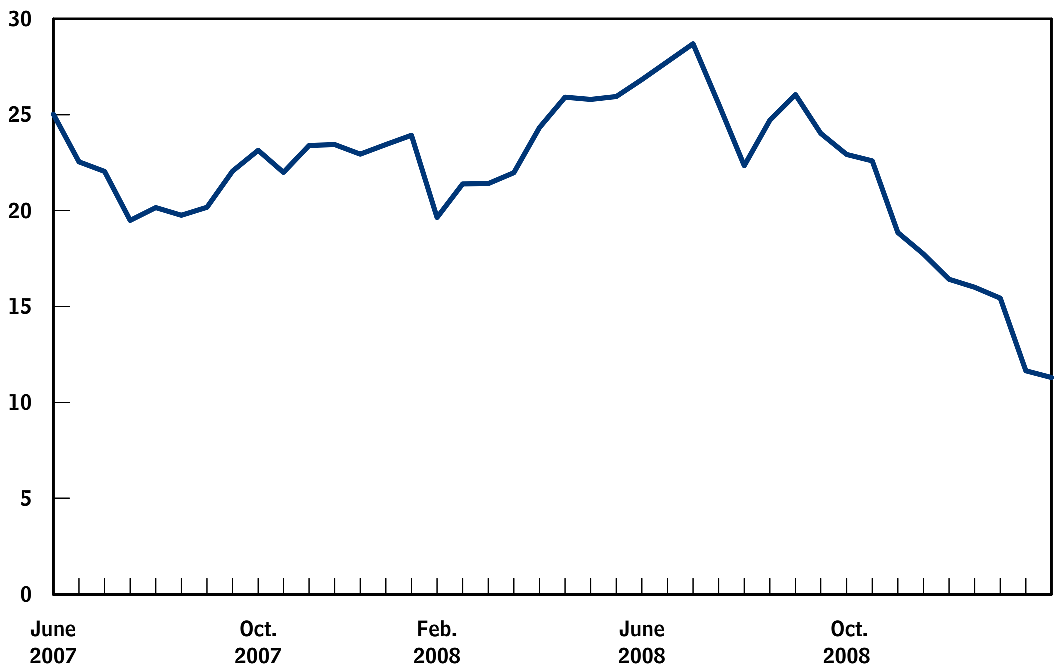

Source: Congressional Budget Office based on data from the European Energy Exchange.

Note: Prices are for "vintage 2009" allowances, which may be used to comply with emission requirements in either 2008 or 2009.

costs over the life of the policy depends on firms' assessment of future costs. There is very little evidence on the effectiveness of borrowing: California's RECLAIM program did not permit borrowing, and the Acid Rain Program does not allow it, in part because of concerns about some firms' borrowing indefinitely and about determining liability if a firm that borrowed allowances was unable to accomplish the necessary emission reductions. As noted above, the cap-and-trade program for carbon dioxide emissions in the European Union permits firms only a very limited form of borrowing by allowing them to comply by using allowances one year in advance of the current year.

Reserve Pool

Policymakers could add additional flexibility to a cap-and-trade system that included banking by allowing firms an opportunity to purchase allowances from a publicly available reserve pool at, or above, a threshold price. If firms did not have a substantial bank of allowances to draw upon and if borrowing were restricted, then the reserve pool would be the primary method of preventing price spikes and relatively costly emission reductions. Thus, the reserve pool could play an especially important role in the initial years of a cap-and-trade program before firms had the chance to build up a bank. The importance of the reserve pool in later years would depend, in part, on the size of the banks that firms had built up.

A spike in compliance costs could make it advantageous for firms to purchase reserve allowances. For example, if allowances in 2020 would cost $40 in the absence of a reserve pool and the threshold price for allowances was $35, then firms would find it cost-effective to purchase reserve allowances. That additional supply of reserve allowances would tend to reduce the price of allowances in 2020 below the $40 it would have otherwise been. Thus, the reserve pool could help limit price increases that might otherwise occur as the result of temporary factors.

