Journal of Economic Perspectives—Volume 27, Number 1—Winter 2013—Pages 87–102

Markets for Pollution Allowances: What Are the (New) Lessons?

Lawrence H. Goulder

About 45 years ago a few economists offered the novel idea of trading pollution rights as a way of meeting environmental goals. Such trading was touted as a more cost-effective alternative to traditional forms of regulation, such as specific technology requirements or performance standards. The principal form of trading in pollution rights is a cap-and-trade system, a system whose essential elements are few and simple. First, the regulatory authority specifies the cap—the total pollution allowed by all of the facilities covered by the regulatory program. Second, the regulatory authority needs to distribute the allowances, either by auction or through free provision. Third, the system provides for trading of allowances.

The idea of cap and trade was implicit in the classic work of Ronald Coase (1960) on how well-defined property rights can assure efficient outcomes despite the presence of externalities. It then took on shape in journal contributions by Crocker (1966), Dales (1968), and Montgomery (1972). The concept materialized into policy starting in 1974, when the US Environmental Protection Agency allowed companies to trade emissions reductions among sources within the firm so long as total, combined emissions did not exceed an aggregate limit (Tietenberg 1985; Hahn and Hester 1989; Foster and Hahn 1995). The EPA's "offset" program, introduced in 1997, went further in allowing for trading across firms. These systems applied to various local pollutants, including volatile organic compounds, carbon monoxide, sulfur dioxide, and nitrogen oxides.

■ *Lawrence H. Goulder is Shuzo Nishihara Professor of Environmental and Resource Economics, Stanford University, Stanford, California. He is also a University Fellow, Resources for the Future, Washington, D.C. and a Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. His email address is goulder@stanford.edu.*

http://dx.doi.org/10.1257/jep.27.1.87. doi=10.1257/jep.27.1.87

Since the 1980s the use of cap and trade has grown substantially. The three other papers in this symposium reveal and assess some of the most important applications. Schmalensee and Stavins indicate that cap and trade has been a principal part of the US Environmental Protection Agency's efforts to reduce US emissions of sulfur dioxide (SO_2) under the Clean Air Act. Newell, Pizer, and Raimi show how cap and trade applied to emissions of greenhouse gases has become an important instrument for climate change policy at the regional (state), national, and international levels.[1] And Fisher-Vanden and Olmstead describe how emissions trading is being used to control water pollution. Cap and trade was also applied to accomplish the phasedown of leaded gasoline in the United States during the 1980s. It has been employed at the municipal level as well, to control a range of pollutants including carbon monoxide, volatile organic compounds, SO_2, and nitrogen oxides (NO_x). An example is the Regional Clean Air Incentives Market (RECLAIM) in the Los Angeles airshed, a program introduced in 1994.[2]

In addition, principles similar to cap and trade have promoted cost-effective environmental protection in programs involving trading of commodities other than pollution. At least 10 nations have implemented programs of individual transferrable fishing rights, in which a limited supply of permits to catch fish is allocated among fishing operators. And some US states have instituted programs involving tradable land-development rights as a way of conserving natural habitats and protecting biodiversity.[3]

The provision for trading of allowances is the key to achieving desired emission reductions at a lower cost than with other, less-flexible, approaches. The separate sources of pollution will tend to have a range of different marginal costs for abating pollution. Facilities with the highest costs of reducing emissions will find it advantageous to reduce their costs by buying additional allowances from other facilities rather than trying to meet the pollution limits given by their original holdings of allowances. Likewise, the facilities for which it is relatively inexpensive to reduce emissions will find it profitable to sell some of their allowances. Even though this obliges them to reduce emissions even more, the returns from the sale of allowances will exceed the additional abatement (pollution-reduction) costs.

[1] Regional programs include the carbon dioxide emissions trading ("cap-and-trade") program in the US Northeast under the nine-state Regional Greenhouse Gas Initiative, which went into effect in 2008. A cap-and-trade program is slotted to go into effect in California in January 2013. National programs include carbon emissions cap-and-trade systems in Australia and New Zealand, and the European Union's 27-country cap-and-trade program. International trading in greenhouse gas emissions is allowed for under the Kyoto Protocol, the international treaty to reduce greenhouse gases.

[2] Cap and trade is not the only form of pollution trading, although it is the one that has gained most attention and been implemented the most. Another trading approach allows firms to receive credits for reducing emissions below some stipulated level, even though they are not penalized if their emissions exceed that level. Here the regulator offers a one-sided option, and there is no cap on aggregate pollution from the covered facilities. This approach has been considered for bringing about greater participation by developing countries in efforts to reduce greenhouse gases (Millard-Ball forthcoming).

[3] For an analysis of a range of issues associated with individual transferable fishing rights and tradable habitats, see Arnason (2012) and Crocker (2005), respectively.

Thus, trading leads to more abatement by those facilities that can reduce emissions most cheaply. It tends to bring marginal abatement costs toward equality, a condition for cost minimization. Regulators do not need to know the marginal abatement costs of individual facilities: they can let the market promote equality in marginal abatement costs. This is a potential advantage over technology requirements or performance standards because regulators generally will not have sufficient information to set the requirements or standards at levels that assure equal marginal abatement costs across the covered entities.

In this overview article, I consider some key lessons about when cap-and-trade programs work well, when they perform less effectively, how they work compared with other policy options, and how they might need to be modified to address issues that had not been anticipated.

I distinguish two types of lessons. The first are, essentially, confirmations of prior theoretical predictions. The second are insights that emerge in response to previously unanticipated circumstances or problems, or as a result of recent analytical contributions. I consider each type of lesson in turn.

