DISCUSSION PAPER

October 2016 ■ RFF DP 16-43

Adding Quantity Certainty to a Carbon Tax

The Role of a Tax Adjustment Mechanism for Policy Pre-Commitment

---

Marc Hafstead, Gilbert E. Metcalf, and
Roberton C. Williams III

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

RESOURCES FOR THE FUTURE

Adding Quantity Certainty to a Carbon Tax: The Role of a Tax Adjustment Mechanism for Policy Pre-Commitment

Marc Hafstead, Gilbert E. Metcalf, and Roberton C. Williams III

Abstract

A concern often raised about a carbon tax is that it does not provide any certainty as to the quantity of emission reductions achieved under the policy. We explore in this Issue Brief how greater emission reduction certainty can be built into a carbon tax. We first define a Tax Adjustment Mechanism for Policy Pre-Commitment (TAMPP). A TAMPP is an adjustment mechanism for the tax rate of a carbon tax to ensure that targeted emission reduction milestones are met over the next few decades. We then provide some guidance based on economic principles related to various design considerations that should be incorporated in a cost-effective and politically realistic TAMPP.

Key Words: carbon tax, emissions uncertainty, climate

JEL Classification Numbers: Q58, H23,Q50

Contents

Adding Quantity Certainty to a Carbon Tax: The Role of a Tax Adjustment Mechanism for Policy Pre-Commitment

Marc Hafstead, Gilbert E. Metcalf, and Roberton C. Williams III*

I. Introduction

In the absence of policy, the social cost of burning fossil fuels exceeds the private cost, because of a “negative externality”: the harm caused by greenhouse gas (GHG) emissions. As a result, greenhouse gas emissions are higher than is socially optimal. Economists agree that incentive-based policy instruments – those that use market forces to influence behavior – can be dramatically less expensive than other regulatory approaches to bring about a given reduction in emissions. This is the concept of "cost effectiveness." Cap-and-trade programs and carbon taxes are the two leading examples of incentive-based policy. These policies put a price on emissions thereby raising the private cost to better reflect social cost. The higher price of fossil fuels induces substitution away from these fuels and thereby a reduction in emissions.

Carbon taxes directly set a price on emissions. Cap and trade programs, in contrast, set a cap on aggregate emissions and then allow market forces to set a price on emissions, either directly through government auctioning of allowances and/or through trades of allowances in organized or informal markets. Allowances are retired as firms burn fossil fuels and release emissions. Just as a tax raises the cost of burning fossil fuels, retiring allowances has an opportunity cost that makes burning fossil fuels more costly.[1]

Where these two market mechanisms differ is in what they control. While obvious, it is worth stating explicitly: a carbon tax sets the price on emissions and market forces determine the quantity of emissions (at the level that equilibrates supply and demand). A cap and trade system, in contrast, sets an aggregate limit on emissions and market forces determine the market clearing price for trading in emissions up to that cap. In a world without any uncertainty, the two

* Hafstead, Resources for the Future; Metcalf, Tufts University and NBER; Williams, University of Maryland, Resources for the Future, and NBER. We wish to thank Ray Kopp for helpful input on earlier drafts, and RFF’s Center for Energy and Climate Economics for financial support for this work.

[1] A carbon tax is an example of Pigouvian pricing, where a tax is levied on pollution causing activities equal to their social marginal damages (Pigou (1932)). In the absence of other distortions, this is socially efficient. Dales (1968) is credited with developing the idea of cap and trade as an alternative to Pigouvian pricing.

approaches lead to equivalent outcomes: the tax rate under the carbon tax equals the equilibrium price on allowances under cap-and-trade, and both produce the same quantity of emissions. However, in a world with uncertainty (such as macroeconomic and cost shocks), setting tax rates or aggregate caps that ex ante lead to the same economic outcomes can lead to very different ex post outcomes.[2]

A sizable economics literature asks under what circumstances one instrument is more efficient than the other in a world with uncertainty. Here, we focus on a different question: how could one design a carbon tax that can provide a level of certainty with respect to quantities as well as price? There are political reasons that quantity certainty is valuable. For example, prevailing climate goals are stated in terms of emission quantities. In 2009 President Obama pledged to reduce US greenhouse emissions to 17% below 2005 levels by 2020[3], then last year pledged to reduce emissions 26-28 percent below 2005 levels by 2025[4].

Reducing uncertainty about emissions reductions also has the potential to improve the efficiency of the carbon tax. Roberts and Spence (1976) first examined hybrid policy instruments – those that combine elements of an emissions tax and a quantity-based policy (such as a cap-and-trade system), and thus provide more emissions certainty than a tax and more price certainty than cap-and-trade – and showed that well-designed hybrid policies are generally more efficient than either pure tax or pure quantity-based policy. But that issue is beyond the scope of this paper.

To be clear at the outset, our focus in this paper is a narrow one. It simply asks how one could design a carbon tax with a mechanism to reduce uncertainty about future emissions, and what tradeoffs different design elements might entail. We do not evaluate whether there is an economic efficiency argument to be made for such a mechanism; we leave that for future research. The normative question of whether a carbon tax *should* include such a mechanism is even further beyond the scope of the paper. If the politics of climate policy are such, however, that adding some mechanism to a carbon tax to provide greater emissions certainty facilitates passage of a carbon tax, then it is worthwhile to examine how that mechanism might be

[2] Weitzman (1974) wrote the seminal paper comparing and contrasting the two instruments in a world with uncertainty.

[3] https://www.whitehouse.gov/sites/default/files/image/president27sclimateactionplan.pdf,

[4] http://www4.unfccc.int/submissions/INDC/Published Documents/United States of America/1/U.S. Cover Note NDC and Accompanying Information.pdf.

designed, what its key elements would be, and what modeling might be undertaken by economists to better understand the implications of such a policy design.

With that as background, we define and discuss the design elements for a Tax Adjustment Mechanism for Policy Pre-Commitment (TAMPP). A TAMPP is an adjustment mechanism for the tax rate of a carbon tax to ensure that targeted emission reduction milestones are met over the next few decades.

