





ISSUE BRIEF 5
**EMISSIONS TRADING
VERSUS CO₂ TAXES VERSUS
STANDARDS**

IAN W.H. PARRY AND WILLIAM A. PIZER



5

EMISSIONS TRADING VERSUS CO₂ TAXES VERSUS STANDARDS

Ian W.H. Parry and
William A. Pizer

SUMMARY

Much attention has focused on the design of a trading program for carbon dioxide (CO₂) emissions, but a more fundamental question is whether emissions trading is really the best regulatory model. In particular, are there potential advantages or disadvantages to a CO₂ tax versus a cap-and-trade program? What about more traditional forms of regulation? This issue brief compares and contrasts these policy approaches, and offers the following observations:

- There are many similarities between CO₂ taxes and tradable allowances or permits. Both reduce emissions by associating a uniform price with emitting activities at any point in time, leading to efficient, low-cost emission reductions. Both can be administered on upstream fossil-fuel producers (based on the carbon content of fuels) to capture economy-wide emissions, or on downstream emitters to capture emissions from large sources. And both can incorporate incentives for carbon sequestration and other offset activities.
- Taxes generally fix the price of emissions, and leave the annual level of emissions uncertain; in contrast, tradable permits generally fix the level of emissions, and leave the price uncertain. Because climate change hinges on the long-term accumulation of global emissions, a predictable price tends to have advantages—for both the environment and the economy—over fixing the level of U.S. emissions for a short time horizon of several years. Over longer horizons, as nations converge on a common target for stabilizing atmospheric greenhouse gas (GHG) concentrations and as international participation in global emission-reduction efforts grows, fixed emissions targets become increasingly advantageous.
- Taxes generally raise government revenue, while tradable permits—at least traditionally—have not. New government revenue, if used to cut other taxes or provide valuable public goods, generates additional economic benefits that are not achieved under a traditional system of tradable permits in which the majority of permits or allowances is allocated for free to regulated entities. On the other hand, the allocation of free permits or allowances under an emissions-trading regime can be tailored to address concerns about an otherwise unequal distribution of regulatory cost burdens across firms and regions.
- These traditional differences between a tax and trading policy are easily blurred in a hybrid emissions trading system where some allowances are auctioned to raise government revenue and where banking and a safety valve (or perhaps borrowing) stabilize prices. Recent proposals for a Federal Reserve-like body to monitor allowance markets address this same issue.
- A few differences between these two types of policies are more immutable. For example, emissions trading does require additional institutions, though experience

suggests that these institutions are likely to arise quickly and for the most part inexpensively. Another difference is that a CO₂ tax tends to reframe the debate in terms of revenue and fiscal policy.

- Traditional forms of regulation—technology and performance standards—represent an alternative to emissions trading or CO₂ taxes, but can be much more costly because they do not allow the flexibility to shift efforts toward the cheapest mitigation opportunities. As a complement to emissions trading or CO₂ taxes, however, flexible standards can address possible additional market failures and potentially lower costs.

Similarities Between CO₂ Taxes and Emissions Trading Programs

A CO₂ tax imposed upstream in the fossil-fuel supply chain (with rates reflecting the amount of CO₂ that will be emitted when the fuel is later combusted in automobiles, during electricity generation, and so on) minimizes the number of entities subject to the tax and therefore has administrative advantages. Roughly speaking, the tax would be passed forward into the price of coal, natural gas, and petroleum products and therefore ultimately into the price of electricity and other energy-intensive goods. These higher energy prices would encourage the adoption of fuel- and energy-saving technologies across the economy and promote switching from carbon-intensive fuels like coal to natural gas and renewable fuels. In these regards, a CO₂ tax closely resembles an upstream emissions-trading system, where the price of allowances is passed forward in the form of higher fuel prices.

Neither policy has to be implemented upstream: CO₂ taxes and emissions-trading programs can be implemented anywhere in the chain from fossil-fuel production (upstream) to ultimate fuel combustion (downstream).¹ Upstream programs, however, are typically more efficient—in the sense that they lead to lower costs per ton of emissions reduced—because they can encompass virtually all emissions sources with minimal administrative burden, thereby maximizing low-cost mitigation opportunities. In contrast, downstream programs necessarily exclude small sources, as does the European Union's Emissions Trading Scheme (EU ETS). And either a tax or tradable-permit program, upstream or downstream, can—via offset and crediting programs—incorporate

incentives for downstream carbon capture and storage at industrial facilities, for forestry expansion on farmland, and for other downstream activities.