Some (Mostly) Reassuring Outcomes

1) In national and subnational cap-and-trade programs applied to local air pollutants, effective monitoring and compliance have enabled cap-and-trade programs to succeed in limiting emissions to specified targets. Difficulties of monitoring have limited the use of cap-and-trade programs aimed at water pollution, and problems of compliance have hampered the effectiveness of cap-and-trade programs under the international Kyoto Protocol.

For the early proponents of cap and trade, one of the touted attractions was that this regulatory approach would establish and maintain clear limits on total emissions of pollution by the covered sectors, with the limit in each period given by the specified cap (or total number of allowances in circulation). The ability to specify an aggregate limit on emissions distinguishes cap and trade from other regulatory approaches: neither limits on the emissions at the firm- or plant-level, nor mandates for the use of certain technologies for pollution abatement, nor sector- or economy-wide pollution taxes specify a total quantity of emissions.

Imposing a limit on total emissions and letting the market determine the price is not necessarily more efficient than imposing a price on emissions and letting the market determine the quantity—as under a pollution tax. Weitzman's (1974) seminal article indicates that the relative advantage of setting the quantity or setting the price depends on the nature of uncertainty about marginal benefits and costs from pollution reductions. But allowing the regulator to choose the quantity of pollution explicitly has considerable practical political appeal.

The promise of keeping aggregate pollution within the stipulated overall cap has been fulfilled in most of the cap-and-trade systems introduced for air pollution control. For example, the US Environmental Protection Agency's programs to reduce sulfur dioxide and nitrogen oxides under the Clean Air Act and the

RECLAIM program for curbing these same pollutants in the Los Angeles region can claim success in reducing emissions to the targeted levels. In addition, the European Union's Emissions Trading Scheme has largely managed to keep greenhouse gas emissions from covered sectors within the levels targeted (although this program is likely to have stimulated a partially offsetting increase in emissions outside of the European Union, a "leakage" phenomenon I discuss below).

Two factors have contributed to these successes. First, emissions of the air pollutants involved have proved relatively easy to monitor, or at least to estimate with some accuracy. In addition, the programs have included strong incentives for compliance. For example, under Europe's Emissions Trading System, the noncompliance penalty is 100 euros per ton, considerably higher than the market price of allowances, which has seldom exceeded 15 euros, and compliance in fact appears to have been very good in all of these programs.[4]

In contrast, under the Kyoto Protocol of 1997, serious problems of compliance have arisen and remain. This largely reflects the lack of significant enforcement capabilities under the Protocol. This is a problem common to many international agreements, rather than any inherent weakness of cap and trade. Under the Protocol, 37 nations committed themselves to maximum levels of emissions of greenhouse gases in the first commitment period, 2008–2012. Parties that did not meet their targets in the first commitment period were required to make up the difference plus 30 percent more in the anticipated second commitment period. However, several parties that are expected to miss their initial targets—including Japan, Canada, and Russia—have simply announced they will not continue to abide by the Protocol in the second commitment period.

In the context of water pollution, the accomplishments are somewhat limited. Cap and trade has enjoyed success in restricting the effluent pollution from regulated point sources. Currently, there are about 13 trading programs, with most of them arising since the turn of the century. As pointed out by Fisher-Vanden and Olmstead, trading of water pollution permits generally has embraced only those sources that are easy to monitor—namely large industrial establishments and municipal sewage treatment plants. The agriculture sector is an important contributor to water pollution, but in general this sector is not covered by enforceable effluent regulations under the Clean Water Act. This reflects the difficulty of monitoring the effluent from these so-called nonpoint sources. It is worth noting that any sort of pollution control, whether via market-based approaches or by way of more conventional approaches, is challenging with nonpoint sources. The absence of cap and trade applied to water pollution from agriculture also reflects the considerable political opposition by the agriculture industry to limits on pollution.[5]

[4] The qualifier "appears" is used because the successful cheaters, by definition, are not observed.

[5] Fisher-Vanden and Olmstead (this issue) point out another important challenge to the application of cap and trade to water pollution: water pollutants often are not uniformly mixed. As discussed by these authors, a simple cap-and-trade system, where given releases of effluent are all traded at the same price, can produce undesirable environmental outcomes.

2) Cap-and-trade programs have brought significant cost reductions relative to conventional regulatory approaches.

The evidence for cost savings from a cap-and-trade policy must always be indirect since researchers never observe the counterfactual world in which an alternative program is introduced under otherwise identical economic and environmental conditions. Moreover, there are not enough instances of cap and trade and other regulatory approaches in roughly similar settings to allow the impact of cap and trade to be identified econometrically.

Still, economists have managed to arrive at plausible estimates of cost savings by estimating the marginal abatement cost curves of the covered facilities, assessing the extent to which marginal abatement costs would differ across facilities under conventional regulation (often the previously prevailing form of regulation), and then calculating the extent to which these differences are eliminated (and total abatement costs reduced) by a cap-and-trade program. The analyses generally rely on the assumption that the market for trading allowances is effective in bringing marginal abatement costs to equality across facilities. Behind this assumption is the implicit assumption that transactions costs are low.

A review by Chan, Stavins, Stowe, and Sweeney (2012) of various analyses using this approach indicates that sulfur dioxide allowance trading under the Clean Air Act yielded cost savings in the range of 15 to 90 percent relative to the costs under conventional forms of regulation. There is some evidence that transactions costs are fairly low (Stavins 1995) and the trading market is fairly fluid, which would support these findings.

