RESOURCES
for the FUTURE

Designing and Evaluating a US Carbon Tax Adjustment Mechanism

Marc A. C. Hafstead and Roberton C. Williams III

Working Paper 20-04
March 2020

About the Authors

Marc Hafstead is an RFF Fellow and Director of RFF's Carbon Pricing Initiative. His research focuses on the evaluation and design of climate and energy policies. With Stanford Professor and RFF University Fellow Lawrence H. Goulder, he wrote Confronting the Climate Challenge: US Policy Options (Columbia University Press) to evaluate the environmental and economic impacts of carbon taxes, cap-and-trade programs, clean energy standards, and gasoline taxes using a sophisticated multi-sector model of the United States. His research also analyzes the distributional and employment impacts of carbon pricing and the design of tax adjustment mechanisms to reduce the emissions uncertainty of carbon tax policies. His work has been featured in the popular press, including the Wall Street Journal, the Washington Post, Axios, and CNNMoney. He joined RFF in 2013 from Stanford University.

Email: hafstead@rff.org

Roberton Williams studies both environmental policy and tax policy, with a particular focus on interactions between the two. In addition to his role as a university fellow at RFF, he is a professor at the University of Maryland, Chief Economist for the Climate Leadership Council, and a research associate of the National Bureau of Economic Research. He was previously an associate professor at the University of Texas, Austin; a visiting research scholar at the Stanford Institute for Economic Policy Research; and an Andrew W. Mellon Fellow at the Brookings Institution. Williams has served as a coeditor of both the Journal of Public Economics and the Journal of Environmental Economics and Management.

Email: roberton@umd.edu

Acknowledgements

We thank Resources for the Future's Carbon Pricing Initiative for financial support and we thank Gib Metcalf, Wesley Look, Susanne Brooks and seminar participants at the Environmental Defense Fund for helpful discussions on tax adjustment mechanisms.

The views expressed herein are those of the authors and do not necessarily reflect the views of Resources for the Future, the University of Maryland, the Climate Leadership Council, or the National Bureau of Economic Research.

About RFF

Resources for the Future (RFF) is an independent, nonprofit research institution in Washington, DC. Its mission is to improve environmental, energy, and natural resource decisions through impartial economic research and policy engagement. RFF is committed to being the most widely trusted source of research insights and policy solutions leading to a healthy environment and a thriving economy. Working papers are research materials circulated by their authors for purposes of information and discussion. They have not necessarily undergone formal peer review. The views expressed here are those of the individual authors and may differ from those of other RFF experts, its officers, or its directors.

Sharing Our Work

Abstract

There is increasing political interest in the United States in an economy-wide carbon tax. Notably, two center-right Republican proposals are attracting attention: the Climate Leadership Council's Carbon Dividends plan, and Rep. Carlos Curbelo's (R-FL) MARKET CHOICE Act. But the emissions uncertainty under a carbon tax is a significant disadvantage, especially since these proposals eliminate or suspend existing regulations, leaving many environmental groups reluctant to embrace them.

This creates a demand for ways to reduce a carbon tax's inherent emissions uncertainty. One leading approach is a tax adjustment mechanism (TAM), which automatically adjusts the tax rate based on where actual emissions are relative to a legislated target.

This paper discusses the role for TAMs in carbon tax design and trade-offs across alternative designs. We use the TAM in the Curbelo bill and a TAM proposed by Metcalf (2018) for the CLC proposal as illustrative examples to show that TAMs can substantially reduce emissions uncertainty. We then consider variations on these mechanisms, illustrating how different design choices affect expected costs and emissions outcomes. Most design dimensions show clear tradeoffs, improving some outcomes while worsening others. Thus, the optimal design varies substantially based on what goals the TAM is intended to achieve.

Key Words: carbon tax, uncertainty, hybrid instruments, emissions tax

JEL Classification Numbers: Q58, Q52, H23, D8

This article has been accepted for publication in Review of Environmental Economics and Policy published by Oxford University Press (doi: 10.1093/reep/rez018).

Contents

1. Introduction

There is increasing political interest in the United States in carbon taxes, which have long been popular among economists as a means to address climate change. Notably, two recent proposals from center-right Republicans are attracting interest: the Climate Leadership Council's Baker-Shultz Carbon Dividends plan, and Rep. Carlos Curbelo's (R-FL) MARKET CHOICE Act. But the emissions uncertainty under a carbon tax is a disadvantage both from an economic perspective (if marginal damages are increasing in emissions, uncertainty raises expected damage) and, perhaps more importantly, from a political perspective. And because center-right carbon tax proposals eliminate (as in Baker-Shultz) or suspend (as in the Curbelo bill) existing or new regulations on carbon dioxide, uncertain emissions reductions have left many environmental groups reluctant to embrace these carbon tax proposals.

One promising approach for reducing this uncertainty (and thus winning support from environmental groups) is to include an automatic tax adjustment mechanism (TAM), which would adjust the path of the carbon tax rate based on where actual emissions are relative to a legislated emissions path.[1] On a pre-specified schedule, actual emissions would be compared to that prespecified emissions path, and if emissions exceed the specified level, then the tax rate automatically adjusts upward (the mechanism might also allow the tax rate to adjust downward if emissions are below the legislated path). These automatic adjustments would thus push emissions back toward the path.

The Swiss Carbon Tax Law is (to our knowledge) the first enacted carbon tax to include such a TAM, and it provides a simple example of how a TAM works. The tax, which covers CO_2 emissions from electricity and heating, started in 2008 at a rate of 12 Swiss francs ("CHF") per ton of CO_2, rising to 36 CHF by 2012, with the increase after that dependent on emissions: if emissions in 2012 exceeded 79 percent of 1990 emissions, then the tax rate would increase to 60 CHF on Jan. 1, 2014. The law included similar adjustments based on emissions in 2016 and 2018.[2]

In the United States, Curbelo's 2018 MARKET CHOICE Act was the first carbon tax bill introduced in Congress to include a TAM. The bill would implement a tax of $24/ton, with a default annual increase of 2% over inflation. And in a review occurring every two years, if realized emissions reductions are behind the goals specified in the

[1] There is no established naming convention for these mechanisms. Hafstead et al. (2017) use the term Tax Adjustment Mechanism for Policy Pre-Commitment (TAMPP); here we shorten that to tax adjustment mechanism (TAM). Metcalf (2018) refers to his proposal as an Emissions Assurance Mechanism. We have also heard the terms Emissions Certainty Mechanisms (ECM) and Environmental Integrity Mechanisms (EIM). However, these latter two terms are less specific: they could refer either to mechanisms that adjust tax levels or to other mechanisms designed to increase emissions certainty.

[2] The bill specified a tax increase to 96 CHF in 2018 if 2016 emissions exceeded 73 percent of 1990 emissions and another tax increase to 120 CHF in 2020 if 2018 emissions exceeded 76 percent of 1990.

bill, the TAM automatically increases the tax rate by an additional $2/ton. The Baker-Shultz Carbon Dividends plan doesn't yet include a TAM, but Metcalf (2018) proposes an adjustment mechanism that could easily be incorporated into the Baker-Shultz proposal.

A small but growing handful of papers look at the problem of reducing emissions uncertainty under an emissions tax (especially a carbon tax), and the role that a TAM could play. Metcalf (2009) proposes a carbon tax with a TAM, and provides a simple numerical illustration of how that TAM could affect emissions outcomes. Hafstead et al. (2017) discusses TAMs in more detail, outlining a variety of choices policymakers will confront in designing a TAM and providing an intuitive discussion of the tradeoffs such choices entail. And recent working papers by Harris and Pizer (2018) and Hafstead and Williams (2019) explicitly model CO_2 emissions uncertainty and quantify how tax adjustments mitigate that uncertainty.[3]

In this paper, we aim to provide an accessible, intuitive, policy-oriented discussion of TAMs, drawing on results from the simulation model from Hafstead and Williams (2019). We discuss the potential roles that a TAM can play and key elements of the design of a TAM. We use simulation results for the Curbelo and Metcalf (2018) TAMs as illustrative examples to show the extent to which proposed adjustment mechanisms can reduce emissions uncertainty. To show how design choices influence emissions and cost outcomes, we then simulate modified versions of each policy. These results largely mirror previous results from Hafstead and Williams (2019).

