ECONOMIC AND BUDGET ISSUE BRIEF

CBO

A series of issue summaries from the Congressional Budget Office

NOVEMBER 25, 2003

Issues in the Design of a Cap-and-Trade Program for Carbon Emissions

If the United States acts to limit climate change, a key step may be reducing the use of fossil fuels (oil, natural gas, and coal) and thus decreasing the amount of carbon dioxide that human activities release into the atmosphere. Even policymakers who think such action is inevitable are far from agreeing about the timing or the size of cuts in carbon emissions because of uncertain science and complex economic considerations. However, there is a growing consensus that if cuts are to be made, economic incentives are a better method than direct government controls. Such controls might assign specific emission rates to individual factories, vehicles, and other emitters or require the use of particular equipment or production processes. By contrast, incentive-based approaches—such as taxes or the "cap-and-trade" program discussed below—would discourage emissions by raising their price. As a result, those approaches would generally be more cost-effective than direct controls.

Under a cap-and-trade program, policymakers would set a limit (the "cap") on total carbon emissions during some period—presumably at a level below what would otherwise occur—and require certain entities to hold rights, or allowances, to the emissions permitted under that cap. Each allowance would entitle the holder to emit some amount of carbon, such as a metric ton. After the allowances were initially distributed, entities would be free to buy and sell them (the "trade" part of the program).

No effort to reduce carbon emissions would be free. Fossil fuels are key inputs in the U.S. economy, so restricting their use would reduce incomes. The appeal of a cap-and-trade program is that, if properly designed, it could motivate firms and households to cut emissions in ways that would minimize the cost of achieving the reductions. Because entities could buy and sell allowances, those that could take advantage of the lowest-cost opportunities to reduce emissions would make the deepest cuts. (They would then sell their excess allowances to entities for whom reducing emissions would cost more.) A cap-and-trade program could also be designed to include the sequestering of carbon—for example, by planting trees, which absorb carbon dioxide from the atmosphere. In that case, regulated entities could have the option of fulfilling part of their allowance requirement by paying other entities (such as farmers) to sequester carbon.

Governments have successfully used cap-and-trade programs to reduce emissions of other air pollutants, such as lead, nitrogen oxides, and sulfur dioxide (an ingredient of acid rain). In the case of sulfur dioxide, researchers estimate that the cap-and-trade program lowered the cost of meeting the emissions target by between 43 percent and 55 percent compared with the cost of requiring all regulated sources of sulfur dioxide to meet a uniform emission rate. Substantial cost savings would also be likely to occur under a cap-and-trade program for carbon dioxide.

Various bills that would limit carbon emissions have been introduced in the 108th Congress. For example, the Climate Stewardship Act of 2003 (S. 139) would cap carbon emissions from several sectors of the economy. The Clean Air Planning Act of 2003 (S. 843) and the Clean Smokestacks Act of 2003 (H.R. 2042) would cap carbon emissions (and other pollutants) from the electricity-generating sector.

Choices about the design of a cap-and-trade program would determine whether reductions in carbon emissions occurred throughout the economy or only in certain sectors. Those choices would also affect the cost of achieving a given cut in emissions as well as the distribution of that cost among shareholders, workers, and consumers. Finally, those choices would determine whether the emissions cap would be met with certainty—but at an uncertain cost—or whether the cost would be capped but the actual level of reductions would remain uncertain.

Who Should Hold Allowances? The "Upstream" Versus "Downstream" Debate

A key decision in designing a cap-and-trade program is whether to implement it "upstream," where carbon enters the economy (when fossil fuels are imported or produced domestically), or farther "downstream," closer to the point where fossil fuels are combusted and the carbon enters the atmosphere. An upstream program would require producers and importers of fossil fuels to hold allowances for the fuel they sold. Their allowance requirements would be based on the carbon emissions released when their fuel was combusted. A downstream program would require users of fossil fuels to hold allowances.

An Upstream Allowance Requirement

An upstream design offers significant advantages in efficiency and could have modest implementation costs. Although carbon is ultimately emitted by hundreds of millions of fossil-fuel users—including vehicles, buildings, and factories—it enters the economy through a relatively small number of fossil-fuel suppliers. By placing the allowance requirement upstream on those suppliers, policymakers could cap virtually all fossil-fuel-based carbon emissions in the United States and could ensure that those emissions were cut at the lowest cost.

An upstream cap would limit fossil-fuel production, leading to higher prices for those fuels as well as for goods and services that they are used to produce (such as gasoline and electricity). The higher prices would give firms and households throughout the nation an incentive to use less fossil fuel—for example, by installing more insulation, using energy-efficient lighting, or driving less—and thus reduce carbon emissions. In addition, higher fossil-fuel prices would encourage the development of new technologies for improving energy efficiency. Finally, an upstream cap would keep the government's administrative costs and the private sector's reporting costs low because only a relatively small number of firms would be regulated.

