



Testimony

**Statement of
Peter R. Orszag
Director**

Containing the Cost of a Cap-and-Trade Program for Carbon Dioxide Emissions

**before the
Committee on Energy and Natural Resources
United States Senate**

May 20, 2008

This document is embargoed until it is delivered at 10:00 a.m. (EDT) on Tuesday, May 20, 2008. The contents may not be published, transmitted, or otherwise communicated by any print, broadcast, or electronic media before that time.

Chairman Bingaman, Senator Domenici, and Members of the Committee, thank you for the invitation to discuss the implications of cap-and-trade programs that are designed to reduce U.S. emissions of greenhouse gases, most prominently carbon dioxide (CO₂). Under a cap-and-trade program, policymakers would set a limit on emissions and allow entities to buy and sell permits (or allowances) to emit CO₂ and other greenhouse gases.

Global climate change is one of the nation's most significant long-term policy challenges. Human activities are producing increasingly large quantities of greenhouse gases, particularly CO₂. The accumulation of those gases in the atmosphere is expected to have potentially serious and costly effects on regional climates throughout the world. Although the magnitude of such damage remains highly uncertain, there is growing recognition of the risk that the damage could prove extensive and perhaps even catastrophic. The risk of potentially catastrophic damage associated with climate change can justify actions to reduce that possible harm in much the same way that the hazards we all face as individuals motivate us to buy insurance.

Reducing greenhouse-gas emissions would provide benefits to society by helping to limit the damage associated with climate change, especially the risk of significant damage. However, decreasing those emissions would also impose costs on the economy—in the case of CO₂, because much economic activity is based on fossil fuels, which release carbon when they are burned.

Most analyses suggest that an appropriately designed program to begin lowering CO₂ emissions would produce greater benefits than costs. Market-oriented approaches to reducing carbon emissions, such as a cap-and-trade program or a carbon tax, would reduce emissions more cheaply than would command-and-control approaches, such as regulations requiring across-the-board reductions by all firms. Those market-oriented approaches are relatively efficient because they create incentives and flexibility for emission reductions to occur where and how they are least expensive to accomplish.

I will focus today on two key design elements of a cap-and-trade system that could help to improve its efficiency further in terms of reducing the cost of emission reductions: (1) structural features to allow the timing of reducing emissions to respond to year-to-year differences in conditions that affect the cost of doing so and (2) the use of the allowances' value created by a cap-and-trade system to reduce its cost.

The Congress is currently considering a bill, S. 2191, which would reduce emissions by establishing a cap-and-trade program.¹ S. 2191 would also establish a Carbon Market Efficiency Board, which would be authorized to transfer emission allowances

1. The Congressional Budget Office (CBO) reviewed S. 2191 as the bill was ordered reported by the Senate Committee on Environment and Public Works on December 5, 2007. As discussed later, on April 10, 2008, CBO provided a cost estimate for the bill as it was ordered reported and a cost estimate for it with a proposed amendment transmitted to the agency on April 9, 2008.

across years to help minimize the cost of meeting a long-term target for reducing emissions. Other approaches—such as imposing limits on the price of allowances—could also be used to contain the costs that a cap might impose on the economy.

My testimony makes the following key points:

- The cost of meeting an emission target with a cap-and-trade program could be reduced, potentially quite substantially, by providing firms flexibility in the timing of their efforts to reduce emissions. In particular, the most cost-effective cap-and-trade design would encourage firms to make greater reductions when the cost of doing so was low and would allow them leeway to lessen their efforts when the cost was high. Providing firms with such flexibility could also prevent large fluctuations in the price of allowances that could be disruptive to the economy. The reduction in economic burden need not come at the cost of additional environmental risk: The flexibility to shift emission reductions across years could be designed to achieve any given cumulative reduction in emissions over the medium or long term.
- One option for allowing firms flexibility in determining when to reduce emissions while also achieving compliance with a cumulative target would be through setting both a ceiling—typically referred to as a safety valve—and a floor on the allowance prices each year. The price ceiling would allow firms to exceed the annual target when the cost of cutting emissions was high, while the price floor would induce firms to cut emissions more than the annual target in low-cost years. The price ceiling and floor could be adjusted periodically to ensure that emission reductions were on track for achieving the long-run target; such a dynamic price system could substantially reduce the cost of a cumulative target for emissions.
- Another option would be to authorize firms to “borrow” future allowances for use in the current year or to “bank” allowances for use in future years. Firms would have an incentive to borrow allowances, though, only if they expected the price in the future to be sufficiently lower than the current price to make borrowing cost-effective. Similarly, firms would have an incentive to bank allowances only if they expected the price in the future to be sufficiently higher than the current price. Most proposals for borrowing and banking would impose limits on the degree to which they could be undertaken, and partially as a result of those limits, this approach is likely to be less effective at reducing cumulative costs for any given cumulative target for reducing emissions than a dynamic price system would be.
- Under the Carbon Market Efficiency Board described in S. 2191, which would be authorized to transfer emission allowances across time periods, regulators would attempt to shift allowances in a manner that led to more reductions when costs were relatively low and less reductions when costs were high. An alternative approach, which may be easier for regulators to implement efficiently, would be to have the board set a ceiling and floor for allowance prices and be responsible for adjusting those price limits periodically as needed to achieve a long-term target for reducing emissions.

