ISSUE BRIEF 6

ALLOWANCE ALLOCATION

RAYMOND J. KOPP

6

ALLOWANCE ALLOCATION

Raymond J. Kopp

SUMMARY

This issue brief provides an overview of concepts and policy decisions related to the allocation of emissions allowances or permits under a cap-and-trade program for limiting greenhouse gas (GHG) emissions. Allocation decisions distribute the wealth embodied in emissions allowances and therefore have economic impacts that can affect the net cost of the program to individual stakeholders and to society as a whole. Allocation decisions do not, however, affect the environmental performance of the program—that is, they do not change the overall level of emissions reductions achieved by the policy.

- Allowances associated with a cap-and-trade system represent an asset with potentially considerable monetary value, perhaps $100 billion or more annually. The value of these allowances or permits is not a measure of the cost associated with meeting the cap, but rather a wealth transfer from those who pay higher energy or emissions prices under the cap-and-trade program to those who hold allowances.

- While the U.S. Acid Rain program allocated sulfur dioxide (SO_2) allowances gratis (for free) to regulated entities (in that case, electric utilities), cap-and-trade systems need not adopt this approach. Permits can be allocated gratis to entities other than those that are directly regulated under the program (including, for example, households or state government). Moreover, allowances need not be allocated gratis: they can be sold by the government, which can retain resulting revenues for other purposes.

- Allocation decisions can affect both the efficiency (overall cost of meeting the cap) and the equity (distribution of the cost) of a cap-and-trade program. Generally, auctioning allowances and using the revenues to lower taxes, or offset particularly distorting taxes, increases efficiency and lowers the overall cost of the program to society. Awarding free allowances to certain stakeholders can address distributional concerns, but can sacrifice some efficiency.

- Allocation can alter economic incentives and the behavior of firms. For example, an output-based, updating approach could award free allowances to firms on the basis of output. For example, free allowances could be distributed to firms within the electric-power sector on the basis of their share of total sector-wide electricity output. Because this approach rewards firms for producing a larger share of output, free allowances will act as an output subsidy, effectively incentivizing firms to produce more. This outcome may or may not be desirable depending on the sector and the policy goals being pursued.

- Allocation to new entrants and retiring sources can be dealt with in a number of ways. However, care must be taken to ensure that the allocation methods used do not alter forward incentives for investment and retirement in ways that may

not be immediately obvious but that lead to suboptimal technology choices (either in terms of encouraging new investments in carbon-intensive technologies or delaying the retirement of uneconomic facilities).

- Arguments for free allocation are typically rooted in equity concerns: the desire to compensate sectors or regions that will otherwise bear a disproportionate share of the cost of regulation, or to blunt immediate impacts on the competitiveness of U.S. firms. As the economy adjusts to GHG constraints over time, these arguments become less compelling while the potential for economic distortions as a result of free allocation tends to grow, making it prudent to phase out free allocation in favor of auctioning allowances.

Overview of Discussion

While many important design features must be addressed in setting up a cap-and-trade system for greenhouse gas emissions, allocation has emerged as a critical challenge in the policy debate. This is unquestionably due to the enormity of the financial assets at stake: under current proposals, tens of billions of dollars per year—perhaps $100 billion dollars or more per year—could be divided up and given away under an emissions trading program. While allocation decisions are first and foremost distributional decisions (who gets what), two key economic concerns are relevant: (1) the risk of unintended consequences from tying allocations to some change in behavior, and (2) using allocation to mitigate costs imposed on particularly vulnerable sectors, households, or regions.

Cap-and-Trade Systems Change Prices and Create Wealth and Obligations

Cap-and-trade systems simultaneously change prices and create assets and liabilities. Entities that are directly regulated under the cap—including producers and processors of fossil fuels in an upstream system—face new liabilities in the form of the obligation to surrender allowances. Matching those liabilities are the new assets created in the form of emissions allowances. These allowances can be given to entities at no charge (whether those entities are directly regulated or not) or held by the government and auctioned. Energy prices downstream of regulated entities will typically adjust to reflect the opportunity cost of surrendering associated allowances, which in turn is a function of carbon dioxide (CO_2) content. Importantly, however, the method by which allowances are allocated will have no impact on the performance of the cap-and-trade system in terms of its ability to achieve targeted emissions reductions.

The wealth embodied in allowances can be substantial. If an economywide cap-and-trade program were instituted in the United States and allowance prices were in the range of $10 per ton of CO_2-equivalent (CO_2e), the total value of allowances circulating under the program would be approximately $50 billion dollars annually. At higher prices on the order of $25/ton CO_2e (akin to expected prices on the European Union CO_2 market for 2008–2012), the value of allowances would be more than $100 billion dollars annually, or slightly less than 1 percent of U.S. GDP.

