

NBER WORKING PAPER SERIES

TAXING EXTERNALITIES:
REVENUE VS. WELFARE GAINS WITH AN APPLICATION TO U.S. CARBON TAXES

Matthew Kotchen

Working Paper 30321
<http://www.nber.org/papers/w30321>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2022

I am grateful to Spencer Banzhaf, Garth Heutel, and Stefano Carattini for helpful comments based on presentation of this research at an early stage. I also appreciate generous comments on an earlier draft from Joe Aldy, Robert Hahn, Arik Levinson, Gib Metcalf, Kyle Meng, and Andrew Plantinga. Any remaining errors or misstatements are nevertheless my own. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Matthew Kotchen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Taxing Externalities: Revenue vs. Welfare Gains with an Application to U.S. Carbon Taxes
Matthew Kotchen
NBER Working Paper No. 30321
August 2022
JEL No. H2,H21,H23,Q38,Q4

ABSTRACT

This paper asserts that reporting of the ratio of welfare gains to tax revenue should be standard protocol in economic analyses of externality correcting taxes. That this comparison might matter is somewhat of a “blind spot” in most economic analyses, for it plays virtually no role when economists recommend taxes to internalize externalities. A simple model illustrates how the ratio of welfare gains to tax revenue plays a central role in a political economy and efficiency framing of Pigouvian type taxes. The analysis also shows intuitive results about how the ratio is increasing in the marginal external costs and the equilibrium elasticity to a tax. The second part of the paper illustrates the wide range of potential results with application of carbon taxes to different fuels in the United States. For example, assuming a social cost of carbon (SCC) and a carbon tax equal to \$50 per tonne, the central estimates imply ratios of 12.1 for coal, 0.36 for natural gas, and very close to zero for diesel and gasoline. When all four fuels are combined, the ratios indicate a more proportional balance between welfare gains and tax revenue, with overall estimates ranging between 0.7 and 2.8. The paper concludes with a general appeal for economists to pay more attention to the relative magnitudes of efficiency gains and tax revenue when analyzing and advocating for externality correcting taxes.

Matthew Kotchen
School of the Environment,
School of Management,
and Department of Economics
Yale University
195 Prospect Street
New Haven, CT 06511
and NBER
matthew.kotchen@yale.edu

1 Introduction

There is perhaps no concept more fundamental to the field of environmental economics than that of Pigouvian taxes. When confronted with an externality problem—regardless of whether it is local in nature (e.g., a nuisance between neighbors) or global in scope (e.g., climate change)—economists can be relied upon deliver a policy recommendation about the need to get prices right, typically through Pigouvian taxes. By internalizing marginal external costs, Pigouvian taxes calibrate private incentives to implement the socially efficient level of market activity (Pigou 1920). Yet, despite the primacy of externality correcting taxes among economists, there are relatively few instances of their implementation in practice. Why? Political scientist Barry Rabe (2018) puts it succinctly, “...compelling ideas from economics do not necessarily suspend the laws of politics.”

People—and therefore political leaders—simply dislike taxes. In response, economists often express frustration and assert that the problem is merely one of communication. They emphasize that tax revenue is itself not a cost: it can be returned lump-sum or used to provide public goods, and unlike revenue raised from other taxes, it is distortion free when correcting an externality. While theoretically well-founded, these arguments often fail to resonate in practice because they sidestep potentially legitimate questions and concerns about tax revenue and the role of government, focusing instead on the more comfortable territory of economic efficiency.

The separation of questions about economic efficiency and tax revenue produces a foundation of theoretical results, yet the partition is difficult to maintain when it comes to policymaking in the real world.¹ This observation provides the starting point for the present paper that aims to highlight the importance of comparing the magnitude of welfare gains and tax revenue when it comes to advocating for, and understanding opposition to, externality correcting taxes. Beyond making a general, conceptual point, the paper provides an empirical application to carbon taxes in the United States, where comparisons are made across fossil fuels and economy wide for purposes of illustrating a range of potential results.

The fact is that economists rarely make comparisons between these two measures—welfare gains and tax revenue—when analyzing and recommending externality correcting taxes.