- Policymakers' choices about whether to distribute the allowances without charge or to auction them—and if they are auctioned, how to use the proceeds—could also have a significant effect on the overall economic cost of capping emissions. Evidence suggests that the cost to the economy of a 15 percent cut in U.S. emissions (not counting any benefits from mitigating climate change) might be half as large if policymakers sold the allowances and used the revenue to lower current taxes on capital that discourage economic activity, rather than giving the allowances away to energy suppliers and energy-intensive firms or using the auction proceeds to reduce the costs that the policy could impose on low-income households. Using the allowances' value to lower the total economic cost could, however, exacerbate the regressivity of the cap-and-trade program.

Containing Costs by Providing Flexibility in the Timing of Emission Reductions

A cap-and-trade program, which creates financial incentives for firms and households to cut their greenhouse-gas emissions, is a lower-cost approach to reducing emissions than more restrictive command-and-control approaches, which mandate how much those entities can emit or what emission-reduction technologies they should use. The lower cost of a cap-and-trade program stems from the flexibility it provides as to where and how emission reductions are to be achieved.

Under a cap-and-trade program for CO₂, policymakers would set a limit on total emissions during some period and would require regulated entities to hold allowances for the emissions permitted under that cap. (Each allowance would entitle companies to emit one ton of CO₂ or to have one ton of carbon in the fuel that they sold.) After the allowances for a given period were distributed, entities would be free to buy and sell them. The trading aspect of the program could lead to substantial cost savings relative to command-and-control approaches: Firms that were able to reduce emissions most cheaply could profit from selling allowances to firms that had relatively high abatement costs. The cost-effectiveness of a cap-and-trade program could be further improved by providing firms with flexibility in determining when to reduce their emissions.

The Importance of Flexibility in the Timing of Emission Reductions

In its most inflexible form, a cap-and-trade program would require that a specified cap on emissions was met each year. That lack of flexibility would increase the cost of achieving any long-term goal because it would prevent firms from responding to year-to-year differences in conditions that affected costs for reducing emissions, such as fluctuations in economic activity, energy markets, and the weather (for example, an exceptionally cold winter would increase the demand for energy and make meeting a cap more expensive), and the technologies available for reducing emissions.

In contrast, because of the long-term nature of climate change, the key issue from an environmental perspective involves emissions over the long term and concentration paths of greenhouse gases, not the year-to-year fluctuations in emissions. In other

words, limiting global climate change will entail substantially reducing the amount of greenhouse gases that accumulate in the atmosphere over the next several decades, but the benefits of doing so are largely independent of the annual pattern of those reductions.² Consequently, a cap-and-trade program could achieve roughly the same level of benefits at a significantly lower cost if it provided firms with an incentive to make greater reductions in emissions at times when the cost of doing so was low and allowed them leeway to lessen their efforts when the cost was high.

Including features in a cap-and-trade program that enabled to firms to reduce emissions less when costs were high and more when costs were low could also reduce the volatility of allowance prices. Experience with cap-and-trade programs has shown that price volatility can be a major concern when a program's design does not include provisions to adjust for unexpectedly high costs and to prevent price spikes. For example, one researcher found that the price of sulfur dioxide allowances under the U.S. Acid Rain Program was significantly more volatile than stock prices between 1995 and 2006 (see Figure 1).³

Price volatility could be particularly problematic with CO₂ allowances because fossil fuels play such an important role in the U.S. economy. In 2006, fossil fuels accounted for 85 percent of the energy consumed in the United States. CO₂ allowance prices could affect energy prices, inflation rates, and the value of imports and exports. If those prices were volatile, they could have disruptive effects on markets for energy and energy-intensive goods and services and could make investment planning difficult.

Design Features Providing Flexibility in the Timing of Emission Reductions

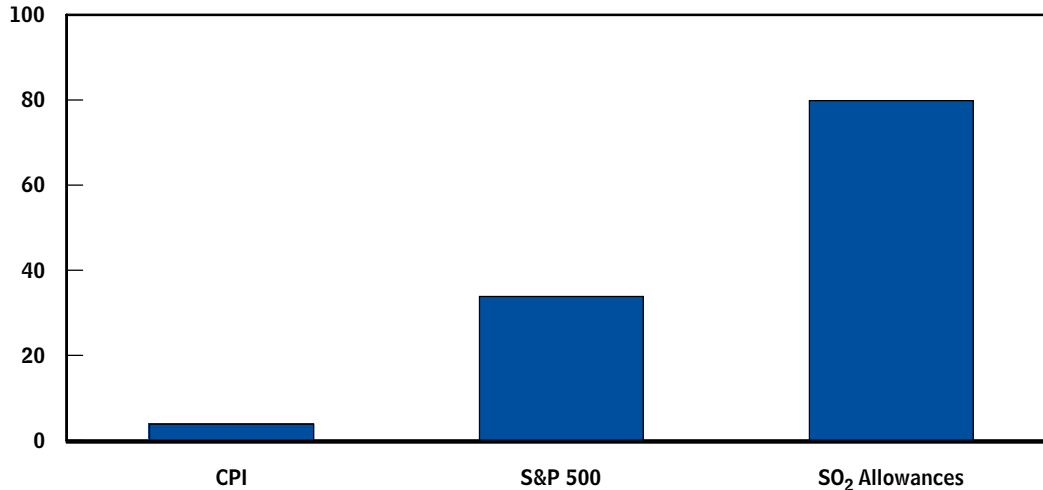
Recent proposals for cap-and-trade proposals include a variety of design features that would provide firms or regulators with flexibility in the timing of emission reductions, thereby reducing the economic costs of the effort to limit greenhouse gas emissions.