Potential Advantages of a CO₂ Tax

Carbon taxes have several advantages over traditional emissions-trading systems, but as discussed later, some of these advantages can be partly captured through modifications to the cap-and-trade approach.

One potentially important advantage of a CO₂ tax is that it establishes a well-defined price for emissions of CO₂ and other greenhouse gases. The price may rise over time, but it is known. In contrast, allowance or permit prices under a cap-and-trade system can be volatile because the supply of allowances is fixed, whereas demand will vary considerably at different points in time. Changes in energy demand, fuel-price fluctuations (like, spikes in natural gas prices), and a variety of other factors can cause demand for allowances to fluctuate significantly. Price volatility in allowance markets may in turn deter both long-term capital and R&D investments in low-carbon technologies that have high up-front costs. The long-term payoffs of making such investments will be very uncertain if the future price of CO₂ is unknown.

Moreover, it typically makes economic sense to allow nation-wide emissions to vary on a year-to-year basis because prevailing economic conditions affect the costs of emissions abatement. This flexibility is inherent in a CO₂ tax because firms can choose to abate less and pay more tax in periods when abatement costs are unusually high, and vice versa in periods when abatement costs are low. Traditional cap-and-trade systems do not provide this flexibility because the cap on economy-wide emissions has to be met, whatever the prevailing abatement cost. Intuitively, imposing strict limits makes economic sense only if (1) we are rapidly approaching a threshold in atmospheric greenhouse gas concentrations beyond which there is a risk of dangerous and extremely costly climate change impacts and (2) strict emissions limits can be globally enforced. It is worth noting that most trading programs do allow banking—that is, firms can save unused allowances for use in future compliance periods—and thus provide some flexibility, especially if initial targets are sufficiently generous for a long enough time to allow a bank to emerge. The topic of borrowing is discussed below.

Another potentially important advantage of CO₂ taxes is that they directly raise revenues for the government, whereas under past emissions trading systems, the government has

¹ See Issue Brief #4 on scope and point of regulation. That issue brief discusses issues related to upstream versus downstream regulation.

traditionally given away free allowances. At current emissions levels, for example, a tax of \$10 per ton of CO₂ on all greenhouse gases would raise about \$70 billion of revenue per year for the federal government, or about 9 percent of federal personal income taxes. This extra revenue could be used to lower the rates of other taxes, such as those on individual income, thereby producing important benefits for the economy. (Such a “tax swap” was implemented by the United Kingdom in conjunction with its 2001 climate change levy, and was proposed in the 110th Congress by Representative John Larson (D-CT) in H.R. 3416.) Income taxes cause a variety of distortions in the economy. For example, by taxing away some of the returns to working and saving, income taxes deter some people from joining the labor force and encourage others to consume too much of their income. Income taxes also induce a bias away from ordinary spending towards items that are deductible from taxes (owner-occupied housing and employer-provided medical insurance, for example). These economic distortions could be reduced if CO₂ tax revenues were used to lower income taxes.

Even if the revenue from a CO₂ tax is not used to cut other taxes, it could still flow to a variety of important uses—including to fund energy R&D; support climate-change adaptation efforts; or provide assistance to stakeholders, communities, and/or low income families adversely affected by the policy. Weighing against this revenue-raising advantage is the risk that the government will spend the additional revenue on programs that cost more than the benefits they provide, thereby in effect increasing the societal cost of the CO₂ tax relative to the cost of a comparable cap-and-trade program with free allowance allocation.

Aside from possible differences in economic efficiency, revenue to the government (and its potential uses) is likely to have different distributional consequences—in terms of costs and benefits to various individuals and firms in the economy—than a free distribution of allowances. Under the latter approach, benefits flow primarily to the recipients of free allowances—typically businesses and their shareholders and/or regions of the country with higher emissions. Revenues that flow to government as the result of a tax can be redistributed more broadly across the population: for example, to lower tax rates for all income groups. Critical questions, therefore, include the degree to which the burden of a market-based CO₂ program is broadly spread across society (or, conversely, concentrated among a particular group of carbon-intensive businesses or regions) and how the government could, and would, spend any tax revenues.