¹ The basic idea here is consistent with the caricature that economists tend to care about triangles, whereas everyone else cares about rectangles. The shape distinction is based on the way that welfare gains are usually measured as triangles in linear graphical analyses, and tax revenue is usually measured as rectangles.

The reason is that tax revenue is typically considered welfare neutral (in both partial and general equilibrium analyses), so comparing revenue to welfare gains is an “apples-to-oranges” comparison from a standard welfare economics perspective. But it is precisely this “apples-to-oranges” comparison that drives much of the debate about externality-motivated taxes in the real world. Those with a primary concern about the externality focus on the welfare gains of imposing a tax, whereas those categorically opposed to new taxes often mount the opposition. The latter group rejects the basic assumption that tax revenue is welfare neutral; they view it more as a cost or mechanism for unwelcome transfers, and they often focus on government inefficiency and the incidence on producers and consumers.

In recognition of these tradeoffs, the fundamental assertion in this paper is that the ratio of the welfare gain to the tax revenue should be an object of interest when evaluating the desirability and political feasibility of an externality correcting tax. This holds when evaluating a policy at the optimal tax rate or at any rate that seeks to address some portion of an externality. At the most basic level, the ratio provides useful information about whether a particular policy generates welfare gains that might be a tiny fraction of the tax revenue, or whether the welfare gains might be orders of magnitude larger. While clearly important from a political economy perspective, the fact that this consideration might matter is a “blind spot” in typical economic analyses, for it plays virtually no role when economists recommend taxes to internalize externalities.

One might argue that literature on “recycling” revenue from environmental taxes provides an exception (see for example Goulder 1995; Parry et al. 1999; Parry and Bento 2000). These papers focus on the potential welfare gains from using externality correcting tax revenue to reduce (or avoid increasing) other distortionary taxes. It is true that tax revenue plays an important role in these studies, and their contributions are central to the literature on environmental taxation. But even in these cases, the revenue itself is assumed welfare neutral and serves only as a go-between to focus on welfare effects and economic efficiency in other markets. Its magnitude is rarely an object of interest on its own.

Section 2 of this paper briefly describes selected examples of Pigouvian taxes from the literature. A general observation is that basic information about tax revenue is rarely even reported because of the sole emphasis on efficiency. It is nevertheless possible to back out estimates in many cases, and doing so indicates that economists sometimes advocate for policies

where the ratio of the welfare gain to tax revenue is exceedingly small. Moreover, because the relative magnitudes of welfare gains to tax revenue are not considered, economists often miss opportunities to provide information that can inform choices among policy alternatives.

Section 3 develops a simple conceptual framework to illustrate how the ratio of welfare gain to tax revenue plays a central role in a political economy and efficiency framing of Pigouvian type taxes. It shows that, absent a need to raise revenue, the ratio plays a pivotal role when evaluating the political feasibility of an externality correcting tax in the presence of tax aversion. In particular, a policy with a lower ratio is less feasible. And even under the standard assumptions, with the focus is on efficiency alone, the ratio is critical for determining whether marginal or non-marginal corrective-tax policies promote efficiency in the presence of administrative, compliance, or corruption costs. Finally, the framework is useful to show that when raising revenue is an objective, but within limits, the ratio is informative about which externalities could be targeted to produce the greatest welfare gains (i.e., those with greater ratios). In each of these cases, linkages are made between the ratio of interest and other measures commonly employed in public economics, including the marginal cost of funds (MCF), the marginal excess burden (MEB), the marginal value of public funds (MVPF), and their average (rather than marginal) analogs.²

Section 4 defines the ratio of interest in more detail in a simplified, partial equilibrium setting with no income effects. It is immediately clear how ratios are increasing in the size of the marginal external costs and the market responsiveness to implementation of a tax, which is itself increasing in both the supply and demand elasticities.

Sections 5 and 6 provide empirical estimates of ratios for implementation of a carbon tax in the United States. While much attention has focused on carbon taxes to avoid the costs of climate change and to raise tax revenue, little attention has focused on the relative magnitudes of these two impacts. The analysis builds on the model first developed in Kotchen (2021) and considers four fuels—natural gas, coal, gasoline and diesel—where the tax is considered to apply separately and combined. While all ratios are sensitive to parameter assumptions, of particular

² As discussed later in the paper, some standard references for each of the concepts are Stiglitz and Dasgupta (1971) and Atkinson and Stern (1974) for the MCF, Harberger (1964) and Auerbach and Hines (2002) for the MEB, and Hendren (2016) and Hendren and Sprung-Keyser (2020) for the MVPF.

interest here is how heterogeneity of the results across fuels exemplifies the different possibilities that the general framework is intended to illustrate.