A Downstream Allowance Requirement

Depending on the design, moving the allowance requirement downstream to fossil-fuel users would raise either the cost of achieving a given reduction in emissions or the cost and difficulty of implementation. A downstream program that was restricted to a subset of emitters—such as electricity generators—would be feasible to implement. But it would not cut carbon emissions at the lowest cost because it would not provide an incentive for low-cost reductions in sectors not covered by the cap. (For example, it would not give motorists an incentive to drive less.) In addition, a cap that covered only part of the economy would not ensure a limit on U.S. carbon emissions overall, because emissions from uncapped sectors could grow.

Alternatively, imposing an allowance requirement on the vast number of vehicles, factories, and electricity generators that burn fossil fuels could create an economywide incentive for low-cost emission reductions. It would also provide incentives for carbon emissions to be captured from smokestacks and tailpipes; however, technologies to do that are not yet economically viable. Moreover, such a system could entail prohibitively high administrative costs for the government and reporting costs for the private sector.

Costs Would Be Felt Upstream and Downstream Under Either Placement of Allowances

The costs of the cap would not be felt only at the point where allowances were required. Regardless of whether the cap was placed upstream or downstream, costs would be distributed throughout the supply chain for carbon. For example, placing the cap on electricity generators would impose costs on coal suppliers (by reducing the demand for coal, which is highly carbon-intensive, and lowering its price) and on electricity consumers (by increasing the cost of producing electricity and thus its price). Likewise, placing the cap on fossil-fuel suppliers would raise costs for electricity producers (by increasing the price of fossil fuels and thus the cost of generating electricity) and for electricity consumers (who would face higher electricity prices resulting from higher production costs).

How Should Allowances Be Allocated Initially? The Free Versus Auction Debate

The decision of how to allocate the allowances permitted under the cap—by selling them or giving them away—could have important implications for both the overall cost of the policy and the distribution of that cost among U.S. households. The importance of that decision stems from the fact that carbon allowances could be worth billions —perhaps even hundreds of billions—of dollars each year.

Selling Allowances in an Auction

Proposals that involve selling carbon allowances generally envision auctioning them off, as the government has done with licenses to use the electromagnetic spectrum. Such auctions could raise a substantial amount of revenue that policymakers could use for various purposes.

Policymakers could lower the overall cost of a cap-and-trade program if they used the auction revenue to offset an unintended effect of the program. That effect—the "tax-interaction effect"—would occur because the price increases resulting from a cap-and-trade program would tend to lower the real (inflation-adjusted) income that people received from working and investing. (In other words, the dollars they earned would buy them less.) Thus, the higher prices would discourage people from pursuing those activities, compounding the fact that existing taxes on labor and capital already discourage work and investment.

That unintended effect may sound minor, but researchers have concluded that its cost could be substantial. Policymakers could offset at least part of that cost by using the revenue from an allowance auction to cut existing marginal taxes on income from capital and labor. Research suggests that using auction revenue to lower those taxes could reduce the economywide cost of a cap-and-trade program by more than 30 percent.

Alternatively, policymakers could use auction revenue to compensate shareholders, consumers, or workers for the costs they would incur because of the cap-and-trade program. Most of the costs of a limit on carbon emissions—perhaps 80 percent or more—would be passed on to consumers through higher prices. The share not passed on to consumers would be borne by investors and workers in industries that supply fossil fuels or use them intensively (such as the electricity-generating sector).

The costs borne by consumers would most likely be widely disbursed among U.S. households. However, those costs would impose a relatively greater burden on lower-income households than on higher-income ones, both because lower-income households consume a larger share of their income and because energy products make up a bigger fraction of their expenditures. The costs borne by shareholders would also be widely disbursed (assuming that stock in affected companies is broadly distributed because shareholders have diversified portfolios). The costs borne by workers, in contrast, would probably be concentrated on relatively few households, imposing particular economic hardships on them and their communities.

Allocating Allowances for Free

Giving allowances away—as occurred in the cap-and-trade program for sulfur dioxide—is another method by which lawmakers could provide compensation for policy-induced costs. Generally, free allocations are seen as a way to compensate producers rather than consumers or workers. Producers would have to receive only a modest portion of the allowances to offset their costs from a cap on carbon emissions, because they would be expected to pass a large share of those costs on to consumers. Thus, a decision to give all of the allowances to selected firms (such as fossil-fuel suppliers or electricity generators) would more than compensate them for their costs and could provide them with substantial profits. Those profits would ultimately benefit shareholders rather than consumers in general.