Finally, emissions trading systems require new institutions to function effectively; that is, they require smoothly running markets where firms can buy and sell permits or allowances and obtain information about permit prices now and in the future. Experience with existing trading programs, such as the U.S. SO₂ trading program, has shown that these institutions can arise quickly and for the most part inexpensively. Some emissions-trading markets have witnessed exceptional volatility during their inception. For example, allowance prices in the U.S. NO_x budget program skyrocketed in the wake of uncertainty about whether Maryland, a net supplier of allowances, would enter the program on time. In the EU ETS, permit prices crashed spectacularly after emissions data pointed to an excess of CO₂ permits rather than the expected shortage.

Hybrid Trading Schemes

The problem of allowance price volatility under a cap-and-trade system can be partly addressed by cost-containment mechanisms, such as a “safety valve,” coupled with allowance banking. With a safety valve, firms can buy an unlimited number of additional permits from the government at a pre-determined, possibly escalating price. The safety valve essentially functions as a cap on permit prices; it is most likely to be triggered when demand for permits and abatement costs are high. Allowance banking allows firms to hold over some allowances, in periods when the demand for permits is slack because abatement costs are low, for use in future periods when permit prices are expected to be higher again. In effect, this mechanism creates a floor under permit prices.

As an alternative to a safety valve or price cap mechanism, allowance borrowing has recently entered the U.S. policy debate. Legislation introduced in the 110th Congress by Senators McCain and Lieberman would allow firms to borrow up to 25 percent of their allowance obligation in a given year for up to five years (paying 10 percent interest annually). Borrowed allowances would be deducted from the allowance pool available in future years. Coupled with somewhat clear expectations about future prices, this mechanism could provide flexibility similar to a tax. Without clear expectations about future prices, however, borrowing would tend to dampen short-term volatility while leaving the market open to fluctuations based on longer-term expectations about the cap and prices.

The second potential advantage of a tax—that it raises revenues for government—can also be achieved by a cap-and-trade program if allowances are auctioned instead of

being distributed for free (conversely, the revenue-generating properties of a CO₂ tax could be offset by including rebates or exemptions). Although auctioning 100 percent of allowances would mimic the revenue advantages of a CO₂ tax, partial auctions or—as suggested by the recent U.S. Climate Action Partnership proposal—a gradual transition to auctions, offer a spectrum of possibilities.

Potential Disadvantages of a CO₂ Tax

CO₂ taxes have several practical disadvantages. One is simply political resistance to new taxes; for example, despite a major effort, the first Clinton administration failed to enact an energy tax motivated on environmental grounds. Nonetheless, a CO₂ tax should not be ruled out entirely on this basis: it is always difficult to predict what policies may or may not be viable in the future, especially under different political leadership and likely greater public awareness of, and concern about, both global warming and the federal debt.

Another concern (noted earlier) is that revenues from a CO₂ tax (or auctioned allowances) might be spent inefficiently or even wasted. This could occur, for example, if revenues go toward special interests, rather than substituting for other taxes or addressing important social needs. In principle, legislation accompanying a CO₂ tax could specify how the new revenue must be used, thereby avoiding the risk that it would be dissipated among competing special interest groups. This approach would require political will, as would—more generally—any effort to pursue a fiscally focused climate policy in which environmental objectives are pursued in a manner that maximizes broader public-good objectives. A shift in focus to a policy approach motivated by revenue and fiscal considerations, as well as by environmental concerns, could have important implications—not only in terms of the jurisdiction of agencies and Congressional committees, but also in terms of the broader debate. At first blush, it might appear that such a shift could increase the political difficulty of achieving desired environmental objectives. On the other hand, a more transparent airing of the energy price implications of a trading program or carbon tax—and of the offsetting social benefits that could be achieved by re-directing revenues raised by the policy for other public purposes—could help to build better understanding of, and deeper support for, the policy among the public and some private-sector stakeholders.