The first set of results relate to the ratios of the welfare gain to tax revenue based on implementing a carbon tax at a range of levels between \$5 and \$125 per tonne in the market for each fuel separately. For example, assuming a social cost of carbon (SCC) and a carbon tax equal to \$50 per tonne, the central estimates imply ratios of 12.1 for coal, 0.36 for natural gas, 0.05 for diesel, and -0.05 for gasoline. Coal has the highest ratio by a wide margin because of its substantially larger marginal external costs and greater equilibrium elasticity. Natural gas represents an intermediate case, while the magnitudes for diesel and gasoline are exceedingly small and even negative. The latter results are driven by the relatively small equilibrium elasticities and the presence of preexisting excise taxes. In the case of gasoline, imposing a carbon tax over and above existing federal and state excise taxes results in a marginal tax rate that exceeds the marginal social costs due to climate damages and local pollution.

When considering all four fuels together with a uniform carbon tax, the ratios indicate a more proportional balance between welfare gains and tax revenue. Continuing to assume \$50 per tonne for both the SCC and the carbon tax, the overall ratio ranges between 0.7 and 1.03 depending on the elasticity scenario. This means that somewhere between seventy cents and one dollar of social net benefits are created for each dollar of tax revenue. Importantly, this is not a standard benefit-cost comparison precisely because the tax revenue is not a cost to society, yet not everyone will perceive it as welfare neutral either. A striking interpretation, potentially in support of the policy, is that a ratio greater than one implies that imposing the tax would pass a standard benefit-cost test even if all the tax revenue were “burned,” that is, not used to generate social benefits. The ratios are even greater if the carbon taxes for each fuel are calibrated to only take effect when they exceed preexisting excise taxes, in which case the overall ratio ranges between 0.88 and 2.82, with a central estimate of 1.84. The final part of Section 5 compares results of the model to those of more detailed and computationally intensive models that are part of the Stanford Energy Modeling Forum, and these comparisons show a high degree of externality validity. This suggests that the simple model employed here produces magnitudes in line with more complicated and computationally intensive approaches.

Section 7 concludes with a discussion of the key findings and an appeal to economists to pay more attention to revenue when analyzing and advocating for externality correcting taxes.

The virtue of Pigouvian-type taxes hinges critically on the assumption of tax revenue’s welfare neutrality—both real and perceived—and more transparency about the relative magnitudes of welfare gains and tax revenue will help clarify how important these assumptions are in particular applications. A focus on the comparison will also make many economists more aware of whether their tax-policy recommendations should be thought of first and foremost as promoting economic efficiency or as a mechanism for raising and dispensing revenue, where different views about this question often frames debate. Providing such basic information will help bridge the gap between economic analyses and the political economy of implementation. For it is surely correct that “compelling ideas from economics do not necessarily suspend the laws of politics” Rabe (2018), yet economic analysis can be conducted in more compelling ways to better inform decision-making that is subject to the laws of politics.

2 Motivating Examples

Let us now turn to selected examples of externality correcting taxes from the literature. Each one is intended to illustrate different ways in which policy evaluation would be aided with additional information about the ratio of the welfare gain to tax revenue.

2.1 Airplane noise

The 1990 Airport Noise and Capacity Act (ANCA) in the Unites States mandated the elimination of certain aircraft because of the noise externality their engines imposed on home owners near airports. Morrison et al. (1999) provide a benefit-cost analysis of the regulatory requirement and find that the present value benefits of \$5 billion fall short of the costs of \$10 billion. The authors then analyze and advocate for what would have been an alternative regulatory approach, implementation of a Pigouvian tax. They estimate that an efficient tax would have generated a net welfare gain of \$0.013 billion per year, which translates into present value net benefits of \$0.184 billion.³ What is not reported, or even mentioned, in the paper is the revenue that imposing the tax would generate. It is nevertheless possible to make the calculations

³ Consistent with other calculations in Morrison et al (1999), this assumes a 7-percent discount rate and a flow of benefits and costs in perpetuity.

that yield a present value of tax revenue equal to \$12.15 billion.⁴ The result is a ratio of the welfare gain to tax revenue that is exceedingly small: 0.015, which implies an average of 1.5 cents of benefits for each dollar of tax revenue.

