NBER WORKING PAPER SERIES

USING TAXES TO MEET AN EMISSION TARGET

Robert I. Harris
William A. Pizer

Working Paper 27781
http://www.nber.org/papers/w27781

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2020

Invaluable comments were received on several earlier drafts from Larry Goulder, Marc Hafstead, Rob Williams, Gib Metcalf, and Rick Van der Ploeg, along with participants at the AERE 2020 Summer Conference, the Harvard-Berkeley-Yale Seminar on the Economics of Climate Change and Energy Transition, the UMD Environmental Taxation Workshop, and the World Congress of Environment and Resource Economists. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Using Taxes to Meet an Emission Target
Robert I. Harris and William A. Pizer
NBER Working Paper No. 27781
September 2020
JEL No. H23,Q54,Q58

ABSTRACT

A sizeable number of papers beginning with Roberts and Spence (1976) have studied the use of price floors and ceilings (or "collars") to manage prices in tradable permit markets. In contrast, economists have only recently begun examining polices to manage quantities under a pollution tax. Importantly, it can be difficult to know how to evaluate these policies, as papers dating back to Pizer (2002) suggest welfare is maximized by not focusing on quantities in the first place. In this paper, we propose an objective function to evaluate these alternative "carbon tax policies to meet an emission target." The objective function includes a discrete jump in marginal emission consequences at the target, where the discontinuity can be interpreted as a true benefit measure or a necessary political constraint. We parameterize these emission consequences using recent legislative proposals, coupling this function with mitigation cost estimates to define the complete objective. This objective identifies the first-best tax policy design, one that requires relatively complex adjustments to mimic a tradable permit system. Turning to simpler, practical rules, we find that such rules achieve much of the difference in expected net benefits between an ordinary, exogenous tax and the first-best tax policy design. However, the ranking among simple rules depends on the interpretation of the higher, above-target emission penalty as a political constraint or a true benefit measure. We find that making these views explicit could facilitate billions of dollars per year in welfare gains.

Robert I. Harris
Duke University
201 Science Drive
Durham, NC 27708
bobby.harris@duke.edu

William A. Pizer
Sanford School of Public Policy
Duke University
Box 90312
Durham, NC 27708
and NBER
billy.pizer@duke.edu

1 Introduction

Beginning with Weitzman (1974), economists have long appreciated the welfare difference between price and quantity regulation, e.g., taxes and tradable permits, when costs are uncertain. In particular, Roberts and Spence (1976) demonstrated that a tradable permit system with a price floor and ceiling could improve welfare beyond taxes or permits alone. Intuitively, taxes and permits alone are nested within the choice set of the blended hybrid. Other papers have extended these results in different directions (Pizer, 2002; Burtraw, Palmer and Kahn, 2010).

Until recently, however, there has been no work examining the reciprocal idea of tax policy that attempts to to achieve an emission target. One could imagine a number of reasons for this omission. The political space, particularly in the United States, has been dominated by tradable permits—though that is not true elsewhere (World Bank and Ecofys 2018). It is trickier, as we discuss, to operationalize hybrid tax policies. Finally, as an empirical matter, price regulation is often preferred on welfare grounds based on relatively flat marginal benefits from abatement. This is particularly true for carbon dioxide and other stock pollutants (Newell and Pizer, 2003), but also air pollution (Pope et al., 2015). Therefore, welfare gains ought to come from adding price-like elements to a quantity regulation, not the other way around.

Regardless of the reason(s) for this gap and motivated by a half-dozen proposals in the 116th Congress, our paper responds by suggesting an objective function to first rationalize this type of policy. Specifically, we are interested in a carbon tax that adjusts based on observable information in order to try to achieve a given cumulative emission target. The objective function we suggest is simply a minimization of the expected cost of reducing emissions, plus a low penalty per ton for expected emissions below the target and a high penalty per ton for expected emissions above the target.

We consider two alternative interpretations for this form of the emission penalty. One views the low penalty as the actual emission damage (e.g., the social cost of carbon) and the high penalty for above-target emissions as a political constraint imposed by stakeholders.[1] The other interpretation is that such stakeholder preferences actually define societal benefits. That is, the high penalty is not a political constraint but an actual, revealed societal benefit placed on mitigation of above-target emissions. In both cases, the objective is the same; the

[1] In a non-stochastic framework, van der Ploeg (2018) considers how to maximize welfare with both a social cost of carbon (SCC) and a cumulative emission constraint. He shows that the result involves pricing carbon at a combination of terms reflecting the SCC and the cumulative emission target. He does not consider the question of uncertainty in the cost of achieving the target or the related policy question of how to adjust a carbon tax as emissions are observed.

difference is how the objective might or might not be separated into welfare and constraint.

We glean stakeholder preferences from these motivating proposals introduced in the 116th Congress while also considering estimates of climate damages (the "social cost of carbon" or SCC). Costs are estimated from available modeling. We then use the identified objective function to find the best policy both within and across a defined set of policy designs. These designs include carbon taxes with different types of adjustments based on observed emissions. The policy parameters within a particular design include alternative carbon tax levels and/or growth rates as well as cumulative emission thresholds that select among the alternatives.

We find, first, that a policy mimicking cap-and-trade with a price floor and ceiling performs the best, as we would expect from the objective function. While the preference for this policy does not depend on how we interpret the objective, the magnitude does. Viewed as a politically-constrained welfare function, the exogenous, simple tax offers a $34 billion per year improvement over no policy. Cap-and-trade offers a 9% or $3 billion further improvement. If we instead view the higher value placed on above-target mitigation as a true benefit, rather than a constraint, the gain from the simple tax jumps to $660 billion per year versus no policy; the cap-and-trade is 2% or $14 billion per year higher. This follows from our calibrated above-target penalty mitigation benefit, which is more than four times higher per ton than the below-target mitigation benefit.

When we turn to taxes with different types of "practical" dynamic adjustments, we find that they achieve between 37-96% of the simple tax-ETS[2] difference, depending on the welfare metric and adjustment. Generally, those that allow larger discrete price jumps do better when the above-target mitigation value is interpreted as a welfare benefit. The more gradual adjustment mechanism does better when the above-target mitigation value is considered a political constraint. These differences add up to $2-3 billion per year.

We recognize there could be considerable disagreement over the form of the objective function, its interpretation, and its parameterization. However, yet another concern is how to recognize policies that are implementable in practice. For example, observed policy proposals have focused on a carbon tax that rises by one of two possible growth rates each period, depending on observed emissions. That is, a relatively gradual tax change. The choice variables are the two growth rates as well as the emission threshold that triggers the higher rate. We have interpreted this as a constraint that practical policies must be simple. With that in mind, we also consider policies that involve two distinct price paths with an emission threshold that triggers jumps between one path and the other. While this is similarly simple, the price changes we identify can be a factor of 30% or much more when cumulative emissions trigger a change. Is this implementable, or are there yet other political constraints on how

[2]We will interchangeably refer to a cap-and-trade policy as an emissions trading scheme or ETS.

much the price can change?

Whether one agrees with our choice of objective and space of policies or not, the larger purpose of this paper is to emphasize the value of (a) thinking carefully about the space of practical policies in order to inform the carbon tax debate, and (b) considering the right objective(s) to identify better and worse policies within that practical policy space. Here, we suggest there is an implied objective defined by the trade-offs stakeholders are (or are not) willing to make. Being explicit about such an objective then allows society to both reduce costs and improve environmental outcomes. Such improvements may come from both a better choice of parameters or an expansion of the policy space. While one can always consider an unconstrained policy space and/or science-based assessments of benefits and costs, those may not speak to stakeholders' revealed preferences and/or constraints associated with the policy process.

Our interest in this topic follows from from policy discussions that began in 2015. In that year, U.S. Senator Bernie Sanders introduced S. 2399, The Climate Protection and Justice Act. This bill would have implemented a carbon tax and, if the emissions exceeded certain emission targets, authorized further regulations to achieve those targets. This approach reflects an emerging, two-part dynamic in the domestic U.S. climate policy debate. First, carbon taxes are increasingly seen as a viable policy option (discussed below). Second, among some stakeholders, this is assumed to go hand-in-hand with the removal of some or all regulatory authority under the Clean Air Act. This, in turn, has prompted other stakeholders to seek assurances—one way or another—that a focus on emissions will not be sacrificed. The idea of using a carbon tax to target emissions is one of those ideas (Murray, Pizer and Reichert, 2017).

Yet, in attempting to use a carbon tax to target emissions it is unclear how to define success. Is exceeding the target 5% of the time success, but 5.01% failure? How do we compare large exceedances 5% of the time with small exceedances 10% of the time? How much are we willing to spend to reduce exceedances further? Is there any benefit to further, below-target mitigation? One solution would be to represent the distribution of emission and cost outcomes under alternative options and allow stakeholders to choose. This approach, however, does not lend itself to easily identifying yet further improvements.

With the introduction of a half-dozen proposals in the 116th Congress that are similar to S. 2399, another approach suggests itself. Use these policies to identify the underlying preferences lurking behind the proposals. Such preferences arguably represent the views of at least the stakeholders supporting and promoting the proposals if not of society as a whole. We believe this may be a helpful exercise, particularly when the parameters relate to familiar concepts in other policies—e.g., emission targets, price floors, and price ceilings.

By presenting the underlying preferences associated with these policies, and estimating the relative advantage, we believe stakeholders and society can make more informed decisions.

The remainder of the paper proceeds as follows. Section 2 reviews recent policy experience and the relevant literature on hybrid instruments, including very recent papers on hybrid price policies. In Section 3, we define our objective based on avoided climate damages (mitigation benefits), mitigation costs, and an above-target emission penalty with competing interpretations. We then discuss our parameter choices based on recent policy proposals and estimates of the social cost of carbon. We also describe our mitigation cost model. Finally, we use this objective to consider the effects of different policies and optimize over policy parameters in Section 4. Section 5 concludes.

2 The path to hybrid (or dynamically-adjusted) taxes

Until very recently, both the theory and practice of environmental regulation focused on what we would call hybrid quantity mechanisms. That is, regulation that targets the emission levels based on a cap-and-trade program. Additional provisions are then added to limit prices, typically a price floor introduced when allowances are auctioned and either a limited or unlimited supply of additional, above-target allowances available at a higher price. Both California and the Regional Greenhouse Gas Initiative operate in this way.

2.1 Hybrid quantities

Both economic theory and political convenience have arguably played a role in this hybrid quantity policy design. On the theory side, economists have provided evidence that price policies will yield higher expected welfare. This follows early work by Weitzman (1974) who highlighted the welfare difference between price- and quantity-based policies when costs are uncertain. He showed that prices are favored when the slope of marginal benefits is relatively flat compared to the slope of marginal costs.[3] Meanwhile, significant evidence points to flat marginal benefits related to climate change mitigation (Pizer, 2002; Newell and Pizer, 2003).

There are a number of ways to understand this result that flat marginal benefits favor price policies. The simplest is that benefits represent a public good being demanded by

[3]Recent work by Pizer and Prest (2016) suggests that impending or expected policy updates can alter the traditional Weitzman result. In particular, a welfare difference between taxes and tradable permits will arise based on how the incentives they create diverge in response to expected policy changes. Real-time incentives created by a tax policy do not change until the tax itself changes, regardless of expectations. But real-time incentives created by a tradable permit policy can fluctuate based on market speculation. These fluctuations can raise or lower welfare depending on whether price changes are driven by a genuine evolution of costs and benefits, or not.

the government. A tax amounts to a flat demand curve; a tradable permit system amounts to vertical demand. Which better approximates marginal benefits when marginal costs (abatement supply) is uncertain? It depends on the slope: Flat marginal benefits are better matched by a tax.

It did not take long for others to realize that a slightly more sophisticated policy could generally raise expected welfare compared to either a tax or a tradable permit policy. Both Roberts and Spence (1976) and Weitzman (1978) suggest the idea of a tradable permit system with a fee for emissions above the quantity target and a subsidy (or minimum price) if emissions fall below the target. In the climate change context, this was noted by a number of authors (McKibbin and Wilcoxen, 1997; Pizer, 2002; Burtraw, Palmer and Kahn, 2010).

Of course, the choice between taxes and tradable permits has never been entirely (or even largely) about expected welfare based on empirical evidence. Generally, tradable permits have been favored in the United States for reasons ranging from Grover Norquist's demand for no new taxes to environmental organizations preferring a policy focus on environmental outcomes. Meanwhile, stakeholders on both sides have not been content with a permit market where the price is not bounded. Businesses tend to be concerned about high prices and environmental groups low prices. Thus, political interests have aligned to favor hybrid quantity policies as well.

2.2 Hybrid taxes

The path to thinking about hybrid taxes began when carbon taxes themselves moved out of the shadows in 2010. Several developments arguably conspired to make this happen. First, the failure of a major effort to enact tradable permit legislation in 2009-2010 led some proponents to reconsider the future prospects of such legislation (Lizza, 2010). Second, there was considerable interest in comprehensive tax reform, manifest in two bipartisan reports on the topic (National Commission on Fiscal Responsibility and Reform, 2010; Debt Reduction Task Force, 2010). Neither report explicitly mentions a carbon tax, but many viewed that as a strategic rather than substantive omission (Bledsoe, 2012). Finally and more recently, conservative supporters of climate action have more boldly articulated carbon taxes as an alternative to regulation under the Clean Air Act (Baker et al., 2017). There was certainly discussion of carbon taxes prior to 2010, but it did not garner the more serious political and stakeholder attention of the last decade.