There are a variety of potential goals that a TAM could be intended to achieve as part of a carbon tax policy. These could include a) providing insurance against very-high-emissions outcomes, b) reducing emissions uncertainty more generally, c) increasing the likelihood of hitting specific emissions targets, or d) managing cost uncertainty. How one ought to design the TAM – and what metrics one might want to use to evaluate outcomes under different designs – will vary substantially based on what goal(s) the TAM is intended to achieve.

The simulation results show that the inclusion of a TAM can substantially reduce emissions uncertainty and increase the probability of hitting a particular emissions target. For example, the simulations indicate that Metcalf's TAM proposal raises the probability of hitting its target for cumulative emissions out to 2035 from 43.4% to 55.6%, and modified versions could push that probability substantially higher still (e.g., adjusting if either annual or cumulative emissions are above benchmarks would raise the probability to 67.7%). This illustrates that when emissions are exceeding benchmarks, adjusting the tax rate – even by relatively small amounts – can often push

[3] Aldy (2017) and Murray et al. (2017) discuss approaches other than a TAM for reducing emissions uncertainty under a carbon tax, such as a legislative review process or regulatory backstop. Aldy et al. (2017) provides a broad introduction to the problem and an overview of the range of potential policy elements to address it, including TAMs and alternatives.

emissions back to the desired path. But while these mechanisms reduce uncertainty, they don't eliminate it: there is still a substantial chance that emissions will be above the target even with a TAM (and while more aggressive adjustments lower the chance of missing a target, they still don't reduce it to zero).

The economic tradeoff is that a TAM will typically increase the cost of the policy, though the effect on costs varies widely based on the design of the TAM, and one could design a TAM that would lower expected costs. A TAM like the one in the Curbelo bill – which can only adjust the tax rate upwards – can only increase costs (since a higher tax rate implies higher costs, all else equal). But that cost increase can be quite small. For example, the TAM in the Curbelo bill increases the expected cost of CO_2 reductions by less than $1/ton; adjusting no more every other year, and by only $2 each time, keeps the expected cost down. In general, the larger and more frequent the upward adjustments, the more the TAM will increase costs. A TAM could also be designed to adjust the tax rate downward if emissions are lower than expected, and those downward adjustments will lower costs. Thus, all else equal, a two-sided TAM (one that can adjust the tax rate either up or down) will have lower expected costs than a one-sided TAM that can only adjust the rate upward, and a TAM that can only adjust down would reduce costs relative to a tax without a TAM.

When we move to evaluating TAM design elements, a few seem to be broadly beneficial. For example, smaller, more-frequent adjustments do better on most metrics we consider compared to larger, less-frequent adjustments (e.g., adjusting by $1/ton each year rather than $2/ton every other year). But even these design elements aren't consistently better on every metric. And for most dimensions of TAM design, there are clear tradeoffs. Changes that reduce emissions uncertainty (such as making larger adjustments) generally increase costs. And there are tradeoffs among different emissions goals as well: for example, a two-sided TAM can cut costs relative to a one-sided TAM without noticeably reducing the probability of hitting the emissions target (or could be combined with slightly larger adjustments to yield a higher probability of hitting the target without increasing cost), but the tradeoff is that it raises expected emissions relative to the one-sided TAM.

The next section of the paper discusses different goals one might have for a TAM, how some simple TAMs would affect the distribution of potential emissions outcomes, and what metrics one might want to use to gauge the performance of a TAM. Section 3 looks at TAM design elements, providing details on the TAMs in the Curbelo bill and Metcalf's (2018) proposal, and going over a range of dimensions along which TAM designs can vary. Section 4 provides a brief overview of the model we use to simulate the effects of different TAMs. Section 5 presents simulation results for the two proposed TAMs and a range of possible modifications to those proposals. The final section concludes and discusses implications of these results.

2. Potential Goals for a Tax Adjustment Mechanism

This section reviews the various goals one might have when including a TAM in a carbon tax policy, looks at how different simple TAMs affect the distribution of possible emissions outcomes, and discusses what metrics one might want to use to evaluate a TAM's effectiveness. Metcalf (2009) introduced the REACT proposal, a simple TAM, as a way to reduce uncertainty about emissions under a tax. But there are a variety of ways to reduce uncertainty: for example, reducing emissions in unexpectedly-high-emissions cases and raising emissions in unexpectedly-low-emissions cases each reduce emissions uncertainty, but they have potentially quite different economic and political implications. Thus, simply saying the goal is to reduce uncertainty isn't sufficient: it matters exactly how that uncertainty is reduced. Moreover, there are other effects of a TAM – beyond reducing emissions uncertainty – that policymakers might find attractive.

The first potential goal one might have for a TAM would be to avoid (or at least reduce the likelihood of) especially high-emissions outcomes. Here, the role of the TAM would be like an insurance policy against unexpectedly high emissions. This could be desirable either on economic grounds (if one thinks that damages are substantially nonlinear, with high-emissions cases causing much more damage) or political grounds (e.g., if environmental groups are especially concerned about emissions above a certain threshold). In this case, it would make sense to have a TAM that is triggered if emissions are higher than expected, raising the tax rate and thus reducing emissions. Such a TAM would affect the upper tail of the distribution of emissions (reducing emissions in high-emissions cases).