The extent to which the reserve pool would be successful in limiting temporary price increases, however, would depend on policymakers' decisions about the level of the threshold price and the size of the reserve pool. A higher threshold price would allow for a larger price spike because it would permit allowance prices to climb higher before firms found it cost-effective to purchase allowances from the reserve pool. In contrast, a lower threshold price would increase the likelihood that firms would want to purchase reserve allowances: The reserve pool would prevent allowance prices from climbing above the threshold price only if the number of allowances that firms wanted to buy at the threshold price was less than the quantity in the pool. That condition would be more likely to hold when the demand for reserve pool allowances was driven by the expectation that compliance costs were *temporarily* high.

Although the reserve pool could help limit temporary spikes in allowance prices, it would not serve as a ceiling for the price of allowances. In particular, the price of allowances might exceed the threshold price if firms anticipated that the cost of meeting annual caps was likely to increase significantly in the future and therefore wanted to bank allowances for future use. Indeed, demand for reserve allowances induced by higher expected prices in future years would come not only from firms subject to the cap but also from third-party traders who might choose to buy reserve allowances, bank them, and sell them at a profit in the future. The size of the allowance market anticipated under a cap-and-trade program for greenhouse-gas emissions would ensure that third-party traders would be interested in taking advantage of any such opportunities for intertemporal arbitrage. For example, the Congressional Budget Office estimated that the value of the allowances created under the cap-and-trade legislation order reported by the Senate Committee on the Environment and Public Works on December 5, 2007, would total nearly $1.2 trillion from 2012 to 2018.[6]

The demand for allowances for banking purposes would increase their prices until they equaled the present value of expected future prices, regardless of the current threshold price. Still, intertemporal arbitrage could lower compliance costs in the aggregate: If firms' expectations about future prices were correct, such arbitrage would

6. Congressional Budget Office, cost estimate for S. 2191, America's Climate Security Act of 2007 (April 10, 2008).

shift the price path for allowances in a manner that lowered total compliance costs over time—with allowance prices lower in the future than they otherwise would have been.

Depending on the design of the reserve pool, firms' purchases of reserve allowances might or might not affect cumulative emissions over time. If the reserve pool was composed of allowances that otherwise would not be allocated, then drawing on it would lead to higher total emissions than would otherwise occur. In contrast, if the reserve pool was composed of allowances that were borrowed from future periods and returned if not used, then using the pool would not lead to higher total emissions.

Finally, if firms expected lower allowance prices in the future (because of the introduction of a new technology, for example), they would want to defer some of their emission reductions. In that case, the reserve pool would be unlikely to help them realize cost savings: Firms would not buy reserve allowances at a threshold price that was higher than the allowance prices they anticipated in the future. As a result, the reserve pool would not facilitate a downward shift of the price path.[7]

Price Floor and Price Ceiling

Policymakers could help reduce fluctuations in allowance prices by setting a floor and a ceiling (often referred to as a safety valve) for the price of allowances. The price floor would induce firms to make more emission reductions than would be necessary to meet the cap in low-cost years. Creating such a price floor would be fairly straightforward if the government chose to sell a significant share of the allowances rather than giving them free of charge to affected businesses: Policymakers could specify a reserve auction price and restrict the supply of allowances to maintain that price. The price ceiling would allow firms to make fewer emission reductions in high-cost years, thereby exceeding the annual cap. Creating such a price ceiling simply would require the government to sell as many allowances as buyers desired at that price.[8]

7. However, if firms had been banking in expectation of higher future prices, news of lower compliance costs in the future could motivate them to cut back on their banking, which, in turn, would shift the price path down.