-
2. Although costs and benefits are difficult to measure, the long-term cumulative nature of climate change implies that the *benefit* of emitting fewer less ton of CO₂ in a given year—referred to as the marginal benefit—is roughly constant. In other words, the benefit in terms of averted climate damage from each additional ton of emissions reduced is roughly the same as the benefit from the previous ton of emissions reduced, and shifting the reductions from one year to another does not materially affect the ultimate impact on the climate. In contrast, the *cost* of emitting one fewer ton of CO₂ in a given year—the marginal cost—tends to increase with successive emission reductions. The reason is that the least expensive reductions are made first and progressively more-expensive cuts would then have to be made to meet increasingly ambitious targets for emission reductions.
 3. See William D. Nordhaus, “To Tax or Not to Tax: Alternative Approaches to Slowing Global Warming,” *Review of Environmental Economics and Policy*, vol. 1, no. 1 (Winter 2007), pp. 37–39.

Figure 1.

Volatility of SO₂ Allowance Prices and Selected Other Prices, 1995 to 2006

(Average annual percentage rate of volatility)



Source: Congressional Budget Office based on William D. Nordhaus, "To Tax or Not to Tax: Alternative Approaches to Slowing Global Warming," *Review of Environmental Economics and Policy*, vol. 1, no. 1 (Winter 2007), pp. 26–44.

Note: Volatility is calculated as the annualized absolute logarithmic month-to-month change in the consumer price index (CPI), the stock price index for the Standard & Poor's 500 (S&P 500), and the price of sulphur dioxide (SO₂) allowances under the U.S. Acid Rain Program.

A Price Ceiling and a Price Floor. The combination of a price ceiling and a price floor offers one method of allowing timing flexibility and thereby reducing the economic burden of achieving any desired cumulative target for reducing emissions:

- Setting a ceiling, or safety valve, for the price of allowances could prevent the cost of reducing emissions from exceeding either the best available estimate of the environmental benefits or the cost that policymakers considered acceptable. The government could maintain a price ceiling by selling companies as many allowances as they would like to buy at the safety-valve price.
- Similarly, policymakers could prevent the price of allowances from falling too low by setting a price floor. If the government chose to auction a significant share of the allowances, it could specify a so-called reserve price and withhold allowances from the auction as needed to maintain that price. The efficiency advantage of a price floor would stem from the fact that it could prevent the cost of emission reductions from falling below the expected benefits or below the level of effort that policymakers intended.

A cap-and-trade program that included both a ceiling and a floor for allowance prices could achieve a long-term target for emissions while minimizing both the overall cost of achieving the target and price volatility. Under such a program, policymakers would specify annual emission targets as well as a ceiling and a floor for the price of allowances for each year. Regulators could adjust the levels of the price ceiling and floor periodically (for example, every five years) to ensure that emission reductions were on track for achieving the long-term target. For example, the rate at which the price floor or ceiling rose over time could be increased if regulators determined that the reductions in the previous five-year period were significantly lower than the amount needed to achieve the long-term target. Alternatively, policymakers could include provisions in a cap-and-trade program that would automatically trigger adjustments in the price ceiling and floor. For example, the rate at which the price ceiling and floor rose could be based on the percentage gap between anticipated and actual emissions in the previous five-year period.

Figures 2 and 3 illustrate the effects of price ceilings and floors. The figures present a simple example of an inflexible cap each year relative to a system involving price ceilings and floors. In Figure 2, the results illustrate what happens in 2018 if the costs of reducing emissions by roughly 15 percent are twice as high or 50 percent lower than expected. Under an inflexible cap, the emission reductions are unaffected. Under a price ceiling, fewer emission reductions are undertaken when costs are high; the result is lower economic costs that year but also less of a reduction in emissions. Under a price floor, more emission reductions are undertaken when costs are low.