Figure 1: Distributions of Cumulative Emissions

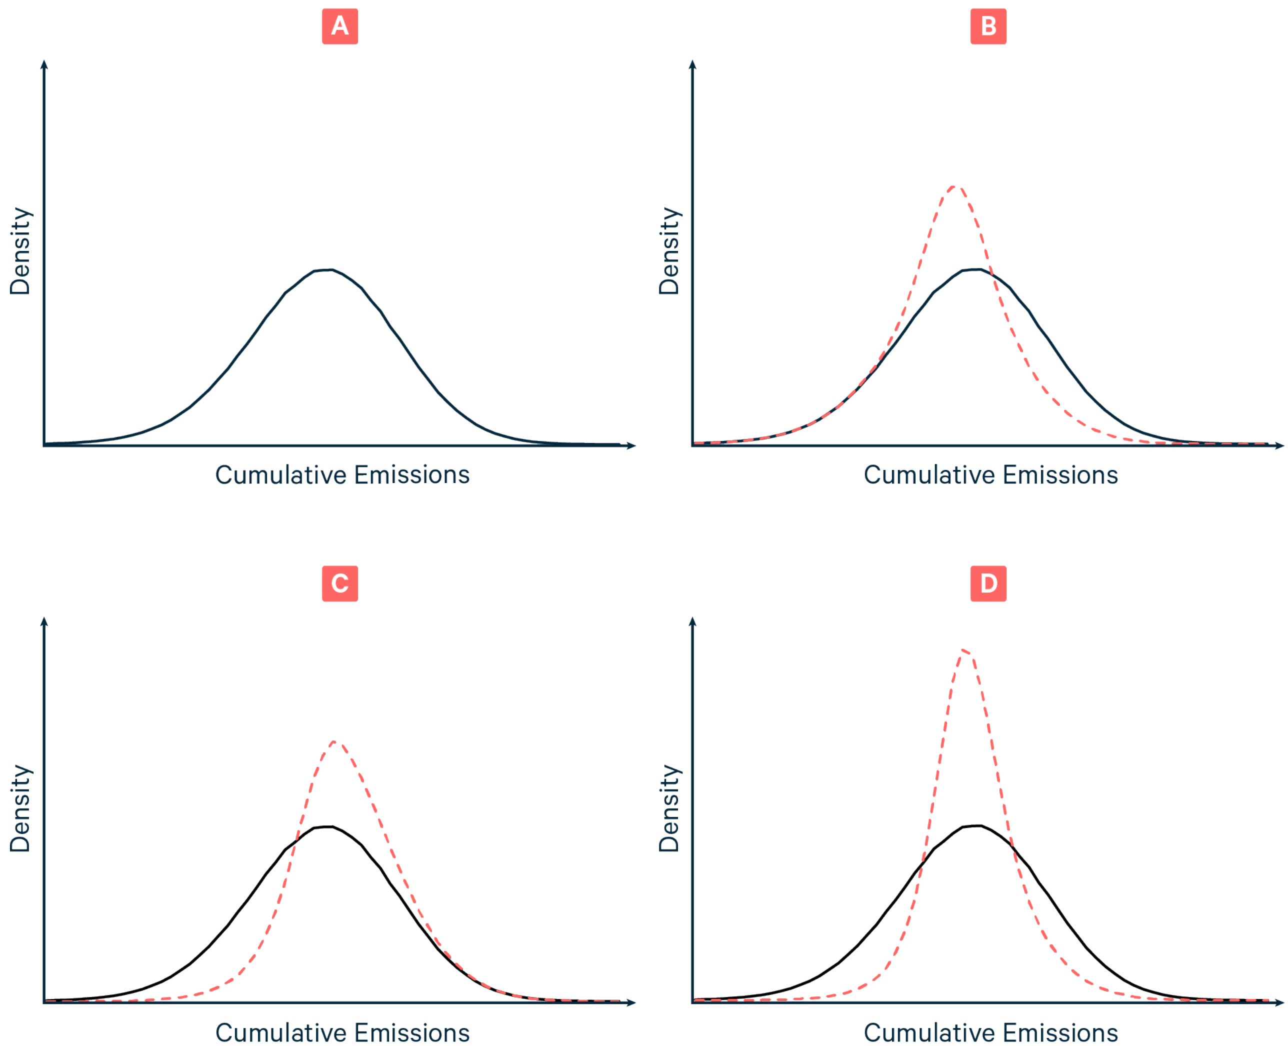

Let's look at how such a TAM would affect emissions. Start by considering the range of possibilities for a carbon tax without a TAM. Figure 1a shows a stylized probability distribution under that policy. The horizontal axis shows the range of possible outcomes for cumulative US carbon emissions over some period of time (now to 2035, for example), and the vertical axis shows how likely any particular outcome is. The curve is bell-shaped: emissions outcomes in the middle are most likely, with probability falling as one moves to lower-emissions (to the left) or higher-emissions (to the right) cases.

This reflects a wide variety of sources of uncertainty in factors which come together to determine carbon emissions: uncertainty about economic growth, technological progress, fossil fuel prices, how emissions respond to a carbon tax, etc. If each of these factors pushes emissions to be higher than expected, then we'll be in a case on the far right of the figure, and if each pushes emissions lower, then we'll be in a case on the far left. But because those outcomes require everything to come together in a particular direction, they're not very likely; the more likely outcome is that some factors will push emissions higher than expected and others will push emissions lower, and we'll wind up somewhere in the middle.

Now consider the effect of adding a TAM designed to reduce the likelihood of high-emissions cases: one that triggers a tax rate increase if emissions are higher than expected. Figure 1b shows how this TAM changes the distribution. The solid black line is the distribution for the carbon tax without the TAM (the same as in Figure 1a). The dashed gray line shows the distribution with the TAM.

Adding the TAM means that when emissions are higher than expected, the tax rate adjusts upwards, reducing emissions. Those high-emissions cases are the ones on the right-hand-side of Figure 1b, so reducing emissions in those cases pulls the right-hand-side of the distribution toward the middle. This shifts the right-hand-side of the dashed gray line (the distribution with the TAM) to the left (lower emissions) relative to the solid line. This doesn't eliminate high-emissions cases: if emissions would have been extremely high without a TAM, then emissions could still be somewhat high even after the TAM pulls them down substantially. But the TAM does substantially lower the probability of high-emissions cases.

Because this TAM is only triggered when emissions are higher than expected, it only matters in those cases.[4] Thus, as one moves to the left on the figure, the two distributions (with and without the TAM) look more and more alike, and indeed become essentially identical as one gets out toward the left tail. This type of TAM has a substantial effect on the right tail of the distribution, but it has little to no effect on the left tail.

A second potential goal one might have for a TAM would be to lower the cost of the policy. Simply setting a lower carbon tax rate would lower costs, but with the tradeoff of higher emissions. Suppose that one is OK with raising emissions in low-emissions cases, but worried about raising emissions in high-emissions cases. Then simply setting a lower rate in all cases could be undesirable, but one could use a TAM to lower the rate only if emissions turn out to be lower than expected. This TAM would be the mirror image of the TAM we just discussed: instead of raising the tax rate when emissions are high, it would lower the rate when emissions are low.

[4] Note, though, that because the TAM is evaluated each year, and the diagram shows cumulative emissions over a longer period, it's possible for the TAM to matter even in cases where cumulative emissions over the entire period are lower than expected (e.g., a case in which emissions are higher than expected for the first few years, triggering the TAM, but then in subsequent years would have been lower than expected even in the absence of the TAM; in this case, the TAM makes a difference even though cumulative emissions over the whole period would have been lower than expected). Thus, the TAM can matter even for the left-hand side of the distribution. But such cases are relatively rare.

Figure 1c shows the effect of such a TAM on the distribution of possible emissions outcomes. Again, the distribution without a TAM is in black, and the distribution with the TAM is the dashed gray line. Because it's triggered only when emissions are low, his TAM raises emissions in the low-emissions cases (the left-hand side of the figure) and has little or no effect in high-emissions cases (the right-hand side). This TAM is the mirror image of the previous one discussed, and thus its effect is also a mirror image, pulling in the left-hand tail of the distribution, whereas the previous TAM pulled in the right-hand tail.

Now think about a combination of these goals: suppose one wants insurance against high-emissions cases, but without increasing costs much, and is OK with raising emissions somewhat in low-emissions cases. Again, this could be either for economic reasons (if marginal damage increases with emissions, it's efficient to put more effort into reducing emissions in the high-emissions, high-marginal-damage cases than in the low-emissions, low-marginal-damage cases) or for political reasons (if there's political pressure to minimize high-emissions cases and to control costs, but little political benefit to especially low-emissions outcomes). With those goals, a combination of the two TAMs already considered could make sense: a TAM that would adjust the carbon tax rate upward if emissions are higher than expected and down if emissions are lower than expected. We refer to this as a two-sided TAM (in contrast to the previous TAMs discussed, each of which is one-sided).

Figure 1d shows how a two-sided TAM would affect the emissions distribution. Like the first one-sided TAM, it raises the rate when emissions are high, thus pulling in the right tail of the distribution. And, like the second one-sided TAM, it lowers the rate when emissions are low, thus pulling in the left tail of the distribution. The net effect is to squeeze the whole distribution, making it more likely that emissions will be close to the expected level and reducing the probability of emissions being substantially higher or lower than expected. Thus, in this case, the TAM acts as an automatic stabilizer, reducing uncertainty in both directions.

A more extreme version of this might make sense when the goal is to be below a particular target level of emissions. That target could come from some known emissions threshold that will cause substantially greater damage if exceeded, or could simply be a country has promised to be below a particular target and wants to keep that promise.[5] If being below the target is important, but there's no difference between being slightly below the target and being well below it, then one would want a carbon tax rate that reduces expected emissions below that target, plus a two-sided TAM to pull in the tails of the emissions distribution (since pulling down emissions in the high-emissions tail raises the probability of hitting the target and pulling up emissions in the low-emissions tail reduces costs). Indeed, if all one cares about is costs and

---

[5] We mention the damage-threshold argument here because it's potentially relevant for cases other than climate change, but it doesn't really apply for a single country's climate policy: while there certainly are such threshold effects, we don't have enough certainty about where they are to target them accurately, and even if we did, what matters is total global emissions, so the target any single country would need to hit depends on what every other country does.

the probability of hitting the target, one would want a very aggressive TAM (one that raises the tax rate drastically when emissions are high, and reduces the rate drastically when emissions are low).[6]

All of the goals we've discussed so far imply reducing emissions uncertainty, and that's certainly the focus of this paper. But TAMs also affect uncertainty about costs. And one can imagine that policymakers might care not just about the expected cost of reducing emissions, but also about the risk of high-cost cases. Thus, while the analysis in this paper focuses primarily on emissions risk and on expected cost, we also briefly discuss effects on uncertainty about abatement costs.