The next section places the TAMPP in the context of a rich literature on price, quantity, and hybrid instruments in greenhouse gas emissions policy design. Our focus is not the standard Weitzman (1974) focus on ex ante efficiency but rather a focus on how policy may play out dynamically over time with either a carbon tax or a cap and trade system. In section 3, we then turn to the TAMPP concept and enumerate key design questions that have to be considered when designing a TAMPP. Section 4 focuses on how the economics modeling community might design new or adapt existing models to assess a carbon tax with a TAMPP feature. A short concluding section summarizes.

II. Policy Design for Carbon Pricing

While much of the discussion over instrument choice for carbon policy has been over the relative merits of price (e.g. tax) or quantity (e.g. allowance) instruments, hybrid instruments are certainly possible. A hybrid instrument adds elements of a price instrument to a quantity instrument or vice versa. A price collar is the archetypal hybridization of a cap and trade system. A price collar combines a price ceiling in a cap and trade system with a price floor, thus limiting the magnitude of price increases or decreases.

With a price collar or a variant in which permits could be sold from a reserve (see, for example, Murray, Newell and Pizer, 2009), we no longer have certainty over cumulative emissions.[5] The hybrid system now adds some elements of a price system to the existing cap and trade system.

Given the focus on design elements to reduce price volatility in a cap and trade system, an obvious question is whether an analogous hybrid is possible for a price instrument: is it possible to reduce *ex post* uncertainty over emissions under a carbon tax? Surprisingly little

[5] We note in passing that the argument that cap and trade provides certainty over emissions is somewhat illusory. Even in the absence of a price collar or some similar mechanism, Congress serves as the ultimate implicit price ceiling. Were prices to rise to levels unanticipated and unacceptable to Congress, they could simply legislate a relaxation of the cap to bring prices down to more politically and economically acceptable levels.

research has been undertaken on this question. We are aware of only one paper on this topic written by Metcalf (2009). His Responsive Emissions Autonomous Carbon Tax (REACT) has the following features:

- An initial tax rate and standard rate of growth for the tax is set at the outset;
- Benchmark targets for cumulative emissions are set for a control period which could be one, five, ten-year or some other time interval;
- If cumulative emissions exceed the benchmark targets at the specified interval, the growth rate of the tax is increased to a higher rate until cumulative emissions fall to or below their benchmark targets in subsequent years.

Metcalf runs some simple simulations to illustrate how the mechanism could operate but does not do an in-depth assessment of the mechanism. Nor does he discuss design principles or possible variations in design for the consideration of policy makers. We turn to such a discussion in the next section. But before doing so, we pause to consider what sorts of "uncertainty" are relevant for the analysis.

At its most basic level, uncertainty refers to the deviation of some quantity of interest from the level that was anticipated when the policy was put in place.[6] That quantity of interest could be an outcome (such as allowance prices or emission levels) or something that influences those outcomes (such as the overall level of economic activity).

In thinking about this problem, it is useful to distinguish between uncertainty about emissions and uncertainty about abatement (i.e., emissions reductions). Abatement in any given year will be determined by the intersection of the tax rate and the marginal abatement curve (see figure 1). Note that figure 1 has abatement along the horizontal axis. Observed emissions equal business as usual (BAU) emissions, i.e. emissions in the absence of a carbon tax, less abatement.[7] This implies that any "certainty" mechanism can provide certainty over emissions but not over emissions reductions – even in the case of a pure cap and trade system with no

[6] We are abstracting away from volatility. While price volatility is of particular concern with a cap and trade system, emissions volatility is less of a concern given the stock nature of greenhouse gas pollution.

[7] We implicitly define abatement relative to BAU abatement activity. That is why the marginal abatement curve goes through the origin in Figure 1. Note, however, that this abstracts from scenarios where a carbon tax bill includes the removal of other climate policies (e.g. the Clean Power Plan).

safety valve. [8] We can track and limit emissions but without knowing BAU emissions, we cannot say what the actual emission reductions were.[9]

Figure 1. Marginal Abatement Costs

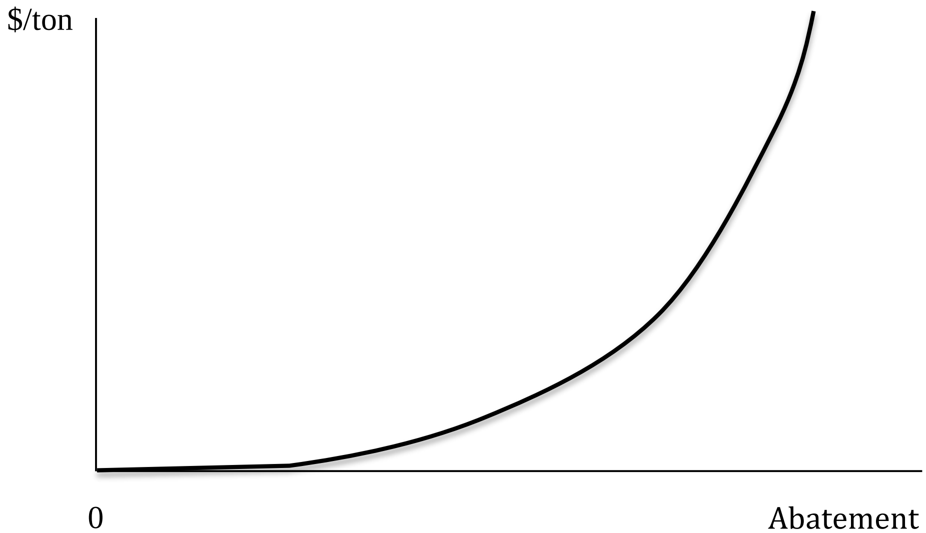

Uncertainty in predicting emissions arises from three sources: 1) unexpected shifts in BAU emissions over time; 2) errors in estimates of the marginal abatement curve at the start of the policy; 3) unexpected shifts in the marginal abatement cost curve over time. Errors in predicting BAU emissions can arise from unexpected changes in the growth rate of the economy, perhaps as part of business cycle impacts; changes in relative fuel prices that might, for example, shift the mix of coal versus natural gas electricity power production; and unforeseen changes in the mix of economic activity in the United States, among other forces.