Using a similar approach, an analysis of sulfur dioxide and nitrogen oxides trading in the Los Angeles area RECLAIM market claimed cost savings of 46 percent relative to the costs of achieving the same aggregate reductions under the prior air quality management program, which involved fixed emissions caps and no trades. The estimates for recent savings may be overestimated, however, as various restrictions on trades have been introduced since the analysis was performed. In addition, some analyses suggest that the efficiency of the trading equilibrium was compromised as a result of interactions between cap-and-trade systems and rate-of-return regulation faced by utilities, an issue to which I return below. Ellerman, Convery, and de Perthius (2010) estimate that Europe's Emissions Trading System achieved cost reductions in the range of 2–5 percent. For other pollution trading markets, the quantitative evidence for cost savings is limited. However, even in these other markets the qualitative conclusion that cap and trade has lowered costs is tacitly supported by the mere existence of trading, as trading shifts responsibility for pollution reduction to facilities that can do so relatively cheaply.

Overall, these considerations suggest some success for many of the cap-and-trade systems that have been introduced. But some important qualifications are in order. To a large extent, these empirical studies show the cost savings compared to a relatively inflexible form of conventional regulation—fixed emissions caps. They show the savings from trading relative to the same regulation without trading. They do not assess cost-savings relative to other, more flexible, nonmarket instruments (such as performance standards) or relative to an alternative market-based instrument: namely, a pollution tax. In addition, the initial assessments of cost savings ignore factors whose importance has only recently come to light. I address these issues below.

Surprises, Challenges, and New Lessons

3) The environmental effectiveness and cost-effectiveness of cap and trade can be significantly compromised by interactions with other regulations.

Virtually all analyses of environmental policies have ignored interactions with other policies. This is particularly important in the case of cap and trade. Economic theory as well as recent experience shows that these interactions can significantly reduce both environmental effectiveness and cost-effectiveness.

One difficulty arises when regulations in one jurisdiction are "nested" within a cap-and-trade system introduced in a higher-level jurisdiction. Suppose, for example, a cap-and-trade system was introduced at the national level in the United States with a national emissions cap. Now suppose that a given state desires further emissions reductions by firms within its boundaries, beyond those that would result from the federal program: through cap and trade or some other instrument, the state prompts further reductions by facilities within its borders. As a result of this state's action, firms within this state will now have excess federal allowances, which they will sell to firms in other states that do not have tougher standards. Since nationwide emissions continue to be determined by the unchanged national cap, the one state's imposition of tougher environmental rules leads to no overall reduction for the nation: it just causes "emissions leakage"—offsetting increases in emissions elsewhere. By affecting the distribution of emissions, these adjustments can raise or lower aggregate environmental damage, depending on how they alter the geographical pattern of pollution concentrations. The national cap effectively prevents lower-level jurisdictions from eliciting further emissions reductions.[6]

The issue came to life when the United Kingdom recently decided to impose a tax on carbon dioxide emissions by electric power generators in the country. For each unit of emissions, these generators will need to pay this tax in addition to the price that they pay for emissions allowances from the EU Emissions Trading System (ETS). Although the tax will likely cause greater abatement by generators within the United Kingdom, it will not cause greater overall abatement in Europe, since overall European abatement is determined by the Europe-wide cap under the ETS. The UK initiative will reduce the UK's demands for emissions allowances from the ETS,

[6] For further discussion of these issues, see Fankhauser, Hepburn, and Park (2010), Burtraw and Shobe (forthcoming), and Goulder and Stavins (2012). The same issue can arise within a single jurisdiction. For example, California introduced a cap-and-trade system as part of its Global Warming Solutions Act. To the extent that other regulations such as a standard for low-carbon fuel aim to achieve further reductions, the affected firms will have excess allowances, and these allowances will be sold to other covered entities. Statewide emissions from the covered sectors will not be reduced further, as they are determined by the state's cap. For discussion of other interactions within a single jurisdiction, see Levinsohn (2012).

putting downward pressure on allowance prices and prompting increased emissions in the rest of Europe. CDC Climate Research (2011) offers a quantitative assessment of the impacts.

The issue also arose when 14 US states attempted to impose tighter limits on greenhouse gases per mile from automobiles below the level implied by existing federal Corporate Average Fuel Economy standards. The 14-state initiative would have caused automobile manufacturers in those states to more than meet the federal corporate auto fuel economy (CAFE) standards, allowing them to sell less-fuel-efficient cars in other states and still remain within the national standard. In Goulder, Jacobsen, and von Benthem (2012), my coauthors and I estimate that about 75 percent of reduction in greenhouse gases achieved in the 14 states would have been offset by increased emissions in other states. As it turned out, the 14-state initiative helped put pressure on automobile manufacturers to accept tighter requirements at the federal level in exchange for elimination of the tougher action by these states.

These difficulties are relevant to recent US initiatives to institute a federal-level tradable clean electricity standard, since some states may wish to impose standards tougher than the federal one.

A second problem arises when firms within the cap-and-trade system are subject to other *non*environmental regulations that affect demands for allowances and the distribution of emissions-abatement effort across firms. This issue arose in the South Coast Air Quality Management District's RECLAIM program to reduce emissions of sulfur dioxide and nitrogen oxides in the Los Angeles area. Electric power generators were important contributors to these emissions: however, these generators were also subject to rate-of-return regulation under the local public utilities commission. As shown by Kolstad and Wolak (2003), these vertically integrated firms *benefited* from higher allowance prices, because the higher prices could be incorporated in the rate base determining the prices that could be charged to consumers. The higher rate base implied higher prices for electricity, which yielded increments to profits despite the higher prices of allowances. These interactions implied a shift in the distribution of wealth from ratepayers to owners of utilities. They also implied a shift in ownership of allowances and abatement effort toward utilities and away from other emitters. This shift compromised cost-effectiveness, as some low-cost abatement by entities other than utilities was crowded out.