A TAM could be designed to serve other purposes as well. For example, suppose a policymaker wants to aggressively reduce carbon emissions, but it's politically difficult to enact a high carbon tax rate. In that case, a TAM could serve as a back-door way to get a higher tax: including a TAM with very low emissions benchmarks (low enough that emissions are almost certain to exceed those benchmarks, thus almost guaranteeing upward adjustments in the tax rate) would mean that even a tax rate that starts low could quickly rise much higher, and in a way that's less obvious than simply specifying a big annual increase.

The different potential goals – and the different ways in which a TAM can change the shape of the emissions distribution – make it hard to generalize about both TAM design and the metrics used to evaluate a TAM. If the goal is to hit a particular target, then it's natural to evaluate the effectiveness of a given TAM based on the probability of hitting that target. If the goal is to reduce emissions uncertainty, but without a specific target, then it's natural to use a metric that measures uncertainty, such as the standard deviation of emissions. If the goal is to limit especially high-emissions outcomes, then it's natural to use a metric focused on the high-emissions tail of the distribution, such as the top end of the 95% confidence interval for emissions (i.e., the 97.5th percentile of emissions). Because different TAM designs affect the shape of the emissions distribution differently, they'll affect those different metrics differently. Thus, in modeling the effects of TAMs and evaluating how effective a given TAM is, we'll look at a range of different metrics: all of those mentioned in this paragraph, plus expected emissions and expected costs.

[6] Of course, for this goal, a tradable permit system could do even better, with the number of permits set at a level just below the target. That system would ensure emissions hit the target, and would mean the permit price adjust upward or downward as much as necessary to hit the target – making it similar to a very aggressively adjusting TAM. But even in such a case, there may be other reasons – particularly political reasons – to prefer a tax over a tradable permit system.

3. Tax Adjustment Mechanism Design

Adding a TAM to a carbon tax requires specifying a variety of design elements that control when and how the tax rate should adjust to realized emissions. In this section of the paper, we use Metcalf (2018) and the Curbelo TAM as two concrete examples that emphasize the diversity in design choices, and provide a brief discussion of each of the key design elements. For a more extensive list of TAM design choices and discussion of the effects of each choice, see Hafstead et al. (2017).

Table 1: Metcalf (2018) and Curbelo TAM Design Elements

Design Element	Metcalf (2018)	Curbelo
Control Period	2021-2035	2020-2032
Benchmark Path	Arbitrary	Model-Based
Benchmark Comparison	Cumulative	Cumulative
Type of Adjustment	Growth Rate	Discrete
Frequency of Adjustment	Annual after 2024	Biennial after 2021
Adjustment Trigger	Asymmetric Two-Sided	One-Sided
Adjustment Size	+5% (-5%)	$2 ($0)

In Metcalf (2018), the author proposes an "emissions assurance mechanism" (EAM) to add emissions certainty to the Baker-Shultz Carbon Dividend Plan. The Carbon Dividend Plan as discussed in Metcalf (2018) begins in 2021 with a $43 tax on carbon dioxide emissions (in $2021) and rises 5 percent above the rate of inflation annually.[7] Metcalf's EAM specifies an emissions pathway (which we term the *benchmark path*) for the years 2021 to 2035 (the control period) to meet an annual emissions target of 45 percent below 2005 emissions in 2035. If, in any year (*frequency of adjustment*) after 2025, cumulative emissions (from the first year of the policy) through the

[7] The Climate Leadership Council is considering growth rates of 3-6 percent above inflation annually for its Carbon Dividends Plan (See Halstead et al. 2018) but has not endorsed a specific growth rate. Metcalf (2018) assumes a 5 percent growth rate.

previous year exceed the cumulative benchmark path emissions (*benchmark comparison*), the annual growth rate increases from 5 to 10 percent (*size and type of adjustment*) in the following year.[8] Alternatively, if cumulative emissions are 90 percent or less of the cumulative benchmark path, the tax rate is held constant in real terms (an example of a two-sided *adjustment trigger*).

Curbelo's bill introduces a carbon tax of $24 per ton (in $2020) beginning in 2020 with an annual growth rate of two percent above inflation. Section 9901(a)(3) defines a "Rate adjustment based on emission levels". In subsection (A), it identifies a cumulative emissions pathway through 2030; this benchmark path is based on a CGE analysis of the expected emissions pathway under the act.[9] In subsection (B), the bill requires the Secretary of the Treasury to, beginning in 2022 and every two years after (*frequency of adjustment*), determine if a price adjustment is required. If cumulative emissions through the previous year exceed the cumulative benchmark path (*benchmark comparison*) the tax rate adjusts upward by $2 per metric ton (*size and type of adjustment*) for the following calendar year. No adjustments are required if emissions are below the benchmark (one-sided *adjustment trigger*).

3.1. Comparison of Alternative Design Elements

3.1.1 Control Period

The Metcalf EAM is specifically designed to help the policy achieve the 2035 target of emissions levels 45 percent below 2005 levels. The Curbelo bill focuses on emissions outcomes through 2030. The designers of these policies could alternatively have specified longer control periods to meet longer-term climate goals. But the large uncertainties in modeling carbon prices past 2035 make it hard to specify policy that far into the future. One can argue that it's better for the legislature to revisit the issue than to try to specify a TAM over that long a period.

---

[8] Due to a lag in the publication of annual emissions levels, emissions through year t-1 are evaluated in year t and any price change would occur in year t+1. For example, in Metcalf (2015) cumulative emissions through 2028 are evaluated in year 2029 and an announcement to increase (or decrease) the price for 2030 is made based on the outcome of the comparison.

[9] The emissions pathway in the Curbelo bill is defined over greenhouse gas emissions. The pathway was defined from model-based analysis of energy-related carbon dioxide emissions (Hafstead, 2018) and adjusted for reductions in other greenhouse gases estimated outside of model-based analysis.

3.1.2. Benchmark Path

As HW2019 show, the benchmark path is one of the most important design elements for any TAM. Metcalf's EAM specifies a straight-line path from a specified initial level of reductions to an annual target in 2035; the path uses limited information from modeling analysis to choose a path, and thus the path is somewhat arbitrary. In such a benchmark path, the choices of initial reductions and endpoints are vital to the performance – in terms of both costs and emissions outcomes – of the TAM. Alternatively, the Curbelo TAM uses model-based estimates (predicted emissions under the carbon tax) to define its emissions path. Based on the analysis in HW2019, we expect TAMs with model-based benchmark paths to perform better (i.e., have lower costs and better outcomes on most emissions metrics) than otherwise similar TAMs that use arbitrary paths.

3.1.3 Benchmark Comparisons

Both the Metcalf and the Curbelo TAM proposals compare cumulative emissions to cumulative benchmarks. Alternatively, the trigger could be based on emissions in a single year or an average over a few years (or a combination of annual and cumulative benchmarks). Metcalf argues that a cumulative comparison provides more flexibility and insulates against cyclical fluctuations in emissions. That flexibility and insulation means the TAM will be triggered less frequently, reducing expected costs but also leading to higher emissions.

3.1.4 Frequency, Type, and Size of Adjustment

The two policies differ in the frequency, type and size of the adjustment. The Metcalf proposal includes both an extended grace period and annual adjustments to the growth rate of the policy. The Curbelo TAM, on the other hand, uses discrete adjustments (where the tax change is a fixed dollar amount) biennially. Frequency and type of adjustments are closely related. With annual adjustments, there is little or no difference between growth-rate and discrete adjustments (e.g., if the tax rate at the time of adjustment is $20/ton, adjusting upward by 5% is identical to adjusting upward by $1/ton). With longer adjustment intervals, however, the type of adjustment matters much more (e.g., with two years between adjustments, discrete adjustments mean the entire tax change happens for the next year, followed by no change in year 2, whereas with a growth rate adjustment, the tax changes in both years).

All else equal, increasing the adjustment size will increase costs and decrease uncertainty. But adjustment size also interacts importantly with other parameters: for example, a 2% adjustment biennially has a somewhat similar effect to a 1% adjustment annually (because the total potential adjustment over two years is the same). Thus, it's important to look at the TAM as a whole. A well-designed TAM with small adjustments might well achieve comparable emissions outcomes (at lower cost) than a poorly-designed policy with larger adjustments.