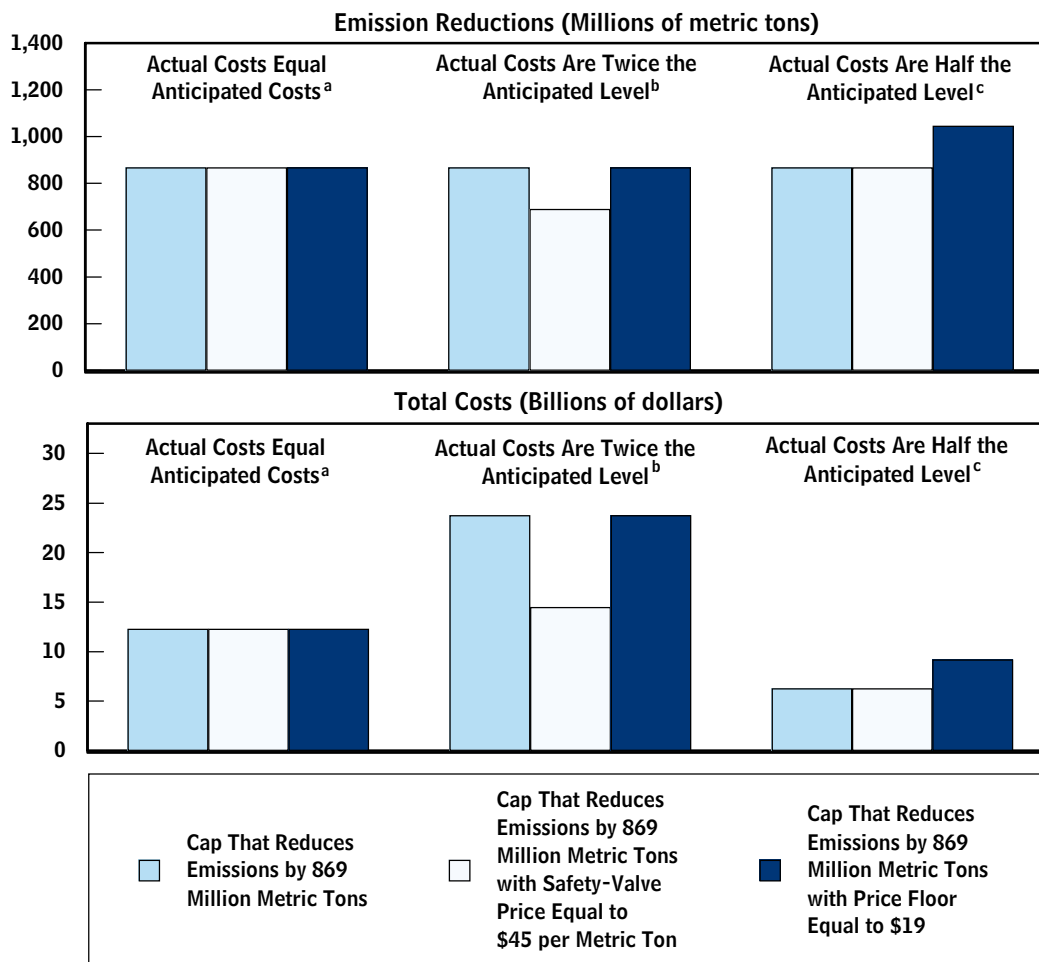
Figure 3 shows the results after one high-cost year and one low-cost year. Cumulative reductions of emissions are the same under the inflexible cap and the combined price ceiling-and-floor system, but costs are more than 20 percent lower under the latter approach. The reason, again, is that more of the emission reductions are undertaken in the low-cost year under that approach.

Borrowing and Banking Allowances. An alternative but generally somewhat less effective approach to reducing economic costs involves allowing companies to borrow future allowances in high-cost years, thereby deferring emission reductions to later years. Borrowing allowances from future years would tend to reduce allowance prices in the current year but then raise prices in the future (because borrowing would allow smaller reductions now but require greater reductions later). Firms would want to borrow allowances only if they expected the price of allowances in the future to be sufficiently below the current price as to make deferring reductions profitable. Most proposals would impose limits on borrowing, furthermore, in part because of concerns about enforcement and questions about who would be liable if the firm that borrowed future allowances was unable to pay them back (if it declared bankruptcy, for example).

Similarly, policymakers could help keep the price of allowances from falling too low by allowing companies to exceed their required emission reductions in low-cost years in order to bank allowances for use in future high-cost years. The additional emission

Figure 2.

Illustrative Comparison of Various Cap-and-Trade Policies to Reduce CO₂ Emissions by Roughly 15 Percent Under Different Cost Conditions in 2018



Source: Congressional Budget Office.

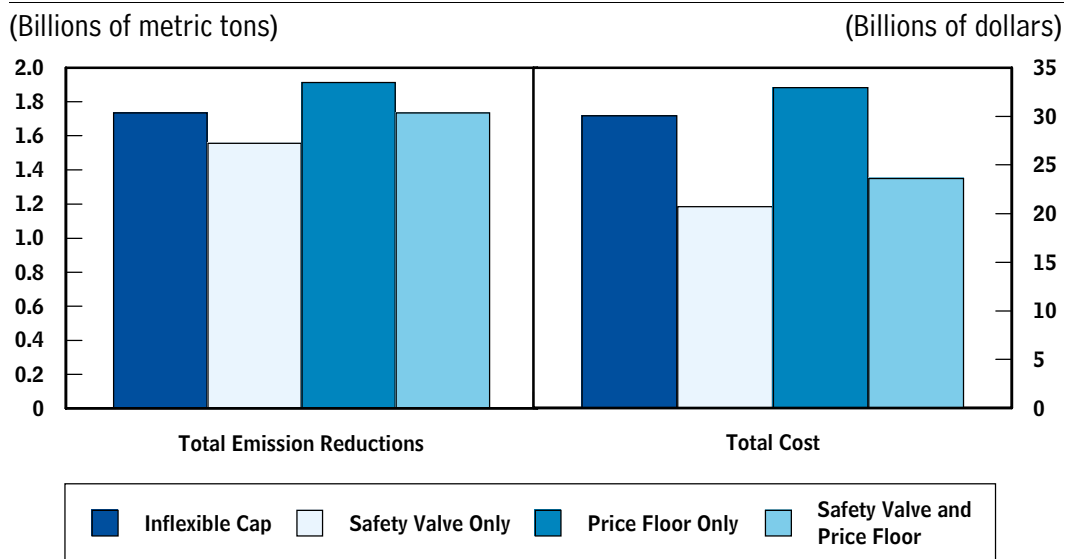
Notes: This example examines the emission reductions and total costs that would result in 2018, assuming that the policy covered only the United States. The cost of firms' emission reductions is derived from Mark Lasky, *The Economic Costs of Reducing Emissions of Greenhouse Gases: A Survey of Economic Models*, Congressional Budget Office Technical Paper No. 2003-03 (May 2003).