Ex ante estimates of the marginal abatement curve could be wrong due to inaccurate estimates of the ease of substituting away from fossil fuels in production or errors in estimating the elasticity of consumer demand for carbon-intensive goods. Unexpected shifts in the marginal abatement curve over time can arise due to changes in the cost and availability of low and zero

[8] Alternatively, one could define emissions reductions relative to some fixed and known level of emissions (e.g., emissions in a particular year in the past). Under that alternative definition, uncertainty about emissions and uncertainty about emissions reductions are the same. But that is a less useful definition for evaluating policy, because under that definition, measured "emissions reductions" will be some mix of unexpected changes in BAU emissions and actual emissions reductions caused by the policy.

[9] To be precise, we are only limiting emissions in covered sectors; important sectors including agriculture and land use are typically excluded from any cap and trade or carbon tax proposal.

carbon fuel technologies. In the next section we discuss design considerations for a TAMPP that increase confidence in a carbon tax's ability to hit emission targets in an economically efficient fashion.

III. Design Considerations for a TAMPP

The basic structure of a TAMPP is straightforward. A time profile of tax rates is set over a control period. At the same time, a final emissions target and intermediate benchmarks are set. If, at specified times during the control period, emissions deviate sufficiently from the intermediate benchmarks, the tax rate changes in order to bring emissions back toward the benchmarks (for example, if emissions exceed the benchmark target, the tax rate would adjust upward).

Figure 2 provides a schematic for a TAMPP. The tax is enacted at time 0 and a final target is set for some designated future date T. Interim benchmarks are set where emissions (annual or cumulative; or emission reductions) are compared to the benchmark and the tax rate adjusted as needed. These adjustments can be designed so that, the final target has a high probability of being achieved.

Figure 2. Schematic of a Tax Adjustment Mechanism for Policy Pre-Commitment

The Swiss Carbon Tax Law provides a reasonably simple example of a TAMPP policy. The tax, which covers emissions from electricity and heating, had an initial rate of 12 CHF per metric ton of CO_2. By 2012, the tax rate had been raised to 36 CHF.[10] The law specified that if emissions in 2012 were greater than 79 percent of 1990 emissions, the tax rate would increase to 60 CHF as of January 1, 2014. The law specifies two additional milestone years (2014 and 2016) with tax rates to adjust (in 2016 and 2018 respectively) if the milestones were not met. The law put in place two different higher tax levels for 2016 and 2018 depending on the level of emissions. The tax would rise to 96 CHF in 2018, for example, if emissions exceeded 73 percent of 1990 emissions. But the tax would rise to 120 CHF if emissions exceeded 76 percent of 1990 emissions.[11]

The Swiss Carbon Tax Law is only one example of the structure that a TAMPP could take. Current policy proposals often include TAMPP-like elements, such as the Whitehouse-Schatz American Opportunity Carbon Fee Act of 2015 (S. 1548) that specified an annual two percent (over inflation) increase in the tax rate until emissions fall to 80 percent below 2005 levels, at which point the tax rate would be held constant in real terms. Policymakers face a number of key design choices in adding a TAMPP to a carbon tax. We review these key design choices next.

Rules v. Discretion

Changes to the carbon tax rate to ensure that targets are met during the compliance period could be spelled out in legislation or left to Congress to periodically make as needed. One might contemplate Congress delegating the authority to change tax rates to an executive branch agency such as the Treasury or EPA or to a new quasi-independent agency similar to the Federal Reserve. But Congressional delegation of tax setting authority to another branch of government might not be constitutional; and it is hard to imagine that Congress would willingly delegate such authority even if it were constitutional.

---

[10] The carbon tax was initially enacted as part of the 1999 Act on the Reduction of CO_2 Emissions and covered emissions between 2008 – 2012. It was subsequently reauthorized to cover emissions through 2020. Firms could opt out of the carbon tax by participating in the Swiss Emissions Trading System. See Sopher and Mansell (2013) for further information on the initial enactment of this law.

[11] *Ordonnance sur la Reduction des Emissions de CO_2*, Le Conseil Federal Suisse, enacted on Dec. 23, 2011 (RS 641.71), available at http://www.news.admin.ch/NSBSubscriber/message/attachments/31399.pdf accessed on Aug. 13, 2016.

If delegation is out of the question, then either the changes to the tax rate must be specified in Congressional legislation (as in the Swiss carbon tax example) or periodically enacted by Congress in response to new information about emissions. The latter approach would make the TAMPP simply guidance for future Congresses; it is unlikely that this would provide the sort of assurance to constituents that want guarantees that a carbon tax can achieve certain emissions targets.

Specifying the changes in legislation would be somewhat unusual, because future tax rates would be dependent on future emissions, and there are relatively few cases in the United States in which legislation specifies changes in future quantities in the tax system based on events that are not specific to the taxpayer being affected. But some examples do exist. Many dollar amounts in the tax code are automatically adjusted for inflation without the need for new legislation, applicable federal rates (interest rates used in the tax system, which determine, for example, the interest charged on tax underpayments) are set based on market rates for Treasury bills and bonds, and the tax credit available to a hybrid car buyer phased out based on how many hybrid vehicles the car's manufacturer had previously sold.

TAMPP Control Period

Over what period should the TAMPP apply? Most policy discussions of climate policy focus on the near term (e.g. 2025 to 2030 for most of the Intended Nationally Determined Contributions (INDC's) submitted to the U.N. Framework Convention on Climate Change (UNFCCC) in the run up to the 2015 Conference of the Parties (COP) in Paris) while also articulating longer term goals (e.g. 80 percent reduction in emissions by 2050). The United States, for example, committed in its INDC to "an economy-wide target of reducing its greenhouse gas emissions by 26-28 per cent below its 2005 level in 2025" while noting that its target "…is consistent with a straight line emission reduction pathway from 2020 to deep, economy-wide emission reductions of 80% or more by 2050."[12]

The control period for a carbon tax and for the TAMPP need not (and should not) be coterminous. While it would be possible to enact a temporary carbon tax, doing so would be inadvisable given the policy uncertainty from which a temporary tax would suffer. While

[12] U.S. Cover Note, INDC and Accompanying Information submitted to the UNFCCC on March 31, 2015 and available at http://www4.unfccc.int/submissions/indc/Submission%20Pages/submissions.aspx. Accessed on Aug. 14, 2016.

temporary taxes may be desirable in certain circumstances, it is inadvisable in the climate context where we face a long-lived pollutant and an on-going environmental concern. Making the tax temporary or subject to periodic reauthorization would only add uncertainty that would complicate decision making for businesses trying to do investment planning for very long lived capital equipment and production processes such as power plants.[13]

While a permanent carbon tax is desirable, the control period for the TAMPP – the length of time over which emissions targets are set - should be finite in duration. The final target year needs to be sufficiently far into the future that meaningful long-run investments can be justified that can contribute to lower emissions but not so distant that any conjectures about the state of technology and energy networks – or about our future understanding of what emissions reductions will be needed to avert serious damage from climate change – become completely speculative. Interim targets will have to be set throughout the control period and the further out in time the control period extends, the more difficult it is to know how to set those targets.