3.1.5 Adjustment Trigger

One of the key design decisions is whether to allow for both upward and downward adjustments in the price path. The Metcalf proposal is an asymmetric two-sided policy: it allows for upward and downward adjustments to the growth rate, but it requires emissions levels to be much lower than the benchmark path for a downward adjustment. The Curbelo TAM, on the other hand, does not allow for downward adjustments: it is a one-sided trigger. As discussed in the previous section, this choice determines whether the TAM affects one or both tails of the emissions distribution.

4. Modeling Approach

To evaluate how TAMs alter the distribution of expected emissions, we utilize the reduced-form model from Hafstead and Williams (2019) (hereafter shortened to HW2019). The simple model projects both GDP and emissions intensity (emissions per unit of GDP) from 2017-2035 and is well-suited to project distributions of emissions over time. While this reduced-form model lacks microfoundations and forward-looking investment behavior, when appropriately calibrated it generates emissions paths similar to those from far more complex environmental-energy CGE models. This reduced-form approach also easily incorporates multiple forms of uncertainty in future output and emissions intensity. Finally, the simplicity of the model lends itself to a Monte Carlo approach: the model can be solved several thousand times in mere seconds.

HW2019 describes the full reduced-form model with uncertainty and readers should refer to that paper for a full description of the model and its calibration. Here, we briefly outline its key features.

4.1. Modeling Uncertainty

The model includes six sources of uncertainty in future emissions levels. The first two are trend uncertainty (uncertainty about the long-run average rate of change) in both GDP and emissions intensity.[10] Because the trend is persistent, this trend uncertainty is more important for long-run emissions predictions than for short-run emissions predictions: small changes in long-run growth rates make little difference in the short run, but compound over time to have big effects on future emissions.

The model also includes cyclical uncertainty in output and emissions intensity, the annual deviations in output or emissions intensity around their long-term trends.[11] Cyclical uncertainty in output represents business-cycle fluctuations and cyclical uncertainty in emissions intensity reflects short-run changes in the relative prices of fuels or emissions-reducing technologies.[12] This cyclical uncertainty is important in the short run, but the cyclical fluctuations tend to average out over time, making them less important in the longer run.

[10] Trend uncertainty is modeled as parameter uncertainty. That is, in the model the underlying trend (i.e., long-run) growth rates for GDP and emissions intensity don't change over time, but policymakers don't know those trend growth rates – only their probability distributions – when setting the policy.

[11] This cyclical uncertainty is modeled as AR(1) processes with annual stochastic shocks to output and emissions intensity.

[12] For example, a short-run increase in the price of natural gas relative to coal (caused by an extremely cold winter) would increase emissions intensity relative to trend. Alternatively, a short-run increase in price of solar panels (caused by tariffs on imported panels) would also increase emissions intensity relative to trend.

Finally, the HW model includes uncertainty in the both the initial elasticity of emissions intensity with respect to an emissions price (how much emissions intensity decreases in response to a given carbon price) and the rate of change in the elasticity over time.[13] As shown in HW2019, this price elasticity uncertainty is the dominant source of overall emissions uncertainty under a carbon tax.[14]

4.2. CGE Calibration

Central case estimates for trends can be calibrated to any projection of GDP and emissions (we use Energy Information Administration's (EIA) Annual Energy Outlook (AEO) 2017 reference case (without Clean Power Plan) projections). To generate emissions path in the HW model that approximate paths from complex traditional CGE models, the HW price elasticities are calibrated to average and marginal price elasticities in the Goulder-Hafstead E3 model (any CGE model could be used). The full procedure, described in HW2019, yields emissions levels in each model that differ by no more than +/- 3-4% per year when holding fixed trend parameters.

4.3. Uncertainty Calibration

Estimated emissions uncertainty under a carbon tax crucially depends on the calibration of the uncertainty parameters in the HW model. Cyclical uncertainty is estimated from historical fluctuations around long-run output and emissions intensity trends. For trend uncertainty, we introduce normal distributions around the central case Annual Energy Outlook (AEO) 2017 estimates; the standard deviations are calibrated such that the 95% confidence interval matches the range of estimates from AEO's alternative reference case scenarios. For price elasticity uncertainty, we use log-normal distributions around the estimated values from the CGE calibration. The size of the confidence interval for the initial elasticity is approximately +/- 50% of the point estimate to allow for large uncertainty in price elasticities.

[13] Price elasticity uncertainty is also modeled as parameter uncertainty and we do not include stochastic shocks to the elasticity; we do not allow for technology breakthroughs to randomly increase the price elasticity in some periods.

[14] The price elasticity uncertainty is a key difference between the HW and Harris and Pizer models; Harris and Pizer assume price elasticities are known.

Because there is relatively little hard evidence to go on, uncertainty calibration is more art than science. As a verification process, HW2019 compared the HW model distribution of emissions reductions to a range of reductions generated by 11 different models that participated in the Stanford Energy Modeling Forum 32 (EMF32). Without specifically calibrating to the EMF32 results, HW2019 found that the range of emissions reductions in the EMF32 models and the within-model variation in the HW model were similar across the set of pricing scenarios considered in EMF32.[15]

[15] See Hafstead and Williams (2018) for details.

5. Quantitative Analysis of Tax Adjustment Mechanisms

This section uses simulation results from the HW model for the CLC/Metcalf and Curbelo proposals (and variants of those proposals) to illustrate how the carbon taxes in these proposals affect emissions, how the TAM components of the proposals reduce emissions uncertainty, and how changing design choices can alter the effects of those TAMs. These results echo many of the results from HW2019.

While the purpose of this section isn't to compare the two proposals, it's nonetheless useful to make comparisons between them in order to show how policy differences affect costs and emissions. Comparisons are somewhat complicated, however, because the timing of the two proposals differs significantly on two dimensions: the year of implementation and the control period. The CLC/Metcalf proposal introduces an economy-wide carbon tax in 2021, Curbelo's bill in 2020.[16] The TAM control period in Metcalf runs through 2035 whereas the control period in the Curbelo TAM runs through 2030. To make comparisons simpler, we extend the Curbelo control period through 2035 by linearly extending the emissions benchmark paths. We still allow the implementation year to differ between policies, but we report cumulative emissions for the years 2020 to 2035 for each policy (applying reference case emissions – that is, projected emissions in the absence of a carbon tax – in 2020 for the Metcalf proposal).

5.1. Emissions Uncertainty of a Carbon Tax

We start by looking at emissions outcomes and expected costs of the Carbon Dividends Plan (with 5% real growth rate) and the Curbelo bill in the absence of TAMs. This provides an indication of what each carbon tax does without adjustments, and then we can compare to results for the policies with TAMs to show the effects of the TAMs relative to the same tax without adjustments.

Figure 2 displays the distribution of cumulative emissions from 2020 to 2035 under business as usual (BAU) and the two non-TAM tax plans. The carbon tax policies significantly shift the distribution of emissions. The obvious (and unsurprising) change in the distribution is that imposing a carbon tax lowers cumulative emissions, causing the distribution to shift to the left. But the carbon tax also widens the distributions: that is, emissions uncertainty increases. This arises because of the importance of uncertainty about the elasticity of emissions intensity with respect to the carbon price. If that elasticity is high, then the carbon tax reduces emissions a lot, but if it's small, the

[16] The Curbelo bill also eliminates the federal excise tax on motor vehicle fuels. Hafstead (2018) indicates that eliminating this tax has minimal effects on CO2 emissions.

emissions reduction is much smaller. And the uncertainty caused by that elasticity gets multiplied by the carbon price, which is why the Carbon Dividends Plan (with a higher baseline tax carbon price) has a wider emissions distribution.

