AEA Papers and Proceedings 2020, 110: 113–118
https://doi.org/10.1257/pandp.20201083

The Cost-Effectiveness Implications of Carbon Price Certainty†

By Joseph E. Aldy and Sarah Armitage*

One hundred years ago, Professor Arthur C. Pigou published his key insight that taxing a negative technological externality would improve social welfare. This has motivated research and policy efforts in pricing pollution—directly through emissions taxes and implicitly through pollution markets such as cap-and-trade programs and tradable performance standards (Aldy et al. 2010). To combat climate change, national and subnational governments around the world have implemented pollution markets (primarily cap-and-trade programs) covering about 15 percent of global carbon dioxide (CO_2) emissions and carbon taxes on another 5 percent of these emissions (World Bank and Ecofys 2018).

Given uncertainty over the costs and efficacy of CO_2 emissions abatement investment, the choice and design of policy instruments can have significant welfare consequences (Weitzman 1974, Stavins 2019). The extensive "prices versus quantities" literature building on Weitzman (1974) has typically focused on the problem confronting the regulator: how to choose between pricing pollution with a tax or setting a quantity target (implemented through a pollution market) when the regulator has incomplete information about firms' pollution abatement opportunities. We focus on the problem facing firms covered by an emissions mitigation policy: how to choose abatement investments in the presence of uncertainty in pollution markets, in contrast to the price certainty under a carbon tax. How firms resolve this problem can in turn influence and inform the regulator's consideration of its problem.

In the context of climate policy, firms make investment decisions subject to two types of policy uncertainty. First, firms operate under uncertainty over policy choice. Will the regulator choose a tax or a quantity-based instrument?[1] This is similar to other economic contexts in which firm investment decisions are influenced by policy uncertainty, such as future monetary policy or tax rates (Bernanke 1983; Baker, Bloom, and Davis 2016). However, once this policy choice uncertainty has been resolved, firms may face residual uncertainty inherent to the selected instrument. This second type of policy uncertainty is significantly more pronounced under quantity-based approaches, such as cap-and-trade, than under an emissions tax.

I. Firm Uncertainty and Forecast Errors in Pollution Markets

To understand the impact of policy uncertainty under an emissions tax versus a cap-and-trade instrument, consider the problem a firm faces in complying with either policy. Under a tax, the firm learns the tax rate from the regulator, determines the marginal abatement costs of its emissions mitigation options, and then makes investments to minimize the sum of its tax bill and abatement costs. The profit-maximizing firm would equate the cost of its marginal investment in pollution abatement with the tax, and in its doing so, the aggregate social costs of reducing emissions are minimized.

---

*Aldy: Harvard Kennedy School, Resources for the Future, NBER, and Center for Strategic and International Studies (email: joseph_aldy@hks.harvard.edu); Armitage: Department of Economics, Harvard University (email: saraharmitage@g.harvard.edu). Ken Norris provided excellent research assistance. We thank Alex Barron, Ed Glaeser, Jerry Green, Sam Hanson, Garth Heutel, Joshua Pfeffer, Robert Pindyck, Jim Stock, and participants at the AEA Annual Conference, AERE Summer Conference, Northeast Workshop on Energy Policy and Environmental Economics, and Harvard Kennedy School energy policy seminar for helpful suggestions and information. Aldy acknowledges financial support from Resources for the Future, BP, and the Taubman Center for State and Local Government; Armitage acknowledges financial support from the Norberg–Bohm Fellowship and the Crump Fellowship.

† Go to https://doi.org/10.1257/pandp.20201083 to visit the article page for additional materials and author disclosure statement(s).

[1] This binary framing abstracts from real-world policy decisions that include hybrid price–quantity approaches and pairing a pollution pricing instrument with technology subsidies, regulatory mandates, etc.