The value of all allowances is *not* a measure of the economic cost of the regulatory program. Rather, allowance value reflects a *transfer* from those paying higher energy or emissions costs as a result of the cap-and-trade program to whatever entities initially receive the allowances (note that the receiving "entity" can be U.S. taxpayers, if allowances are auctioned to raise money for the federal treasury). What, then, is the cost of the regulatory program itself? It is the sum of the cost associated with each ton that has to be reduced to meet the emissions cap. In turn, the price of allowances depends on the cost of the marginal—or last, most expensive—ton reduced. A quick numerical example may be helpful here: suppose the economy generates ten tons of emissions before we impose a cap of seven tons. The three tons that must be reduced cost $1, $5, and $10, respectively. Here the *cost* of the program is $16 ($1 + $5 + $10). The *marginal* cost of the last, most expensive ton is $10; this sets the market price of allowances in our cap-and-trade program. Finally, the total value of the seven allowances will be $70: 7 tons x $10/ton. There is generally no simple relationship between program costs and the value of the allowances, though for the CO_2 policies currently under consideration in the U.S. Congress, costs are significantly smaller than the value of the allowances.

Allowance Allocation Options

Allowance allocation can affect two important economic dimensions of a cap-and-trade program: efficiency and equity. Efficiency refers to the overall economic cost of meeting the emissions cap, while equity refers the distribution of that cost across all sectors and households in the economy. Generally, pursuing equity objectives means sacrificing some efficiency. Several approaches can be used to determine the initial allocation of allowances under an emissions trading program.

Allowances can be given for free to entities that are especially affected by the policy—whether those entities are directly regulated (that is, required to surrender allowances) or not. The entities most burdened by the trading program will be those that are least able to pass associated costs—either the direct cost of surrendering allowances or the higher cost of energy under a system that regulates emissions upstream—through to their customers. These issues of cost "pass-through" are discussed further in Issue Briefs #7 and #8, which examine concerns about competitiveness, and in Issue Brief #11, which addresses cost and allocation issues specific to the electricity sector.

Allowances can be distributed to individual entities on the basis of past or current behavior. Alternatively, allowances may be simply auctioned and the revenue retained (and ultimately re-distributed) by the government. Any combination of the above methods can be employed.

In the case of the national SO_2 trading program established under the acid rain provisions of the Clean Air Act, the vast majority of allowances were given for free to those entities with emissions that were regulated under the cap. This same model was used in the eastern states' nitrogen oxides (NOx) trading program and in the first phase of the European Union Emissions Trading Scheme (EU ETS). Nonetheless, there is no economic reason why the question of how allowances should be allocated cannot be separated from the question of how compliance obligations should be assigned—that is, there is no reason why allowances cannot or should not be provided to entities other than those directly regulated under an emissions trading program. In fact, where most of the costs of compliance are passed through to entities that are *not* directly regulated (typically in the form of higher energy prices), equity considerations may argue for an allocation focused on compensating downstream energy users.

In that case, recipients of free allowances would sell those allowances to entities that do face a direct compliance obligation. An emissions allowance can be thought of as just another input—like capital or labor—that the regulated entity needs to produce its intended product. Regardless of how allocation occurs, the allowances must eventually find their way into the hands of the regulated entities.

In the simplest case, the government may give allowances free of charge (gratis) to regulated or unregulated entities, or sell allowances to the regulated entities. To date, most existing trading programs—including the U.S. SO_2 and NOx programs as well as the EU ETS—have allocated most allowances for free to regulated entities. This gratis allocation transferred the wealth represented by the permits from the government to regulated entities, thereby affecting the equity of the program. Yet economists regularly point out that selling allowances and using the revenue to cut other taxes (or avoid tax increases) can substantially lower overall program costs. Thus, in the case of the U.S. SO_2 and NOx programs as well as the EU ETS, concerns about compensating regulated industry appear to have trumped efficiency considerations.

In the case of the national SO_2 trading program established under the acid rain provisions of the Clean Air Act, the vast majority of allowances were given for free to those entities with emissions that were regulated under the cap.

Interestingly, the allowance allocation plans that have been announced for Phase 2 of the EU ETS, as well as the allocation approaches that have been proposed for the northeastern states' Regional Greenhouse Gas Initiative and in several draft bills introduced in the 110th Congress, rely on a mix of gratis allocation to different entities and allowance sales (auctions). Perhaps even more interesting, some Congressional proposals feature gratis allocations to entities such as states and energy-intensive commercial enterprises that are not directly regulated under the proposed policy. This change in thinking with regard to allocation policy might be taken to reflect a greater preference for efficiency. But since these proposals generally do not propose to use allowance-auction revenues to reduce taxes or displace existing distortionary taxes, their break with past allocation precedents is more likely to reflect different equity priorities.