A safety valve is a ceiling on the price of emission allowances.

- Assumes that the actual marginal cost of reducing emissions by 869 million metric tons is \$45 per metric ton, the cost that policymakers anticipated when they set the cap.
- Assumes that the actual marginal cost of reducing emissions by 869 million tons is \$60 per metric ton but that the safety valve induces less reductions (691 million tons instead of 869 million), up to a marginal cost of \$45 per metric ton.
- Assumes that the actual marginal cost of reducing emissions by 869 million tons is \$15 per metric ton but that the price floor induces more reductions (1,047 million tons instead of 869 million) at a marginal cost of \$19 per metric ton.

Figure 3.

Illustrative Comparison of Total Emission Reductions and Total Costs and After One High-Cost and One Low-Cost Year



Source: Congressional Budget Office.

Notes: This example represents the cumulative emission reductions and costs over two years of a cap-and-trade policy that would reduce emissions of CO₂ by 869 million tons in each year (roughly a 15 percent reduction in 2018). The cost of firms' emission reductions is derived from Mark Lasky, *The Economic Costs of Reducing Emissions of Greenhouse Gases: A Survey of Economic Models*, Congressional Budget Office Technical Paper No. 2003-03 (May 2003).

A safety valve is a ceiling on the price of emission allowances.

For the high-cost year, CBO assumes that the marginal cost of reducing emissions by 869 million tons is \$60 per metric ton but that the safety valve induces less reductions (691 million tons instead of 869 million), up to a marginal cost of \$45 per metric ton.

For the low-cost year, CBO assumes that the marginal cost of reducing emissions by 869 million tons is \$15 per metric ton but that the price floor induces more reductions (1,047 million tons instead of 869 million) at a marginal cost of \$19 per metric ton.

reductions motivated by banking in low-cost years would put upward pressure on the price of allowances in those years.

S. 2191 and the Carbon Market Efficiency Board. S. 2191 would address sustained high prices for allowances by allowing an administrative board, the Carbon Market Efficiency Board, to transfer future allowances to the current year. That action could be viewed as a form of forced borrowing—that is, it would require firms to trade lower reductions today for higher reductions in the future, even if they would not have found it profitable to do so voluntarily. Such transfers could ultimately raise or lower the overall cost of meeting a long-run target depending on how the price of allowances changed over time. For example, if a low-cost, low-carbon energy technology became available in the future, transferring future allowances to the current period would have successfully shifted emission reductions to a time when the cost of

achieving them was lower. Alternatively, if policymakers borrowed future allowances with the expectation that such a technology would become available, but it did not, then the transfer could cause even more reductions to be made at a relatively high-cost time. (An alternative approach to the one embodied in S. 2191, which may be easier for regulators to implement efficiently, would be to have the board be the entity responsible for setting a ceiling and a floor for allowance prices and for adjusting those price limits periodically as needed to achieve a long-term target for reducing emissions.)

Using the Value of Allowances to Reduce Economic Costs

In establishing a cap-and-trade program, policymakers would create a new commodity: the right to emit CO₂. The emission allowances would have substantial value. For example, on April 10, 2008, CBO estimated that the value of the allowances created under S. 2191 (as order reported) would be roughly \$145 billion (in 2006 dollars) once the proposed program took effect in 2012; in subsequent years, the aggregate value of the allowances would be even greater. (See Box 1 for a short description of CBO's cost estimate for S. 2191.)

Options for Distributing Emission Allowances

Policymakers would need to decide how to allocate the allowances that corresponded to each year's CO₂ cap. One option would be to have the government capture their value by selling the allowances, as it does with licenses to use the electromagnetic spectrum. Another possibility would be to give the allowances to energy producers or some energy users at no charge. The European Union has used that second approach in its two-year-old cap-and-trade program for CO₂ emissions, and in the United States, the federal government has distributed nearly all of the allowances issued under the 13-year-old U.S. cap-and-trade program for sulfur dioxide emissions (which contribute to acid rain) that way.

Selling the allowances would provide lawmakers with an opportunity to reduce the overall economic impact of a CO₂ cap. For instance, the government could use the revenue from auctioning allowances to reduce existing taxes that tend to dampen economic activity—primarily, taxes on labor, capital, or personal income. As research indicates, a CO₂ cap would exacerbate the economic effects of such taxes: The higher prices caused by the cap would lower real (inflation-adjusted) wages and real returns on capital, which would be equivalent to raising marginal tax rates on those sources of income. Using the value of the allowances to reduce such taxes could help mitigate that adverse effect of the cap. Alternatively, policymakers could choose to use the revenue from auctioning allowances to reduce the federal deficit. If doing so lessened the need for future tax increases, the end result could be similar to dedicating the revenue to cuts in existing taxes.

The decision about whether or not to sell the allowances and how to use the proceeds could have a significant impact on the overall cost. For example, researchers have

Box 1.**CBO's Cost Estimate for S. 2191**

On April 10, 2008, the Congressional Budget Office (CBO) issued a cost estimate for S. 2191, the America's Climate Security Act of 2007, as ordered reported by the Senate Committee on Environment and Public Works in December 2007. CBO also issued a cost estimate for a slightly amended version of the legislation that was transmitted by the committee on April 9.