Metcalf (2009) was the first to propose the more specific idea of a hybrid carbon tax policy, e.g., a tax that would adjust in some way to achieve an emission target. This idea also remained largely in the shadows until 2015, when two carbon tax bills were introduced with reference to an emission target (Sanders, 2015; Whitehouse, 2015). The Sanders bill proposed a \$15 per ton CO_2 tax, rising at roughly 5% annually above inflation. In addition, it specified a set of targets, declining to 1.3 billion tons in 2050 (17% of 2005 emissions). If emissions were deemed to not be meeting the specified targets, additional regulations were authorized and required. The Whitehouse bill proposed a \$45 per ton CO_2 tax rising at 2% annually above inflation until such point as emissions fell below 20% of 2005 emissions. At that point, the carbon tax would rise just at the inflation rate.

These bills, particularly the Sanders bill, reflected underlying sentiments that both a carbon tax was increasingly possible and that it would need to be responsive to observed emissions. Such sentiments spurred additional research on various ways to increase emissions certainty under a tax. An online forum of the *Harvard Environmental Law Review* conveniently compiled many ideas in one place (Aldy, 2017; Aldy et al., 2017; Hafstead, Metcalf and Williams, 2017; Murray, Pizer and Reichert, 2017). One observation across these papers is that such policies can be organized by how they function in two key dimensions: One dimension is whether the carbon tax dynamically adjusts or whether other tools (such as regulation) are introduced. The other is whether such adjustments occur automatically, triggered by observable events, or whether they are discretionary based on a procedural determination. Among them, the idea of an automatically-adjusting (or "rules-based") carbon tax figured prominently. Two authors went on to draft a more expansive paper on the topic (Hafstead and Williams, 2019).

More recently, in 2017 a coalition of largely conservative stakeholders put together a comprehensive carbon tax proposal that would also seek to eliminate any further regulation carbon regulation under the Clean Air Act. This was followed by expression, of environmental advocates, that such an elimination would require an adjustment mechanism that is sensitive to environmental outcomes. This was followed, in 2019, by the introduction of eight carbon pricing bills in Congress. Seven of the bills are carbon taxes and four increase the tax rate faster when cumulative emissions exceed certain targets.

Discussion and research on increasing emissions certainty under a tax continued in a very recent symposium of the *Review of Environmental Economics and Policy* (Aldy, 2020; Brooks and Keohane, 2020; Hafstead and Williams, 2020; Metcalf, 2020). The symposium papers, Brooks and Keohane (2020) in particular, note that cumulative emissions uncertainty poses a political challenge—the support of the environmental community is contingent on ensuring desired emission outcomes. For a tax to garner broad enough political support, it will need to be responsive to emissions. In addition to the political economy motivations, two themes from this symposium are (1) how to design a tax adjustment mechanism and (2) how to evaluate the performance of such a mechanism. Assuming the mechanism is automatic, design considerations include the frequency, type, and size of adjustment, the target emissions path, and whether to allow downward price adjustments (Hafstead and Williams, 2020; Metcalf, 2020).[4] Evaluating the performance of a tax mechanism is viewed as more of a question: How should the mechanism trade off costs and emissions certainty?[5] Then, what is the right emissions measure: average cumulative emissions, the probability of achieving the target, the 97.5-percentile of the emissions distribution, the magnitude of expected exceedances, or something else?

While our paper also considers different design options, our main contribution is to make this second theme explicit in our objective function. The symposium papers note that a singular focus on achieving a particular emissions target may come at the expense of increased compliance costs—that is, measuring the success of the policy based on environmental outcomes is distinct from a traditional focus on welfare. Our objective function makes this explicit by building in the societal (or political) value of achieving the emission target that is implicitly revealed by the proposals in the 116th Congress.

2.3 Distinctions between hybrid taxes and hybrid quantities

The different ways quantity and price hybrid policies seek to control uncertainty in price and emissions, respectively, also raises an important distinction that has been somewhat overlooked in the literature. In particular, intervention to manage prices under a quantity policy is almost always forward looking, while intervention to manage quantities under a price policy is almost always backward looking. Moreover, managing prices "works" in the sense of directly limiting prices, while managing emissions (as we shall see) requires trial and error.

Why? Part of the reason is that for most emission problems, and certainly climate change, the quantity we are trying to manage is really the cumulative stock of emissions over some time period. At a minimum, we are thinking about a year; but perhaps we have in mind cumulative emissions over many decades. Meanwhile, the more frequent emission flow (over a day, a month or even year) can be quite cyclical and noisy. Conveniently, a well-functioning market associated with a quantity policy provides information about prices that is both almost continuous and forward-looking. That is, when firms buy and sell emission allowances, they are thinking about the cumulative supply and demand over the current compliance period, responding to expectations about both future emission limits and abatement costs. Why sell an allowance now that you think will fetch a higher price

[4] Aldy (2020) again discusses a more discretionary review and updating process.

[5] In addition to these dimensions, Aldy (2020) also posits other metrics for judging discretionary adjustments, including temperature, competitiveness, and comparable international action.

later in the same compliance period? Likewise, why buy one if you think the price will fall? With banking (and perhaps borrowing), this reasoning can project far into the future. The conclusion in the literature is that expected prices in an allowance market should and do rise at the interest rate in order to equate cumulative supply and demand over time (Yates and Cronshaw, 2001; Schennach, 2000). Moreover, when we alter the allowance supply based on a price floor or ceiling, the price (path) adjusts almost immediately. This allows policies to easily manage prices: Price floors and ceilings operate by simply subtracting or adding allowances until the price limit is met. This often happens all at once, when an auction is held with a price floor or ceiling.

The reverse is not true with a price policy. Continuous emission monitoring may allow relatively frequent observations of emissions. However, current, observed emissions generally reflect only the current emission price (the tax), not any expectation about the future. That is, while arbitrage opportunities link prices over time in an allowance market, there is no similar link among emissions over time under a carbon tax.[6] The exception would be the extent to which long-term capital investments, rather than frequent operational choices, determine emissions.

More to the point, as we adjust the price we only see the impact on emissions as they occur over time. There is no sense, comparable to the above quantity management, where we could adjust prices and quickly see a transparent outcome reflecting expected emissions through the end of a policy horizon. Rather, we must adjust the price, see what happens, and then perhaps adjust the price again. Viewed this way, adjusting the price to control emissions may be better thought of as a dynamic adjustment process, a name which reflects their characteristic of updating over time to control cumulative emissions (and that we will use for the remainder of this paper).

Unless, that is, we use a sophisticated emission and mitigation model to project effects into the future and determine the necessary tax adjustment each period. This would not be transparent. Rather, it would attempt to replicate what an allowance market with banking (and borrowing) does each period: Determine a market price to balance expectations about supply. Setting a price this way would necessarily involve tax adjustments over time in response to new information.

[6]We note that Heutel (2020) considers a policy where firms can opt to pay a tax in the current period or in a future period. He shows that firms can be induced to pay all the tax in the future period, once uncertainty is revealed, by implementing an arbitrarily high tax in the first period. This has important implications for the choice of tax or trading policies based on whether they effectively make use of new information over time. However, the particular point here is that trading policies reveal future market expectations over time (through observed prices) in a way that tax policies do not (through observed emissions). Price policies will require government modeling and expertise to construct that information. We discuss how a tax policy might mimic emissions trading below and in Section 4.1.

This leads to the question that motivates this paper: How do we best design a dynamic adjustment process? Indeed, what defines "best"?

3 Defining and evaluating an objective function for carbon tax adjustments

The traditional metric for normative policy analysis is welfare—benefits minus costs—and we continue to consider that our guiding principle. In particular, we are going to suggest as an objective to maximize the function,

$$- S^{-1} \sum_{s=1}^{S} \left(aE_T^s + (b-a)(E_T^s - \bar{E})\mathbb{1}(E_T^s > \bar{E}) + \sum_{t=t_0}^{T} e^{-\delta(t-t_0)} C_t^s(e_t^s) \right) \tag{1}$$

where $s \in \{1, ..., S\}$ indexes uncertain states of nature, $t \in \{t_0, ..., T\}$ indexes time, e_t^s is annual emissions in period t and state s, $E_t^s = \sum_{r=1}^{t} e_r^s$ or cumulative emissions, δ is a discount rate, $\{a, b, \bar{E}\}$ are climate damage parameters, and $C_t^s(e_t^s)$ is a state- and time-dependent cost function. In this way, the first two terms represent damages and the third term costs.

Alternatively, we can rearrange Equation (1),

$$- S^{-1} \sum_{s=1}^{S} \left(aE_T^s + \sum_{t=t_0}^{T} e^{-\delta(t-t_0)} C_t^s(e_t^s) \right) - S^{-1} \sum_{s=1}^{S} (b-a)(E_T^s - \bar{E})\mathbb{1}(E_T^s > \bar{E}) \tag{1a}$$

Here, we interpret the first two terms as true damages and costs. The third term represents a political constraint that above-target emissions be penalized additionally relative to the true benefit parameter a. Regardless of the interpretation of the third term—true damages or a political constraint—the objective function and our calibration remains the same. We discuss additional details of this interpretation as a political constraint in Appendix A.

It is useful to note that the climate damage/penalty terms reflect an underlying marginal damage/penalty function of the form

$$MD(E_T^s) = a + (b-a)\mathbb{1}\left[E_T^s > \bar{E}\right] \tag{2}$$

$\bar{E}$ is the cumulative emissions target over the period of interest and we assume $b > a$, implying higher marginal damages/penalty when the target is exceeded.

For our exercise, we choose the discount rate $\delta = 4\%$. This is halfway between the benefit discount rate (3%) and cost discount rate (5%) used in the 2015 Regulatory Impact Analysis of the Clean Power Plan (U.S. EPA, 2015). In this analysis, it principally guides how we trade off mitigation costs over time and how quickly optimal carbon prices will grow. A low rate will encourage relatively flat prices over time while a higher rate implies lower prices now and higher prices in the future.

We also focus on the 2020-2050 time period, which has been most relevant in recent policy discussions. Thus t runs from $t_0 = 2020$ to $T = 2050$. As a final preliminary, we consider $S = 10,000$ uncertain future outcomes.

We now discuss the remainder of the terms and parameters in more detail.

3.1 Revealed marginal damages and/or political pressure

While evidence suggests relatively flat marginal benefits (Pizer, 2002; Newell and Pizer, 2003), considerable stakeholder attention has focused on emissions or cumulative emissions. This includes previous global calls for emissions to be cut by 50% by 2050 (or 80% by developed countries) (Wintour and Elliott, 2009), more recent calls for climate neutrality by 2050 (Climate Home News, 2019), as well as setting the social price of carbon based on an emission target (United Kingdom, 2009).

This is also manifest in the noted tax proposals that indicate one tax rate if an emission objective is achieved, and another tax rate if it is exceeded. One interpretation is a revealed marginal damage function with a discontinuous jump at the cumulative emissions target—i.e., Equation (2). Marginal damage functions of this form are notionally implied by any policy defined by an emissions goal and, hence, by the recent literature on hybrid taxes.

Alternatively, we can view this policy feature as a necessary political compromise. Environmental advocates will require some attention to a cumulative emission target and, if a low carbon tax is insufficient, a higher carbon tax must be applied. From this perspective, the attention to a particular cumulative emission target does not reflect true damages, but a feature of the political landscape.

In either case, the target and above-target damage penalty point to some type of *policy-revealed preference*, either among stakeholders influencing the policies or, in a welfare sense, for society as a whole. That is, hybrid quantity policies with specific emission targets developed over the past two decades and more recent carbon tax proposals that refer to a set of emission targets both reveal a preference for quantifying success in terms of achievement of a particular emissions outcome. There is a revealed notion that we should differentially value emissions below the cumulative target from emissions exceeding the cumulative target.

While a strict emphasis on emissions would indicate a "vertical" damage function or political constraint on emissions, this discontinuous horizontal function would seem to better reflect actual preferences or constraints. That is, there is value to reducing emissions below the target and some willingness to exceed the target if at least some significantly higher effort (e.g., higher tax) is applied. This most explicitly began with Metcalf (2009).

Metcalf's proposal implicitly assumes such an objective given his choice of a penalty growth rate for years in which emissions exceed a threshold level. Similarly, Hafstead and Williams (2019) evaluate various choices in designing an updating mechanism to achieve emissions certainty objectives, all of which suggest that achieving a particular target is important. Despite this type of feature being implicitly revealed through policy proposals and recent literature, we are the first to explicitly include it in the policymaker's objective function.

In order to calibrate this model and determine a, b, and $\bar{E}$, we turn to these recent proposals in the 116th Congress. In particular, we examine the NPV carbon prices when the emission target is met to determine a, and the NPV carbon prices when the emission target is exceeded to determine b. Averaging across all years and proposals (see Appendix C) we find $a = \$34.1$ per ton and $b = \$142.5$. Similarly, we cumulate (where necessary interpolating or extrapolating) the triggering annual emission targets to yield a 2020-2050 cumulative target. Averaging across the proposals, and converting to a CO_2 target, we find $\bar{E} = 78.5$ billion metric tons of CO_2 (see Appendix C).

The preceding discussion looks to policy proposals to reveal climate damage parameters. If we view the above-target penalty as a political constraint, we should instead look to objective estimates of climate damages. For example, IWG (2016) reports central estimates of \$42 per ton in 2020 rising to \$69 in 2050. In NPV terms at our 4% discount rate, \$69 in 2050 would be \$28 in 2020. Averaging these numbers, \$42 and \$28, we find an approximate mitigation benefit of \$35 per ton over cumulative emissions to 2050. Hence, we view our parameter estimate of $a = 34.1$ as representative of both revealed preference *and* an objective estimates of climate damages.