Of course, policymakers may wish to compensate the industries most affected by the carbon regime or ease the transition for firms and workers facing adjustments. Compensation can be provided in a straightforward way under an emissions-trading regime by granting free allowances to particular firms or groups.² Compensation can also be provided under a CO₂ tax regime, although legislatively, this is more complex.

Finally, policymakers may wish to reduce emissions in a gradual fashion by setting progressively more stringent targets each year, perhaps because atmospheric CO₂ concentrations are already judged to be dangerously high, or because steady progress on emissions reductions more effectively communicates America's seriousness about tackling climate change to the international community. A traditional cap-and-trade system with no safety valve is best tailored to achieving defined emissions targets; in contrast, progress on emissions reductions is less certain under a CO₂ tax because emissions will vary from year to year with economic conditions. A cap-and-trade program with a safety valve represents a potential compromise between these approaches: the safety valve limits allowance prices and emissions-abatement costs, but the trigger price for the safety valve can be steadily increased over time, providing more certainty about emissions levels over the longer term.

What About Recent Proposals for Federal Reserve-like Oversight of Carbon Markets?

In July 2007, a new proposal emerged in Congress for government oversight of carbon markets via a new body, much like the Federal Reserve. Like the Fed, this body could intervene in response to unexpectedly high (or low) prices or to curb excessive price volatility.³ The basic idea is that this type of oversight would deliver some of the market-stabilizing benefits of a safety valve while providing greater confidence in the achievement of longer-run emissions goals. Although this proposal does not eliminate the trade-off between price and emissions certainty, it introduces an additional nuance into the current debate about cap-and-trade proposals with and without explicit cost caps.

At the same time, empowering an outside agency to intervene in the market poses risks. Designed or operated poorly, such oversight could exacerbate volatility. For

² See Issue Brief #6 for a longer discussion of the allocation issue.

³ See <http://www.senate.gov/-warner/pressoffice/statements/20070802a.htm> and <http://www.nicholas.duke.edu/institute/carboncosts/>.

example, consider a provision that requires prices to remain above a particular threshold for a period of time before intervention occurs. As permits trade above the threshold, it becomes increasingly likely that the government will intervene to lower prices. At that point, allowance buyers begin waiting for the intervention—who wants to buy now if prices are going to be lower in the future? As demand falls, prices drop, the likelihood of intervention recedes, and therefore prices begin to rise again. In this scenario, the prospect of intervention could have the perverse effect of *increasing* price volatility and market instability. Alternatively, if interventions are quantitatively limited, the most valuable role of the outside agency—addressing a truly exceptional shortage—is compromised.

In sum, the idea of an independent oversight body for future carbon markets is likely to be the subject of additional discussion and elaboration as Congress debates different climate policy proposals going forward. On the one hand, this approach may provide additional opportunities to fine tune the balance between emissions and cost uncertainty in a tradable permit program. At the same time, however, the implications of such a mechanism must be carefully evaluated and important design questions considered in terms of minimizing any additional political or market risks associated with potential intervention.

Is There Any Role for Traditional Regulation?

From a cost-effectiveness standpoint (that is, in terms of minimizing cost per ton of emissions reduction achieved), market-based instruments like CO₂ taxes and emissions trading systems, applied to all emissions sources, are typically superior to traditional regulation. (Examples of traditional regulation include facility-specific pollution-control requirements, limits on emissions per kilowatt-hour of electricity generation, fuel economy requirements imposed on new vehicles, or regulations on fuels). Under market-based policies, the marginal cost of abatement is equalized across all sectors of the economy, across all firms within a sector, and across all opportunities for abatement. The least expensive abatement options are implemented first, such as substituting less carbon-intensive fuels for more carbon-intensive fuels, adopting energy-efficient technologies, and conserving energy at the household level by, for example, driving less and reducing residential heating and cooling loads.