The legislation would create a cap-and-trade system for carbon dioxide and other greenhouse gases. (The bill actually calls for two separate cap-and-trade programs—a bigger one covering most types of greenhouse gases and a smaller one covering hydrofluorocarbons.) Some of the emission allowances would be auctioned—through a new entity, the Climate Change Credit Corporation; the remaining allowances would be distributed at no charge to states and other recipients. Over the roughly 40 years that the proposed cap-and-trade programs would be in effect, the number of allowances—and thus the emissions of relevant gases—would be reduced each year.

On the basis of an analysis of the results of several economic models, CBO estimates that if the legislation was enacted, the auction price of emission allowances for those gases would rise from about \$23 per metric ton of carbon-dioxide-equivalent (mt CO₂e) emissions in 2009 to about \$44 per mt CO₂e in 2018.¹ (In 2006 dollars, the auction price per metric ton of

1. A carbon dioxide equivalent is defined for each greenhouse gas as the quantity of that gas that makes the same contribution to global warming as one metric ton of carbon dioxide, as determined by the Environmental Protection Agency.

estimated that the efficiency cost of a 15 percent cut in emissions could be reduced by more than half if the government sold allowances and used the revenue to lower corporate income taxes, rather than devoting the revenue to providing lump-sum rebates to households or giving the allowances away (see the top panel of Figure 4).⁴

4. The efficiency cost of a policy reflects the economic losses that occur because prices are distorted so that they do not reflect the nonenvironmental resources used in their production. That cost includes decreases in the productive use of labor and capital as well as costs (both monetary and nonmonetary) associated with reducing emissions. To provide perspective on the magnitude of such efficiency costs, they are depicted as a share of gross domestic product.

CBO's Cost Estimate for S. 2191

CO₂e would rise from about \$21 in 2009 to \$35 in 2018.) Measured relative to base-case emissions (that is, those that would occur under current law), emissions of the main greenhouse gases covered by the programs would decline by 7 percent in 2012 and by 17 percent in 2018; over the 2012–2050 period, emissions would decline by a total of 42 percent relative to the base case.

Enacting S. 2191 as it was ordered reported would increase revenues by about \$1.19 trillion over the 2009–2018 period, CBO estimates. Direct spending from distributing those proceeds would total about \$1.21 trillion over the period. The net effect of the original legislation (as ordered reported) would be to increase the deficit (excluding any effects on future discretionary spending) by an estimated \$15 billion over the next 10 years. The effect of the amended version, in contrast, would be to reduce the deficit (again excluding any effects on future discretionary spending) by roughly \$80 billion over the same period. In addition, if policymakers appropriated the amounts necessary to implement S. 2191, discretionary spending would increase over the 2009–2018 period, CBO estimates, by about \$4 billion under the original legislation and by about \$80 billion under the amended version.

The cost estimates for the two versions of the bill differ because the amendment would increase the proportion of allowances that were auctioned, deposit some of the auction proceeds in a Climate Change Deficit Reduction Fund, and make spending from that fund subject to appropriation.

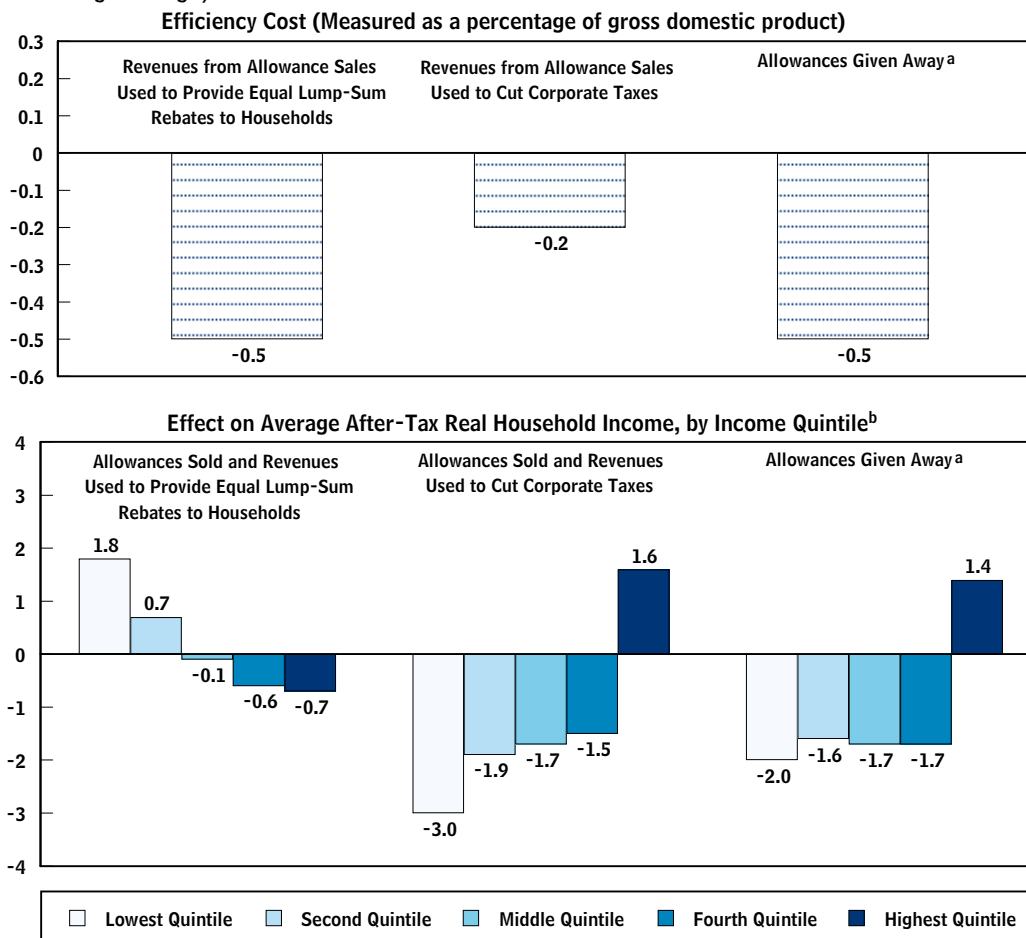
The Distributional Consequences of Different Approaches