The Clean Air Act's sulfur dioxide allowance trading market offers yet another case where the cap-and-trade system was vulnerable to other regulations, as detailed in the accompanying article by Schmalensee and Stavins. In this case, the other regulation was the Clean Air Interstate Rule, which was promulgated in 2005, well after the cap-and-trade program's implementation in 1990. This rule imposed stringent emissions-reduction requirements that eventually led to significant reductions in the demand for sulfur dioxide allowances in the trading market. As a result, the cap in the sulfur dioxide trading program became no longer binding, and allowances prices subsequently have collapsed. Although the Clean Air Interstate Rule accomplished significant reductions (which many might applaud), the neutering of the cap-and-trade program suggests that the reductions were not accomplished as cost-effectively as would have been the case if instead the reductions had been achieved by a tightening of the cap (which would have required Congressional action).

Schmalensee and Stavins (this issue), along with Burtraw (forthcoming), claim that a key lesson from this episode is the importance of building flexibility into cap-and-trade systems. The absence of institutional rules permitting adjustments of the cap in the face of new information contributed to the need to invoke different, potentially less-efficient, regulations. Making it easier to adjust the cap might have some drawbacks, however. Greater flexibility could adversely affect the credibility of the government's commitment to a given time profile for the emissions cap and introduce new uncertainties into the system.

In sum, interactions with other regulations can compromise cap-and-trade's environmental effectiveness, distort the demands for allowances, or make a cap-and-trade program irrelevant. Ignoring regulatory interactions can be imprudent, just as a doctor in prescribing a medication without knowing what other medications the patient is taking would be reckless.

4) Volatility of allowance prices has been a significant concern.

Under cap and trade, the supply of allowances is highly inelastic in the short term, changing only as a result of government policy decisions (that one hopes are predictable). With highly inelastic supply, shifts in demand can cause significant price changes, and irregular shifts in demand can produce price volatility.

Some existing cap-and-trade systems have displayed considerable allowance price volatility. The energy supply crisis in California in the summer of 2000 gave power companies incentives to bring online some older power generators in the Los Angeles region. This led to a significant increase in the demand for emissions allowances for nitrogen oxides under the RECLAIM program, since allowances were needed to validate the emissions produced by these generators. As a consequence, NO_x allowance prices rose from about $400 per ton to an average in the year 2000 of over $40,000 per ton—with the average allowance price reaching $70,000 in the peak month of 2000 (Ellerman, Joskow, and Harrison 2003).

There was also significant price volatility in the first (that is, the pilot) phase of cap and trade under the European Union's Emissions Trading Scheme. About a year after its implementation, emissions allowance prices dropped dramatically with the release of information that indicated that the Phase I permit allocations were generous in the sense that they barely constrained the covered sources. The December 2008 futures prices fell from 32.25 euros to 17.80 euros between April 19 and May 12, 2006. There was even greater volatility for the Phase I permit prices contained in December 2007 contracts. These prices dropped from 31.65 euros on April 19, 2006, to 11.95 euros on May 3, 2006. When Phase II of the program began in 2008, allowance prices rose to more than 20 euros in the first half of 2008 and averaged 22 euros in the second half of 2008. In the first half of 2009, they fell to 13 euros. Since then, allowance prices have remained below 13 euros.

Is price volatility a problem? Critics of cap and trade point out that it is hard for producers to make sound investment decisions when the prices of allowances (and associated costs of production) fluctuate and are subject to uncertainty. Others claim that unstable allowance prices can produce macroeconomic disruptions. On the other hand, the ups and downs of allowance prices can play a beneficial countercyclical role. During economic downturns, the demand for allowances will fall, putting downward pressure on allowance prices. Lower allowance prices soften the impact of the pollution regulation on firms during the difficult economic times.

Reflecting the idea that significant swings in allowance prices should be avoided, policymakers have come up with ways to limit price volatility. One is to incorporate within the trading system an allowance price floor, price ceiling, or both. To impose a ceiling, the regulator can make available for sale additional allowances once the price reaches a given level. This prevents allowance prices from rising further. To enforce a price floor, the regulator buys allowances (and removes them from circulation) whenever the floor price is reached, thereby preventing prices from falling further.

The presence of a price ceiling implies that once the ceiling is reached, overall emissions no longer are constrained to the level of the original cap, because new allowances are being introduced to maintain the ceiling price. Thus, certainty about the total level of emissions is sacrificed for the sake of reduced uncertainty about allowance prices. Some interested parties have questioned whether this swap is worthwhile.

Another way to reduce potential price volatility is to allow for intertemporal banking and borrowing of allowances. With intertemporal borrowing, firms can credit toward present emissions the allowances allocated to them for future time periods. With intertemporal banking, firms can apply to future periods the allowances they do not use in the current period. Such intertemporal flexibility makes the current supply of allowances more elastic in any given period, which helps dampen price volatility. Of the major tradable allowance systems tried in the United States, RECLAIM offered the fewest opportunities for banking allowances. Stavins (2007) and Ellerman and Joskow (2008) suggest that much of the allowance price volatility experienced by RECLAIM was due to the absence of provisions for banking. Similarly, volatility in allowance prices for Phase I of Europe's Emissions Trading system has been attributed in part to the fact that the program prevented banking of allowances from the first phase to the second (Market Advisory Committee 2007; Schmalensee and Stavins, this issue).