The Pigouvian tax is by definition efficient assuming welfare neutrality of the tax revenue, and on these terms, the tax may be preferred to the ANCA mandate. But the fact that only 1.5 cents of benefits are created for each dollar of tax revenue should certainly play a role in debate about whether the tax policy is a preferred approach. It would only take a small amount of administrative cost (just over 1.5 percent of the revenue) for the tax to fail a benefit-cost test. Moreover, rather than a comparison to the inefficient mandate of the ANCA, the more relevant comparison is likely to be an efficient mandate that induces behavior consistent with the Pigouvian tax. This would produce the same welfare gain without implementing what is effectively a revenue raising policy in comparison to the relatively small net benefits. While the preferred policy approach is certainly open for debate, the objective here is to make two observations: First, economists, perhaps even unknowingly, sometimes advocate for externality correcting taxes when the ratio of benefits to tax revenue is exceedingly small. Second, the information necessary to make such a comparison transparent is often not reported because of the sole focus on efficiency.

2.2 Pollution taxes in the electricity sector

Griffin (1974) provides an early analysis of pollution taxes in the U.S. electricity sector that illustrates a further way in which the ratio of welfare gains to tax revenue can inform choices among policy alternatives. He considers a tax on sulfur oxide emissions consistent with proposed legislation at the time. The analysis considers three different tax rates (10, 15, and 20 cents per pound of sulfur) and a range of estimated control costs across 9 difference scenarios. While the net welfare effects are explicitly reported and discussed across scenarios as being the “critical variable to economists (p. 683),” there is also mention of how “the politician is likely to be more

⁴ Using the approach and estimates reported in Morrison et al. (1999), the procedure is as follows. Average industry revenue is \$88.7 billion, and private marginal cost is \$13,145, which they assume is equal to price. The marginal external cost equal to the tax is \$127, and this implies an assumed mark-up of one percent. The initial quantity is 6,747,813 flights (equal to $\$88.7B/\$13,145$). Then using the assumed demand elasticity of -0.7, the after-tax quantity of flights decreases to 6,700,578. Multiplying this quantity by the tax yields annual revenue of \$0.851 billion, which translates into a comparable present value of \$12.15 billion.

concerned with the distributional effects of the tax (p. 683).” However, not mentioned or reported is the tax revenue itself.

Backing out estimates of the tax revenue based on information provided, one finds that the ratios vary dramatically across scenarios: from of a low of 0.25 to a high of 9.96.⁵ That is, depending on the scenario, the tax policy produces a welfare gain of between 25 cents and \$10 per dollar of tax revenue. But perhaps of even more interest is a comparison across scenarios with the same compliance cost assumptions and different tax rates. In the three low-cost scenarios, the net welfare gain is roughly constant across the three tax levels, ranging between \$42 and \$43 billion over a 10-year period, but the tax revenue increases with the tax rate. In particular, the estimates over the 10-year period are \$4.22 billion at 10 cents, \$5.58 billion at 15 cents, and \$7.4 billion at 20 cents, and the corresponding ratios are 9.95, 7.63, and 5.76. While all the ratios might seem relatively favorable, the differences among them is policy relevant. With welfare benefits roughly equal across scenarios, the preferable option is likely to depend on whether raising revenue is considered desirable or not, yet this observation is not even apparent with standard reporting that does not report tax revenue.

2.3 A congestion charge

In 2003, the city of London began implementing a daily charge on driving and parking in designated zones and at designated times to reduce traffic congestion. An additional aim of the policy was to establish a source of tax revenue earmarked for spending on London’s public transportation. Unlike the previous examples, this is a policy with the dual aims of reducing an externality and raising tax revenue. The information is therefore more readily available for comparing the magnitude of these two impacts.