As soon as allowances are seen as representing wealth—perhaps a considerable amount of wealth—it becomes obvious that how this wealth is distributed via the allowance allocation method will alter the relative well-being of individual firms and stakeholder groups in the economy. Allocation can also, however, alter the behavior of the aggregate economy and the pattern of GHG emissions going forward *if* the allocation is dependent on current or future behavior (in contrast to an allocation based entirely on historic behavior). This is because an allocation based on future actions or behavior inevitably creates incentives for those actions or behaviors. Since those actions or behaviors in turn can affect the efficiency of the cap-and-trade program, it is imperative that the incentive properties of any updating allocation method be well thought through as later discussion of an example from the EU ETS illustrates.

Using Gratis Allocation to Mitigate the Costs of the Emissions Reduction Program to Individual Entities

As noted above, equity and other distributional objectives can be achieved through the allocation of allowances. An example is provided by draft legislation (S. 1766) introduced in the 110th Congress by Senators Bingaman and Specter. This proposal would allocate a portion of the permits for free to both regulated and unregulated entities in the energy and manufacturing sectors, as well as to states. In addition, it would steadily increase the portion of allowances auctioned relative to the portion being distributed gratis (specifically, the portion of allowances auctioned increases from 12 percent of the total allowance pool in 2012 to 26 percent by 2030). Revenues from auctioning allowances would be used to fund technology development, climate-change adaptation, and assistance to low-income households. Other legislative proposals in both the House and Senate follow the Bingaman/Specter approach and use allowance allocation for a variety of purposes besides compensating regulated industry, including to provide credits for early reductions, to promote CO_2 sequestration on agriculture lands, to provide adaptation assistance to communities and ecosystems that are particularly vulnerable to the effects of climate change, to subsidize energy costs for low-income households via a direct allocation to states, and to establish a dedicated source of funding for low-carbon technology R&D and commercialization activities.

While it is feasible to use allocation as the Bingaman/Specter bill proposes (that is, to distribute the cost burden imposed by the cap more equitably), doing so effectively requires good information about which sectors, households, and regions of the country will bear the cost of meeting the emissions cap. Unfortunately, this information is not readily available in a reliable and objective form; moreover, due to the magnitude of the wealth embodied in allowances, there are massive incentives for sectors, households, and regions to claim significant costs in an attempt to capture a larger share of the available allowance pool.

Gratis Allocation: Grandfathering Based on Emissions

Suppose a decision has been made to allocate allowances for free to a particular sector. How might allowances be allocated within that sector? As has already been noted, gratis allocation to regulated entities has been the norm in emissions trading programs to date, and the simple method applied to distribute allowances to individual firms has usually involved the concept of "grandfathering." Each regulated entity receives a share of the total allowance pool that is equal to its share of total emissions from all regulated entities in a defined baseline year (equivalently, the emissions of each regulated entity in the baseline year are multiplied by the ratio of the emissions cap to total emissions in the baseline year).

Gratis Allocation: Grandfathering Based on Output

Grandfathering is a straightforward allocation method, but it relies on past behavior, thereby granting the greatest number of allowances to the historically largest emitters. Grandfathering can also be used in an allocation method that does not reward past emissions but is instead based on past output. That is, each regulated entity within a sector receives a share of the total allowance pool that is equal to its share of total sector-wide output (rather than emissions) in a given baseline year. Thus, the entity with the highest historic output captures the largest share of allowances, not necessarily the entity with the highest emissions.

To date, grandfathering allocations has awarded free allowances only or primarily to regulated entities, but the grandfathering approach can also be applied more broadly to distribute allowances to entities that are not directly regulated. For example, allowances could be awarded to large energy consumers to offset the impacts of higher energy

prices. In such cases, allowances might be allocated on the basis of historical output or labor input or some other metric related to the entity's ability to pass along higher energy costs.

Gratis Allocation: Output-Based Updating

Any grandfathering approach to allocation is based on past behavior and therefore generally does not take into account changes that occur in a sector over time. A method that does take change into account is output-based updating, which is the dynamic analog to output-based grandfathering. In the updating case, output shares are recalculated over time, and successive allocations are revised to reflect each entity's changing share of sector-wide output.