On the other hand, setting a control period of only a few years reduces incentives for the long-lived energy investments needed to move towards a zero carbon economy. The Swiss carbon tax is a good example. With a control period that extends only to 2020, it is difficult to see how the law will provide incentives for significant additional reductions in the post-2020 era.

In general, the longer-lived the relevant investments are, the longer the control period should be. And the more uncertainty we have about how our understanding of the costs and benefits of carbon mitigation will evolve over time, the shorter the control period should be. Pinning down a specific number is difficult, but it should clearly be at least 10-15 years to overlap with the period in which the U.S. has made international commitments to meet certain emissions targets.[14] However, it is hard to imagine how to credibly set targets more than 35-40 years into the future.

This discussion has assumed a specific end date for the control period. One could also imagine an endogenous length of control whereby the control period ends once an emissions threshold has been achieved. After that date, the tax rate might be fixed in real terms or increase

[13] Metcalf (2010) finds that the need to reauthorize the production tax credit for renewable electricity production every two years has impacted wind electricity investments.

[14] We are not suggesting that the TAMPP should be designed specifically to reflect Administration commitments made through the UNFCCC process. There are good reasons to enact a carbon tax and include a TAMPP component regardless of the state of UNFCCC negotiations. But the TAMPP targets should not be inconsistent with any timetable developed as part of the UNFCCC negotiating process.

at a relatively slow real rate. Under this approach, it would still be useful to have at least an approximate target end date in mind for transparency and clarity in business planning.

The importance of the length of the control period also depends on the likely timing of future legislation. If we were certain that an updated law will be enacted 10 years from now, then any targets beyond 10 years would matter only as a guideline for what that updated law might look like, so the distinction between a 20-year versus 50-year control period would be relatively unimportant. In contrast, if the law is infrequently revisited, then targets further into the future become more important.

Targets and Interim Benchmarks

Emissions targets and interim benchmarks can be set in terms of emissions relative to some base year, some absolute emissions cap, or emission reductions relative to a BAU emissions baseline. The first type of target is consistent with the 2025 target articulated in the U.S. INDC as noted above. It is a percentage reduction relative to emissions in 2005 and so (indirectly) sets an absolute cap on emissions in 2025. The second approach would simply make the cap explicit and is consistent with the approach taken in the annual allowance allocations in the Waxman-Markey cap and trade bill passed by the U.S. House of Representatives in 2009.[15]

In addition, interim benchmarks could be set in terms of annual emissions, cumulative emissions, or some moving average of emissions. Given the stock nature of the pollutant, a longer-term target (cumulative emissions or a moving average over a relatively long period) more closely corresponds to what determines damages. And using a moving average rather than an annual snapshot would smooth out short-term fluctuations. Setting benchmarks in terms of cumulative emissions provides even greater smoothing.

How much smoothing is desirable depends on how persistent unexpected shocks are likely to be. If all shocks persist forever (i.e., if the underlying quantities follow random walks), then smoothing provides no advantage, and simply causes adjustments to lag unnecessarily behind the shocks, so adjusting based on the most recent emissions would be ideal. In contrast, if all shocks are very transitory, then smoothing is desirable: adjusting based on annual emissions won't help at all in hitting the final target, and will introduce unnecessary price volatility as the tax rate adjusts based on every little fluctuation.

[15] The Waxman-Markey annual declining allowance allocation is not the same as an absolute cap given the ability to bank allowances from previous years.

In practice, there will be a mix of transitory and persistent shocks: for example, business-cycle and weather shocks are transitory, whereas technology shocks will be more persistent. The ideal would be for benchmark quantities to cover a long enough period to smooth out the transitory shocks, but a short enough period to respond quickly to persistent shocks.

Targets and benchmarks specified as reductions from a BAU emissions path are also possible but raise a number of difficult questions. How is the BAU path set? How often is it updated to reflect changes in the BAU economy? Who updates the path and according to what model? Should the BAU path closely reflect what actual BAU emissions would be or is it simply a benchmark against which to assess reductions? We see little if any advantage to setting targets or benchmarks in terms of emission reductions from a hypothetical BAU path, no matter how accurately one believes such a path could be estimated.

Despite the problems of setting targets and benchmarks relative to a BAU path, it is important to emphasize that the costs of reducing emissions depends on the legislated emissions trajectory (however specified) relative to the BAU emissions pathway. For example, the costs of reducing emissions to a given target would depend on whether, in the absence of policy, the world experienced breakthroughs in low or zero-cost carbon technologies.

Regardless of how targets and benchmarks are defined, the number of interim benchmarks should be set to reasonably ensure stakeholders that the final target will be achieved. However, as we note below, while frequent adjustments of tax rates are beneficial, such adjustments could occur between benchmark years and in that case there would be no strong economic rationale for choosing the number of interim benchmarks. If adjustments do not occur between benchmark years, then that would be a reason for more frequent interim benchmarks.

Types of Adjustments

Adjustments to the tax rate could take a variety of forms. One approach would be to specify tax rates (dollars per ton of CO_2) in the legislation with schedules contingent on whether interim benchmarks have been met or not. This is the approach taken in the Swiss Carbon Tax Law.

Metcalf (2009) suggests an alternative approach. A tax rate is specified for the first year along with an annual percentage increase in the tax rate. If cumulative emissions at the intermediate milestones exceed the target specified for that year, the annual rate of increase in the tax rate jumps to a higher level and stays at that level until cumulative emissions at a future benchmark no long exceed the benchmark target. For illustrative purposes, Metcalf suggested a "standard" growth rate of the tax of 4 percent in real terms and a higher "catch-up" rate of 10 percent (real).

One might combine elements of the Swiss approach and the approach suggested by Metcalf (2009). The TAMPP might specify two or more percentage increases in the tax rate for the next control period (two years in the Swiss case but possibly longer) that depend on the amount by which emissions exceed some targeted emissions level.