3.2 Model of costs

Having defined the climate damage/penalty parameters in (1), our other task is to define costs or $C_t^s(e_t^s)$. This amounts to building a simulation procedure, or "test track," that allows us to simulate the economy's response to a carbon tax with various dynamic adjustment mechanisms. Such a test track should reveal both costs and emissions outcomes. The goal is to have a tractable procedure that remains consistent with historical trends and is informative about future emissions.

We break task this into two parts by defining

$$C_t^s(e_t^s) = \tilde{e}_t^s \left(\frac{1}{2}\beta_{1,t}(r_t^s)^2 + \frac{1}{3}\beta_{2,t}(r_t^s)^3 \right) \tag{3}$$

where $\tilde{e}_t^s$ is baseline emissions in state s and period t and $r_t^s = (\tilde{e}_t^s - e_t^s)/(\tilde{e}_t^s)$ is the fraction of emissions reduced. We define costs this way in order to derive marginal costs of abatement (the negative of the derivative of costs with respect to emissions) as

$$MC_t^s(r_t^s) = -\frac{dC_t^s(e_t^s)}{de_t^s} = \beta_{1,t}r_t^s + \beta_{2,t}r_t^{s2} \tag{4}$$

That is, marginal costs expressed in terms of reduction fraction are quadratic, with the restriction that $MC_t^s = 0$ when $r_t^s = 0$.

We approach costs this way based on the idea that estimated marginal mitigation costs tend to be more consistent when viewed in terms of fractional reductions (Weyant and Hill, 1999). We then intend to choose the parameters β based on modeling estimates.

In particular, we use the Dynamic Integrated Economy/Energy/Emissions Model (DIEM-CGE) (Ross, 2014), a dynamic computable general equilibrium model, to derive the β's. We simulate a number of different price paths and emission reduction in that model. We then fit the results in each time step of the model to Equation (4). Again, $MC_t(\cdot)$ is the economy-wide marginal cost (\$/MtCO_2) for reducing emissions by the proportion r_t below baseline emissions in year t. Here we have dropped indicator for state s, as the DIEM-CGE model is not stochastic. That is, we allow this function to evolve over time t, as calibrated by DIEM, but focus our uncertainty analysis on baseline emissions not marginal costs in terms of percent reductions r_t^s. Total costs are then given by (3) for each year with stochastic emissions e_t^s described next. We report annualized net present value costs (using our 4% discount rate) in our results. Appendix D contains more detail on this procedure.

3.3 Model of emissions

Rather than using the single emission forecast in DIEM-CGE, we construct stochastic emission forecast based on two sources of uncertainty: (1) underlying economic growth, and (2) aggregate emission intensity. Using U.S. carbon dioxide (CO_2) emissions data from the World Resources Institute's Climate Data Explorer Tool (2015) and real GDP time series data in chained 2009 dollars from the U.S. Bureau of Economic Analysis (2018), we estimate the following models:

$$\Delta \log\left(\mathrm{GDP}_t\right) = \phi \Delta \log\left(\mathrm{GDP}_{t-1}\right) + \mu_g\left(1-\phi\right) + \varepsilon_t \tag{5}$$

and

$$\Delta \log\left(\frac{\mathrm{CO}_2}{\mathrm{GDP}}\right)_t = \alpha_1 \Delta \log\left(\frac{\mathrm{CO}_2}{\mathrm{GDP}}\right)_{t-1} + \alpha_2 \Delta \log\left(\frac{\mathrm{CO}_2}{\mathrm{GDP}}\right)_{t-2} + \mu_r\left(1-\left(\alpha_1+\alpha_2\right)\right) + \eta_t \tag{6}$$

where ϕ, α_1, and α_2 are autoregressive coefficients, ε_t and η_t are shocks identically and independently distributed $N(0,\sigma_g)$ and $N(0,\sigma_r)$, respectively, where σ_g and σ_r are estimated from the data, and μ_g and μ_r are drift coefficients. That is, we estimate an ARIMA(1,1,0) process with drift of the log of GDP and an ARIMA(2,1,0) process with drift of the log of the ratio of emissions to GDP (these models being chosen based on the Akaiki information criteria). Following this estimation, we use these models to forecast baseline emissions from 2016 through 2050. We draw 10,000 sequences s for $\{\varepsilon_t, \eta_t\}$ and use equations (5) and (6) to generate 10,000 sequences of baseline emissions, $\tilde{e}_t^s$. For more details on our estimation and forecasting procedure, see Appendix B.

We graph historical emissions through 2015 along with our projections for 2016 through 2050 in Figure 1. We also plot the annual target path that we derive from proposals of the 116th Congress following the procedure described in Appendix C. Our cumulative target for 2050, $\bar{E}$, is defined as the sum of these annual targets. Figure 3 in the Appendix shows that business-as-usual (BAU) emissions fail to meet the cumulative target for 2050 in all 10,000 runs of our forecast. In fact, mean BAU cumulative emissions are 181 billion metric tons while our cumulative target is 78.5 billion metric tons (the vertical line in Figure 3).

4 Policies and results

In this section we develop and discuss the different types of policies that we will analyze using the model developed in the preceding section. We begin by discussing the bounding cases of a simple tax and ETS, and then introduce hybrid policies. In all cases, we imagine a cumulative target of 79 billion tons CO_2 for emissions over the 31 years from $t_0 = 2020$ to $T = 2050$. We then turn to the numerical analysis to understand how different policies achieve different outcomes.

Figure 1: Historical (1929-2015) and Projected (2016-2050) Emissions

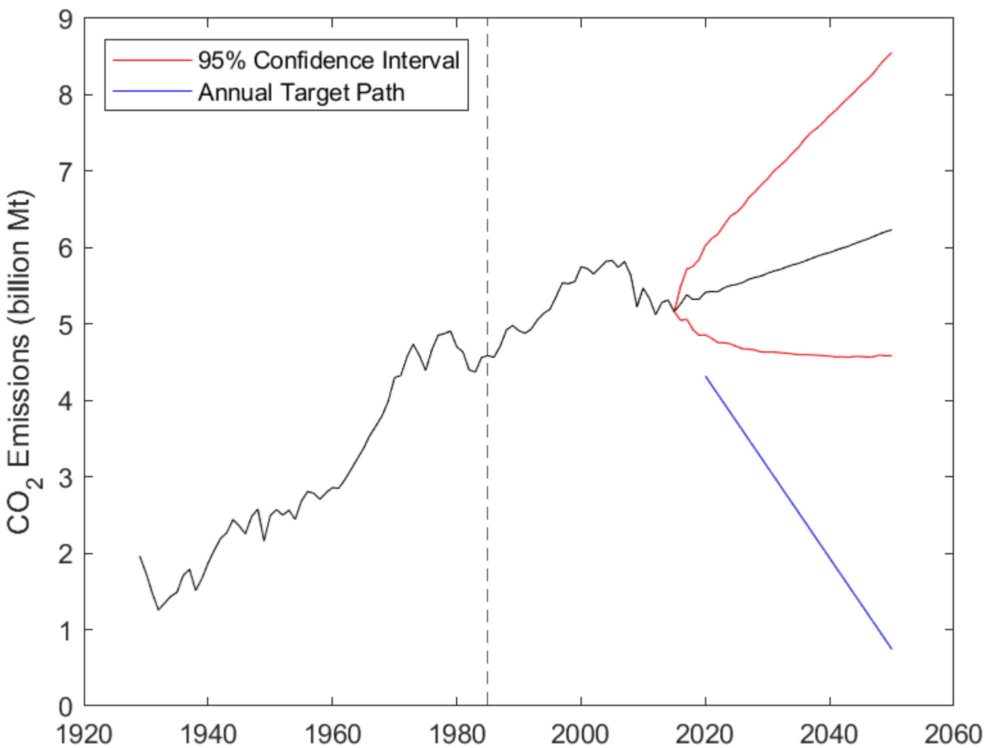

4.1 Approaches to hybrid (or dynamically-adjusted) tax policies

All of the policies that we consider involve a price being set each period t, in many cases differentiated by observable emissions in each state, e_t^s, but in all cases being equated by firms to marginal costs. This price most often reflects a tax, but sometimes can be interpreted as a market price arising among clever forward-looking agents.

We start by noting that, based on the objective (1), there is a first-order condition for an optimum

$$MC(e_t^s)e^{-\delta(t-t_0)} + a + (b-a)\mathbb{E}[\mathbb{1}(E_T^s > \bar{E})|t,s] = 0$$

where $\mathbb{E}[()|s,t]$ is a shorthand for the expectation of () at time t in state s. We further note that at time t in state s, future values e_{t+u}^s remain uncertain (both because of uncertain future baseline emissions $\tilde{e}_{t+u}^s$ and the choice of e_{t+u}^s). That is, the marginal cost should equal the expected marginal damage from the constraint on cumulative emissions in period T, given what we know in period t and state s. This condition also must also be met in expectation as we think about future marginal costs,

$$\mathbb{E}[MC(e^s_{t+u})e^{-\delta(t+u-t_0)}|t,s] + a + (b-a)\mathbb{E}[\mathbb{1}(E^s_T > \bar{E})|t,s] = 0$$

The first-order condition reflects the desire (or political need) to set the NPV marginal cost in each period and state of nature to equal expected marginal damages and/or political constraint. Expressed in terms of NPV dollars today (2020), this marginal value lies in the interval $[a,b]$ by construction. There is also an expectation that welfare (constrained)-maximizing prices will rise at the discount rate in the future.

We can imagine this price arising from a market of clever agents facing an emission trading scheme (ETS) policy with cumulative cap $\bar{E}$ along with price floor $ae^{\delta(T-t_0)}$ and price ceiling $be^{\delta(T-t_0)}$ in the last period T. These agents therefore attempt to glean $a + (b-a)\mathbb{E}[\mathbb{1}(E^s_T > \bar{E})|t,s]$ in each period t and state s in order to figure out if the cumulative cap will be exactly binding in the last period or if price floor or ceiling will be triggered.

The first policy (a) that we consider is exactly that—a policy that sets the price in each period t and state s based on observed emissions up to that period, knowledge of the (deterministic) cost function (4), and knowing the stochastic process for future $\tilde{e}^s_{t+u}$. This information is coupled with goal of maximizing the expected objective (1) or, equivalently, minimizing expected costs and the firm's payments in the last period to purchase allowances (the price depending, in part, on whether the ceiling or floor is triggered).

Alongside this policy, for comparison, we present (b) a similar policy that sets the price each state s and period t in order to achieve the target $\bar{E}$ in expectation without a price ceiling or floor (equivalent to setting $a = 0$ and b higher than any expected marginal cost). That is, policy (b) approximates a pure ETS. For convenience, we call policy (a) a price collar and (b) an ETS.

Such forward-looking policies (a-b) are likely too complex to implement as actual tax policies because they do not involve automatic, transparent rules that can be easily written into law. Instead, they rely on the regulator to update the price by inputting realized emissions into a forward-looking emission/cost model to re-solve for the cost-minimizing price path. However, we have already noted these policies are an approximation to actual cap-and-trade policies. Policies (a-b) only allow prices to change at fixed intervals and based on specific information; market prices under cap-and-trade can change at any frequency and may use additional information.

Next we consider a policy (c) that fixes the price path (rising at discount rate δ) across all states to achieve the target $\bar{E}$ on average across all states. Note this does not consider the proposed objective, replacing the first two terms in (1) with an actual quantity constraint that $\mathbb{E}[E_s^T] = \bar{E}$ exactly. Separately, we consider a fixed price path policy (d) where the starting price is chosen to maximize the objective (1). Policies (c) and (d) are equivalent to a simple tax that rises at the discount rate, with the starting prices chosen slightly differently.

Policies (a-d) represent a hybrid ETS, an ordinary ETS, and alternative simple tax policies that rise over time. We now consider dynamic adjustment process for taxes based on simple rules. In particular, we consider a set of simple rules of the form $\tau_t^s = f_t(E_t^s, \tau_{t-1}^s)$, where τ_t^s is the tax in period t and state s and $f_t()$ is a simple time-dependent function. That is, we allow the tax rate in a given period and state to depend on the previous-period tax rate, cumulative emissions, and a time dependent function.

The three policies we consider are defined by different functions f_t:

(e) $f_t^{\text{Path-level}} = (\tau_0 + (\Delta \times \mathbb{1}[E_t^s > \hat{E}_t]))e^{\delta t}$. Here the tax rate falls on one of two paths, starting at two different points τ_0 and $\tau_0 + \Delta$ with both paths rising at the discount rate δ. Each period the choice of which tax rate depends on whether cumulative emissions exceed a threshold $\hat{E}_t$. The parameters $\{\tau_0, \Delta, \hat{E}_1, ...\hat{E}_T\}$ are chosen to maximize (1).

(f) $f_t^{\text{Path-growth}} = \tau_0 e^{(\gamma + (\Delta \times \mathbb{1}[E_t^s > \hat{E}_t]))t}$. There are again two paths for the tax rate, but now they start at the same point τ_0 and rise at different rates γ and $\gamma + \Delta$. The parameters $\{\gamma, \Delta, \hat{E}_1, ...\hat{E}_T\}$ are chosen to maximize (1).