Under cap-and-trade, the firm learns the aggregate emissions cap and allowance allocation from the regulator and determines the marginal costs of its abatement options and must then form an expectation over allowance prices that will emerge in the secondary allowance market. Given these forecast market prices, the firm then makes investments to minimize the sum of expected allowance purchases and abatement costs. *If* every firm forms identical allowance price expectations that equal the realized price, then cap-and-trade can replicate the cost effectiveness of the emissions tax. The more likely outcome, however, is that uncertainty over expected allowance prices could result in variation in firm expectations. This is the inherent policy instrument uncertainty of cap-and-trade, which persists even after questions around policy design and stringency have been resolved.

Figure 1 illustrates how such inherent policy uncertainty could increase the cost of abating emissions. Suppose that there are N firms that have identical marginal abatement cost functions and identical levels of emissions absent investment in emissions abatement, e_0. The horizontal axis of Figure 1 presents firm-level quantity of emissions abatement, where e_0 corresponds to the origin. The aggregate emissions level of these N firms is $E_0 = N \cdot e_0$. The socially optimal policy would have each firm reduce its emissions to e^*, reflecting the quantity of emissions abatement $q_A^* = e_0 - e^*$ that equates marginal abatement costs and marginal benefits. If the regulator knew each firm's abatement cost function, then the regulator could simply assign e^* to each firm. Since this information is generally private, however, the regulator has two policy options: (i) set a tax on each unit of emissions, τ^*, equal to the marginal benefits of reducing emissions, or (ii) set an aggregate emissions cap E^*, which would require aggregate emissions abatement of $Q_A^* = E_0 - E^*$, and implement it through a cap-and-trade program.

In the tax case, so long as the regulator has an appropriately estimated social cost of carbon, this approach will deliver both cost-effective emissions abatement (q_A^*) and the socially optimal outcome. The per-firm resource cost under this policy is represented by the area $A + B$ in the figure, and the social resource costs across N firms would be $N \cdot (A + B)$.

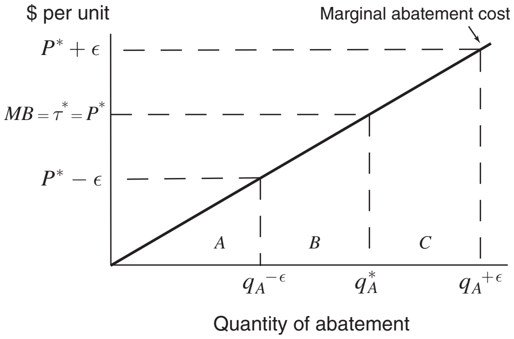

FIGURE 1. INVESTMENT ERRORS: TAX VERSUS CAP-AND-TRADE

Under an emissions cap, the regulator creates emissions allowances (the sum of which equals the cap) granting the holder the right to emit one unit of emissions and allocates these allowances to firms. Firms then trade allowances on a secondary market so that firms hold sufficient allowances to cover their emissions. Analogous to the firm's problem under a tax, the firm under cap-and-trade must identify the least-cost combination of emissions abatement and net allowance purchases. If every firm in this pollution market expects the allowance price to be P^*, then each undertakes the socially optimal level of abatement q_A^*.

However, firms may err in their expectations of allowance prices. Suppose that one-half of the firms expect allowance prices to be high, say $P^* + \epsilon$, and undertake investment in emissions abatement capital with marginal costs consistent with this expectation, but the other one-half of the firms expect lower allowance prices, say $P^* - \epsilon$, and undertake investment in abatement capital consistent with this lower price. In this case, cap-and-trade does not realize a cost-effective outcome. One-half of the firms will each bear resource costs represented by the area $A + B + C$, and the other one-half will each bear resource costs represented by the area A. The social resource costs would be $(N/2) \cdot (C - B)$ higher under this scenario than under the carbon tax; this increase in social resource costs reflects both the size of the forecast errors (the ϵ deviations from P^*) and the slope of the marginal abatement cost function. Note that any allowance price clearing in this market between $P^* - \epsilon$ and $P^* + \epsilon$ would be a Nash equilibrium. In this highly stylized illustration, it is conceivable that a cap-and-trade program could include firm-level forecast errors at the abatement investment stage, resulting in a market-clearing price that is equal to the optimal tax but incurs higher abatement costs than under the tax.