Leape (2006) provides an *ex post* analysis of the congestion charge that draws on other studies and provides new estimates. Annual net benefits in the early years were estimated at £67 million, and the revenues were less than expected at £63 million and £97 million in the first and second year, respectively. These estimates imply annual ratios of 1.06 and 0.69, and this provides a sense of the relative magnitudes in a case where real policy was focused on both the

⁵ Table 2 in Griffin (1974) reports the average annual emissions over the a 10-year period. These emissions multiplied by the tax rate and 10 years in each scenario produces estimates of tax revenue that are comparable with the 10-year estimates of the net welfare gain.

numerator and denominator. Interestingly, an unexpected feature of the policy was a significant degree of noncompliance, and enforcement through penalty notices was found to raise additional revenue of £70 million, resulting in lower ratios of 0.5 and 0.4. In this case, lower ratios are not necessarily a disadvantage given the aim of raising revenue, but the importance of noncompliance and enforcement does highlight the need to consider administrative burdens of collecting revenue that will ultimately affect both the numerator and denominator.

2.4 A soda tax

The final example is a sin tax on sugar-sweetened beverages in the United States. Allcott et al. (2019) solve for the optimal tax that accounts for both externalities and internalities, and their estimates range between 1 and 2.1 cents per ounce, depending on the specification. The net welfare gain for their baseline scenario is \$2.4 billion per year. Once again, the tax revenue is not explicitly reported in the paper, but making the calculation for the baseline scenario yields an estimate of \$6.7 billion per year.⁶ The ratio of the welfare gain to tax revenue is therefore 0.36, and it is worth noting that the externality (increased public health care costs) and the internality (consumer psychological biases) contribute to the welfare gain in roughly equal proportions.

Beyond providing a non-environmental example, the soda tax is useful for illustrating an additional point: the increased focus on revenue as part of distributional analyses. Allcott et al. (2019) consider the welfare consequences of the policy at different levels of income assuming equal lump-sum redistribution of the tax revenue. This type of analysis is increasingly considered in research on externality correcting taxes as interest has grown in whether such policies are regressive or progressive, or have other distributional impacts of potential concern.⁷ While this exemplifies one way in which greater attention is focused on tax revenue, it does not address the more overarching concern about the relative magnitudes of the overall welfare gain to the tax revenue. Yet such a comparison itself highlights how important redistribution assumptions might be for the regressivity or progressivity of an externality correcting tax.

⁶ The per capita tax revenue is shown in Figure IX of Allcott et al. (2019), although the number itself is not reported. Through personal communication, the estimate was found to be \$21.52 per capita, which is then multiplied by 311 million adult equivalents.

⁷ Examples of papers that focus on the distributional impacts of carbon taxes (which are the focus later in this paper) combined with revenue redistribution include Grainger and Kolstad (2010), Metcalf (2019), Williams et al. (2014) and Goulder et al. (2019). Other studies also consider how the expenditure of tax revenue affects political support (e.g., Kotchen et al. 2017; Klenert et al. 2018; Carattini et al 2018).

3 Why ratios matter

This section has two objectives. The first is to show in the simplest possible way how the ratio of the welfare to gain to tax revenue plays a central role in a political economy and efficiency framing of externality correcting taxes. The second is to discuss how the ratio in this setting relates to other measures commonly employed in public economics, particularly the MCF, the MEB, and their analogous measures based on averages.

At the most general level of abstraction, we can express the level of welfare in an economy as a function of an excise tax imposed on a particular good. Specifically, we can write $W(\lambda)$, where λ is the level of the tax, and it is assumed that the associated tax revenue, denoted $TR(\lambda)$, is returned lump sum. To simplify notation, the measures of welfare and tax revenue are normalized such that $W(0) = TR(0) = 0$. If the market for the good subject to the tax is perfectly competitive and there are no other distortions, welfare is maximized at $\lambda = 0$. If the good is associated with a negative externality, then welfare is maximized at some $\lambda > 0$, which is set to internalize the social marginal costs at the efficient quantity (Pigou 1932).⁸