That higher baseline price also means that the Carbon Dividends Plan achieves greater emissions reductions, despite the one-year lag in implementation relative to Curbelo's bill. [17]

Figure 2: Distribution of Cumulative Emissions Under Alternative Scenarios, 2020-2035

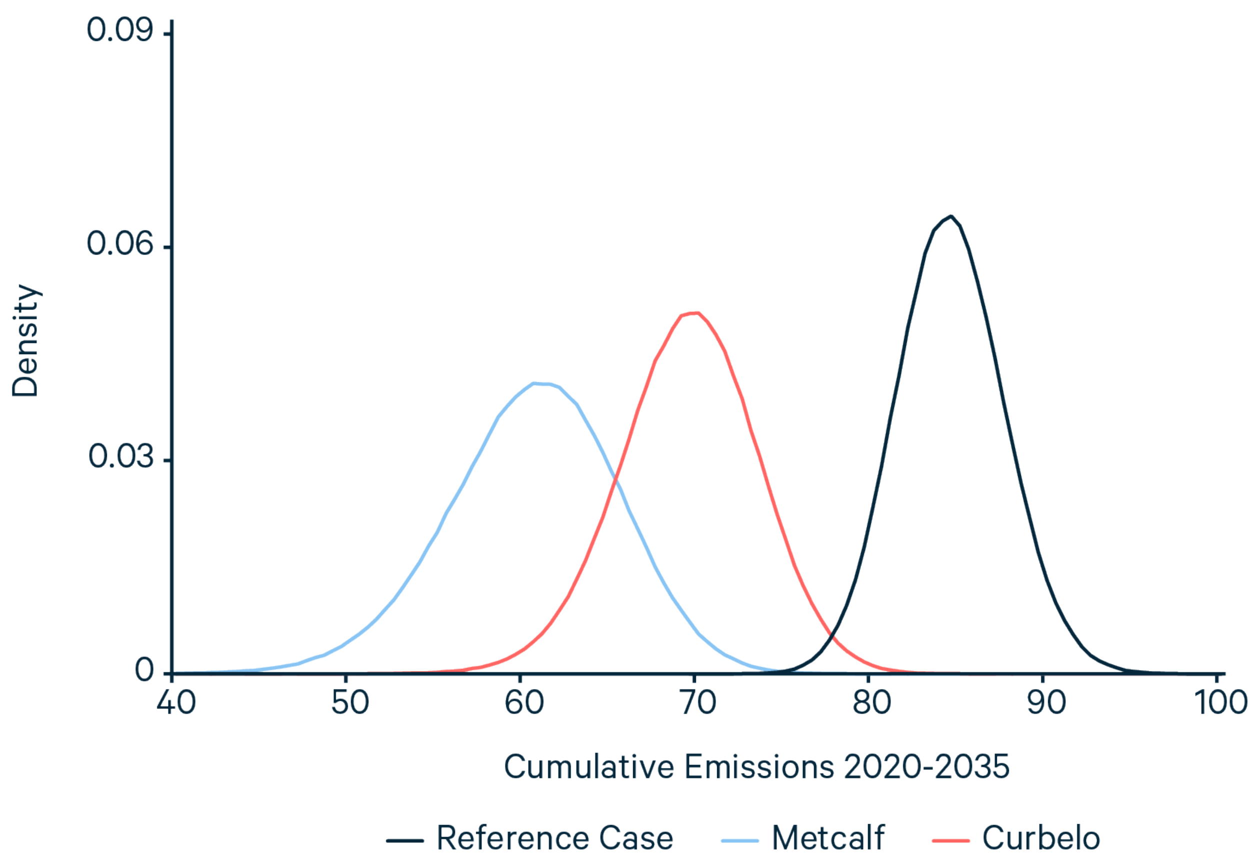

[17] Note that while the distributions under the two policies overlap, one should not conclude that a lower tax rate could achieve less emissions (or more emissions reductions) than a tax with a higher rate. The left-hand tail of each distribution reflects cases in which multiple uncertainty elements (economic growth, technological changes, world fossil fuel prices) come together to reduce emissions, whereas the right-hand tail reflects the opposite, cases in which those elements come together to boost emissions. Thus, the overlapping distributions mean that differences in those uncertain elements can matter more than the difference in tax rates. But it is still true that for any given case (i.e., holding those elements constant), a higher carbon tax will result in lower emissions. Or, put more concisely, one can't just look to see if the distributions overlap because the two distributions are highly correlated.

5.2. Emissions Uncertainty of a Carbon Tax under Alternative Tax Adjustment Mechanisms

Figure 3 displays the distribution of cumulative emissions from 2020 to 2035 for the CLC/Metcalf and the Curbelo policies with and without TAMs. As described in Section 2, the TAMs primarily affect uncertainty (pulling in one or both tails of the distribution), rather than shifting the whole distribution. The baseline price path largely determines the level of emissions, and the TAMs role is to change the shape of the distribution by changing the tail(s).

Figure 3: Distribution of Cumulative Emissions with and without Tax Adjustment Mechanisms, 2020-2035

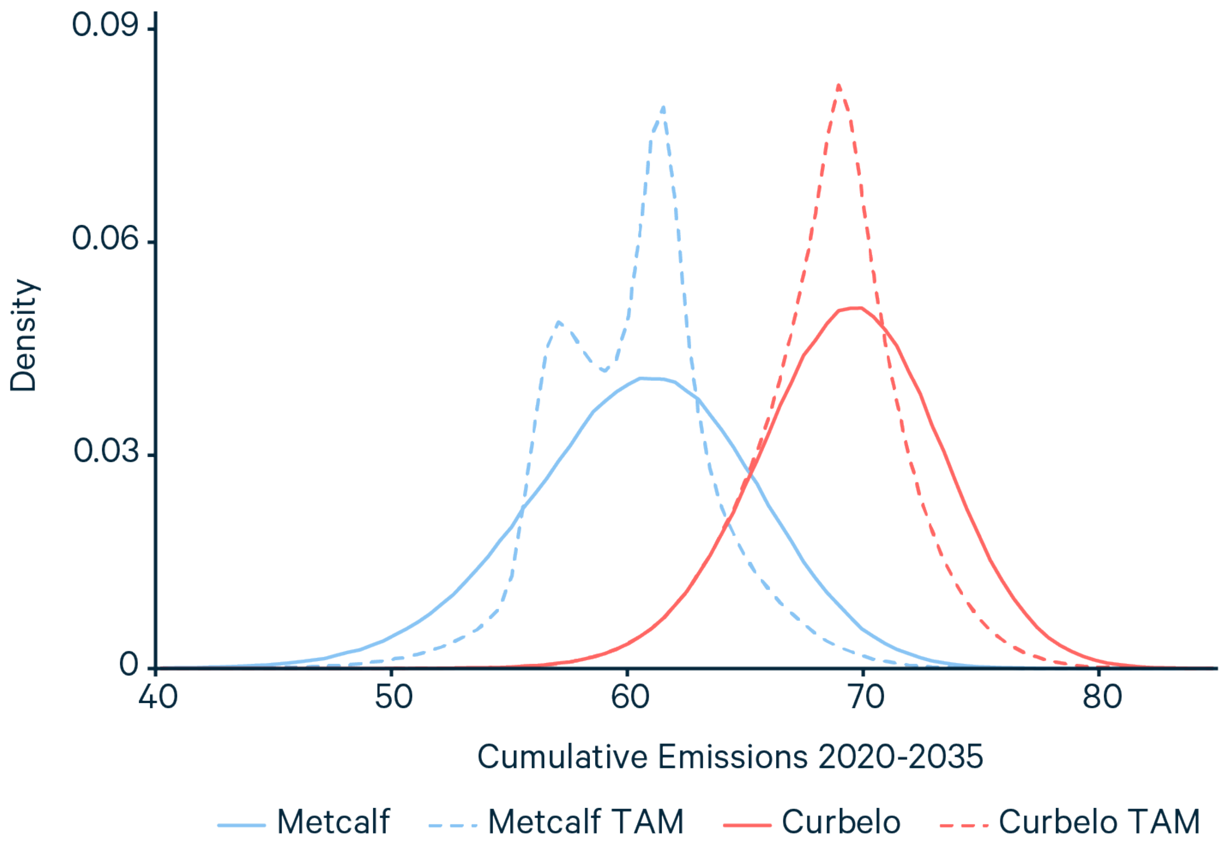

The Curbelo TAM is a one-sided mechanism and thus pulls in only the right tail of the distribution (reducing emissions in high-emissions cases), leaving the left tail (low-emissions cases) unaffected. The Metcalf TAM is an asymmetric two-sided mechanism and trims both the right and the left tails. Interestingly, the asymmetry in the mechanism (the growth rate is reduced to 0 percent only if cumulative emissions are well below the path) introduces a bimodal (two-peaked) distribution for cumulative emissions. If instead the two-sided mechanism were symmetric, the distribution would be unimodal (have only a single peak). In each case, adding a TAM substantially changes the shape of the emissions distribution, reducing emissions uncertainty, with that effect concentrated in the affected tail(s).