A key assumption in this graphical analysis is that the investment in pollution abatement includes substantial fixed costs that are irreversible. Firms cannot instantaneously adjust their pollution abatement—and hence marginal costs and net purchase of emissions allowances—in response to realized allowance prices in a pollution market. The types of investment necessary to reduce CO_2 emissions—building wind farms, retrofitting existing building shells, installing carbon capture and storage equipment with fossil fuel power plants, constructing biorefineries, and others—are long lived and irreversible. In practice, firms making irreversible emissions abatement investments may fail to minimize compliance costs under cap-and-trade.

II. Cost-Effectiveness Anomalies in Real-World Pollution Markets

Since the 1980s, policymakers have implemented numerous quantity-oriented policies to address environmental and energy objectives. In addition to cap-and-trade markets for CO_2, sulfur dioxide (SO_2), and nitrogen oxides (NO_x), policymakers have designed tradable performance standards to reduce lead in gasoline, promote fuel economy among vehicle manufacturers, increase the renewable share of power generation, and raise the biofuel share of transportation fuels.[2] Indeed, virtually all transportation fuel and electricity consumption in the United States today is subject to at least one cap-and-trade program or tradable performance standard.

The theoretical appeal of cap-and-trade programs and tradable performance standards lies in the potential for the market to allocate effort in a cost-effective manner, just as in any other efficient market. Montgomery (1972) formally shows how firms operating under a cap-and-trade program each have an incentive to equate their marginal abatement costs with the allowance price and, as a result, marginal abatement costs are equalized among all firms in the market. Complementing this static cost effectiveness across firms, Rubin (1996) demonstrates the potential for dynamic cost effectiveness in cap-and-trade programs that permit intertemporal trading (banking allowances for future compliance purposes, or borrowing future vintage allowances for contemporary compliance purposes). In practice, however, firm behavior in a variety of cap-and-trade programs and tradable performance standards has deviated from the theoretical conditions for cost effectiveness.

Consider the California Low Carbon Fuel Standard, which requires refineries to satisfy a performance benchmark based on the carbon content of transportation fuels and allows refineries to buy and sell compliance credits. Since 2016, 83 percent of trading days have had multiple transactions in which the state of California reported credit prices. On only 15 multitransaction trading days did the within-day transactions share a common credit price. The within-trading-day credit price standard deviation averaged about \$10 per ton of CO_2. The maximum credit price exceeded the minimum credit price by more than 20 percent, on average, within multitransaction trading days. There were as many credit days in which the maximum price paid was double the minimum price paid as there were days with a common price across transactions. If buying firms are equating their marginal costs of compliance with the price paid for credits, then firms are not equating their marginal costs of compliance in this market.

Dynamic cost effectiveness for cap-and-trade programs with intertemporal trading calls for allowance prices to increase at the rate of interest, following a Hotelling-style result. As Figure 2 illustrates for the US SO_2, EU Emissions Trading System (ETS) CO_2, and Regional Greenhouse Gas Initiative (RGGI) CO_2 markets, prices are volatile, reveal occasional spikes and troughs, and do not follow a price path that increases with the interest rate. This extreme volatility suggests that uncertainty about future prices may be substantial.

Price volatility comes from many sources, which is part of the reason why firms have difficulty predicting realized prices. For example, energy market demand shocks can influence allowance prices. The Southern California RECLAIM NO_x market experienced a 100-fold increase in allowance prices during the

---

[2] Tradable performance standards establish a quantitative benchmark that firms must meet. If a firm beats the benchmark, its overcompliance generates a credit that may be traded to another firm, such as one that fails to meet the benchmark.

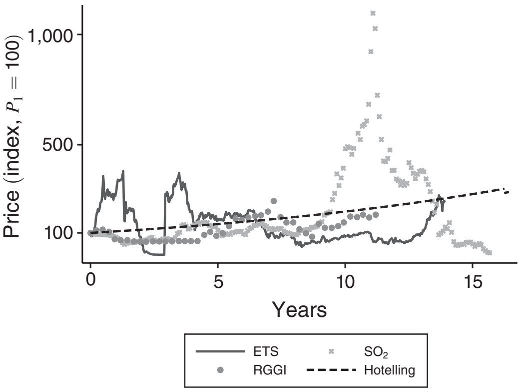

Figure 2. Allowance Prices over Time, Various Pollution Markets

Notes: This figure depicts nominal allowance prices—indexed to 100 for the initial date in each series—for the SO_2 cap-and-trade program, the RGGI CO_2 cap-and-trade program, and the EU ETS CO_2 cap-and-trade program. It also illustrates a Hotelling price trajectory consistent with a 7 percent annual growth rate.