(g) $f_t^{\text{Metcalf EIM}} = \tau_{t-1}^s e^{\gamma + (\Delta \times \mathbb{1}[E_t^s > \hat{E}_t]))}$. Now the tax rate bifurcates in each period, growing either by γ or $\gamma + \Delta$. The parameters $\{\tau_0, \gamma, \Delta, \hat{E}_1, ...\hat{E}_T\}$ are chosen to maximize (1). This is called "Metcalf EIM" (Emissions Integrity Mechanism) because it mimics Metcalf (2009) and (largely) Metcalf (2020).

Policies (e-g) include automatic tax adjustment mechanisms in which the adjustment rules are based on a pre-defined sense of "reasonableness"—we can think of policies in this class as *rules of thumb*. But rather than choosing parameters in an ad hoc way, adjustment parameters are optimally chosen to improve the objective developed in Section 3. Hafstead and Williams (2019) begin to think about this in noting the importance of using an expected emissions path compared to an uninformed target path as the basis for $\hat{E}_t$ and making adjustments. However, that may or may not be optimal and the other adjustment parameters could also be selected to improve the objective function. We now turn to our results.

4.2 Results

Based on the way we have characterized our objective, we care about both mitigation costs and emission outcomes, which can be translated into benefits (avoided damages) or responses

Figure 2: Cumulative Emissions and Cost Distributions

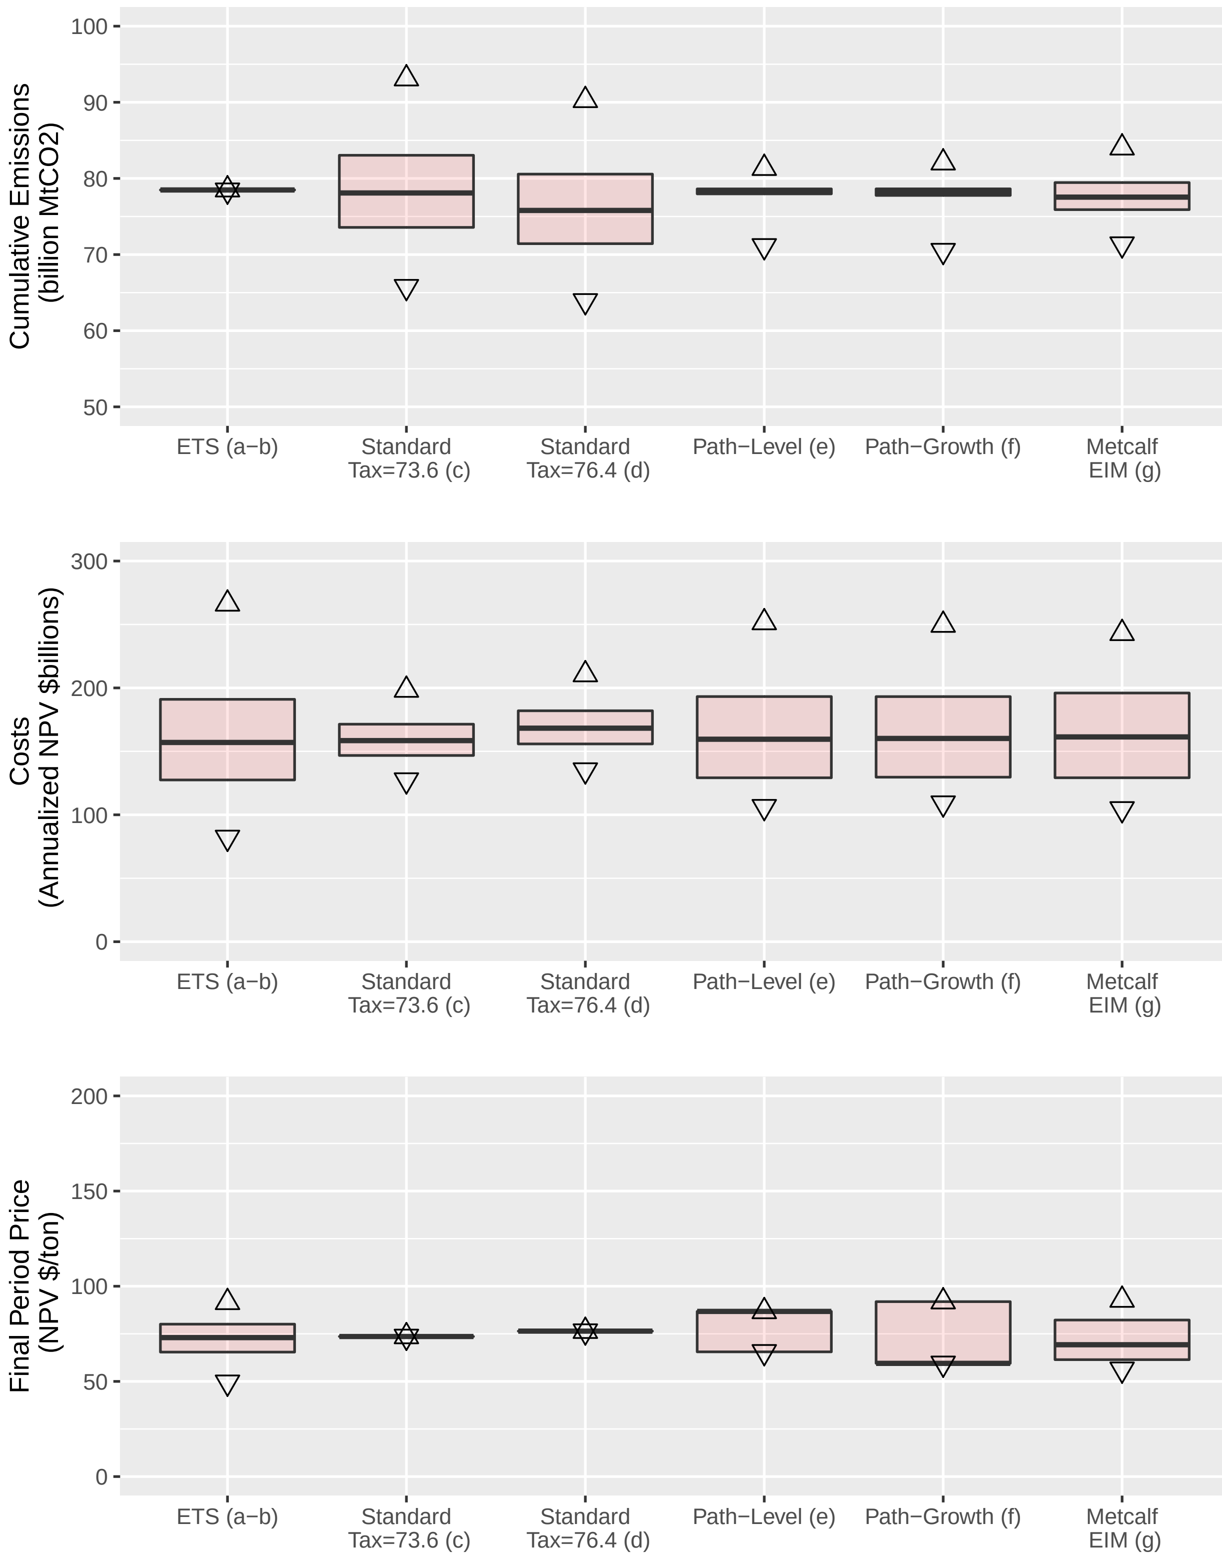

Triangles represent the 2.5 and 97.5 percentiles of the distributions. Boxplots represent the interquartile range and the median.

to political pressure. With this in mind, we first present the results for all seven policies in Figure 2 before turning to numerical results in Table 1.

Figure 2 shows the distribution of cumulative emissions and costs (in billions of dollars, annualized over 2021-2050) for all seven policies (a-g). The boxes indicate the 25-75% interquartile range across the simulated outcomes, the thicker line in the middle indicates the median, and the small trianges indicate the 2.5-97.5% range. We first notice that the ETS policies (a) and (b) are identical (so much so, that we only provide one boxplot for both). They both exactly achieve the cumulative emission target of 79 billion tons but have a relatively wide range of cost outcomes. It is notable that policy (a), the ETS with a price floor and ceiling achieves the emission outcome exactly: It implies that that the price ceiling (\$142/ton CO_2) is sufficiently high to allow the target to met in all scenarios based on our emission and cost estimates. It also implies that the price floor (\$34/ton CO_2) is *not* high enough to encourage any reductions below the target. This is reflected in both the top panel, showing identical emissions in every scenario, and in the bottom panel, showing prices that range from about \$50 to \$100 (e.g., well within the \$34 to \$142 range)

The second related outcome that we notice is the similarity between policies (c) and (d), our two tax policies that rise over time but are otherwise fixed over uncertainty. Recall that policy (c) was designed to match the cumulative emission target on average. Policy (d) was chosen to maximize objective (1). Why is policy (d) different at all? If marginal costs and emission consequences were linear, they would be the same (Weitzman, 1974). However, we know that benefits (avoided climate damages and/or political constraints) are very *non-linear* based on the discontinuity at $\bar{E} = 79$ billion tons. In particular, under the tax (c) of \$73 (see Table 2), roughly half the emissions will be above the target and face the higher marginal consequences (\$142) and roughly half will be below the target and face the lower marginal damages (\$34). Compared to the average emission consequence of $\$88 = (\$34 + \$142)/2$, it suggests a tax optimized using (1) will be higher than the \$73 tax that matches the target on average. On the other hand, as the tax rises above \$73, emissions fall, more scenarios fall below the target, and the expected marginal emission consequence in (1) falls below \$88. The equilibrium ends up being \$76.[7]

Finally, we turn to the main focus of the paper, policies (e-g). These policies adjust the carbon tax periodically using relatively simple rules—a binary choice each period determined by whether observed cumulative emissions up to that period exceed a threshold. The binary choice is either a high or low growth rate applied to the previous period tax rate (g - "Metcalf

[7]As the tax rises, we also do not expect the same emission response in all scenarios. In particular, it is typically harder to reduce emissions in the high-emission scenarios. This will further depress the optimized tax (d).

EIM"), or a fixed high or low path (e-f). The two paths in policy (e - "Path-level") are both growing at the discount rate but start at different prices; the two paths in policy (f - "Path-rate") start at the same price but grow at different rates. In all cases (e-g), the parameters of the policy and threshold are chosen to maximize the objective.

We make two observations about policies (e-g). First, they all do relatively well (visually) in terms of hitting the cumulative emission target and generating broadly similar costs. Generally, the interquartile range is tightly around the cumulative target, with only the tails falling below or above. They all look much closer in outcomes to the ETS policies (a-b) than standard tax policies (c-d). Second, policies (e-f), the policies that jump between two otherwise fixed price paths, do look slightly closer to the ETS than the Metcalf EIM policy (g). The Metcalf EIM has slightly lower high-cost tail. Given the style of policy included in most proposals in the 116th Congress match the Metcalf EIM, this is an interesting result.

To put these results in more quantiative terms, and see the parameter outcomes, we turn to Tables 1 and 2. Note that the top panel A of Table 1 treats the above-target emission penalty $b - a$ as a political constraint, and does not include that in the damage or net benefit calculation. The bottom panel B takes the revealed preference approach, and includes the above-target emission penalty in climate damages. Therefore panel B represents how the objective itself compares policies, while panel A reflects how someone skeptical of the emission target would interpret the results of necessary political compromise.

The tax and ETS comparison in Table 1 (policies a-d) add to the visual discussion above. In particular, we can see more clearly that the two ETS outcomes are identical. Moreover, we can see that as we raise the tax slightly from policy (c), which matches expected emissions, to policy (d), which maximizes the objective, we achieve a slight ($1.1 billion/year) improvement in the objective in Panel B, column (3). However, if we do not view the reduction in above-target emissions as a real welfare gain, in Panel A, column (3), this lowers net benefits by $5.38 billion/year. This is the underlying cost of political compromise to the emission target skeptic.

We can also see the standard Weitzman result in Panel A: With flat marginal benefits, the equivalent tax (policy (c)) achieves a $2.02 billion/year gain over an ETS (policy (a-b)) in column (3).

Focusing on Panel B for a moment, and the revealed benefit interpretation, we can think about policies (a) and (d) as defining a range from the "first-best" ETS with a price ceiling and floor to the "worst" simple approach of a fixed tax schedule that rises at the discount rate. We can then make several important observations. First, comparing other policies as deviations from the first-best ETS in column (4), the fixed tax schedule (policy (d)) achieves 97.9% of the welfare benefits associated with the ETS. In dollar terms, the best policy

Table 1: Policy Comparisons (all dollar values are $2009)

A. Traditional Mitigation Benefits and Political Constraint

	Mitigation Costs ($B) (1)	"Traditional" Climate Damages ($B/yr) (2)	Net benefits relative to no policy ($B/yr) (3)	Net benefits relative to ETS (%) (4)	Net benefits relative to tax, ÷ ETS relative to tax (%) (5)
No policy	0.00	351.66	0.00	-100	NA
(a-b) ETS & price collar	161.72	152.16	37.67	0.00	100
(c) Tax to hit target	159.71	152.16	39.69	5.35	159.87
(d) Tax to max objective	169.57	147.68	34.31	-8.93	0.00
(e) Path-level	164.80	150.97	35.79	-4.99	44.11
(f) Path-growth	165.24	150.76	35.55	-5.62	37.09
(g) Metcalf EIM	164.72	149.29	37.55	-0.33	96.35

B. Revealed Climate Mitigation Benefits

	Mitigation Costs ($B) (1)	"Revealed" Climate Damages ($B/yr) (2)	Net benefits relative to no policy ($B/yr) (3)	Net benefits relative to ETS (%) (4)	Net benefits relative to tax, ÷ ETS relative to tax (%) (5)
No policy	0.00	985.40	0.00	-100	NA
(a-b) ETS & price collar	161.72	152.21	671.47	0.00	100
(c) Tax to hit target	159.71	169.46	656.24	-2.27	-8.06
(d) Tax to max objective	169.57	158.46	657.37	-2.10	0.00
(e) Path-level	164.80	152.70	667.90	-0.53	74.68
(f) Path-growth	165.24	152.71	667.44	-0.60	71.45
(g) Metcalf EIM	164.72	155.51	665.17	-0.94	55.35

Notes: The letters in parentheses correspond to the policy descriptions in the text.

achieves \$671 billion in annualized net benefits, while the worst \$657. This should not be surprising. With a cumulative emission target of 79 billion tons, emissions are being reduced by more than 50%. The benefits of these reductions are large compared to the variation in Figure 2.