To summarize the effects of the TAM on the distribution of emissions, we focus on five specific emissions metrics: expected emissions, the standard deviation of emissions, the right-tail of the distribution (summarized by the 97.5 percentile), and the probability that the policy exceeds either the annual or cumulative benchmark at the end of the control period. As discussed in Section 2, which of these metrics to focus on and how to weight them will depend on the goal(s) the TAM is intended to achieve.

Tables 2 and 3 display expected costs and the emissions metrics for the Metcalf and Curbelo TAMs, respectively. The tables include the policies without TAMs, the central case TAMs (we use “central case” to refer to the policy designs as proposed) and alternative mechanism designs. These alternative mechanism designs highlight the importance of design choices and illustrate that there are important tradeoffs across different metrics and thus the optimal policy design varies substantially based on the goal(s) of the TAM.

Table 2: Emissions and Costs, Metcalf Tax Adjustment Mechanism

	Cumulative Costs 2021-2035 ($billion)	Costs per Ton Reduced ($2017)	Cumulative Emissions (2020-2035) (bmt)			Targets	
Central Case	Mean	Mean	Mean	Std Dev	97.5	Annual	Cumulative
Without TAM	$582.96	$24.37	60.74	4.90	69.79	54.9%	43.4%
With TAM	$620.91	$25.81	60.41	3.44	67.54	62.3%	55.6%
Benchmark Paths							
Arbitrary - Metcalf (CC)	$620.91	$25.81	60.41	3.44	67.54	62.3%	55.6%
Arbitrary - Straight Line	$440.55	$20.54	63.17	4.34	70.51	32.9%	6.9%
Model-Based	$671.35	$27.13	59.72	3.59	67.54	73.9%	74.2%
Benchmark Comparisons							
Cumulative (CC)	$620.91	$25.81	60.41	3.44	67.54	62.3%	55.6%
Annual	$634.20	$26.14	60.20	3.49	67.54	67.0%	67.6%
Annual and Cumulative	$635.31	$26.21	60.23	3.34	67.54	70.8%	67.7%
Frequency of Adjustment							
Annual, Grace Period (CC)	$620.91	$25.81	60.41	3.44	67.54	62.3%	55.6%
Annual, No Grace Period	$634.26	$26.23	60.34	2.81	66.13	66.4%	57.5%
Biennial, Grace Period	$618.99	$25.75	60.44	3.44	67.54	61.9%	54.4%
Biennial, No Grace Period	$627.87	$26.05	60.39	3.01	66.62	64.8%	55.4%
Adjustment Trigger							
Asymmetric Two-Sided (CC)	$620.91	$25.81	60.41	3.44	67.54	62.3%	55.6%
Symmetric Two-Sided	$546.75	$23.85	61.65	3.22	67.54	38.6%	23.8%
One-Sided	$647.89	$26.40	59.95	4.15	67.54	62.3%	61.7%

Table 3: Emissions and Costs, Curbelo Tax Adjustment Mechanism

	Cumulative Costs 2021-2035 ($billion)	Costs per Ton Reduced ($2017)	Cumulative Emissions (2020-2035) (bmt)			Targets	
Central Case	Mean	Mean	Mean	Std Dev	97.5	Annual	Cumulative
Without TAM	$181.34	$12.04	69.66	3.95	77.17	45.9%	40.5%
With TAM	$210.01	$13.24	68.78	3.25	75.04	55.3%	58.7%
Benchmark Paths							
Model-Based (CC)	$210.01	$13.24	68.78	3.25	75.04	55.3%	58.7%
Arbitrary - Straight Line	$181.59	$12.06	69.65	3.93	77.07	45.9%	40.5%
Benchmark Comparisons							
Cumulative (CC)	$210.01	$13.24	68.78	3.25	75.04	55.3%	58.7%
Annual	$214.02	$13.38	68.64	3.24	75.00	59.1%	62.7%
Annual and Cumulative	$217.59	$13.52	68.54	3.21	75.00	62.5%	66.3%
Frequency and Size of Adjustment							
Biennial, $2 (CC)	$210.01	$13.24	68.78	3.25	75.04	55.3%	58.7%
Annual, $1	$208.63	$13.19	68.82	3.28	75.16	54.7%	57.5%
Annual, $2	$234.80	$14.21	68.17	2.81	73.54	67.6%	71.3%
Biennial, $4	$236.93	$14.28	68.12	2.77	73.31	68.6%	72.7%
Adjustment Trigger							
One-Sided (CC)	$210.01	$13.24	68.78	3.25	75.04	55.3%	58.7%
Asymmetric Two-Sided (90%, -$2)	$207.06	$13.16	68.90	2.99	75.04	55.3%	57.9%

The central case TAMs reduce expected emissions, the standard deviation of emissions, and the right-tail of emissions outcomes while increasing the probability that annual or cumulative targets are met. The Metcalf TAM, as a two-sided mechanism, is more effective than the Curbelo TAM at reducing the standard deviation of cumulative emissions but less effective at increasing the probability of meeting emissions targets. On the other hand, the Metcalf TAM has a smaller increase in total costs relative to a policy without the TAM; the Metcalf TAM raises expected costs by 7 percent, the Curbelo TAM raises expected costs by 16 percent.

As described in Section 3, there are many alternative TAM design options. In tables 2 and 3, we quantitatively evaluate a number of alternative options to each adjustment mechanism. These results generally echo those from HW2019.

5.2.1. Benchmark Path

The Metcalf and Curbelo TAM benchmark paths differ substantially. Metcalf uses an arbitrary path (a path not based on predicted emissions under the tax) whereas Curbelo uses a path informed by predictions from a CGE model. For Metcalf, we compare the central case path to two alternatives: an even more arbitrary path – a straight-line from reference case emissions to the 2035 annual target[18] – and a model-based path. For Curbelo, we compare the central case model-based path to an arbitrary straight-line path.

In both cases, the model-based TAM significantly reduces uncertainty relative to the straight-line paths but the expected costs are also higher under the model-based paths. Imposing a meaningful carbon tax causes a substantial drop in emissions right away, so emissions are almost always below the straight-line path early on (even if the tax causes a smaller drop in emissions than expected), so adjustments start later with the straight-line path (resulting in higher emissions, but lower costs). The straight-line path starting from current emissions performs particularly poorly on the target metrics under the Metcalf policy; straight-line paths and two-sided mechanisms are a particularly poor combination of policy designs – prices decline in the early years of the policy even if the policy is significantly underperforming relative to expectations. Metcalf's arbitrary straight-line policy performance is more similar to the policy with a model-based path than the policy with the straight-line path because the Metcalf's benchmark path begins in 2021 at a level that is roughly consistent with model-based expectations of emissions levels in the first year of a carbon tax.

[18] Several early TAM proposals used this kind of straight-line-from-reference-case-emissions benchmark. For example, the REACT proposal (Metcalf, 2009) used targets that matched Waxman-Markey permit allocations; these allocations were not particularly stringent in the early years of the policy. More recent proposals have moved away from that, partly because early versions of HW2019's modeling showed the problems with this design.

5.2.2. Benchmark Comparisons

Both the Metcalf and Curbelo TAMs compare cumulative emissions to the cumulative benchmark path. Here, we compare that to the performance of TAMs that utilize annual emissions comparisons or annual/cumulative comparisons (i.e., those that adjust if either annual or cumulative emissions exceed the relevant benchmark). Both policy alternatives perform at least as well on all emissions outcomes as the cumulative comparison for both the Metcalf and Curbelo TAMs. In particular, the annual/cumulative comparison significantly improves the probability of hitting both the annual and cumulative targets. The increased emissions performance increases expected costs, but the cost increases are relatively small: 2-3 percent above the central case expected costs. Intuitively, the annual comparison reacts more quickly when emissions start to go off track, and the cumulative comparison continues to react if emissions remain off target, so the combination makes a big difference for the likelihood of hitting the target.