There is no good economic argument to pick one approach over another. What matters is clarity and certainty in the rules for tax rate adjustment so that the business community can plan with reasonable certainty (given their private forecasts of emission trajectories).

Frequency and Size of Adjustments

How often the tax rate adjusts and how large those adjustments should be are linked: the more frequent the adjustments are, the smaller each adjustment should be, with the size of each adjustment being roughly proportional to the period between adjustments: several small adjustments will add up to the same overall effect as one big adjustment, so if adjustments occur half as often, each one will need to be roughly twice as big.

Smaller and more frequent adjustments will tend to provide both lower costs and better environmental outcomes. Infrequent big tax adjustments aren't cost effective (consider a big tax increase: one could achieve the same total abatement more cost-effectively by doing more abatement before the increase and doing less afterwards, which is exactly what a series of smaller adjustments would do). And from an environmental standpoint, less frequent adjustments make it harder to hit a given target. Larger adjustments also raise the risk of substantial price movements as well as political push-back.

The frequency of adjustment need not be fixed to the period defined by interim benchmarks in the legislation (if any). If, for example, the policy specifies five or ten year interim benchmarks, the legislation could define more frequent adjustments (annual or bi-annual) to allow for responses to shocks that occur during years between benchmarks. Adjustments occurring between benchmark years would then be based on smoothing between the benchmarks.[16]

[16] The initial response of the economy to a tax could make smoothing between benchmark years problematic in the first few years of the policy. This could be an argument for not making any adjustments in those first few years.

However, there are practical limits to how often adjustments can be made. For example, adjustments cannot be made more frequently than emissions data is updated. More frequent updating could also raise enforcement and compliance costs, though this seems like it would be a significant issue only with very frequent updating (e.g., monthly or quarterly).

For a given frequency of adjustment, the size of adjustments pose a tradeoff: larger adjustments make it more likely that emissions will stay close to the target quantities, but this would likely also imply higher costs.[17,18] This affects the credibility of the policy along two dimensions. If more sizable rate changes are mandated for larger shortfalls from the target, credibility in the mechanism is enhanced. However, very large tax rate adjustments can undermine political credibility as they increase the chances that Congress could intervene to prevent very large tax rate changes. In the end, a balance must be struck between an adjustment process that provides credibility in the environmental outcomes and a process that does not lead to abrupt and large economic costs.

Given the above, we would expect adjustments to be no more frequent than annual. At the other end of the continuum, decadal adjustments seem too infrequent. The Swiss model of two-year interim targets is not unreasonable. In the end, this comes down to a tradeoff between economic and environmental factors, which favor more frequent adjustment, versus practical and political considerations that could push in the opposite direction.

Adjustment Trigger

Design considerations involving the adjustment trigger include, among other things, whether the trigger is one or two-sided; whether it is discrete or continuous; and whether there is a target range for deviations from the target. A one-sided target only responds to undershooting the target (e.g. cumulative emissions in excess of allowed cumulative emissions) by raising the tax rate. A two-sided target would provide for reductions in the tax rate (or rate of growth of the tax rate) in the case of overshooting the target (as in the case of a technology shock that significantly reduces abatement costs). We see no particular reason for choosing a one-sided

[17] This tradeoff depends on how elastic emissions are with respect to the tax rate, with a lower elasticity implying a need for larger adjustments. This means that uncertainty about that elasticity is especially problematic when designing the policy.

[18] Large adjustments could also lead to overshooting and oscillations where we alternately fall short of and exceed the target. An adjustment that is a function of the gap between emissions and the target as discussed in the next sub-section could reduce the potential for overshooting.

target over a two-sided target though we would be concerned about transitory declines in emissions (for example due to a recession) that lead to a tax rate reduction that then requires a larger tax rate increase in the future as emissions return to their long-run trend.[19]

Triggers could be discrete or continuous. Metcalf (2009) proposes a discrete two-sided trigger where the tax rate grows at a standard rate of 4 percent real but then jumps to 10 percent real if emissions exceed the target. It then reverts to the 4 percent growth rate when emissions fall below subsequent targets. The Swiss ordinance is an example of a one-sided target as there is no provision for lowering the rate at any adjustment period. The Swiss ordinance also illustrates the possibility of multiple discrete adjustments (in the Swiss case, two) depending on how far emissions deviate from the target.

Adjustments could also be continuous. The tax rate changes could be a function of the deviation of emissions from the target. As a simple example, the percentage change in the tax rate (Δ) might equal

$$\Delta = max\left(1.0, \alpha\left(\frac{E-T}{T}\right)\right),$$

where α is some positive constant, *E* is the measure being tracked, and *T* is the target. This example puts a cap of 100 percent on the tax rate increase at any given adjustment benchmark. For example, if α were set to 10, an overshoot of 3 percent in the measure *E* relative to its target would lead to an increase in the tax rate of 30 percent. Note that the formula could be made symmetric for undershooting the target (e.g. *E* is 2 percent lower than *T* leads to a 20 percent decline in the tax rate), asymmetric (the value of α could depend on whether $E > T$ or $E < T$), or one-sided (α equals zero when $E < T$). One could also use a more complex formula (making the function nonlinear, for example).

A continuous adjustment will generally be more cost-effective and better from an environmental standpoint, since it will imply smoother and more predictable changes in tax rates over time (if emissions are near the cutoff for a discrete adjustment, a small change in emissions could produce a big change in the tax rate). But this advantage might be small, especially if the adjustments are frequent and relatively small in magnitude. Further, politicians may view such a continuous adjustment mechanism as too complex and opaque.

[19] This problem could be mitigated, however, by smoothing the emissions targets, as noted earlier (e.g., using a weighted average of several years' emissions rather than a single year).

The threshold for triggering a tax rate change could be based on the target itself or a band around the target. Above, we have described thresholds where the tax adjusts if emissions exceed the target. An alternative might trigger adjustments if emissions exceed some band around the target (e.g., exceeding the target by more than 2%). One might even envision a tiered threshold, based on color coded bands: a narrow green band around the target requires no action; a wider yellow band serves warning that the target is being exceeded and action may be required in the future (or that leads to an immediate but modest tax change). One might even require action if too much time (e.g. 2 years) is spent in the yellow band. Finally emissions in an even wider red band require an immediate increase in the tax rate (or a larger tax change than would occur in the yellow band).