Second, among our optimized dynamic adjustment rules (e-g), they all obtain the majority of the \$14 billion in annualized difference between the ETS (a) and tax (d), highlighted in column (5). The Path-level based policy (e) does the best, at only \$3.6 billion worse than the ETS (75% of the ETS-tax difference). The Path-growth policy (f) is close at only \$4 billion worse. The Metcalf EIM policy drops a bit more, at \$6.3 billion. The advantage of both (e) and (f) is that they allow bigger price adjustments to hit the target—jumping between two distinct price paths—while policy (g) only allows relatively small growth differences from the previous period price. Whether \$2-3 billion in additional net benefits is worth pushing for the design changes suggested by Table 1 would be something stakeholders might consider.

The discussion changes a bit as we turn back to Panel A, where welfare no longer includes the additional, \$142 per ton above-target penalty that instead becomes a necessary political constraint. The ETS (a-b) still provides a better way to meet the political constraint and improve welfare compared to a tax subject to the same constraint (d). This is evidenced by the \$3.4 billion decline in welfare from row (a-b) to row (d) in column (3). This is still a relatively small difference (8.9%) compared to the overall welfare gain of the ETS, shown in column (4). But, it is a larger % difference and smaller \$ difference compared to Panel B, because reductions in the immense (\$142 per ton) above-target penalty are no longer considered real benefits.

Turning to the simple tax rules (e-g) in Panel A, they again obtain benefits compared to the simple tax (d). But now the Metcalf EIM is the best and almost the same as the ETS. That is, it obtains 96% of the tax-ETS difference in column (5). Its response to political pressure leads to similar costs as Path-level and Path-growth, but it is less able to effectively reduce above-target emissions and, instead, ends up slightly reducing overall emissions. This leads to lower damages in Panel A (column 2, comparing row g to e/f), and higher damages in Panel B (column 2, comparing row g to e/f). So, the Metcalf EIM is not as good for environmental advocates compared to Path-level and Path-growth policies, but better for stakeholders who view the revealed preferences from recent policy proposals as a political constraint. But, again, Path-level and Path-growth are both better than the simple tax and much better than no policy at all.

4.3 Discussion and next steps

We can summarize our main simulation results as follows:

1. With high levels of mitigation, the welfare difference between price and quantity policies is small relative to the gain compared to no policy. While the scale varies somewhat depending on how we view the above-target penalty, the worst policy still obtains at least 90% of the best policy gains.

2. An ETS is preferred to a fixed tax, regardless of whether we view the high, above-target penalty as a constraint or true damages. An ETS with a price floor and ceiling offers little improvement over an ordinary ETS. An exogenous tax schedule is not as cost-effective.

3. There can be a difference between our objective-maximizing deterministic tax and the deterministic tax that matches expected emissions under an objective-maximizing ETS. This difference is driven by the non-linearity in the consequences of marginal emissions.

4. When we turn to taxes with simple dynamic adjustment mechanisms, these policies are able to achieve the intended target with relative accuracy, but do so at higher cost than an ETS. Altogether, they can achieve a large portion of the welfare difference between the deterministic tax policy and an ETS.

5. Among the dynamic adjustment mechanisms that we consider, one's preferences will depend on how one views the above-target emission penalty with a difference of between $2 and 3 billion per year. The Path-level and Path-growth are better able to reduce that penalty and achieve the target because they tend to allow larger price adjustments. Thus they are preferred when the above-target penalty is considered a real benefit.

6. The Metcalf-EIM has more trouble making those targeted adjustments and instead leads to more expected emissions above the target but lower overall expected emissions (compared to Path-level and Path-growth). It would be preferred when the above-target penalty is considered a political constraint.

These types of observations demonstrate the value of our proposed approach even as we readily note many open questions. That is, it is possible that Figure 2 or similarly designed graphics could be used directly by stakeholders to make normative judgements about policies. Hafstead and Williams (2020) and others implicitly suggest such an approach as they present a range of figures and summary statistics on costs and emissions. When we began this exercise, we also constructed a large number of figures that presented different features of the emission and cost outcomes (ultimately we found Figure 2 to be the most informative).

However, even the best-designed graphical comparison of emission and cost outcomes allows, at best, a discrete ranking of proposals. It does not easily point to where improvements exist in terms of parameter choice or further adjustments to policy design. Forcing stakeholders to contemplate the preferences underlying their proposals, one way or another, allows us to write out an objective. Ultimately, it should be relatively uncontroversial if we can achieve better environmental and cost outcomes, which an explicit objective allows us to do.

Of course, there are many open questions. At the top of the list is the sensitivity of our results to different assumptions about how we value and/or constrain emission outcomes. As discussed in Appendix C, we have taken a particular approach to using policy proposals in the 116th Congress to inform selection of a, b, and $\bar{E}$. We also allow a to be informed by 2016 estimates of the social cost of carbon. How do our results change when those parameters change? Alternatively, are there different forms of a political constraint that could be applied?

For our cost model and target, the required marginal cost to hit the target on average falls close to the middle of the interval $[a, b]$. What happens when the required marginal cost is at or below the lower bound a (e.g., the target is too easy) or above the upper bound b (the target is too hard)? And what happens when the range $[a, b]$ is smaller compared to the uncertainty about the price required to hit the target?

Regarding the price required to hit the target, we have included one source of cost uncertainty—baseline emissions. A more realistic model would include uncertainty in the marginal cost equation (4). It might also link costs over time, so that costs in one period reflect mitigation (and likely capital investment) choices in previous periods.

Finally, there are a variety of questions about policy design. We have not, for example, considered how the frequency of adjustment matters. One could also imagine more complex policies, with more than two choices each period, or rules that depend on more than cumulative emissions. Our view is that this is best informed by (a) examining whether and what kinds of changes are likely to have a large welfare effect and (b) what kinds of designs are likely practical from a stakeholder perspective.

5 Conclusion

Recent interest in tax policies that adjust dynamically to achieve an emission target raises a number of previously unexamined, or barely examined, questions. Such questioning begins with the basics: how such policies might work and how different choices deliver a range of emission and cost outcomes. These policies are fundamentally different than hybrid ETS policies studied in the past, in that hybrid tax policies will generally feature a cruder, periodically backward-looking adjustment. In contrast, hybrid ETS policies can adjust as quickly as markets can adjust and will be as forward-looking as market agents' cleverness allows.

Soon after one examines how they might work, a natural next question is what makes one policy choice better than another? A fundamental problem for any welfare analysis is that most economics and climate science suggests that taxes should not adjust in the first place; such evidence argues that an exogenous tax leads to higher expected welfare than a tax that adjusts to achieve an emission target. To create a metric for better and worse tax adjustment policies that does not lead us back towards exogenous taxes, we use recent proposals to reveal specific preferences for a discontinuity in the marginal value placed on increasing lower emission levels. These proposals place a higher tax on emissions when emission targets are being missed versus when they are being met.

We might view the high willingness to pay for above-target reductions as a political constraint imposed by the particular stakeholders behind these proposals, or as the true benefit value. Either way, it defines an objective. Here, we have used recent proposals in the 116th Congress to calibrate such a function: Reductions down to 79 billion tons (cumulative emissions 2020-2050) are valued at \$142 per ton, while further reductions are valued at only \$34 per ton. The lower value also corresponds to recent estimates of the social cost of carbon.

Combining this damage/political constraint function with a mitigation cost function, we can define a complete objective function. In turn, we can use this function to compare across policy designs as well as to adjust parameters within such designs. With our optimized parameters and policies, we find first that the ETS and simple, exogenous tax both offer dramatic welfare gains, with a difference of only 10% between the two. This follows from our target, which represents more than 50% reduction in emissions, and our estimated emission uncertainty, which is much smaller than this 50% reduction.

Turning to our taxes with dynamic adjustment mechanisms, they all achieve much of the tax-ETS difference. The difference among the mechanisms that we consider amounts to \$1-3 billion per year. The ranking among these alternative tax rules, however, depends on how one views the above-target penalty and importance of the target—is it a political constraint or an indication of real welfare? If the former, the Metcalf EIM responds to the constraint with better welfare outcomes than the two policies that switch between two otherwise fixed paths (Path-level and Path growth). If the latter, the more gradual adjustment in the Metcalf EIM does worse.

Beyond these specific conclusions, we have sought to make two general points. One is that policy proposals can be a means to reveal preferences useful for policy analysis. Some might view these as societal preferences and "benefits;" others as a political constraint. Yet others might prefer to argue against any consideration of these preferences in the first place, given these preferences are at odds with other evidence from economics and climate science. Our view is that when a persistent position exists in the policy debate, it makes sense to at least consider how policymakers might usefully work with that position.

The second, more important point, is that there is significant value in translating otherwise abstract or discrete preferences (or constraints) into a continuous numerical objective. This allows us to identify adjustments to policy parameters and design that have the potential to make all stakeholders better off. Without such an effort, policymakers (and society) may miss out on billions of dollars per year in achievable welfare gains.

References

Aldy, Joseph E. 2017. "Designing and Updating a US Carbon Tax in an Uncertain World." *Harv. Envtl. L. Rev. F.*, 41: 28–40.

Aldy, Joseph E. 2020. "Carbon Tax Review and Updating: Institutionalizing an Act-Learn-Act Approach to U.S. Climate Policy." *Review of Environmental Economics and Policy*, 14(1): 76–94.

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

Baker, James A, Martin Feldstein, Ted Halstead, N Gregory Mankiw, Henry M Paulson, George P Shultz, Thomas Stephenson, and Rob Walton. 2017. "The Conservative Case for Carbon Dividends." 8.

Bledsoe, Paul. 2012. "Carbon tax could be part of eventual tax reform package."

Brooks, Susanne A, and Nathaniel O Keohane. 2020. "The Political Economy of Hybrid Approaches to a U.S. Carbon Tax: A Perspective from the Policy World." *Review of Environmental Economics and Policy*, 14(1): 67–75.

Burtraw, Dallas, Karen Palmer, and Danny Kahn. 2010. "A symmetric safety valve." *Special Section on Carbon Emissions and Carbon Management in Cities with Regular Papers*, 38(9): 4921–4932.

CAIT Climate Data Explorer. 2015. "CAIT Climate Data Explorer." World Resources Institute, Washington, DC.

Climate Home News. 2019. "Guterres asks all countries to plan for carbon neutrality by 2050."

Debt Reduction Task Force. 2010. "Restoring America's future: reviving the economy, cutting spending and debt, and creating a simple, pro-growth tax system." Bipartisan Policy Center, Washington, DC. OCLC: 681901007.

Hafstead, Marc A C, and Roberton C. Williams, III. 2020. "Designing and Evaluating a U.S. Carbon Tax Adjustment Mechanism to Reduce Emissions Uncertainty." *Review of*

Environmental Economics and Policy, 14(1): 95–113.

Hafstead, Marc, and Roberton C. Williams, III. 2019. "Mechanisms to Reduce Emissions Uncertainty under a Carbon Tax." *RFF Working Paper*, 31.

Hafstead, Marc, Gilbert E. Metcalf, and Roberton C. Williams, III. 2017. "Adding Quantity Certainty to a Carbon Tax through a Tax Adjustment Mechanism for Policy Pre-Commitment." *Harv. Envtl. L. Rev. F.*, 41: 41–57.

Heutel, Garth. 2020. "Bankability and Information in Pollution Policy." *Journal of the Association of Environmental and Resource Economists*, 7(4): 779–799. Publisher: The University of Chicago Press.

Hyndman, Rob J., and Yeasmin Khandakar. 2008. "Automatic Time Series Forecasting: The forecast Package for R." *Journal of Statistical Software; Vol 1, Issue 3 (2008)*.

IWG, Interagency Working Group on Social Cost of Greenhouse Gases and United States Government. 2016. "Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866."

Kwiatkowski, Denis, Peter C.B. Phillips, Peter Schmidt, and Yongcheol Shin. 1992. "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?" *Journal of Econometrics*, 54(1): 159–178.

Lizza, Ryan. 2010. "As the World Burns." *The New Yorker*.

McKibbin, Warwick J., and Peter J. Wilcoxen. 1997. "A Better Way to Slow Global Climate Change." Australian National University, Economics and Environment Network.

Metcalf, Gilbert E. 2009. "Cost Containment in Climate Change Policy: Alternative Approaches to Mitigating Price Volatility." National Bureau of Economic Research, Inc 15125.

Metcalf, Gilbert E. 2020. "An Emissions Assurance Mechanism: Adding Environmental Certainty to a U.S. Carbon Tax." *Review of Environmental Economics and Policy*, 14(1): 114–130.