Source: Cantor Fitzgerald SO_2 Monthly Price Index, ICE EU Allowance Futures daily prices, RGGI quarterly auction results

2000–2001 California electricity crisis (Fowlie, Holland, and Mansur 2012). In contrast, slow economic growth over 2012–2013 contributed to low EU ETS allowance prices. However, the observed volatility in allowance markets is not simply a function of the volatility of the underlying energy commodities. The allowance price volatilities observed under the EU ETS and SO_2 programs exceed the volatilities of oil and natural gas futures prices, as well as of S&P 500 index prices, over comparable periods.

The deviations from cost effectiveness under cap-and-trade may also reflect the impact of overlapping policies. If the government imposes a new mandate—such as a renewable power mandate or a performance standard—on top of a cap-and-trade program, then that additional policy influences firms covered by the initial program in two ways. First, the new mandate restricts the investment options under cap-and-trade, which would increase the costs of reducing emissions without changing the level of emissions, so long as the cap in cap-and-trade remains binding (Goulder and Stavins 2011). Second, compliance with the mandate effectively reduces scarcity in the allowance market and thus depresses the equilibrium allowance price. This latter effect does not occur under an emissions tax. If some firms do not anticipate the overlapping mandate—or incorrectly anticipate its impact on the pollution market—then their allowance price forecast errors will result in higher-than-necessary costs for compliance.

III. US CO_2 Tax and Cap-and-Trade Simulations

To illustrate the potential resource costs of forecast errors under cap-and-trade relative to a carbon tax, we undertake simulations of an economy-wide carbon pricing policy for the United States. In these simulations, we consider the abatement decisions of a representative firm. In contrast to many existing models of firm behavior in allowance trading markets, we model abatement as a dynamic, long-lived investment rather than a variable input. Accumulated abatement capital persists across compliance periods, subject to a depreciation rate.

Our simulations also incorporate the key no-arbitrage condition for allowance markets with intertemporal trading, which requires expected future allowance prices to rise at the rate of interest following a Hotelling rule. However, we introduce the possibility that shocks to price *levels* will cause the expected price path to jump from one Hotelling trajectory to the next. As discussed above, this volatility better reflects the evolution of allowance prices in real-world allowance trading markets (Figure 2). In our first set of simulations, the representative firm makes abatement investment decisions given these stochastic allowance prices, where the firm learns period *t*'s allowance price only in period *t* as shocks are realized in each period.

In our second set of simulations, we model a tax that increases at the rate of interest. Firms again make abatement investment decisions in response to this carbon price path, which is known ex ante with certainty. To compare the total resource costs required to achieve a given level of emissions reductions under the stochastic allowance price scenario versus the smoothly increasing carbon tax scenario, we conduct a series of carbon tax simulations that vary over the initial carbon price and then map these to stochastic allowance price simulations that realize the same quantity of emissions reductions.

We calibrate a quadratic abatement cost function with modeling results from a recent Stanford Energy Modeling Forum study of US carbon taxes (Barron et al. 2018). To model the allowance price trajectory, we assume that prices follow geometric Brownian motion and estimate drift and volatility parameters by using historical allowance price data from the EU ETS over 2008–2018. Our estimated drift and volatility parameters correspond to 5.2 percent expected annual real price growth and 42.9 percent annual volatility. Full details of the model calibration and simulation procedure are provided in the online Appendix.