IV. An Agenda for Research on TAMPPs

Research on TAMPPs would be useful, both for evaluating different design choices under a TAMPP and for comparing a TAMPP to other alternative policies (e.g., a carbon tax without any formal adjustment mechanism or a cap-and-trade program). But as noted earlier, we are aware of only one paper that addresses any portion of this issue: Metcalf's (2009) paper about REACT, a specific example of a TAMPP (see Section II for details of the REACT policy). That paper includes some simple simulations, but they are intended to be illustrative, not to provide rigorous modeling of the REACT policy. In this section, we outline some of the useful directions that research on TAMPPs might take, as well as some of the challenges such research would face.

A good starting point would be simple analytical modeling. Weitzman's (1974) paper comparing price and quantity regulations and much of the literature that followed it have used simple analytical models. Such models have major advantages in transparency and generality of results. They can also put a sharp focus on key underlying forces that drive important economic results. Weitzman's simple modeling structure, for example, highlighted the importance of the relative slopes of the marginal damage and marginal benefit curves for emissions in determining whether price or quantity instruments are ex ante more efficient.

But as analytical models become more complex, they soon become intractable. Uncertainty and dynamics are essential for modeling a TAMPP, and those elements together lead to inherently complex models. Moreover, reaching quantitative conclusions is likely to require numerical simulations. Thus, we believe that while research might start with analytical models, numerical simulation will quickly become necessary.[20]

One could attempt to model the underlying structure of the economy and energy sectors in detail, in a manner similar to the computable general equilibrium (CGE) models that are commonly used to model the response of carbon emissions and the broader economy to the introduction of a carbon price. Indeed, an existing CGE model could be the core of a numerical model to evaluate a TAMPP.

The major problem with such an approach is that CGE models are almost all deterministic, and uncertainty is obviously a vital element of any model used to evaluate a TAMPP. The most common approach to uncertainty in CGE modeling is to undertake Monte Carlo analysis with deterministic CGE models. Probability distributions for key parameters are assumed and simulations are then run where parameter draws from those distributions are taken.[21] Monte Carlo approaches are useful for illustrating model sensitivity to key parameters and could be used to estimate uncertainty over the marginal abatement cost curve at the time of policy implementation but fail to address other types of uncertainty arising from unexpected shocks over time. This highlights the internal inconsistency of this approach: the Monte Carlo simulations explicitly incorporate uncertainty but the underlying models have no uncertainty.[22]

---

[20] Papers on dynamic problems in the literature on policy instrument choice under uncertainty typically use numerical simulation (often in addition to analytical models). See, for example, Hoel and Karp (2002) or Pizer (1999). The modeling in Metcalf (2009) was entirely numerical.

[21] Webster et al. (2003) take such an approach in modeling with a global climate model that includes a CGE model of the world economy as one element of the broader model. Abrell and Rausch (2016) use a Monte Carlo experiment to characterize uncertainty in marginal abatement costs curves for ETS and non-ETS sectors in Europe.

[22] A further difficulty with Monte Carlo analysis is determining what probability distributions to use for the key parameters. In many cases, there is no empirical evidence, and so such distributions must rely on ad hoc assumptions. This problem isn't unique to Monte Carlo analysis, though; it (or very similar problems) apply to every method for handling uncertainty discussed in this section.

Adding explicit uncertainty to an existing CGE model or building a new CGE model with explicit uncertainty would be a tremendous undertaking.[23] Thus, using a CGE model directly is likely infeasible, though CGE models could be useful for parameterizing other approaches.

Dynamic stochastic general equilibrium (DSGE) models are another potentially promising approach. They have primarily been used to study macroeconomic problems, but are starting to be used in environmental applications. Such models include dynamics and uncertainty, but the tradeoff is that they have greatly simplified representations of the structure of the economy, typically modeling only a single aggregate sector, and almost never modeling more than 2 or 3 sectors.[24] This greatly limits their ability to represent the range of carbon emissions abatement options needed to provide meaningful insight into the greenhouse gas mitigation problem. Nonetheless, a properly parameterized DSGE model could be very useful for modeling a TAMPP by providing a framework that properly addresses the uncertainty in the business cycle and uncertainty in shocks to future abatement costs.

A simpler approach wouldn't try to model the underlying structure of the economy at all, but would instead take a much more reduced-form approach, in which emissions would be a function of the carbon tax rate (perhaps representing the speed of adjustment to tax changes by also including the rate from one or more previous time periods), with random shocks to the level and slope of that function. This is the approach that Metcalf (2009) took.

A major challenge for either of the latter two approaches – DSGE or reduced-form – is parameterizing the response of emissions to a carbon price. Key elements that would need to be parameterized include how much emissions respond to a given price (i.e., the elasticity of the abatement supply curve), how quickly that response occurs,[25] and how random shocks could

[23] In deterministic models, agents choose actions each period to maximize some objective function. In a stochastic model, agents actions are governed by a decision or policy rule that govern which actions to take conditional on different realizations of a shock. Often, this rule will be non-linear and yet solution methods for non-linear stochastic models almost always use first or second order approximations. Further, most of these approximation methods are only useful if the economy is close to its steady state and are not appropriate for solving transitions from one steady state to another, as would be the case with a carbon tax. Moreover, these approaches are very computationally intensive, and thus other aspects of a CGE model would likely need to be substantially simplified in order to make analysis with explicit uncertainty computationally tractable.

[24] Examples include Heutel (2012), Grodecka and Kuralbayeva (2015), and Dilusio (2016).

[25] Some adjustments will be almost immediate, such as changes in the dispatch order for electric power generation by existing plants, while other responses could take decades, such as retirement of long-lived emissions-intensive capital.

alter that response.[26] We don't have direct empirical estimates of any of those elements, because the US has never imposed a national carbon price (and even if one extrapolates from experience in other countries that have imposed a carbon price, the sample is quite small).