In contrast, unlimited banking in the US Sulfur Dioxide Allowance Trading Program is generally viewed to have been a successful design feature of that program, as it mitigated issues of price volatility and led firms to reduce emissions faster than they would have without banking (Ellerman, Joskow, and Harrison 2003). Banking is also considered responsible for a large share of the gains from trade under the program.

That said, allowing intertemporal banking is not a panacea. Nordhaus (2007) finds that sulfur dioxide allowance prices between 1995 and 2006 were about as volatile as oil prices, and that they were much more volatile than prices of stocks, other assets such as houses, and most consumer goods. Sulfur dioxide allowance prices were particularly volatile in the late 2000s, as a series of court and regulatory decisions changed expectations about the future stringency of the cap (Schmalensee and Stavins, this issue; Palmer and Evans 2009; Bravender 2009).

5) Because of interactions with the fiscal system, certain decisions about the design of a cap-and-trade system—namely, the choice between auctioning and freely allocating allowances, and the way that any auction revenues are returned to the economy—significantly affect policy costs. Indeed, these decisions can determine whether a cap-and-trade program is more cost effective than some more conventional pollution control approaches.

The early assessments of cap and trade tended to be partial equilibrium in nature. Since the early 1990s, however, several studies have examined cap and trade (and other environmental policies) in a general equilibrium framework. These studies reveal that general equilibrium connections between cap and trade and the fiscal system have a first-order impact on the costs of cap and trade.

One of the key findings concerns the method of introducing emissions allowances into circulation. The regulating authority can give out all allowances free, auction them all out, or use a combination of free allocation and auctioning. A time-honored notion in economics is that while this choice affects the distribution of wealth, it does not affect cost-effectiveness because no matter how the allowances are initially distributed, the process of trading will assure that reductions in emissions happen in a cost-effective manner.

In a general equilibrium framework that accounts for interactions with the fiscal system, this logic no longer holds. By yielding government revenue, auctioning has the potential to reduce the government's reliance on distortionary taxes—such as income, sales, and payroll taxes—to finance its expenditures. The implied reductions (or avoided increases) in distortionary taxes can confer a benefit in terms of economic efficiency. In contrast, when allowances are given out free, the government does not receive these revenues, and society does not enjoy this potential benefit. The word "potential" is important here: if the revenues are recycled in ways that do not reduce marginal rates of prior taxes or that do not avoid increases in marginal rates of these taxes, this benefit is not realized.[7]

The potential benefits are substantial. Parry and Williams (2010) provide general formulas suggesting that auctioning can reduce the costs of meeting a given target for emissions reductions by almost half compared to a program with free permits. In a model focusing on the US economy in Goulder, Hafstead, and Dworsky (2010), we find that the costs of achieving a 42 percent reduction in carbon dioxide emissions under cap and trade are about 33 percent lower under 100 percent auctioning with

[7] While the choice between auctioning and free allocation has implications for cost-effectiveness, the choice about how to distribute the allowances *within a program involving free allocation* does not influence the cost-effectiveness of that program. This property was implicit in Coase (1960) and was first emphasized by Montgomery (1972).

recycling of revenues in the form of cuts in distortionary taxes as compared with 100 percent free allocation.

Historically, cap-and-trade policy has relied principally on free allocation. This is changing, however, especially for cap-and-trade programs aiming to cap greenhouse gas emissions. The European Union's Emissions Trading Scheme, the Regional Greenhouse Gas Initiative in the northeastern United States, and the State of California's new climate change policy all are moving toward auctioning more than half of their allowances. This change offers the potential for very large benefits in terms of economic efficiency, although the political motivation for these changes appears to have been a concern about distributional implications—the view that continued reliance on free allocation would generate windfalls to the recipient firms—as well as interest in obtaining funds to support various environmental programs. Economic analysis indicates that the concern about potential windfalls has merit. Studies of nitrogen oxide allowance trading under the US Clean Air Act (Bovenberg, Goulder, and Gurney 2005) and of potential carbon dioxide allowance trading in the United States (Bovenberg and Goulder 2001; Smith, Ross, and Montgomery 2002), suggest that the rents from 100 percent free allocation would substantially overcompensate firms for the costs they would otherwise face under these programs. In fact, these studies show that a fairly small share of the allowances—generally less than 30 percent—needs to be freely allocated to provide sufficient rents to prevent an overall decline in firm equity values.

In fact, the decision about whether to auction or freely allocate emissions allowances can determine whether a cap-and-trade program is more cost effective than certain more conventional regulatory alternatives. As we show in Parry, Williams, and Goulder (1999), to the extent that the cost of environmental policies are shifted forward to consumers (in the form of higher prices paid for pollution-intensive goods and services), the consumer price level will rise, implying a reduction in real factor returns. This depresses factor supply, and the resulting efficiency loss in factor markets (termed the "tax-interaction effect") raises the costs of environmental policies. In Goulder, Parry, Williams, and Burtraw (1999), we show that the tax-interaction effect is larger under emissions-pricing policies like cap and trade than for performance standards or technology mandates, which do not raise consumer prices as much. This potential disadvantage of cap and trade is overcome when cap and trade involves an auction and auction revenues are used to finance cuts in pre-existing distortionary taxes. In that case, cap and trade is more cost effective than these alternatives. But cap and trade can be more costly than the alternatives when allowances are given out free or when auction revenues are not used to finance cuts in prior tax rates.

Thus, the method of introducing allowances and the way that any revenues from the system are recycled importantly influence the cost-effectiveness of a cap-and-trade system. It can determine whether cap and trade is more or less cost effective than more conventional policy instruments. For cost-effectiveness, the design of a cap-and-trade system is of first-order importance.