The ways in which lawmakers allocated the revenue from selling emission allowances would affect not only the total economic cost of a cap-and-trade policy but also its distributional consequences. The ultimate distributional impact of a cap-and-trade program would be the net effect of two distinct components: the distribution of the costs of the program (including the cost of paying for the allowances) and the distribution of the allowances' value. (Because someone would pay for them, someone would benefit from their value.) Market forces would determine who bore the costs of a cap-and-trade program, but policymakers would determine who received the value of the allowances. The ultimate effect could be either progressive or regressive, imposing disproportionately large burdens on high-income or low-income households, respectively.

Figure 4.

Effects of a 15 Percent Cut in CO₂ Emissions, with the Allowances' Value Used in Various Ways

(Percentage change)



Sources: Terry M. Dinan and Diane Lim Rogers (top panel), "Distributional Effects of Carbon Allowance Trading: How Government Decisions Determine Winners and Losers," *National Tax Journal*, vol. 55, no. 2 (June 2002) and Congressional Budget Office (bottom panel).

Notes: These figures do not reflect any of the benefits from reducing climate change.

The policy examined here is a cap-and-trade program designed to reduce carbon dioxide (CO₂) emissions by 15 percent from 1998 levels. (CBO performed the analysis in 2000 and used 1998 emission levels so the distributional effects could be based on actual, rather than projected, data on consumer spending and taxes.) In the bottom panel, the costs of the cap-and-trade policy are shown as decreases in real household income, measured as a percentage of after-tax income before the policy change. Those numbers reflect data on each quintile's cash consumption and estimates of cash income. (A quintile contains one-fifth of U.S. households arrayed by income.) Because of data limitations, those numbers should be viewed as illustrative and broadly supportive of the conclusions in this analysis rather than as precise estimates.

- These estimates assume that the government would use any positive net revenue remaining after accounting for ways in which the policy affected the federal budget to provide equal lump-sum rebates to households. The results would be more regressive if the government used any positive net revenue to decrease corporate taxes or payroll taxes.
- Indicates the net effect of households' increased expenditures because of cap-induced price increases and the income that households would receive as a result of the allowance-allocation strategy.

Market Forces Would Determine Who Bore the Costs of a Cap. Obtaining allowances—or taking steps to cut emissions to avoid the need for such allowances—would become a cost of doing business for firms that were subject to the CO₂ cap. However, those firms would not ultimately bear most of the cost of the allowances. Instead, they would pass the cost along to their customers (and their customers' customers) in the form of higher prices. By attaching a cost to CO₂ emissions, a cap-and-trade program would thus lead to price increases for energy and energy-intensive goods and services. Such price increases would stem from the restriction on emissions and would occur regardless of whether the government sold emission allowances or gave them away. Indeed, the price increases would be essential to the success of a cap-and-trade program because they would be the most important mechanism through which businesses and households were encouraged to make investments and change their behavior to reduce CO₂ emissions. (In regulated electricity industries, distributing the permits at no cost might mitigate or prevent price increases in those markets but only at the cost of requiring even larger price increases in other markets. Ultimately, consumers will, in one way or another, bear costs roughly equal to the value of the permits.)

The rise in prices for energy and energy-intensive goods and services would impose a larger burden, relative to income, on low-income households than on high-income households. For example, without incorporating any benefits to households from lessening climate change, CBO estimated that the price increases resulting from a 15 percent cut in CO₂ emissions would cost the average household in the lowest one-fifth (quintile) of all households arrayed by income slightly more than 3 percent of its income; such increases would cost the average household in the top quintile just under 2 percent of its income (see Table 1).⁵

The higher prices that resulted from a cap on CO₂ emissions would reduce demand for energy and energy-intensive goods and services and thus create losses for some current investors and workers in the sectors of the economy supplying such products. Investors might see the value of their stock decline, and workers could face the risk of unemployment as jobs in those sectors were cut. Stock losses would tend to be widely dispersed among investors, because shareholders typically diversify their portfolios. In contrast, the costs borne by workers would probably be concentrated among relatively few households and, by extension, communities.

Policymakers Would Determine Who Received the Value of the Allowances. Although the price increases triggered by a cap-and-trade program for CO₂ emissions would be regressive, the program's ultimate distributional effects would depend on policymakers' decisions about how to allocate the allowances. As noted above, those allowances would be worth tens or hundreds of billions of dollars per year. Who received that value would depend on how the allowances were distributed.