8. One criticism of a cap-and-trade program that includes only a price ceiling is that it could reduce firms' incentives to replace carbon-intensive capital equipment and to develop new technologies for lowering carbon dioxide emissions. The fact that the range of potential future prices would be truncated at the high end by the price ceiling, but not at the low end, would reduce the expected price for allowances. See Dallas Burtraw and Karen Palmer, "Dynamic Adjustment to Incentive Based Policy to Improve Efficiency and Performance" (draft, Resources for the Future, Washington, D.C., November 30, 2006); and Congressional Budget Office, *Policy Options for Reducing CO_2 Emissions* (February 2008).

Figure 3.

Example of a Floor and Ceiling for Allowance Prices Based on Illustrative Price Fluctuations

(Allowance price)

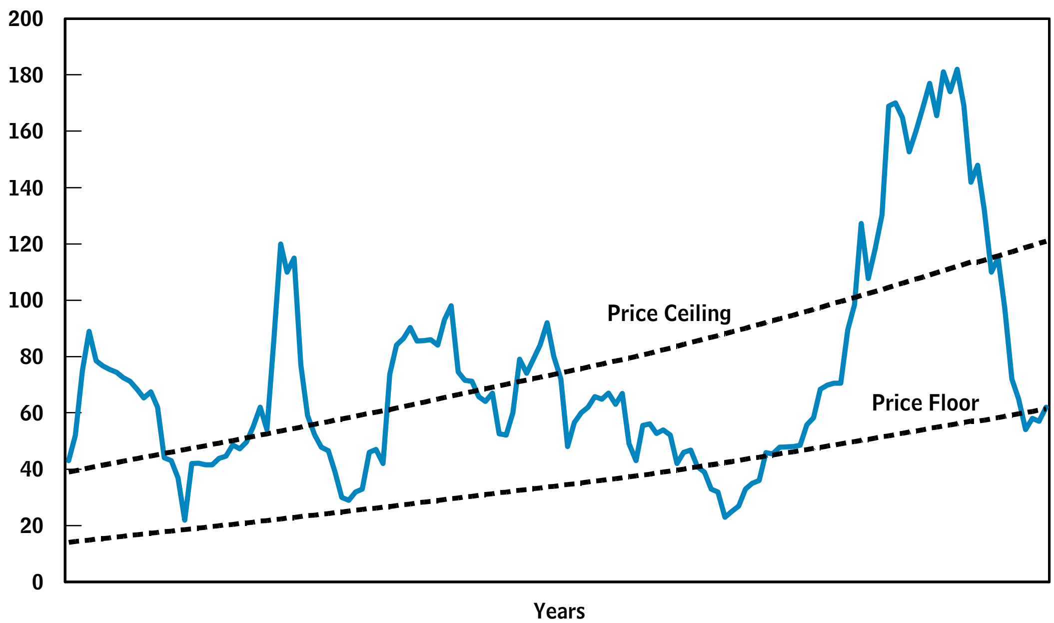

Source: Congressional Budget Office.

Just as firms would wish to allocate their compliance costs over time to match the opportunity costs of other investments, the price floor and ceiling could be increased over time at the same rate that firms would discount future costs. The rising floor and ceiling would, on average, induce greater emission reductions in successive years.

A cap-and-trade program that included both a floor and a ceiling could limit the potential range of allowance prices (see the illustrative example in Figure 3), but it would not ensure any particular outcome for reducing emissions. Emissions would exceed the cumulative limit implied by the sum of annual caps if the additional emissions that occurred when firms purchased allowances at the price ceiling exceeded the additional reductions induced by the price floor. Conversely, emissions would fall short of the cumulative limit if additional emissions when prices hit the ceiling were outweighed by additional reductions when prices reached the floor.

Policymakers could attempt to adjust the price floor and ceiling in order to achieve a desired long-term target—for example, by shifting the "price corridor" upward if reductions were falling below the desired level—but such adjustments could be less effective if firms were allowed to bank allowances. If firms anticipated an upward shift in the price corridor, they would have an incentive to buy large quantities of allowances at the current ceiling price, bank them, and use them in the future after the price ceiling had increased. That action could lead to excess emissions (above the cap) after the price increase occurred: If firms banked the allowances bought at the ceiling price, those allowances would be used later without having brought about actual reductions in emissions.