These considerations do not contradict the idea that cap and trade generally has lowered the costs of pollution control. This is because cap and trade often has substituted for some of the more costly methods of control, such as fixed facility-level caps on emissions. But these broader concerns show that cap and trade needs to be carefully designed to assure lower costs than other regulatory alternatives. Auctioning and judicious revenue-recycling are needed to assure greater cost-effectiveness than some of the relatively flexible alternatives such as performance standards.

6) Should cap and trade displace other approaches?

Cap and trade cannot achieve all the efficiency-related goals of environmental policy. If the concern is economic efficiency, then in many settings it should complement, rather than substitute for, other instruments for environmental protection. The reason is that cap and trade cannot address all of the market failures responsible for pollution that is excessive from an efficiency point of view. And the same point applies to a pollution tax. As a form of emissions pricing, cap and trade addresses the market failure stemming from the emissions-related externality: it establishes a price for the otherwise external costs associated with pollution. But several other important market failures are not confronted by cap and trade (or by a pollution tax).

For example, an "innovation market failure" is associated with the spillover knowledge and the associated external benefits resulting from knowledge-generating activities. Additional measures—for example, a subsidy to research and development—are called for to confront this market failure directly. In its early history, some analysts touted cap and trade as the preferred instrument not only for encouraging conservation by consumers and substitution to cleaner known production processes by firms, but also for stimulating technological change—in particular, the invention of cleaner technologies. By raising the relative price of pollution-intensive production methods, cap and trade can provide incentives for innovation.[8] But efficiency calls for supplementing cap and trade with another instrument that directly addresses the innovation market failure. It is a common principle of policy analysis that multiple market failures generally call for multiple policy instruments.[9] Cap and trade is an excellent instrument for dealing with the externality associated with emissions, yet it should not displace other approaches that address other market failures.

But is cap and trade the best instrument for confronting the emissions externality? The main alternative is a pollution tax. A number of authors have analyzed the relative strengths and limitations of the cap-and-trade and pollution-tax options (for example, Metcalf 2007; Stavins 2007; Metcalf and Weisbach 2009; Goulder and Schein 2012). Although numerous issues are involved, perhaps the first point to emphasize is that both approaches offer similar advantages relative to conventional approaches for curbing emissions. Both approaches effectively impose, at the margin, a price

[8] However, as pointed out by Gans (2012), in general equilibrium, cap and trade (or, more generally, emissions pricing) can sometimes reduce incentives to innovate.

[9] For quantitative assessments of the significance of this principle in the context of environmental regulation, see Goulder and Schneider (1999), Fischer and Newell (2008), and Acemoglu, Aghion, Bursztyn, and Hemous (2011).

for each unit of emissions. This is the case for cap and trade even when allowances are initially given out free to the covered entities. After all, even when allowances are received for free, each additional unit of emissions carries an opportunity cost: one more unit of pollution either reduces the number of allowances the firm can sell, or it raises the number of allowances the firm will need to buy to remain in compliance. By establishing one price for pollution that facilities must face, both approaches encourage equality of abatement costs at the margin across facilities, which works toward cost-effectiveness.

Moreover, there is no inherent difference between the two approaches in terms of the distributional impacts on facilities. Under a cap-and-trade system, free allocation of allowances can cushion the impact of the regulation on covered firms, shifting the burden onto the general public (since more free allocation implies less revenue collected by the auction). Under a pollution tax, offering inframarginal exemptions to the tax yields the same opportunities for altering the distribution of impacts.

The two approaches do differ in some important ways, however. A pollution tax avoids the problem of emissions price volatility. On the other hand, the pollution tax does not impose a predetermined cap on aggregate emissions; some would regard this as a disadvantage.

It has often been suggested that a cap-and-trade system would be more costly to administer than a pollution tax. One claim is that administrative costs are higher because a cap-and-trade program would involve more entities whose emissions must be tracked. This claim is incorrect. The number of covered entities depends on where the cap-and-trade system or pollution tax is imposed—upstream, midstream, or downstream—and both approaches can be introduced at any of these levels. Still, recent experience suggests that a cap-and-trade system might involve somewhat greater administrative challenges for two reasons: 1) there are costs of setting up a market for auctioning and trading allowances (which may be higher than the costs of incorporating a pollution tax within the existing tax-collection institutions), and 2) under a cap-and-trade system, the regulator must not only keep track of the emissions of covered facilities, but also establish a registry to record changes in ownership of allowances as a result of allowance purchases or sales.

At the same time, current policy conditions and political economy considerations might favor cap and trade, at least in the climate policy context. Given the existence of other cap-and-trade systems overseas, it might be easier to achieve international harmonization through a US cap-and-trade program than with a US carbon tax (Jaffe, Ranson, and Stavins 2010; Metcalf and Weisbach 2009). Cap and trade has been an easier political sell than a pollution tax, partly because cap and trade is less costly to the covered firms than a pollution tax would be.[10] It is also partly because the public, often averse to any new tax, has tended to view a cap-and-trade program as something very different from a tax measure. However, this political advantage

[10] This statement assumes that the pollution tax policy does not include inframarginal exemptions. Such exemptions would function much like free allowances under a cap-and-trade system, lowering the costs to the covered firms.

seems to be waning, at least in the United States, where opponents of cap-and-trade policies for limiting carbon emissions have started to refer to them as "cap and tax" policies (for example, "The Cap and Tax Fiction," in the *Wall Street Journal* 2009).

The bottom line is that neither a pollution tax nor a cap-and-trade approach clearly dominates. The degree of efficiency in reducing emissions seems to depend more on the extent of emissions pricing (under either form) and on the particular design of the emissions-pricing instrument (for example, the degree to which a cap-and-trade program relies on auctioning of allowances).