Murray, Brian C., William A. Pizer, and Christina Reichert. 2017. "Increasing Emissions Certainty under a Carbon Tax." *Harv. Envtl. L. Rev. F.*, 41: 14–27.

National Commission on Fiscal Responsibility and Reform. 2010. "The Moment of Truth: Report of the National Commission on Fiscal Responsibility and Reform | National Commission on Fiscal Responsibility and Reform."

Newell, Richard G., and William A. Pizer. 2003. "Regulating stock externalities under uncertainty." *Journal of Environmental Economics and Management*, 45(2): 416–432.

Pizer, William A. 2002. "Combining price and quantity controls to mitigate global climate change." *Journal of Public Economics*, 85(3): 409–434.

Pizer, William A., and Brian Prest. 2016. "Prices versus Quantities with Policy Updating." National Bureau of Economic Research w22379, Cambridge, MA.

Pope, C. Arden, Maureen Cropper, Jay Coggins, and Aaron Cohen. 2015. "Health benefits of air pollution abatement policy: Role of the shape of the concentrationresponse function." *Journal of the Air & Waste Management Association*, 65(5): 516–522.

Roberts, Marc J., and Michael Spence. 1976. "Effluent charges and licenses under uncertainty." *Journal of Public Economics*, 5(3): 193–208.

Ross, Martin. 2014. "Structure of the Dynamic Integrated Economy/Energy/Emissions Model: Computable General Equilbrium Component, DIEM-CGE." Duke University 14-

12, Durham, NC.

Sanders, Bernard. 2015. "S.2399 - 114th Congress (2015-2016): Climate Protection and Justice Act of 2015."

Schennach, Susanne M. 2000. "The Economics of Pollution Permit Banking in the Context of Title IV of the 1990 Clean Air Act Amendments." *Journal of Environmental Economics and Management*, 40(3): 189–210.

United Kingdom. 2009. "Carbon valuation in UK policy appraisal: a revised approach."

U.S. Bureau of Economic Analysis. 2018. "Real Gross Domestic Product."

U.S. EPA. 2015. "Regulatory Impact Analysis for the Clean Power Plan Final Rule." U.S. EPA.

van der Ploeg, Frederick. 2018. "The safe carbon budget." *Climatic Change*, 147(1): 47–59.

Weitzman, Martin. 1974. "Prices vs. Quantities." *Review of Economic Studies*, 41(4): 477–491.

Weitzman, Martin. 1978. "Optimal Rewards for Economic Regulation." *American Economic Review*, September, 68(4): 683–691.

Weyant, John P., and Jennifer N. Hill. 1999. "Introduction and Overview." *The Energy Journal*, 20: vii–xliv.

Whitehouse, Sheldon. 2015. "S.1548 - 114th Congress (2015-2016): American Opportunity Carbon Fee Act of 2015."

Wintour, Patrick, and Larry Elliott. 2009. "G8 agrees to climate targets despite differences with developing nations." *The Guardian*.

Yates, Andrew J., and Mark B. Cronshaw. 2001. "Pollution Permit Markets with Intertemporal Trading and Asymmetric Information." *Journal of Environmental Economics and Management*, 42(1): 104–118.

Ye, Jason. 2019. "Carbon Pricing Proposals in the 116th Congress."

Additional Figures

Figure 3: Baseline Cumulative Emissions 2020-2050

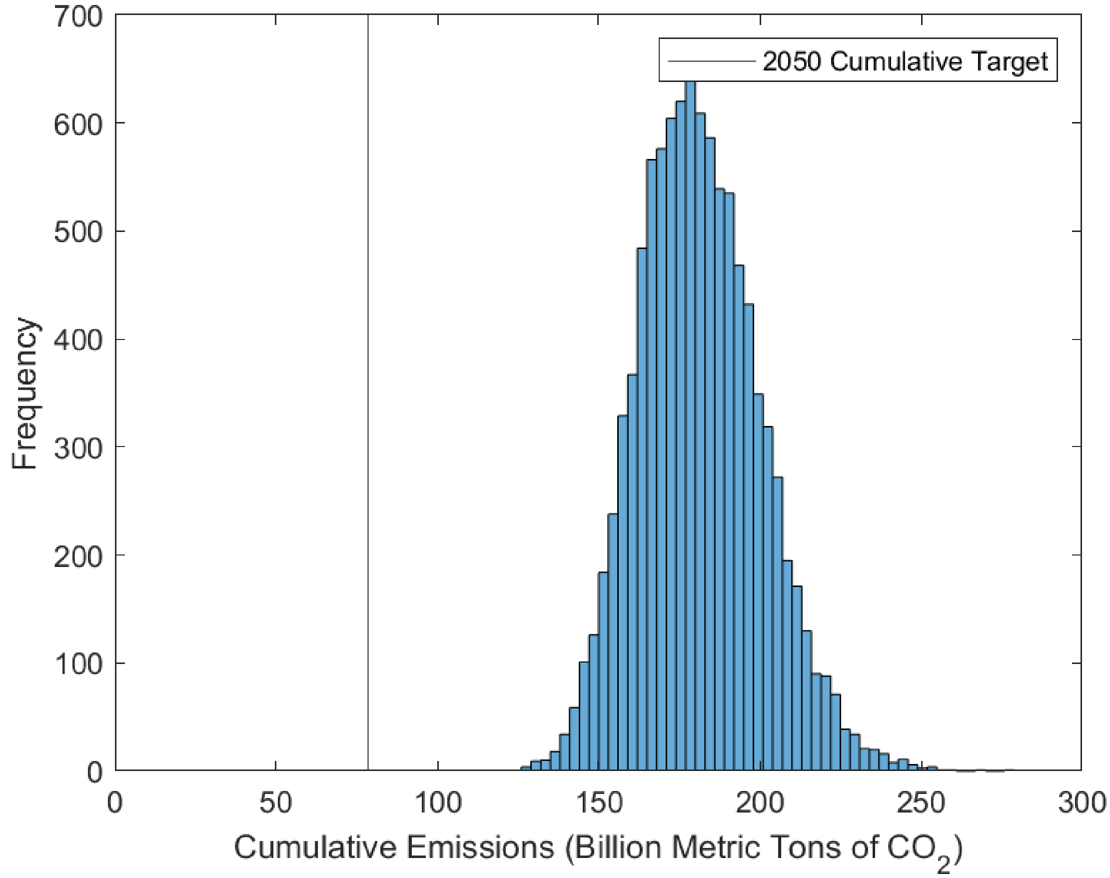

Figure 4: Price Paths

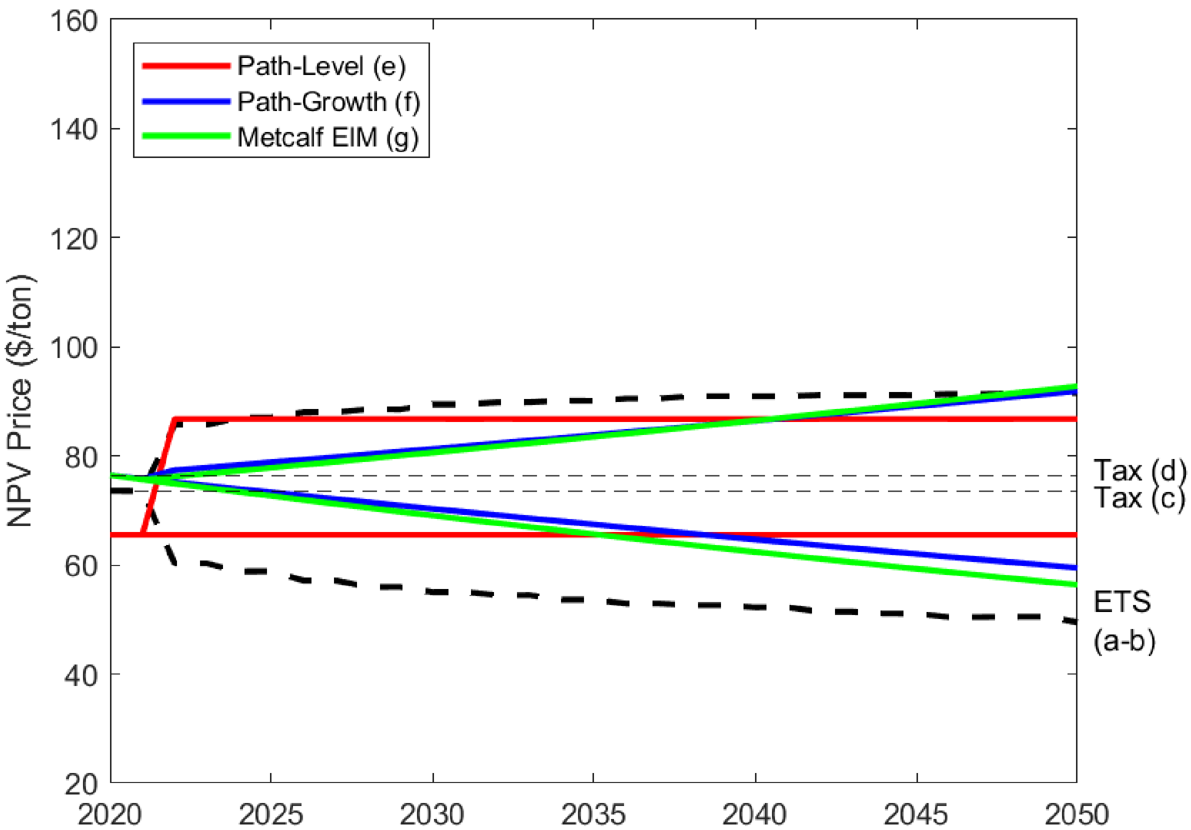

Figure 5: Threshold Paths

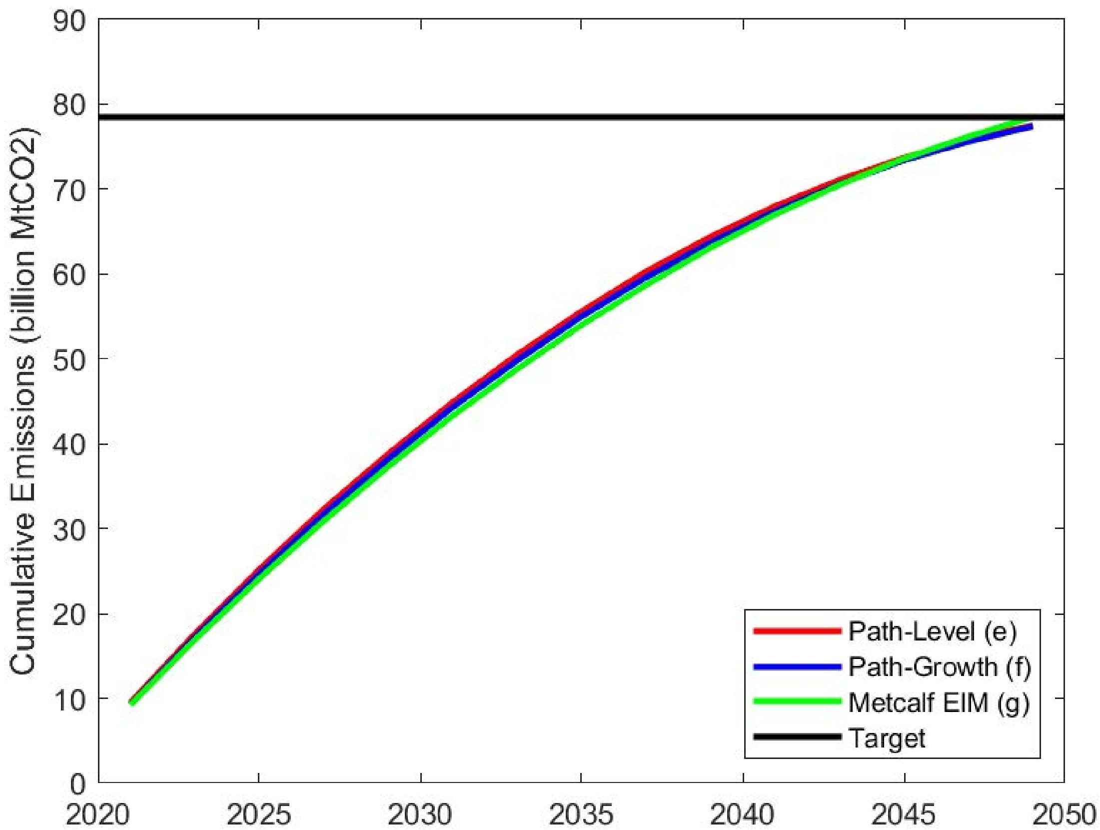

Table 2: Policy Parameters

Policy Parameters				
	Starting Price 1 ($)	Starting Price 2 ($)	Growth Rate 1 (%)	Growth Rate 2 (%)
(c) Tax to hit target	73.56	—	4.0	—
(d) Tax to min objective	76.44	—	4.0	—
(e) Path-level	65.52	86.77	4.0	—
(f) Path-growth	76.44	—	3.1	4.6
(g) Metcalf EIM	76.44	—	2.9	4.7

Notes: The letters in parentheses correspond to the policy descriptions in the text.

Appendices

A Penalizing excess emissions as a political constraint

In our objective function to evaluate alternative carbon tax mechanisms, we argue the emission target can be interpreted as a political constraint rather than part of the welfare measure itself. That is, welfare is determined by mitigation costs coupled with emission reductions valued at estimates of the benefit per ton without regard to any emission target. This "true" benefit per ton is given by parameter a in Equation (1). However, the need to garner certain political support requires a policymaker to penalize emissions above a specified target at per ton rate considerably higher than the benefit estimate. That is, we have to be willing to spend up to that higher per-ton penalty rate to hit the target before exceeding it. This penalization rate is given by $b \gg a$ in (1).