One could use the results from a CGE model (or models) to parameterize that response to emissions to a carbon price. This approach would run the CGE model for a range of different carbon tax rates (to measure the emissions response to the tax), trajectories for the tax (to measure the speed of adjustment), and underlying model parameters (to measure how random shocks could change the response) – in essence, running a Monte Carlo analysis along these dimensions. This still relies upon the CGE model providing a reasonable representation of the emissions responses to carbon pricing, but since the CGE model includes more of the underlying structure of the economy and energy sector, its parameters can be estimated based on a wider range of historical shocks to the economy.[27]

Even if parameterized based on a CGE model, however, the reduced-form approach has the fundamental problem that it can't represent the effects of firms anticipating future carbon tax adjustments. Suppose emissions are well above the target under the TAMPP, and a firm is considering making a long-term investment that will lower its carbon emissions. Because emissions are high, future tax adjustments under the TAMPP will almost certainly raise the tax rate, thus making that long-term investment look more attractive than it would look based just on the current carbon price. That kind of anticipation of tax changes will generally make the TAMPP perform better (more likely to hit emissions targets, and in a more cost-effective way), and thus failing to capture it in a model will bias the results.[28] A DSGE model has the potential to avoid this problem, since it can explicitly capture firms' anticipation of future tax changes.

---

[26] The simulations in Metcalf (2009) use a function that implicitly assumes away the latter two elements. In that model, emissions respond immediately to a change in the carbon tax, and there is no uncertainty about the magnitude of that response; the BAU level of emissions is uncertain, but the reduction a given carbon tax rate will cause from that BAU level is entirely deterministic. Metcalf parameterizes the function based on runs of the EPPA CGE model.

[27] For example, substitution elasticities among different energy sources in a CGE model could be estimated using events that caused exogenous shifts in relative prices of different energy sources (such as shocks to the world oil market). But those same prior events would not be sufficient for directly estimating the reduced-form response of emissions to a carbon tax.

[28] If short-term shifting of emissions is possible, then anticipation could also make the TAMPP perform worse. A firm that anticipates a carbon tax increase at the start of next year and can do short-term shifting of emissions would shift emissions from next year into this year. This would incur some costs, but do nothing to lower cumulative emissions. But because the potential for such shifts seems smaller than the importance of long-lived investments, anticipation seems likely to boost the performance of a TAMPP rather than hurt it.

Under any of these approaches, empirical research on the uncertainty about future emissions paths would be important. As noted earlier (in Section II), we see three key sources of uncertainty: 1) unexpected shifts in BAU emissions; 2) errors in estimates of the marginal abatement curve at the start of the policy; and 3) unexpected shifts in that marginal abatement curve over time. The first of these – shifts in BAU emissions – is straightforward to estimate based on prior data. Nonetheless, we are unaware of empirical work that has explicitly focused on the magnitude and persistence of random shocks to BAU emissions. Such estimates would be valuable for designing and evaluating any policy designed to manage uncertainty about carbon abatement.

The second and third of these are harder to estimate, however, because they can only be directly observed after a carbon pricing policy is in place, and there are relatively few cases of carbon pricing to work with. But those cases might be enough to provide some lessons, or it might be possible to draw information from pricing of emissions other than carbon. Some work in this area already exists,[29] and further research could be highly useful.

Summing up, we see a fruitful research agenda for incorporating uncertainty into CGE modeling. First, we see great value in doing more Monte Carlo simulations with existing CGE models. While this approach has an internal inconsistency in that these models assume economic agents are making decisions in a world without uncertainty, the approach still has value. It sheds light on where reducing parameter uncertainty can be most fruitful in reducing error bars on key model results and can highlight the extent of the uncertainty surrounding the marginal abatement curve at the start of the policy. But it cannot answer many key questions about TAMPP mechanisms because a deterministic CGE model cannot adequately model future shocks that would lead to TAMPP adjustments.[30]

---

[29] For example, Kaufman, Obeiter and Krause (2016) finds that estimates prior to the introduction of carbon pricing tend to overestimate marginal abatement costs, thus leading either to overestimates of permit prices under cap-and-trade or underestimates of abatement under an emissions tax.

[30] Note that agents could still react to anticipated future policy in these models if the models incorporate forward looking behavior. This is distinct from the observation that the CGE models that would be used for these Monte Carlo runs have economic agents that operate as if the world is deterministic. But the models cannot incorporate reactions to shocks. Or, in other words, agents in these models can react to policy changes that are entirely predictable before the policy starts, but not to any other policy changes (such as TAMPP adjustments caused by unexpected changes in BAU emissions or in the abatement cost curve).

Second, building new or adapting existing DSGE models to study climate policy should have a high priority in the research agenda. While the models will need to be simplified in many ways to be computationally tractable, even simple DSGE models have the potential to tell us quite a bit about how adaptive policy (such as a TAMPP) interacts with risk preferences and uncertainty. Simple DSGE models could also provide useful insight as to the size and direction of biases that come from running Monte Carlo simulations with deterministic CGE models.[31]

At the same time that a research program to incorporate uncertainty explicitly into CGE modeling proceeds, there is a need to inform policy makers on near term policy initiatives. Deterministic CGE models can be used to determine what initial tax rate and price path would lead to a given target. This is simply an *ex ante* estimate based on the assumptions in the model and should not be construed as "truth;" in other words, how emissions actually decline for a given ex ante price path will differ due to errors in the estimation of the marginal abatement costs and unexpected shocks and may require the TAMPP to come into play if the emissions path is sufficiently off the ex ante target.

V. Conclusion

Including a Tax Adjustment Mechanism for Policy Pre-Commitment (TAMPP) in a carbon tax could provide some assurance to the public that United States policy is committed to meaningful GHG emission limits (as laid out in the TAMPP). To that end, we have provided a brief review of the literature on TAMPP-type mechanisms and have enumerated a number of policy elements that would go into a TAMPP.

While some design elements are a matter of legislative preference, other design elements are quite important if the policy is to be successful at its goal of providing policy assurance without adding inefficient or other unintended elements to a carbon tax. Table 1 provides a summary of the design elements and our initial thoughts on preferred courses of action (when those exist).

Table 1. TAMPP Design Considerations

Design Feature	Options	Guidance
Rules vs. Discretion	• Congressional delegation for rate setting	• Legislated tax change rules preferable to

[31] Farmer et al. (2015) and Stern (2016) call for a "third wave" of climate modeling including the use of DSGE models.