Conclusions

Trading rights to pollute—which was just an idea in the minds of a few economists 45 years ago—has now taken form in many locales and for many types of pollution. This novel approach has largely lived up to its basic promises: that is, in most places where it has been tried, it has succeeded in bringing down pollution to the targeted levels and has achieved those emissions reductions at lower cost than would have been possible under many of the more conventional forms of regulation. At national and subnational levels, the environmental targets have largely been met under cap-and-trade systems for local pollutants including sulfur dioxide and nitrogen oxide compounds, as well as for carbon dioxide, the principal greenhouse gas.

Important challenges remain, however. The application of cap and trade for control of water pollution has been limited by difficulties of tracking the nonpoint sources, particularly the water pollution generated by the agricultural sector. The international-level use of cap and trade to limit greenhouse gas emissions has been limited by difficulties in enforcement.

We have reached a much deeper understanding of the potential environmental and economic impacts of cap and trade. Research reveals how the simple textbook version of cap-and-trade system can be modified to address potential difficulties such as the problem of price volatility. It also makes clear how the impacts of cap and trade depend on interactions with other regulations and with the existing tax system. These interactions are of first-order importance: they influence whether cap and trade manages to reduce pollution, and they indicate that the particular design of a cap-and-trade system makes a substantial difference to its cost. Indeed, the design can determine whether the program yields efficiency gains.

Cap and trade has some advantages and some drawbacks relative to the chief alternative form of emissions pricing—a pollution tax. Neither approach dominates the other. When well designed, either form of emissions pricing will offer several advantages over conventional forms of regulation. Yet neither cap and trade nor a pollution tax is a cure-all for environmental problems: emissions pricing does not eliminate the need to engage other environmental policy instruments to address environment-related market failures other than the one stemming from the emissions externality.

■ *I am very grateful to Dallas Burtraw, Steven Cliff, Sheila Olmstead, Christian de Perthuis, Robert Stavins, Roberton Williams III, and the editors of this journal for helpful comments and suggestions, and to Santiago Saavedra Pineda for excellent research assistance.*

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous. 2011. "The Environment and Directed Technical Change." *American Economic Review* 102(1): 131–66.

Arnason, Ragnar. 2012. "Property Rights in Fisheries: How Much Can Individual Transferable Quotas Accomplish?" *Review of Environmental Economics and Policy* 6(2): 217–36.

Bovenberg, A. Lans, and Lawrence H. Goulder. 2001. "Neutralizing the Adverse Industry Impacts of CO_2 Abatement Policies: What Does It Cost?" In *Behavioral and Distributional Effects of Environmental Policy*, edited by C. Carraro and G. Metcalf, 45–85. University of Chicago Press.

Bovenberg, A. Lans, Lawrence H. Goulder, and Derek J. Gurney. 2005. "Efficiency Costs of Meeting Industry-Distributional Constraints under Environmental Permits and Taxes." *RAND Journal of Economics* 36(4): 950–70.

Bravender, Robin. 2009. "Acid Rain Credits Nosedive on CAIR Concerns." *Greenwire*, March 27. http://www.eenews.net/public/Greenwire/2009/03/27/4.

Burtraw, Dallas. Forthcoming. "The Institutional Blind Spot in Environmental Economics." *Daedalus*.

Burtraw, Dallas, and William Shobe. Forthcoming. "Rethinking Environmental Federalism in a Warming World." *Climate Change Economics*.

CDC Climate Research. 2011. "Carbon Price Flaw? The Impact of the UK's CO_2 Price Support on the EU ETS." Climate Brief No. 6, June.

Chan, Gabriel, Robert N. Stavins, Robert Stowe, and Richard Sweeney. 2012. "The SO_2 Allowance Trading System and the Clean Air Act Amendments of 1990: Reflections on 20 Years of Policy Innovation." *National Tax Journal* 65(2): 419–52.

Coase, Ronald H. 1960. "The Problem of Social Cost." *Journal of Law and Economics* 3(October): 31–44.

Crocker, Thomas D. 1966. "The Structure of Atmospheric Pollution Control Systems." In *The Economics of Air Pollution*, edited by H. Wolozin, 61–86. New York: W. W. Norton.

Crocker, Thomas D. 2005. "Markets for Conserving Biodiversity Habitat: Principles and Practice." Chap. 7 in *Species at Risk: Using Economic Incentives to Shelter Endangered Species on Private Lands*, edited by Jason F. Shogren. Austin: University of Texas Press.

Dales, John H. 1968. *Pollution, Property and Prices: An Essay in Policy-making and Economics*. University of Toronto Press.

Ellerman, A. Denny., Frank J. Convery, and Christian de Perthuis. 2010. P*ricing Carbon: The European Union Emissions Trading Scheme*, especially chap. 6, "Emissions Abatement." Cambridge University Press.

Ellerman, A. Denny, and Paul L. Joskow. 2008. "The European Union's Emissions Trading System in Perspective." Pew Center on Global Climate Change.

Ellerman, A. Denny, Paul L. Joskow, and David L. Harrison, Jr. 2003. *Emissions Trading: Experience, Lessons, and Considerations for Greenhouse Gases*. Washington, DC: Pew Center on Global Climate Change.

Fankhauser, Samuel, Cameron Hepburn, and Jisung Park. 2010. "Combining Multiple Climate Policy Instruments: How Not to Do It." *Climate Change Economics* 1(3): 1–17.