5.2.3. Frequency and Size of Adjustment

The Metcalf policy allows for annual adjustments to the growth rate of the carbon tax after a five-year grace period. Adjusting every other year (biennial) leads to fewer adjustments, reducing costs and increasing emissions uncertainty, but these changes are relatively small. Removing the grace period and allowing for adjustments immediately means more adjustments on average, increasing costs, but reducing uncertainty.

The Curbelo TAM allows for adjustments of $2 every other year. Adjusting every year with $1 increments (which keeps the total possible adjustment over a two-year period the same) slightly reduces costs but also slightly increases emissions uncertainty. Larger adjustments of either $4 biennially or $2 annually increase costs but significantly reduces emissions uncertainty. All else equal, larger adjustments will always increase costs and reduce uncertainty: the size of the changes will depend on the design specifics of the TAM.

5.2.4. Adjustment Trigger

Two-sided mechanisms reduce expected costs relative to one-sided mechanisms that can only increase the tax rate, because the downward adjustments reduce costs. The change in emissions metrics varies. Two-sided mechanisms have higher expected emissions than one-sided mechanisms, because the downward adjustments increase emissions in low-emissions cases. But the standard deviation is much lower with the two-sided mechanisms, which pull in both tails of the emissions distribution instead of just the right tail. For both the Metcalf/CLC and Curbelo policies, the 97.5 percentile is identical across adjustment triggers: downward adjustments only occur when emissions are low, so adding the potential for those adjustments doesn't affect the high-emissions tail of the distribution.

Adjustment triggers highlight how much the ideal TAM design can vary based on the goal(s) the TAM is intended to achieve. If the goal of a TAM is to reduce right-side risk at a low cost, two-sided mechanisms are optimal. But if the goal is to increase the probability of meeting a cumulative emissions target, one-sided mechanisms do better.

These model results show that adding a TAM to a carbon tax policy can substantially reduce emissions uncertainty, while raising costs slightly. There are substantial tradeoffs across different outcomes, meaning that the optimal TAM design depends strongly on what goals the TAM is intended to achieve.

6. Conclusion

In this paper, we aimed to provide an intuitive discussion of the effects of TAMs and of the important tradeoffs in TAM design. That discussion used simulation results for the Metcalf (2018) and Curbelo TAMs from Hafstead and Williams's (2019) reduced-form model of carbon dioxide emissions with uncertainty to show the effects that a real-world TAM could have on emissions and costs. And we then used simulations of alternative versions of those TAMs to show how design choices affect cost and emissions outcomes, illustrating the tradeoffs inherent in such design choices.

We find that both of these TAMs substantially reduce emissions uncertainty and increase the probability of hitting annual and cumulative emissions targets in 2035. And while the TAMs increase the expected cost of the policy relative to a carbon tax with no adjustments, that cost increase is relatively small for both the Metcalf and Curbelo TAM designs.

Comparing results across a range of alternative design scenarios, a few TAM design elements seem to be broadly beneficial. Smaller, more-frequent adjustments do better than larger, less-frequent adjustments on almost every metric we consider. The same is true for using a model-based benchmark path (i.e., setting the benchmark path based on predicted future emissions under the tax) rather than an arbitrary path (such as a straight-line path to the target). But even these design choices aren't consistently better on every metric one might reasonably care about.

And on most dimensions of TAM design, there are clear tradeoffs. First, there exists a general tradeoff between emissions certainty and expected costs: design changes that generally reduce emissions uncertainty across a range of metrics (such as making larger adjustments) almost always increase costs. Second, there are often also tradeoffs among different emissions certainty goals: the design that reduces the risk of very-high-emissions outcomes at the least cost could perform poorly at hitting a particular target or reducing expected emissions. These tradeoffs demonstrate that the question of how to design a TAM depends on exactly what goals one has for the TAM and how one weights those different (and often conflicting) goals. Indeed, the question of whether to include a TAM at all depends on how much weight one puts on reducing emissions uncertainty versus reducing expected costs for a given level of expected emissions.

These are difficult questions for an economist to answer. If the motivation is purely to maximize economic welfare (i.e., minimize the sum of abatement cost and pollution damage), then the answers depend strongly on how marginal damage changes as cumulative emissions increase. Large discontinuities in the marginal damage curve would suggest large adjustments should be made to avoid significantly increasing overall damages. But it's hard to estimate where those discontinuities are (if they exist at all). And as the social cost of carbon literature has demonstrated, estimating the level of marginal damage from CO2 emissions is challenging enough; estimating the slope is more so. But at least those are questions that economists have the tools to address. In many cases, however, the motivation for controlling emissions uncertainty is political rather than economic and it's hard to know how to weight those motivations (e.g., meeting a pledge to reduce emissions below a given target).

Therefore, rather than trying to define the "optimal" policy design, we have tried in our work on this topic (both in this paper and in Hafstead and Williams (2019)) to measure and explain the tradeoffs, with the idea that policymakers can then make educated decisions, based on their own weighting of the different potential goals. Given the rapidly growing interest in TAMs in the policy community, we hope this will be a valuable contribution.

We can see a number of promising directions for future research in this area. Some would be technical advances on the modeling. The Hafstead and Williams (2019) model doesn't allow for forward-looking behavior by emitters, which might make a significant difference, so incorporating that into the model would be a nice contribution. More research measuring emissions uncertainty – especially uncertainty in how emissions will respond to a carbon tax – would also be valuable: Hafstead and Williams's (2019) calibration of that key elasticity is necessarily somewhat ad hoc, and Harris and Pizer (2018) ignore uncertainty in that elasticity entirely (perhaps because of the calibration difficulty).

Going in a different direction, more research could be valuable about the role of emissions uncertainty in climate policy. For example, how does emissions uncertainty affect environmental agreements? That's a much more vague, open-ended question, but it might help provide a bit more guidance about how to weight different potential goals for if policymakers decide to address emissions uncertainty with a carbon tax adjustment mechanism.

References

Aldy, Joseph E. "Designing and Updating a US Carbon Tax in an Uncertain World." Harvard Environmental. Law Review Forum 41 (2017): 28-40.

Aldy, Joseph E., Marc Hafstead, Gilbert E. Metcalf, Brian C. Murray, William A. Pizer, Christina Reichert, and Roberton C. Williams III “Resolving the Inherent Uncertainty of Carbon Taxes” Harvard Environmental. Law Review Forum 41 (2017): 1-13.

Barron, Alexander R., Allen A. Fawcett, Marc A.C. Hafstead, James R. McFarland, and Adele C. Morris, “Policy Insights from the EMF32 Study on US Carbon Tax Scenarios.” Climate Change Economics 9(1), 1840003 (2018).

Hafstead, Marc, “Considering a Carbon Tax-Gasoline Tax Swap: Projected Energy-Related US CO_2 Emissions Reductions under the MARKET CHOICE Act.” Resources for the Future, Washington DC, Issue Brief 18-08 (2018).

Hafstead, Marc, Gilbert E. Metcalf, and Roberton C. Williams III. "Adding Quantity Certainty to a Carbon Tax through a Tax Adjustment Mechanism for Policy Pre-Commitment." Harvard Environmental. Law Review Forum 41 (2017): 41-57.

Hafstead, Marc A.C. and Roberton C. Williams III. “Mechanisms to Reduce Emissions Uncertainty under a Carbon Tax.” Resources for the Future, Washington DC, DP 18-25 (2018).

Halstead, Ted, George P. Shultz, Lawrence Summers, Rob Walton, Christine Todd Whitman, and Janet Yellen, “Exceed Paris.” Climate Leadership Council, Washington DC (2018).

Harris, Robert and William A. Pizer, “Taxes vs. Tradable Permits: Information Asymmetry and Dynamic Adjustment Mechanisms”, Duke University Working Paper (2018).

Metcalf, Gilbert E. "Cost Containment in Climate Change Policy: Alternative Approaches to Mitigating Price Volatility." Virginia Tax Review 29, no. 2 (2009): 381.

Murray, Brian C., William A. Pizer, and Christina Reichert. "Increasing Emissions Certainty under a Carbon Tax." Harvard Environmental. Law Review Forum 41 (2017): 14-27.