Nonetheless, traditional regulations—such as technology standards that dictate the use of a particular technology or manner of operations, and performance standards that limit emissions generated per unit of economic output or activity—are frequently proposed as alternatives or complements to emissions taxes or tradable permits.⁴ The cost of such regulations is often less visible: emissions control requirements or performance standards raise the cost of certain goods and activities, resulting in price increases and income reductions that are not obviously tied to CO₂ emissions. Traditional regulation can also modify specific behavior directly, without appealing to incentives, and target preferred technologies or mitigation actions. No money is exchanged in the form of taxes paid or allowances traded—changes in behavior are simply required by law. While some view these features of traditional regulation as advantageous, they come at the cost of higher—perhaps much higher—costs.⁵ Thus, while imposing sector- or source-specific requirements might appear to reduce the cost of emissions abatement (by avoiding effects on energy prices completely or by reducing demand for allowances and hence lowering allowance prices), the total cost to society—taking into account the less transparent costs of traditional regulation—is likely to be higher than if the same overall result were achieved with a market-based program only.

Unlike market-based instruments, performance or technology standards typically do not impose an economywide carbon price and therefore fail to meet the conditions for efficiently distributing the burden of emissions reductions across different firms, households, and mitigation options. In contrast to minimum performance standards that must be met by every facility or product, *tradable* performance standards offer some ability to equalize marginal costs. Facilities or products that beat the standard cheaply generate credits used to offset excess emission rates at facilities or by products that miss the standard, achieving the standard on average at a lower cost. However, even tradable performance standards often overlap coverage in some areas, exclude coverage in other areas, and always fail to provide proper incentives for conservation. For example, tradable performance standards for the power sector and efficiency standards for appliances would overlap, as would a tradable fuel-economy standard for cars and a renewable or carbon-based fuel standard for gasoline. In the case of industrial facilities, where facility output is not

⁴ See, for example, recent proposals in California (<http://www.arb.ca.gov/cc/ccea/ccea.htm>) and the bill introduced by Senators Sanders and Boxer (S. 309, <http://www.sanders.senate.gov/news/record.cfm?id=269618>).

⁵ One estimate found that using fuel economy standards for light-duty vehicles and a renewable portfolio standard in the electricity sector would cost ten times as much as an economywide tradable permit system. See Pizer, W. et al., 2006. Modeling Economywide versus Sectoral Climate Policies Using Combined Aggregate-Sectoral Models, *Energy Journal* 27(3), 135-168.

easily defined on a consistent basis, it would be difficult as a practical matter to develop output-based performance standards that could be applied to a diverse population of sources.

Finally, performance standards do less to promote conservation than market-based instruments. Both types of regulation lead to emissions reductions, the cost of which raise the price of emissions-intensive goods, like motor fuel and electricity. Market-based instruments like taxes and emissions trading, however also associate a cost with the remaining emissions that *do* occur, further raising the price of these goods. While this may seem like a bad thing for consumers, it is precisely that price increase that encourages the right amount of conservation—such as driving less or using less electricity. For example, vehicle fuel-economy standards reduce emissions per mile traveled, but do not generate incentives to reduce driving (on the contrary, drivers of more efficient vehicles face lower costs per mile traveled and hence weaker incentives to reduce driving). While avoiding the increase in fuel prices that would accompany a cap-and-trade program or emissions tax might seem desirable on the surface, pursuing the same carbon-reduction objectives via product performance standards means higher costs and lower income somewhere else.

While economic analyses reach uniformly negative conclusions about the cost-effectiveness of traditional regulations as an *alternative* to emissions taxes or tradable permits, for the reasons discussed above, an economic argument can be made for performance standards as a *complement* to a market-based carbon regime, either to address additional market failures and/or because the price incentive for reducing CO₂ emissions under the market-based regime does not reflect the full value of those reductions to society. Examples of market failures that might be amenable to traditional regulatory approaches include the possibility that purchasers may undervalue more energy-efficient vehicles or appliances, or that efforts to develop new technologies may generate substantial public benefits (in the form of new knowledge) that are not appreciated by the firm conducting the research.⁶ Finally, the inability to price greenhouse-gas reductions appropriately may arise from political opposition to higher energy prices and/or concerns about the international competitiveness of energy-intensive industries.

⁶ The latter point is discussed at length in Issue Briefs #9 and #10, concerning technology policy.