Fischer, Carolyn, and Richard G. Newell. 2008. "Environmental and Technology Policies for Climate Mitigation." *Journal of Environmental Economics and Management* 55(2): 142–62.

Foster, Vivien, and Robert W. Hahn. 1995. "Designing More Efficient Markets: Lessons from Los Angeles Smog Control." *Journal of Law and Economics* 38(1): 19–48.

Gans, Joshua. 2012. "Innovation and Climate Change Policy." Unpublished Paper.

Goulder, Lawrence H., Marc A. C. Hafstead, and Michael Dworsky. 2010. "Impacts of Alternative Emissions Allowance Allocation Methods under a Federal Cap-and-Trade Program." *Journal of Environmental Economics and Management* 60(3): 161–81.

Goulder, Lawrence H., Marc R. Jacobsen,

and Arthur van Benthem. 2012. "Unintended Consequences from Nested State and Federal Environmental Regulation: The Case of the Pavley Greenhouse-Gas-Per-Mile Limits." *Journal of Environmental Economics and Management* 63(2): 187–207.

Goulder, Lawrence H., Ian W. H. Parry, Roberton C. Williams III, and Dallas Burtraw. 1999. "The Cost-Effectiveness of Alternative Instruments for Environmental Protection in a Second-Best Setting." *Journal of Public Economics* 72(3): 329–60.

Goulder, Lawrence H., and Andrew Schein. 2012. "Carbon Taxes vs. Cap and Trade." Unpublished paper, Stanford University.

Goulder, Lawrence H., and Stephen H. Schneider. 1999. "Induced Technological Change and the Attractiveness of CO_2 Emissions Abatement Policies." *Resource and Energy Economics* 21(3–4): 211–53.

Goulder, Lawrence H., and Robert N. Stavins. 2012. "Interactions between State and Federal Climate Change Policies." Chap. 7 in *The Design and Implementation of U.S. Climate Policy*, edited by Don Fullerton and Catherine Wolfram. University of Chicago Press.

Hahn, Robert W., and Gordon Hester. 1989. "Where Did All the Markets Go? An Analysis of EPA's Emissions Trading Program." *Yale Journal on Regulation* 6(1): 109–153.

Jaffe, Judson, Matthew Ranson, and Robert N. Stavins. 2010. "Linking Tradable Permit Systems: A Key Element of Emerging International Climate Policy Architecture." *Ecology Law Quarterly* 36(4): 789–808.

Kolstad, Jonathan, and Frank Wolak. 2003. "Using Enivronmental Emissions Permit Prices to Raise Electricity Prices: Evidence from the California Electricity Market." CSEM Working Paper 113, Center for the Study of Energy Markets, University of California, Berkeley.

Levinsohn, Arik. 2012. "Belts and Suspenders: Interactions among Climate Policy Regulations." Chap 8 in *The Design & Implementation of U.S. Climate Policy*, edited by D. Fullerton and C. Wolfram. University of Chicago Press.

Market Advisory Committee to the California Air Resources Board. 2007. *Recommendations for Designing a Greenhouse Gas Cap-and-Trade System for California.* Report to the California Air Resources Board.

Metcalf, Gilbert E. 2007. "A Proposal for a U.S. Carbon Tax Swap: An Equitable Tax Reform to Address Global Climate Change." Discussion Paper 2007-12, The Hamilton Project.

Metcalf, Gilbert E., and David Weisbach. 2009. "The Design of a Carbon Tax." *Harvard Environmental Law Review* 33(2): 499–556.

Millard-Ball, Adam. Forthcoming. "The Trouble with Voluntary Emissions Trading." *Journal of Environmental Economics and Management.*

Montgomery, W. David. 1972. "Markets in Licenses and Efficient Pollution Control Programs." *Journal of Economic Theory* 5(3): 395–418.

Nordhaus, William D. 2007. "To Tax or Not to Tax: Alternative Approaches to Slowing Global Warming." *Review of Environmental Economics and Policy* 1(1): 26–44.

Palmer, Karen, and David A. Evans. 2009. "The Evolving SO_2 Allowance Market: Title IV, Cair, and Beyond." Resources for the Future Weekly Policy Commentary, July 13. www.rff.org/Publications/WPC/Pages/090713-Evolving-SO2-Allowance-Market.aspx.

Parry, Ian W. H., and Roberton C. Williams III. 2010. "What Are the Costs of Meeting Distributional Objectives for Climate Policy?" *B. E. Journal of Economic Analysis & Policy* 10(2).

Parry, Ian W. H., Roberton C. Williams III, and Lawrence H. Goulder. 1999. "When Can Carbon Abatement Policies Increase Welfare? The Fundamental Role of Distorted Factor Markets." *Journal of Environmental Economics and Management* 37(1): 52–84.

Smith, Anne E., Martin E. Ross, and W. David Montgomery. 2002. "Implications of Trading Implementation Design for Equity–Efficiency Tradeoffs in Carbon Permit Allocations." Unpublished paper, Charles River Associates.

Stavins, Robert N. 1995. "Transactions Costs and Tradeable Permits." *Journal of Environmental Economics and Management* 29(2): 133–47.

Stavins, Robert N. 2007. "A US Cap-and-Trade System to Address Global Climate Change." Discussion Paper 2007-13, The Hamilton Project.

Tietenberg, Thomas H. 1985. *Emissions Trading: An Exercise in Reforming Pollution Policy.* Washington, DC: Resources for the Future.

Wall Street Journal. 2009. "The Cap and Tax Fiction." June 26, A12.

Weitzman, Martin L. 1974. "Prices *vs.* Quantities." *Review of Economic Studies* 41(4): 477–91.