The optimization involves the same function but we no longer interpret the portion related to penalizing above-target emissions as welfare. Instead, it is the result of political pressure.

This is not the same as establishing quantitative limit on expected above-target emissions and then otherwise maximizing welfare. To see the difference mathematically, note our optimization (1) can be written as

$$\min_z f(z) - \lambda g(z) \tag{7}$$

where $-\lambda = (b - a)$, $f(z) = \mathbb{E}[aE_T^s + \sum_{t=1}^T e^{-\delta t} C_t^s(e_t^s)]$, $g(z) = \mathbb{E}[(E_T^s - \bar{E})\mathbb{1}(E_T^s > \bar{E})] - g_0$, and z includes various policy parameters that determine e_t^s in each period t and state of nature s (as well as cumulative $E_s = \sum_t e_t^s$ in each state of nature). In (1) we have $g_0 = 0$ but any fixed value of g_0 would yield the same optimization result. A typical constrained optimization would instead (i) pick a fixed quantitative constraint g_0 to be met, (ii) treat λ, the constraint's shadow price, as endogenous, and (ii) add the complimentary slackness condition that $\lambda \leq 0$ or $g(z) \leq 0$ to close the system (note λ is non-positive).

One might ask why not define the political constraint as a constraint on $\mathbb{E}[(E_T^s - \bar{E})\mathbb{1}(E_T^s > \bar{E})]$? For example, we could require the expected above-target emissions to be less than 1% of the target. A similar approach was taken in van der Ploeg (2018) to convert uncertainty about the emissions limit necessary to achieve 2° in an otherwise deterministic model: He picked an emission limit corresponding to a 2/3 probability of achieving 2°. This approach implies that increasingly flexible policies would lower costs as the constraint becomes easier to achieve and λ in Equation (7) wold fall.

We do not believe this is consistent with actual political constraints. There is a certain necessity associated with choosing a probability of 2° in the van der Ploeg example: There may well be little to no additional information about the climate sensitivity until *after* the cumulative emission limit is reached. The same is not true about the price necessary to achieve a given cumulative emission target: Each year will bring some new information about whether a higher or lower price is necessary. Pro-environmental stakeholders may be unwilling to agree to "acceptable above-target emissions" (e.g., g_0). In particular, such an agreement would not allow them to claim further reductions approach fails to further their goal of reducing above-target emissions ($g(z)$) as more flexible policies lower costs. Our view is that efforts to improve flexibility will need to go hand-in-hand with both lowering costs and reducing above-target emissions. That is, there is pressure to "share the flexibility savings with the environment" and not just use such policies to lower costs only.

B Estimation and forecasting

In Section 3.3, we define two autoregressive processes for GDP and CO_2/GDP. While we have data for the period 1929-2015, we estimate our models using only 1985-2015 data. After estimating various models using different lengths of the data, we concluded that there seemed to be state transitions occurring before 1985 that make the model predictions less informative of future emissions than we would like; thus, we use the data following what appears to be the final transition.

We use the 'forecast' package in R to conduct our estimation. The 'auto.arima' function chooses the best autoregressive integrated moving average ARIMA(p, d, q) process for each time series by following the procedure outlined in Hyndman and Khandakar (2008). First, the routine chooses the optimal number of differences d by using repeated KPSS (Kwiatkowski et al., 1992) tests, which test the null hypothesis of stationarity against the alternative of a unit root. Then the program uses a search algorithm to find the autoregressive order p and moving average order q by minimizing the small sample Akaike information criterion (AICc) of the fitted models. We check the procedure's model choices by plotting the auto-correlation and partial autocorrelation functions (ACF and PACF) of both of the differenced series, and verify that AR(1) and AR(2) seem appropriate for the differenced data. (Thus, ARIMA(1,1,0) and ARIMA(2,1,0) seem appropriate for the GDP and CO_2/GDP processes, respectively.) The parameter estimates come from maximizing the log-likelihood function and the covariance matrix comes from the Hessian. The procedure uses a Kalman filter to find the innovations and their variance. After estimating each model, we again check our model choice, by plotting the ACFs of the residuals and conclude, using portmanteau tests, that they look like white noise.[8] Equations 5 and 6 in the text characterize the models. These equations are rewritten here:

$$\Delta \log (\text{GDP}_t) = \phi \Delta \log (\text{GDP}_{t-1}) + \mu_g (1 - \phi) + \varepsilon_t$$

and

$$\Delta \log \left(\frac{\text{CO}_2}{\text{GDP}} \right)_t = \alpha_1 \Delta \log \left(\frac{\text{CO}_2}{\text{GDP}} \right)_{t-1} + \alpha_2 \Delta \log \left(\frac{\text{CO}_2}{\text{GDP}} \right)_{t-2} + \mu_r (1 - (\alpha_1 + \alpha_2)) + \eta_t$$

where ϕ, α_1, and α_2 are autoregressive coefficients, ε_t and η_t are shocks identically and independently distributed according to the Normal distributions $N(0, \sigma_g)$ and $N(0, \sigma_r)$, respectively, where σ_g and σ_r are estimated from the data, and μ_g and μ_r are drift coefficients. Parameter estimates are in Table 3.

Table 3: Model Coefficients

Coefficient	Estimate
ϕ	0.44
	(0.16)
μ_g	0.03
	(0.00)
α_1	-0.42
	(0.15)
α_2	-0.63
	(0.14)
μ_r	-0.02
	(0.00)
σ_g	0.0002
σ_r	0.0002

Standard errors in parentheses.

Then we write our own forecasting and simulation procedure using these parameter estimates. Our forecast procedure is straightforward. We take the coefficient estimates to be fixed and, for 10,000 runs, draw 35 years of error terms for each model from Normal distributions with mean zero and the estimated variances. Using our fitted models and the error draws, we forecast log(GDP) and log(CO_2/GDP) for the period 2016-2050 and then calculate GDP and CO_2 from these projections.[9] We plot our forecasts for CO_2, GDP, and

[8] Plots of the ACF and PACF are available from the authors by request.

[9] We realize that forecasting 35 years into the future using only 30 years of historical data may cause concern about the validity of our forecasts. However, our simulation procedure is for illustrative purposes and we believe it captures the trend in emissions without becoming intractable.

$\log(CO_2/GDP)$ in Figures 1, 6, and 7, respectively.

Figure 6: Historical (1929-2015) and Projected (2016-2050) GDP

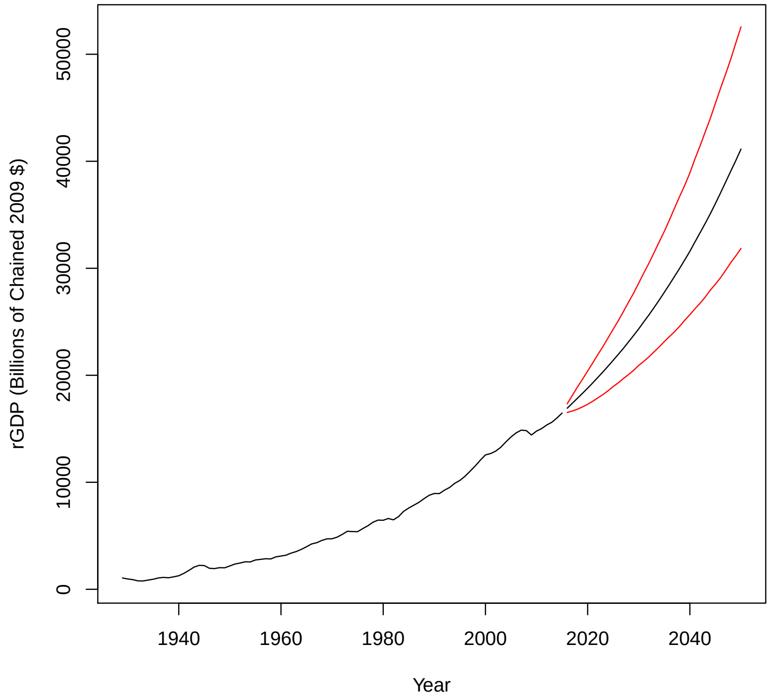

Figure 7: Historical (1929-2015) and Projected (2016-2050) log(CO2/GDP)

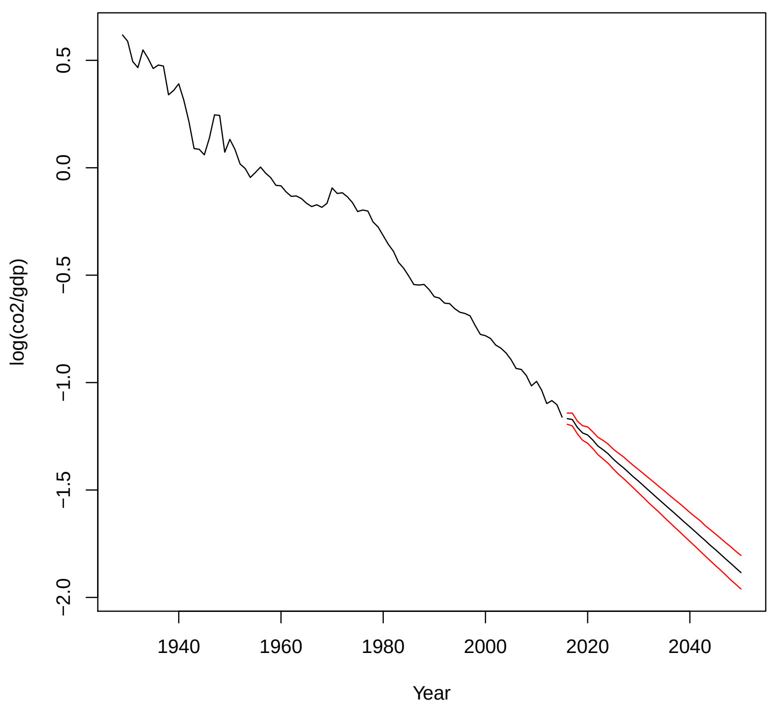

C Setting the target and marginal damage function parameters

As of the writing of this paper, eight carbon pricing proposals had been introduced in the 116th Congress (2019-2020).[10] Seven of these proposals are for a carbon tax while the other is for a cap-and-trade program. The proposals are generally defined by an initial tax rate that grows at some interest rate (or by some level) annually or biennially as shown in Figure 8. Additionally, five proposals have emission reduction targets specified by year while two others define an 80% reduction target not specific to any year. Figure 9 depicts these targets relative to 2005 emission levels. An important distinction between our work and these proposals is the set of greenhouse gas emissions covered. Each proposal differs: some apply to all greenhouse gases while others, for instance, only apply to emissions from fossil fuels and certain industrial products and processes. Our model, meanwhile, is for economy-wide carbon dioxide emissions. To accomodate this, we translate each emission reduction target into a CO_2 reduction target using the same percentages drawn from the proposal relative to 2005 CO_2 emissions. Finally, several of the proposals include built-in environmental integrity mechanisms which aim to provide greater certainty over emissions outcomes (i.e., adjustments similar in spirit to those considered in this paper). These adjustment mechanisms would increase the tax at a rate (or level) higher than the standard increase if certain targets are not met.

We define our emission reduction targets based on the mean of the annual targets defined in the proposals containing targets. For proposals with less frequent reduction targets, we extrapolate to obtain annual targets. Our targets are shown as the "Harris-Pizer" targets in Figure 9. We then convert these into annual emission targets by multiplying by 2005 CO_2 emissions. We obtain the cumulative 2050 target, $\bar{E} = 78.5$ billion metric tons of CO_2, by summing these annual targets.

To obtain a and b, we look at the net present value of the tax rates when the targets from the proposals are achieved and the NPV of the tax rates when the targets are not achieved, respectively. (NPV is calculated with a discount rate of $\delta = 0.04$.) These are depicted in Figures 10 and 11. We take $a = \$34.1$ as the average starting tax rate among all the proposals. For the proposals with integrity mechanisms, we define the maximum tax rate over time to be the tax if the proposal's targets are never met and the tax always grows at the adjustment rate. The NPV of this tax rate over time is depicted in Figure 11. We take $b = \$142.5$ as the average final maximum tax rate (NPV-terms) of this subset of proposals with built-in adjustments.

[10] See Ye (2019) for a summary.