3.1 Political economy

The focus of this paper is not only on efficiency. Let us for the moment consider a positive political economy perspective. In particular, assume there are two groups that may differ in their political influence: one cares about the welfare gains of correcting the externality, the other opposes the implementation of new taxes. If the balance of power between the different interest groups determines the level of an externality correcting tax, one might expect the level of the tax to satisfy

$$\min_{\lambda} \{ \kappa TR(\lambda) - \rho W(\lambda) \}, \quad (1)$$

where $0 \leq \rho, \kappa \leq 1$ are weights reflecting political influence. One could equivalently write the problem as maximizing benefits rather than minimizing costs, but as we will see, stating the

⁸ The Pigouvian tax rate is typically defined as the optimal tax when the externality associated with the taxed good is the only market failure or distortion, in which case the optimal tax is equal to the marginal external costs at the welfare maximizing quantity. There is, however, an extensive literature that considers the optimal tax in second-best settings, where comparisons with the Pigouvian level are often a focal point of analysis. See Bovenberg and Goulder (2002) for a survey of the foundational literature.

problem as a cost minimization is useful for seeing direct analogs to well-established concepts in the theory of taxation. Without loss of generality, we can simplify further by normalizing $\rho = 1$ so that κ has a direct interpretation: it reflects the extent to which a monetary unit of tax revenue is considered a cost compared to the same unit's worth of welfare gain. The special case of $\kappa = 0$ reverts the problem back to that of the standard assumption for maximizing efficiency in the presence of an externality—that is, tax revenue is welfare neutral and there are no administrative costs. In contrast, $\kappa = 1$ treats tax revenue fully as a cost, as if it were “burned.”

Assuming a unique interior solution, the first-order condition that defines the chosen level of the tax is $\kappa TR'(\lambda) - W'(\lambda) = 0$. Dividing by marginal revenue, this condition has a familiar interpretation as the marginal cost of funds (MCF):

$$MCF(\lambda) = \kappa - \frac{W'(\lambda)}{TR'(\lambda)} = 0, \quad (2)$$

where $W'(\lambda)$ is the Harberger (1964) measure of marginal excess burden (MEB), which in this case includes the avoided marginal external costs. Although the MCF is not typically used to motivate externality correcting taxes, the framework can readily accommodate these cases.⁹ The standard externality framework implies that $MCF(\lambda) = -W'(\lambda)$, and it follows that minimizing (1) is equivalent to maximizing the MEB inclusive of the marginal external costs. This is the Pigouvian solution where $W'(\lambda) = 0$. If $\kappa > 0$, however, the tax revenue is itself associated with a cost that comprises part of the MCF.¹⁰

An alternative, and perhaps more realistic, way to view the policymaking process is a bit more crude. It is one where the minimand in (1) is not used to set the level of the tax, but rather to evaluate the political costs and benefits of a particular tax proposal that emerges through what may be a more complicated and obscure process. In this case, political leaders might consider a given λ and ask whether the policy is a net winner or loser from their own political economy

⁹ The term MCF has different interpretations in the literature, but the most common is for cost accounting in a modified version of the Samuelson condition for optimal provision of a public good (Auerbach and Hines 2002). The standard references are Stiglitz and Dasgupta (1971) and Stern and Atkinson (1974), although neither use the term MCF. Ballard and Fullerton (1992) provide a helpful overview and also discuss circumstances where the MEB component of the MCF can be of either sign.

¹⁰ Although not central to the analysis here, it is interesting to note that the tax rate satisfying (1) need not be decreasing in the level of tax aversion. To see this, totally differentiate the first order condition to see that $d\lambda/d\kappa = TR'(\lambda)/\Omega$, where Ω is the second-order condition that is assumed to be negative. Hence the sign of this expression depends on whether marginal revenue is increasing or decreasing, that is, on which side the allocation resides on the Laffer curve.

perspective. The net cost is the minimand in (1), and the question is whether it is positive or negative for a given λ . Diving through by $TR(\lambda)$, we know that sign will be the same as the sign of

$$ACF(\lambda) = \kappa - \frac{W(\lambda)}{TR(\lambda)}. \quad (3)$$

This expression represents the average cost of funds (ACF), which is the difference between the tax aversion parameter and the average excess burden (AEB), or benefit as the case will be for a negative externality, assuming the tax rate is not set too high.