Our simulation results illustrate that the presence of price uncertainty and forecast errors weakly increases the resource costs of achieving a given level of emissions reduction. Figure 3 depicts total emissions reductions associated with ten years of abatement investment on the horizontal axis and the corresponding total abatement investment costs on the vertical axis. The effective abatement cost function—that is, the total resource costs required to achieve some level of emissions reduction over time—is shifted upward for the stochastic price scenarios relative to the tax scenarios. Given our model parameterization, we find that the median percentage increase in abatement investment costs under stochastic allowance prices is 21 percent.

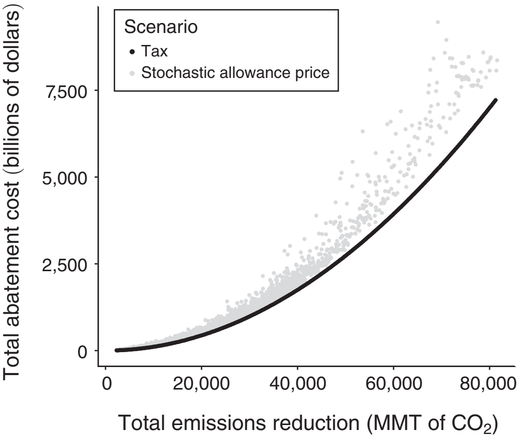

Figure 3. Total Emissions Reduction from Ten Years of Abatement Investment

Notes: Refer to the online Appendix for details on model calibration and simulations. MMT refers to million metric tons.

These simulation results help to illustrate that firm-level price forecast errors create an additional asymmetry between price and quantity instruments, in addition to the standard considerations around the relative slopes of marginal cost and marginal benefit functions that arise from the regulator's uncertainty. The presence of forecast errors drives a wedge between the effective abatement cost functions under the two types of policies, creating additional cost inefficiencies for cap-and-trade programs and tradable performance standards.

IV. Conclusion

While pollution markets have enabled the attainment of important environmental and energy goals, their inherent uncertainty increases the likelihood that firms err in forecasting future allowance (or credit) prices and thus undertake investment that does not equate compliance costs across firms. Failing to address this inherent uncertainty in future climate policy means that carbon pricing through cap-and-trade will likely be less cost effective than carbon pricing through a carbon tax.

Given the scope of the climate challenge—the need to focus trillions of dollars of investment to transform the energy foundation of the modern economy—minimizing costs will be critical from a welfare and political viability perspective. Our findings can inform both the choice of instrument (tax versus quantity-based approaches) and its design (such as hybrid tax–quantity approaches that reduce the inherent uncertainty faced by firms) to promote cost-effective efforts to mitigate climate change risks.

REFERENCES

Aldy, Joseph E., Alan J. Krupnick, Richard G. Newell, Ian W. H. Parry, and William A. Pizer. 2010. "Designing Climate Mitigation Policy." *Journal of Economic Literature* 48 (4): 903–34.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–1636.

Barron, Alexander R., Allen A. Fawcett, Marc A. C. Hafstead, James R. McFarland, and Adele C. Morris. 2018. "Policy Insights from the EMF 32 Study on U.S. Carbon Tax Scenarios."

Climate Change Economics 9 (1): 1–47.

Bernanke, Ben S. 1983. “Irreversibility, Uncertainty, and Cyclical Investment.” *Quarterly Journal of Economics* 98 (1): 85–106.

Fowlie, Meredith, Stephen P. Holland, and Erin T. Mansur. 2012. “What Do Emissions Markets Deliver and to Whom? Evidence from Southern California’s NOx Trading Program.” *American Economic Review* 102 (2): 965–93.

Goulder, Lawrence H., and Robert N. Stavins. 2011. “Challenges from State-Federal Interactions in US Climate Change Policy.” *American Economic Review* 101 (3): 253–57.

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

Rubin, Jonathan D. 1996. “A Model of Intertemporal Emission Trading, Banking, and Borrowing.” *Journal of Environmental Economics and Management* 31 (3): 269–86.

Stavins, Robert. 2019. “The Future of U.S. Carbon-Pricing Policy.” NBER Working Paper 25912.

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

World Bank and Ecofys. 2018. *State and Trends of Carbon Pricing 2018*. Washington, DC: World Bank Group.