Policymakers could attempt to minimize that adverse outcome by limiting the amount of allowances that could be banked or by eliminating banking altogether, but doing so could restore some price variability and reduce some cost savings. In particular, eliminating banking would not let firms smooth allowance prices between the boundaries defined by the floor and ceiling; the greater the distance between the floor and ceiling, the more costly that problem would be. In addition, significant restrictions on banking could lead to volatile allowance prices at the end of each compliance period. For example, if firms anticipated that fewer allowances would be available than required to meet the demand at a price below the ceiling, then allowance prices would jump to the ceiling level. In contrast, if firms anticipated that the supply of allowances was ample to meet the demand at a price below the floor, the price would drop to the level of the floor. The most disruptive outcome would occur if expectations about the supply of and demand for allowances changed. In that case, allowance prices could bounce between the floor and ceiling prices.

The Managed-Price Approach

A managed price for allowances takes the notion of defining a price corridor a step further, in essence eliminating the distance between the floor and the ceiling. Under this approach, legislators would set a cap on cumulative emissions over a period of several decades but would not set annual caps. Regulators, in turn, would be charged with setting allowance prices for each year of the policy—with the objective of choosing prices that would minimize the cost of achieving the multidecade cumulative cap. For example, legislators might specify a limit on cumulative emissions occurring from 2012 to 2050. Regulators would then specify a single allowance price for each year of that span.[9]

To achieve the goal of minimizing total compliance costs, regulators would weigh the cost of reducing emissions today against the cost of reducing them at a future date, just as firms would with banking and borrowing. The specified allowance prices would then rise annually at a rate equal to an estimated average rate of return on investments that firms could have made in lieu of reducing emissions. The key decision for regulators would be, What initial price (which would then rise at the defined rate) would be necessary to set the economy on track for achieving the desired cap on emissions? Regulators would need to adjust the chosen price path periodically to ensure that the level of emissions was on target for achieving the cap.

9. A variation of the concept may be found in Center for Clean Air Policy, *Preventing Market Disruptions in Cap-and-Trade Programs* (October 2008), available at www.ccap.org. Under that approach, regulators would attempt control the price of allowances sold in four separate lots over the course of each year during the initial years of the cap-and-trade program.

Under this approach, firms would be able to purchase allowances from the government in each year at the specified price. Firms would cut emissions if the cost of doing so was less than the price of an allowance. Firms would not be allowed to bank or borrow allowances, but smoothly increasing allowance prices would automatically capture much of the intertemporal cost savings that banking and borrowing were designed to achieve.

This option is similar to a tax on greenhouse gases that regulators would adjust periodically on the basis of new information.[10] But it also has features of a traditional cap-and-trade program, including these:

- It would allow policymakers to distribute a share of the allowances to firms for free if they wished to do so. For example, they could give allowances to coal producers, electricity generators, or to dislocated workers in the coal industry. Firms could sell any allowances that they were given but did not use. However, giving away too many allowances could undermine regulators' ability to maintain the managed price.

- It would allow firms to meet a portion (specified by policymakers) of their requirement for allowances by purchasing "offsets," which are credits for qualifying emission reductions in areas not subject to the cap. For example, individuals or firms not subject to the cap might generate offsets through biological sequestration that stored carbon in plants and soil (by pursuing afforestation or adopting certain agricultural practices, for instance) and then those offsets could be available for purchase by firms subject to the cap.[11]

- It would allow firms to reduce uncertainty about their future compliance costs by entering into futures contracts. For example, firms subject to the cap could agree to buy allowances at fixed prices in the future from traders that were willing to absorb the risk that the price of allowances would turn out to be higher or lower than anticipated.