While updating sounds like an improvement over static allocation, it brings with it new issues. Because regulated entities know their future allowance allocation will be tied to output, and allowances are valuable, this approach creates incentives for firms to increase their share of output so they can increase their share of allowances. Incentives to increase output have two implications. First, as firms compete to increase output and capture a larger share of the allocation, output prices fall (with the allocation acting like a subsidy on output). Second, as prices fall, consumers have a smaller incentive to reduce their consumption of the goods and services produced by the regulated sector. While lower prices may be good thing for consumers, the fact that conservation is not fully incentivized increases the overall cost of the cap-and-trade program.

Gratis Allocation: Changing Incentives

There is no limit to the variety of approaches and methodologies that could be used to distribute free allowances to different entities and stakeholder groups. Many forms of allocation have been and will be proposed to achieve some economic and/or political objectives. From the standpoint of economic efficiency and environmental effectiveness, however, what matters most is the effect the allocation method has on the future behavior of entities in the economy. As should be evident from the foregoing discussion, this effect may not be immediately apparent.

Under the EU ETS, for example, a regulated entity loses its allocation if it closes a regulated facility. This seems like a reasonable rule—no emissions, no allocation. But the effect of this rule is to create forward-looking incentives to keep inefficient and perhaps highly emitting facilities operating just so the parent firm can claim allowances. This outcome would likely not be desirable in the power sector, but could be viewed as advantageous for sectors that are subject to external competitive pressures; in this case, keeping facilities from closing and moving abroad would likely be viewed as a good outcome.

From the standpoint of economic efficiency and environmental effectiveness, however, what matters most is the effect the allocation method has on the future behavior of entities in the economy.

Allocation to New and Retiring Sources

One of the more challenging issues that arises in designing an allocation methodology is how to handle the entry of new sources and the retirement of existing sources. Where will new sources get allowances and what happens to the allowances given to retiring sources that no longer need them? If allowances are auctioned, new entrants and retiring sources pose no special problems—new entrants buy allowances like all existing sources, while retiring sources should be holding no excess allowances.

The problem of accommodating new and retiring sources comes about when some or all allowances are allocated gratis. In this case, the government is transferring wealth to the private sector. If new entrants are not afforded the same wealth transfer as existing sources, they may be disadvantaged. Similarly, retiring sources benefit if they are able to retain their allocations after ceasing operation.

Deciding how to allocate emissions allowances under a CO_2 cap-and-trade program amounts to deciding how to distribute an asset worth, in aggregate, tens (if not hundreds) of billion dollars per year.

There is no single view on how to treat this issue. As noted previously, the EU ETS sets aside allowances for future allocation to new entrants and reclaims allowances from retiring sources. In contrast, the current U.S. SO_2 program has a very limited allowance set-aside for new entrants and allows retiring entities to retain their allowances. In some recent climate-policy proposals in the United States, allocations to new entrants are conditioned on the achievement of certain performance standards. For example, new coal-fired power plants might be required to achieve the same emissions level as integrated gasification combined-cycle plants to qualify for allowances from a reserve pool or set-aside for new entrants.

As already noted, the problem with setting allowances aside for new entrants and reclaiming allowances from retiring sources lies in the incentives this creates for future business behavior. Tying allowances for new entrants to the achievement of certain technology benchmarks can favor technology in unintended ways and on grounds other than curbing GHG emissions. Obviously, the concern about creating incentives that distort future behavior in undesirable ways diminishes in importance over time under a policy that gradually shifts to auctioning all or most allowances, as was recently proposed by a coalition of business and environmental groups known as the U.S. Climate Action Partnership.

Conclusion

Deciding how to allocate emissions allowances under a CO_2 cap-and-trade program amounts to deciding how to distribute an asset worth, in aggregate, tens (if not hundreds) of billion dollars per year. It is a hard distributional question that in some sense begs a legislative answer. Congress has typically been the authority best equipped to adjudicate questions of a fundamentally distributional nature. At the same time, analysis can inform important economic questions. First, the impact of a cap-and-trade program is not as obvious as it might seem: regulated businesses do not necessarily bear the brunt of program costs. More to the point, regulated entities need not be the only entities that receive free allocations. Second, there is growing interest in using auctions to distribute a large share of allowances (and, in some recent proposals, eventually most or nearly all allowances). This change in thinking about allocation has come about for a variety of reasons: one rationale is that using auction revenue to cut other taxes (or to avoid tax increases) can substantially reduce the cost of the climate policy. Finally, it is very important to consider how allocation rules can spur future behavior in possibly unintended ways. Unintended changes in incentives and behavior have the potential to significantly raise the cost of the climate program.