If policymakers decided to give some of the allowances to producers, they would minimize the cost of emission reductions by basing those allocations on historical production or emission levels (a process called grandfathering) rather than on current levels. Tying allocations to current levels of production would create greater incentives for some emission-cutting strategies than for others and would not result in the least costly mix of strategies.

Finally, recent research shows that compensating electricity generators by giving them free allowances—even with grandfathering—could raise the cost of reducing carbon emissions. The reason is that in many parts of the country, electricity production is still dominated by regulated generators. If those generators received allowances for free, regulators would probably not let them reflect the value of the allowances—that is, the amount they would receive if they sold the allowances—in the price of electricity. (Researchers found that to be the case for sulphur dioxide allowances.) Consequently, the rise in electricity prices that would result from free allocation would be lower than the rise that would result from auctioning allowances. Those lower price increases would do less to decrease existing distortions in the pricing of electricity (where the price is often less than the marginal cost of generation) and would not give consumers as much incentive to reduce their electricity use.

Research indicates that the cost of a cap-and-trade program could be more than twice as high if generators were given allowances instead of having to buy them in an auction.

The Bottom Line: Reducing Overall Policy Costs Versus Providing Compensation

In essence, with a cap-and-trade program, policymakers would face a trade-off between using the allowances' value to offset the distributional impact of the cap and using it to reduce the overall cost to the economy. Even if policymakers decided to use the entire value of the allowances to provide compensation, that value would be insufficient to fully compensate all investors, consumers, and workers for their share of the program's costs. Thus, policymakers would have to weigh competing objectives when deciding on the appropriate combination of uses for that value.

How Should the Cap Be Designed? Fixed Caps, Declining Caps, Safety Valves, and Circuit Breakers

The cap-and-trade program for sulfur dioxide emissions set a fixed cap on those emissions, which electricity generators had to meet regardless of the costs to them, coal producers, or electricity customers. If the government set a fixed cap for carbon emissions, the result would be the same: the cap would be met (barring cheating) regardless of the costs that it imposed on shareholders, workers, or consumers. Some policymakers favor the idea of a fixed cap because it sets a maximum level of emissions.

A cap could also be designed to phase in gradually, giving regulated entities time to adjust to the ultimate level of the cap. Like a fixed cap, a declining cap would establish a maximum level of emissions, but the more time the economy had to adapt to the final cap, the less it would cost.

Neither a fixed nor a declining cap would place an upper bound on the potential cost of complying with the program. Policymakers could limit the per-unit cost of reducing emissions to a certain amount (say, $25 per ton) by agreeing to supply an unlimited quantity of allowances at that price, so that competition did not increase the price. If policymakers set such a "safety-valve price," the total number of allowances (and, hence, the level of emissions) could exceed the initial number that firms were either given or required to purchase.

Many analysts favor a safety-valve approach because it would help prevent the U.S. economy from incurring higher-than-expected costs for emission reductions. Predicting the costs of such reductions is difficult, and those costs could rise steeply as deeper cuts were made. (Initial reductions would be less expensive than later ones because low-cost options for cutting emissions would be used first, leaving more expensive options for further decreases.) A safety-valve price would address that uncertainty by setting an upper limit on the cost of emission reductions.

Some proposals for a cap-and-trade program envision setting a "circuit-breaker" cap that would depend on the price of allowances. Under such a policy, the cap would decline gradually as long as the price of allowances stayed below a predetermined trigger price. If the allowance price increased to the level of the trigger, the cap would be frozen. Once that happened, the price of allowances could remain above the trigger price—which would be likely to occur, since the frozen cap would become harder to meet over time as the economy grew. Unlike a safety-valve price, a circuit-breaker policy would not set an upper limit on the cost of emission reductions. If the price of allowances once again fell below the trigger price (because of the introduction of new technologies, for example), the cap would again start to decline.

Regardless of how a cap on carbon emissions was designed, it would be beneficial if policymakers were able to adjust it in the face of new information. Scientific understanding of the effects of man-made emissions on the Earth's climate—and of the potential harm that might result from climate change—is continuing to evolve. Therefore, it could be important to preserve the flexibility to raise the cap (by giving away or selling more allowances) or to lower it (by buying back or requiring firms to surrender allowances that had been distributed).

Related CBO Publications: This brief is based on *An Evaluation of Cap-and-Trade Programs for Reducing U.S. Carbon Emissions* (June 2001) and *Shifting the Cost Burden of a Carbon Cap-and-Trade Program* (July 2003). Those and the following related CBO publications are available at the agency's Web site (www. cbo.gov): *The Economics of Climate Change: A Primer* (April 2003) and *Who Gains and Who Pays Under Carbon-Allowance Trading?* (June 2000).