5. Those numbers are based on an analysis that CBO conducted using 1998 data; see Congressional Budget Office, *Who Gains and Who Pays Under Carbon-Allowance Trading? The Distributional Effects of Alternative Policy Designs* (June 2000). CBO is in the process of updating those figures, using recent data on households' expenditures and income.

Table 1.

Effects on U.S. Households of the Higher Prices Resulting from a 15 Percent Cut in CO₂ Emissions

	Average for Income Quintile				
	Lowest	Second	Middle	Fourth	Highest
Annual Cost Increase in 2006 Dollars	680	880	1,160	1,500	2,180
Annual Cost Increase as a Percentage of Income ^a	3.3	2.9	2.8	2.7	1.7

Source: Congressional Budget Office, *Who Gains and Who Pays Under Carbon-Allowance Trading? The Distributional Effects of Alternative Policy Designs* (June 2000).

Notes: These numbers do not reflect any of the benefits from reducing climate change.

The policy examined here is a cap-and-trade program designed to lower U.S. carbon dioxide (CO₂) emissions by 15 percent from 1998 levels. (CBO performed the analysis in 2000 and used 1998 emission levels so that the distributional effects could be based on actual, rather than projected, data on consumer spending and taxes.) CBO assumed that the full cost of cutting emissions would be passed along to consumers in the form of higher prices and that the price increase for a given product would be proportional to the amount of CO₂ emitted from the fossil fuels used in its production.

These numbers reflect data on each quintile's cash consumption and estimates of cash income. (A quintile contains one-fifth of U.S. households arrayed by income.) Because of data limitations, the numbers should be viewed as illustrative and broadly supportive of the conclusions in this analysis rather than as precise estimates.

a. The cost increases are equivalent to percentage declines in households' after-tax income.

Lawmakers could more than offset the price increases experienced by low-income households or the costs imposed on workers in particular industrial sectors by providing for the sale of some or all of the allowances and using the revenue to pay compensation. For example, when CBO examined the ultimate distributional effects of a cap-and-trade program that would reduce CO₂ emissions in the United States by 15 percent, it concluded that lower-income households could be better off (even without any benefits from reducing climate change considered) as a result of the policy if the government chose to sell the allowances and use the revenue to pay an equal lump-sum rebate to every household in the United States. In that case, the size of the rebate would be larger than the average increase in low-income households' spending on energy and energy-intensive goods.⁶ Such a strategy would, on net, increase average income for households in the lowest income quintile by about 2 percent (see the bottom panel of Figure 4). At the same time, the net average income for households in the top quintile would fall by less than 1 percent, CBO estimated.

6. One researcher has suggested that an environmental tax credit based on earnings could offer another means of reducing the regressive effects of the price increases that would result from a tax or cap on CO₂ emissions. See Gilbert E. Metcalf, *A Proposal for a U.S. Carbon Tax Swap* (Washington, D.C.: Brookings Institution, October 2007).

In contrast, if lawmakers chose to use the allowances to decrease corporate income taxes, the overall cost to the economy would fall but the distributional effects would be significantly more regressive than the initial price increases. Because low-income households pay relatively little in corporate taxes, the cut in corporate tax rates would not offset their increased spending on energy and energy-intensive goods. Households in the top income quintile, however, would experience an increase in after-tax income as a result of the policy. Should policymakers decide to use the revenue from selling allowances to decrease payroll taxes, the effects (not shown in the figure) would be regressive as well, although less so than for a cut in corporate taxes.⁷

Giving all or most of the allowances to energy producers to offset the potential losses of investors in those industries—as was done in the cap-and-trade program for sulfur dioxide emissions—would also exacerbate the regressivity of the price increases. On average, the value of the CO₂ allowances that producers received would more than compensate them for any decline in profits caused by a drop in demand for energy and energy-intensive goods and services. As a result, the companies that received allowances could experience windfall profits.

For example, in 2000, CBO estimated that if emissions were reduced by 15 percent, as in the scenario discussed above, and all of the allowances were distributed free of charge to producers in the oil, natural gas, and coal sectors, the value of the allowances would be 10 times as large as the producers' combined profits in 1998. Profits for those industries have climbed substantially since then, yet the value of the allowances associated with the policy that CBO analyzed would still be large relative to those producers' profits.⁸ Because the additional profits from the allowances' value would not depend on how much a company produced, such profits would be unlikely to prevent the declines in production and resulting job losses that the price increases (and resulting drop in demand) would engender.

In addition, those profits would accrue to shareholders, who typically are from higher-income households, and would more than offset those households' increased spending on energy and energy-intensive goods and services. Low-income households, by contrast, would benefit little if allowances were given to energy producers for free, and they would still bear a disproportionate burden from the price increases that would nonetheless occur. Thus, giving away allowances would be significantly regressive, making higher-income households better off as a result of the cap-and-trade policy while making lower-income households worse off. Further, giving away the allowances would preclude the government from dedicating the value of the allowances to reducing the overall economic impact of the policy.

7. For those results, see Congressional Budget Office, *Trade-Offs in Allocating Allowances for CO₂ Emissions* (April 25, 2007).

8. Specifically, CBO estimated that the value in 1998 of the allowances stemming from the 15 percent reduction in U.S. emissions would total \$155 billion (in 2006 dollars). By comparison, profits for U.S. producers of oil, natural gas, and coal totaled \$13.5 billion in 1998 (in 2006 dollars). Those companies' total profits have grown substantially—for example, in 2006, they totaled \$174 billion.