---

10. For a discussion of a tax on carbon dioxide emissions that could be adjusted over time, see Gilbert Metcalf and David Weisbach, *Design of a Carbon Tax,* Olin Working Paper No. 447 and Public Law Working Paper No. 254 (University of Chicago, December 2008). For an overall discussion of timing flexibility, a tax on carbon dioxide emissions, and an overview of the design features that can be included in a cap-and-trade program to provide timing flexibility, see Congressional Budget Office, *Policy Options for Reducing CO2 Emissions*. For a discussion of the potential advantages of reducing greenhouse-gas emissions by a tax, see Gilbert Metcalf, "An Equitable Tax Reform to Address Climate Change" (discussion paper, Brookings Institution, The Hamilton Project, October 2007), available at www.brookings.edu/papers/2007/10carbontax_metcalf.aspx. For a discussion of the potential advantages of a cap-and-trade program, see Robert Stavins, "A U.S. Cap-and-Trade Program to Address Climate Change" (discussion paper, Brookings Institution, The Hamilton Project, October 2007), available at www.brookings.edu/papers/2007/10climate_stavins.aspx.

11. See Congressional Budget Office, *The Potential for Carbon Sequestration in the United States* (September 2007).

■ It would set a fixed cap on cumulative emissions over a period of several decades.

In order to forecast the price path that would minimize the cost of achieving the desired limit on cumulative emissions, regulators would rely on information about trends in future emissions, firms' and households' responses to higher energy prices, and the availability of technologies in the future. (Analysts would use that same information to predict the allowance prices under cap-and-trade programs that included annual caps and allowed banking.)

By design, the managed-price option would prevent temporary factors—such as fluctuations in the economy, the weather, or energy markets—from affecting allowance prices and from increasing the cost of achieving a multidecade cap. However, the initial price path that regulators set might lead to more or less cumulative emissions than they had anticipated. They would need to make periodic adjustments to the price path to ensure that adequate progress was being made toward the cap and to reflect significant changes in trends in future emissions and new information about future technologies. For example, an increase in underlying emission trends would mean that allowance prices had to be higher than regulators originally foresaw to keep emissions from exceeding the cap; in contrast, the development of an unexpectedly cheap technology for reducing emissions would mean that the target could be met with lower prices and less economic cost. Those adjustments would entail one-time increases or decreases in the current and future prices—that is, upward or downward shifts in the price path.

Allowing a longer time between adjustments to the price path would provide longer periods of certainty about prices for firms that needed to comply with the policy and that wished to invest in new technologies for reducing emissions. In addition, longer intervals would provide more time for the effect of temporary influences on emissions to balance out (for a cold winter to be offset by a mild winter, for example). However, infrequent adjustments to the price path—say, every 8 to 10 years rather than every 3 to 5 years— might also need to be larger: Cumulative emissions would have had more time to get off course if the initial price path was set too high or too low, and more information about new technologies and baseline emission trends would have accumulated.[12]

Lawmakers might decide to delegate responsibility for setting and modifying allowance prices to regulators, such as a regulatory agency, commission, or the executive branch. In that case, the Congress could specify a mandatory cap for period covered (through 2050, for example), specify how frequently prices should be adjusted, and require reports on that updating process. Regulators would then use forecasts of the cost of reducing emissions to estimate the least-cost path for allowance prices and to

12. In contrast, banking would allow firms to immediately adjust to new information suggesting that compliance costs might be higher than previously anticipated. Firms would be able to adjust to new information indicating that compliance costs might be lower than previously anticipated only if they were allowed to borrow future allowances.

make periodic adjustments. If lawmakers did not wish to delegate that responsibility and did not want to rely on additional legislation, which might be quite cumbersome, the Congress could specify how prices should be modified under alternative conditions. A disadvantage of this approach compared with delegating price-setting authority is that it could be less flexible in responding to new information.