	• Congressional discretion • Tax rate changes included in initial legislation	discretion to reduce policy risk • Delegation seems unlikely and provides no clear advantages
TAMPP Control Period	• Some period of time ranging from a few years to the end of the century	• Too short a period increases policy risk • Too long a period brings in too much economic, technological, and environmental uncertainty • A period of 15 to 25 years may be optimal though this is a matter of judgment
Interim Benchmarks	• Some number of interim benchmarks, ranging from zero to annual benchmarks throughout the control period	• No clear economic guidance • Benchmarks spaced in the five to ten year range may be reasonable • If adjustments occur only in benchmark years, see Frequency of Adjustment, below
Types of Adjustments	• Tax rate changes in absolute or percentage terms • Single or multiple possible rate changes at an interim benchmark	• No clear guidance • Room for flexibility
Frequency and Size of Adjustment	• Frequency can be short (annual) to very infrequent (10 years or longer) • Size of adjustment related to	• More frequent adjustments generally preferred, subject to practical and political limitations

	frequency of adjustment (more frequent adjustment means adjustments can be smaller)	• Adjustments can be made between interim benchmark years (and should, unless benchmark years are very frequent)
Adjustment Trigger	• One- or two-sided • Discrete or continuous • Based on a formula • Target bands	• Number of appealing options • What matters in the end is a clear and transparent approach that both provides assurance for environmental goals and reduces uncertainty for investment

In particular, any TAMPP should be built into the legislation rather than left to Congressional discretion. While future Congresses always have the ability to alter previous legislation, policy inertia favors making the TAMPP a default in the carbon tax legislation. How far into the future the target TAMPP emission target is set (and at what level) is a matter of judgment. Setting final target dates too far into the future risks setting targets with – at best – speculative knowledge about the state of the economy or mitigation technologies that will be available at that future date. Setting final target dates just a few years out does not provide sufficient time for meaningful emission reduction targets.

Both final targets and interim benchmarks are best designed either as absolute emission limits or as reductions from a benchmark year. While the true effect of any GHG mitigation policy is the emissions reduction from the business-as-usual (BAU) emissions path, and the true cost depends on reductions from the BAU path, this BAU path cannot be directly observed or determined with certainty (even ex-post). And in the end what matters for measuring damages from GHG emissions is the stock of emissions in the atmosphere resulting from the accumulation of annual emissions over time. So benchmarks that relate to actual emission caps (or reductions from a given historic emissions level) relate more directly to future damages.

Policymakers have considerable discretion in how they design the tax rate adjustment if interim targets are not met. What is most important is clarity and certainty in the rules for the tax rate adjustment so that the businesses and individuals can respond with reasonable confidence to likely future government policy. There is also no set guidance for how frequently interim benchmarks should be assessed. More frequent adjustments will generally lead to lower costs and better environmental outcomes, but practical and political considerations will limit the frequency of adjustment.

A similar tradeoff applies for using a discrete or continuous adjustment. The continuous adjustment will generally be superior on economic and environmental grounds, but those advantages could be small (particularly with frequent adjustments) and the apparent simplicity of discrete adjustments is a political advantage.

Finally, policymakers have considerable flexibility as to how to design other elements of the trigger. It can be one-sided or two-sided; can be designed in absolute or percentage terms; and can include the use of bands (green, yellow, red, for example) with different responses within each band.

Further research can be useful hone the guidance on the optimal design of a TAMPP; we have laid out a research agenda that can contribute to better informed carbon tax design in the face of uncertainty over future emission trajectories, damages, and mitigation technology.

References

Abrell, Jan and Sebastian Rausch. 2016. "Combining Price and Quantity Controls under Partitioned Environmental Regulation," Zurich, Working Paper 16/233.

Dales, J.H. 1968. Pollution, Property, and Prices. Toronto: University of Toronto Press.

Dilusio, Francesca. 2016. "Environmental Policy and the Business Cycle: The Role of Adjustment Costs on Abatement," Rome: Universita degli studi di Roma,

Farmer, J. Doyne; Cameron Hepburn; Penny Mealy and Alexander Teytelboym. 2015. "A Third Wave in the Economics of Climate Change." Environmental and Resource Economics, 62(2), 329-357.

Grodecka, Anna and Karlygash Kuralbayeva. 2015. "The Price Vs Quantity Debate: Climate Policy and the Role of Business Cycles," London: Grantham Research Institute,

Heutel, Garth. 2012. "How Should Environmental Policy Respond to Business Cycles? Optimal Policy under Persistent Productivity Shocks." Review of Economic Dynamics, 15, 244-264.

Hoel, Michael and Larry Karp. 2002. "Taxes Versus Quotas for a Stock Pollutant." Resource and Energy Economics, 24(4), 367-384.

Kaufman, Noah; Michael Obeiter and Eleanor Krause. 2016. "Putting a Price on Carbon: Reducing Emissions," Washington, DC.

Metcalf, Gilbert E. 2009. "Cost Containment in Climate Change Policy: Alternative Approaches to Mitigating Price Volatility." Virginia Tax Review, 29, 381-405.

Metcalf, Gilbert E. 2010. "Investment in Energy Infrastructure and the Tax Code." Tax Policy and the Economy, 24, 1 - 33.

Murray, Brian C.; Richard G. Newell and William A. Pizer. 2009. "Alternative U.S. Climate Policy Instruments: Balancing Cost and Emissions Certainty: An Allowance Reserve for Cap-and-Trade." Review of Environmental Economics and Policy, 3(1), 84-103.

Pigou, Arthur C. 1932. The Economics of Welfare. London: MacMillan and Co.

Pizer, William A. 1999. "The Optimal Choice of Climate Change Policy in the Presence of Uncertainty." Resource and Energy Economics, 21(3), 255-287.

Sopher, Peter and Anthony Mansell. 2013. "Switzerland: The World's Carbon Markets: A Case Study Guide to Emissions Trading," EDF and IETA: New York.

Stern, Nicholas. 2016. "Current Climate Models Are Grossly Misleading." Nature, 530, 407-409.

Webster, Mort; Chris Forest; John Reilly; Mustafa Babiker; David Kicklighter; Monika Mayer; Ronald Prinn; Marcus Sarofim; Andrei Sokolov; Peter Stone, et al. 2003. "Uncertainty Analysis of Climate Change and Policy Response." Climatic Change, 61, 295-320.

Weitzman, Martin L. 1974. "Prices vs. Quantities." Review of Economic Studies, 41(4), 477-491.