Figure 8: 116th Congress Proposals: Tax Rates

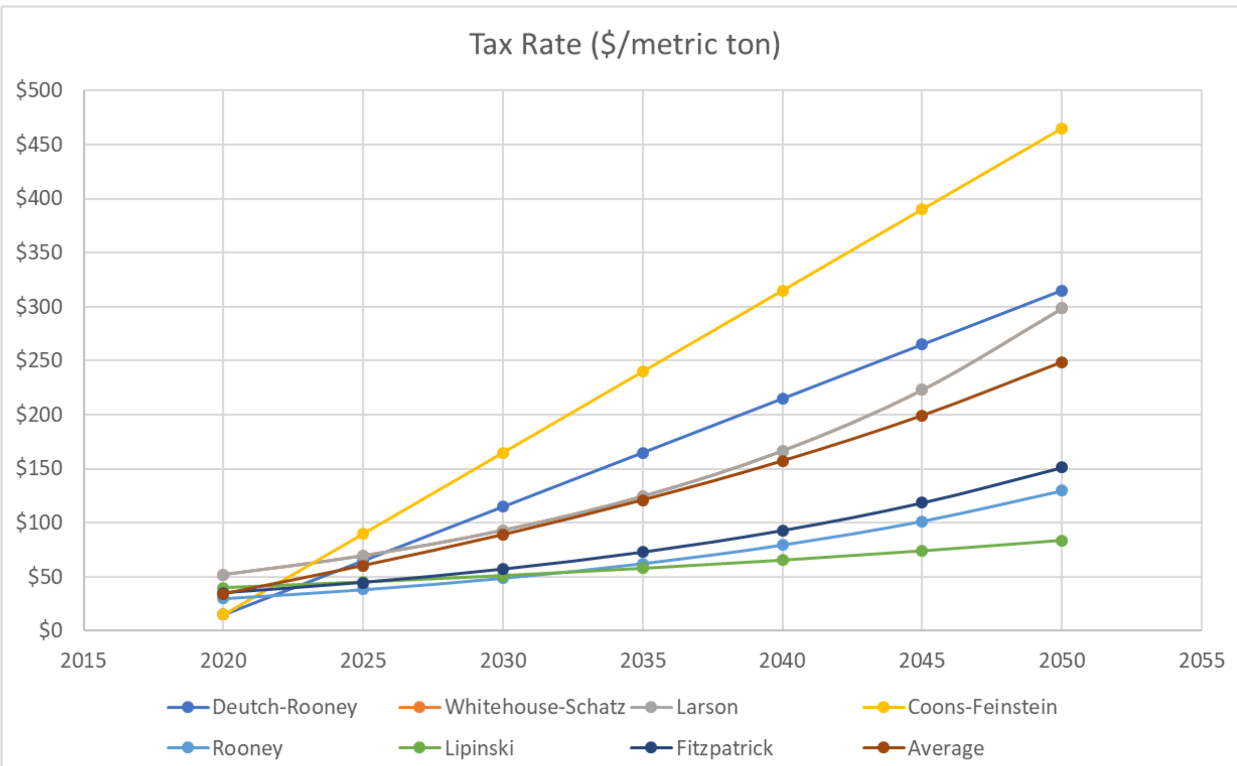

Figure 9: 116th Congress Proposals: Emission Targets

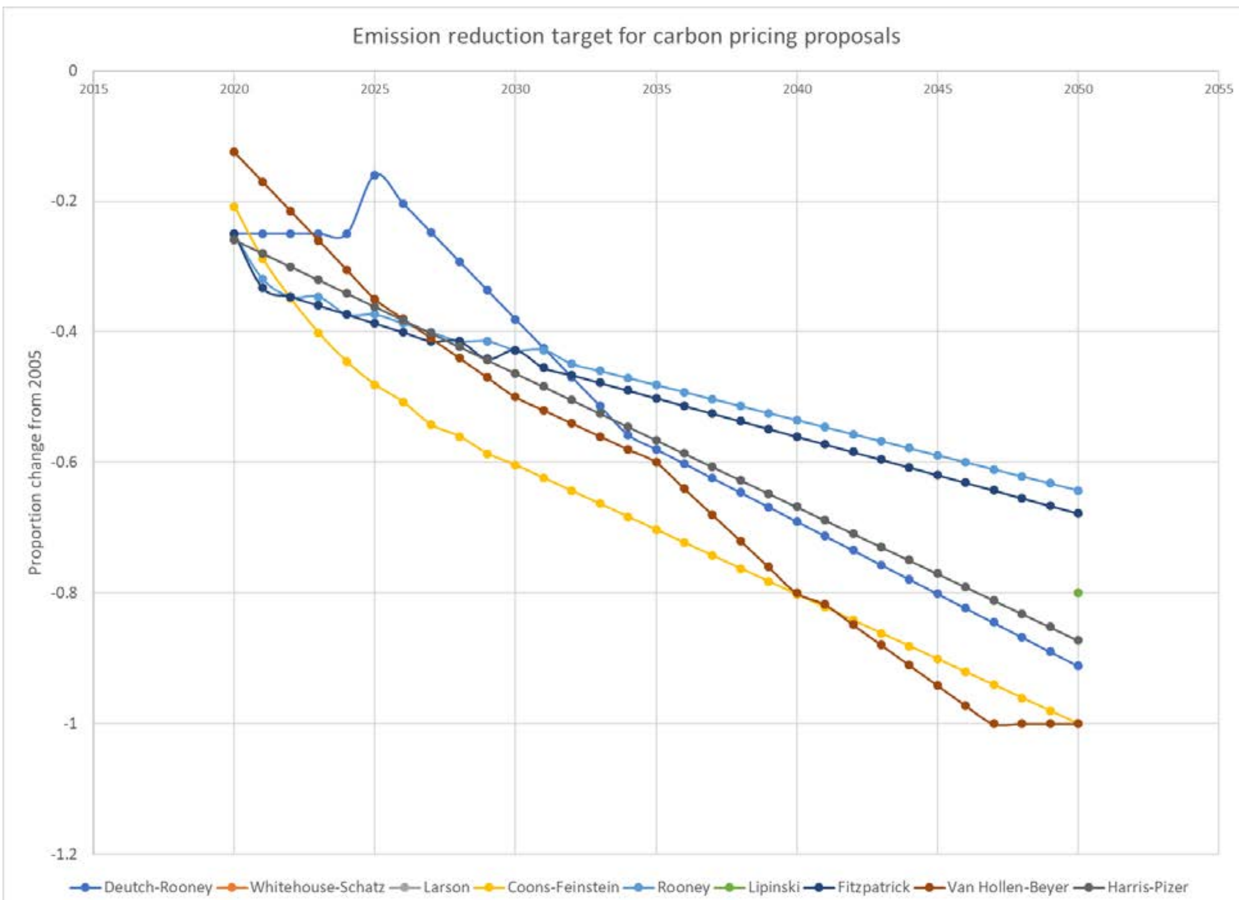

Emission reduction targets are shown relative to 2005 levels. The Harris-Pizer line represents our emission reduction targets. Our targets are defined from the regression of emission targets from the enumerated proposals on the year.

Figure 10: 116th Congress Proposals: NPV Tax Rate

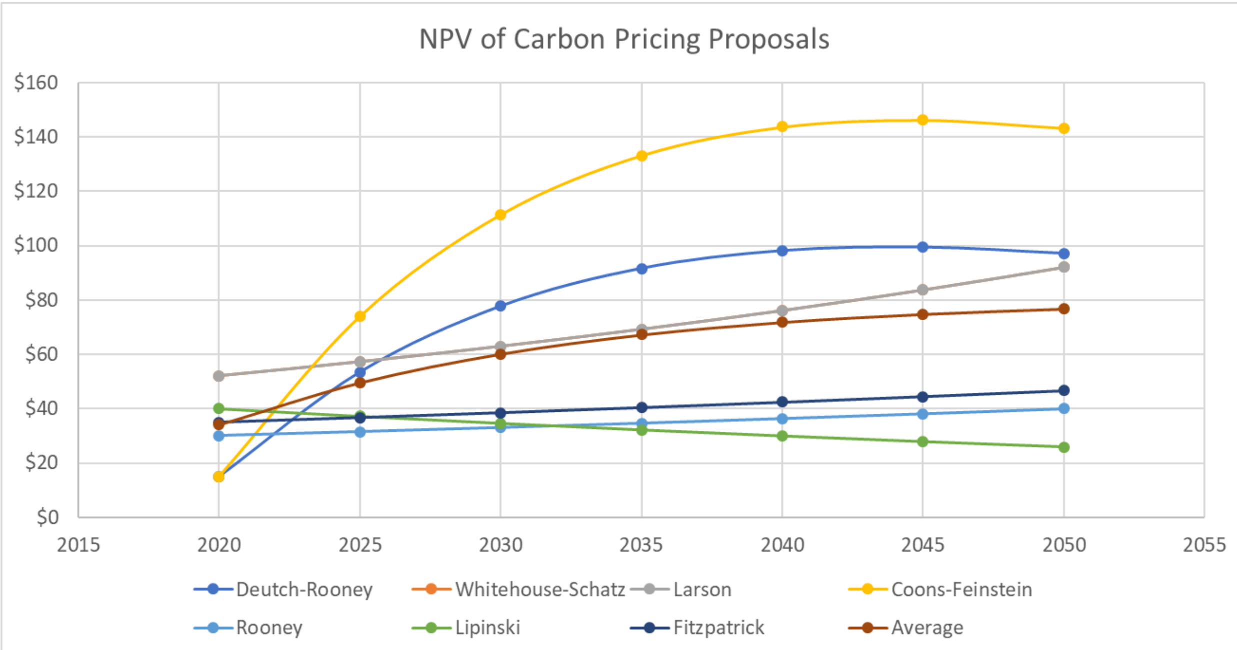

Figure 11: 116th Congress Proposals: NPV Max Tax Rate

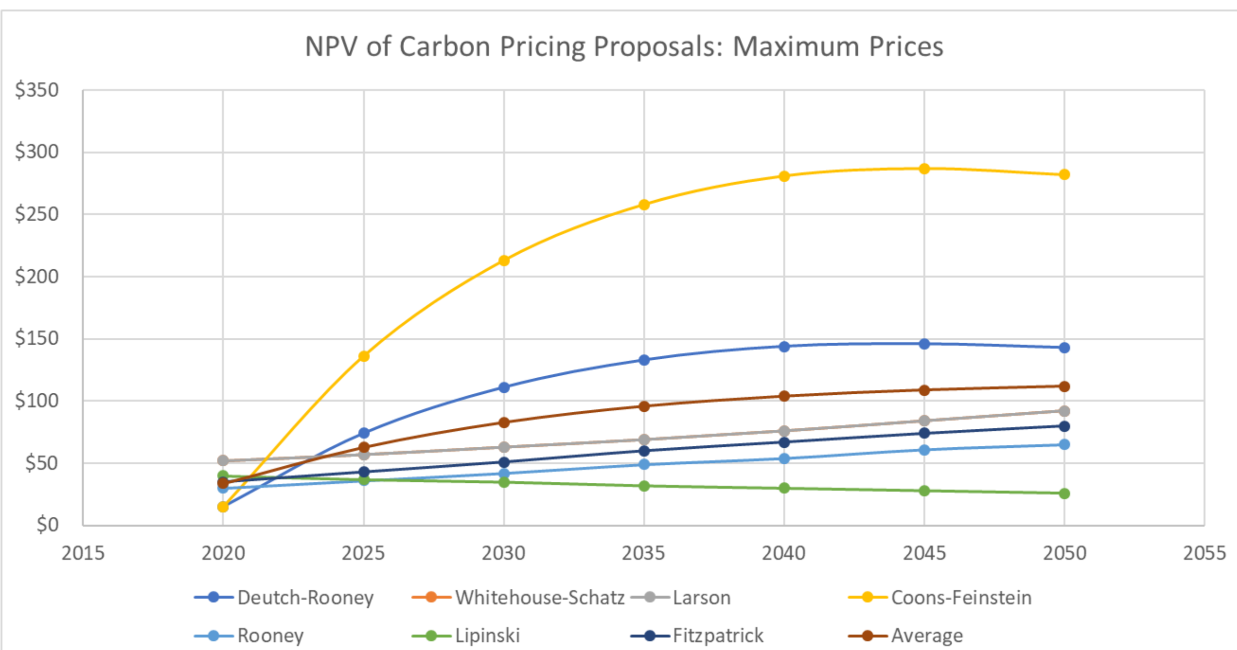

D Deriving the marginal cost curves

The DIEM-CGE model is a dynamic computable general equilibrium model of the global and U.S. economies that is multi-regional and multi-sectoral. It allows us to explore how various policies affect sectors of the U.S. economy, particularly how policies like carbon pricing affect the economy, the energy and electricity sector, and greenhouse gas emissions. We use the model to derive estimates of greenhouse gas reductions in response to carbon prices. Figure 12 shows the carbon prices that correspond to a range of percent reductions in the years 2020, 2030, 2040, 2050. From this, we estimate quadratic marginal cost curves for each decadal year, and then interpolate these coefficients across years to derive curves for every year in our period of interest (2020-2050). That is,

$$MC_t(r_t) = \beta_{1,t} r_t + \beta_{2,t} r_t^2 \tag{8}$$

where MC_t is the economy-wide marginal cost (\$/MtCO$_2$) for reducing emissions by the fraction r_t below baseline emissions in year t. As an example, the marginal cost curve for the year 2025 is $MC_{2025}(r) = 215.0r + 289.3r^2$.

Then we calculate the total cost for each year as the integral of the marginal cost curve from 0 units of reduction up to the amount of reduction relative to business-as-usual emissions. That is simply the total cost function defined in Equation 3. We rewrite it here

$$C_t(e_t) = \tilde{e}_t \left(\frac{1}{2} \beta_{1,t} (r_t)^2 + \frac{1}{3} \beta_{2,t} (r_t)^3 \right) \tag{9}$$

where $\tilde{e}_t$ is baseline emissions in period t and $r_t = (\tilde{e}_t - e_t)/(\tilde{e}_t)$ is the fraction of emissions reduced.

The net present value total cost across all periods is then

$$C^{NPV} = \sum_{t=2020}^{2050} \frac{C_t}{(1+\delta)^{t-2020}} \tag{10}$$

where $\delta = 0.04$ is the discount rate. We report annualized NPV costs (in billions \$), which are

$$C^{NPV,annualized} = \frac{C^{NPV} \cdot \delta}{1 - 1/(1+\delta)^{31}}.$$

Figure 12: Marginal Cost Curves from DIEM-CGE