Conclusion

Significant differences exist between emissions taxes and trading programs. In particular, emissions taxes will generate revenue and set prices, whereas trading programs have traditionally distributed most allowances for free and fixed emissions. Recent proposals for emissions-trading programs with allowance auctions and safety valves (and other mechanisms), however, suggest that many of the key features of a CO₂ tax can be partly included in a trading program. The same is not true for a tax: it is not possible to create fixed emissions limits without resorting to emissions allowances or permits. And whereas tax revenues can be redistributed, industry stakeholders have frequently responded to carbon-tax proposals by seeking exemptions or voluntary agreements in lieu of taxes. Politically, this represents a very different challenge than adjudicating competing claims for allowance allocations. Many other program design questions—such as point of regulation and whether to include offsets and other crediting mechanisms—have always applied equally to emissions taxes and trading systems.

What, then, are the fundamental differences between the major policy options? Emissions-trading programs do require additional institutions: markets, brokers, and information tools to function effectively and manage risk. These institutions tend to arise quickly and inexpensively but there is generally some risk of excess volatility, especially in the early phases of implementation. A tax approach does tend to reframe the traditionally environmental issue as, at least partially, a revenue issue—with attendant political, jurisdictional, and institutional consequences. Of course, similar issues are likely to arise in connection with revenue-generating allowance auctions. All this suggests that designing a CO₂-reduction policy is more usefully viewed as a matter of selecting different program features along a continuum than as a simple dichotomous choice between taxes and tradable allowances. In that selection process, trade-offs must be made between emphasizing certainty about prices versus certainty about emissions and between raising revenue versus compensating some stakeholders through the free distribution of allowances.

The comparison between a market-based approach (whether taxes or tradable allowances) and traditional regulation is much simpler. While there is possibly an economic rationale for traditional regulations as a *complement* to a market-based policy when other market failures exist (or when the emissions price under a market-based system is constrained for political or other reasons to be less than its social value), there is no economic rationale for such regulations as an *alternative* to, or

substitute for, market-based programs. Traditional regulation is always more expensive because it: (1) generally fails to trade low-cost reductions off against high-cost reductions, (2) tends to provide overlapping incentives for reductions from some types of sources while excluding others, and (3) often fails to provide proper incentives for conservation. Nonetheless, the

desire to pursue preferred technologies or mitigation activities and to reduce the obvious price impact on energy end-users (even recognizing that the result is likely to be higher costs elsewhere) often means that substantial support exists for traditional types of regulation in some sectors of the economy.

	CO ₂ tax	Cap and trade	Traditional regulation <i>(e.g., source-specific emissions standards)</i>
Certainty over CO₂ price or cost?	Yes. The tax establishes a well-defined price.	No. But price volatility can be limited by design features, such as a safety valve (price cap) or borrowing.	No.
Certainty over emissions?	No. Emissions vary with prevailing energy demand and fuel prices.	Yes, in its traditional form (over capped emissions sources). No, with the use of additional cost containment mechanisms.	No; regulating the rate of emissions leaves the level uncertain.
Efficiently encourages least-cost emissions reductions?	Yes.	Yes.	No, but tradable standards are more efficient than non-tradable standards.
Ability to raise revenue?	Yes. Results in maximum revenue generation compared to other options (assuming cap-and-trade alternative includes substantial free allocation of allowances).	Traditionally—with a largely free allocation—no. Growing interest in a substantial allowance auction suggests opportunity to raise at least some revenue now and possibly transition to a complete auction that generates maximum revenue in the future.	No.
Incentives for R&D in clean technologies?	Yes. Stable CO ₂ price is needed to induce innovation.	Yes. However, uncertainty over permit prices could weaken innovation incentives.	Yes and no. Standards encourage specific technologies, but not broad innovation.
Harm to competitiveness?	Yes, though if other taxes are reduced through revenue recycling, competitiveness of the broader economy can be improved.	Yes (as with a tax), but giving firms free allowances offsets potentially harmful effects on profitability.	Somewhat. Regulations increase the cost of manufacturing but, unlike taxes or tradable permits, do not raise the price of fossil energy.
Practical or political obstacles to implementation?	Yes. New taxes have been very unpopular.	Yes. Identifying a reasonable allocation and target is difficult.	Yes. Setting the level of the standard is difficult.
New institutional requirements?	Minimal.	Yes, but experience with existing trading programs suggests that markets (for trading permits and exchanging information across firms and time periods) arise quickly and relatively inexpensively.	Minimal (unless tradable).