Equation (3) makes clear the pivotal role that the ratio of the welfare gain to tax revenue plays when evaluating the political feasibility of an externality correcting tax in the presence of tax aversion. With full information, one could calculate the sign and magnitude of (3). But more realistically, leaders are only likely to have a rough sense of κ , in which case (3) is useful to show that a policy with a greater ratio is more likely to be perceived as a net winner from a political economy perspective when there exists aversion to taxes. The same can be said for a marginal increase in the tax from an initial level, where the analog to (3) is simply $\kappa - W'(\lambda)/TR'(\lambda)$ for some λ . We have thus shown the importance of the ratio of welfare gains to tax revenue for both marginal and non-marginal policy changes in a political economy framework.

3.2 Efficiency

The ratio also has importance when the focus is solely on efficiency. Pigou (1947) recognized the potential role of administration and compliance costs, and it is reasonable to assume that any such costs, which could include corruption, will scale in proportion to the tax revenue (Polinsky and Shavell 1982). In such cases, the formulation of problem (1) with $\rho = 1$, is precisely the regulators problem that seeks to maximize efficiency, where κ need only be reinterpreted as the administrative or compliance cost per unit of tax revenue. While Polinsky and Shavell (1982) consider this problem in greater detail, the important observation here is that this provides another circumstance where the ratio of welfare gains to tax revenue is an important object of interest: with administrative, compliance, or corruption costs, the ratio is

critical for determining whether marginal or non-marginal corrective-tax policies promote efficiency.¹¹

The ratio can also play an important role when choices must be made among externality correcting taxes. In particular, the ratios can help set priorities among different externality-correcting taxes.¹² Consider, for example, a scenario where there exists the possibility to tax multiple goods with externalities, but there is a limit to the total revenue that is acceptable to raise. The solution to this problem for all i goods that are positively taxed will satisfy $W'_i(\lambda_i)/TR'_i(\lambda_i) = \theta + \kappa$, where θ is the shadow value on the total revenue constraint.¹³ In this case, the ratio of the marginal welfare gain to marginal tax revenue matters for setting optimal taxes among the taxed goods and for choosing which goods not to tax. Additionally, if the individual tax rates are taken as given among goods, and the regulator were to choose among a menu of which goods to tax, the key input would be the ratio $W_i(\lambda_i)/TR_i(\lambda_i)$ for all i policies (i.e., the AEB for each policy). Use of a greedy algorithm to solve the combinatorial problem, for example, would entail ordering the policies based on their ratio and choosing them in descending order until the revenue constraint binds.

4 Defining ratios

Let us now turn to specifying in more detail the ratio of interest: the welfare gain to tax revenue for an externality correcting tax. The aim is to identify factors that affect its magnitude and to consider a standard special case. To keep things simple, the focus is on a partial equilibrium perspective, where the market has no distortion other than the externality of interest.

¹¹ See Hahn and Hird (1991) for a broad discussion and empirical estimates of different ways to measure the cost of regulation that may affect a policy's overall efficiency.

¹² A series of recent papers (Hendren 2016; Hendren and Sprung-Keyser 2020; Finkelstein and Hendren 2020) advance the notion of the marginal value of public funds (MVPF) for empirical welfare analysis. The MVPF applies to policies that require government expenditures, where it is interpreted as the benefit per dollar of government spending for a particular policy. The MVPF can thus be used to target policies that maximize benefits per unit of government expenditure. In the case of an externality correcting tax, however, the welfare gain is generally associated with an increase in tax revenue rather than an expenditure. Hendren and Sprung-Keyser (2000) note the possibility for policies to have a negative net cost to the government, yet they interpret all such cases as having an infinite MVPF, with the rationale that these policies “pay for themselves.” Part of the contribution here is to show that with externalities, lumping all policies that pay for themselves into the same category can overlook important information.

¹³ For simplicity, interactions across markets for goods are not considered, so the cross price effects are assumed to be zero. The optimality condition is consistent with the solution to the following cost minimization problem: $\min_{\{\lambda_i\}} \{\kappa \sum_i TR_i(\lambda_i) - \sum_i W_i(\lambda_i)\}$ subject to $M \geq \sum_i TR_i(\lambda_i)$, where M is an overall revenue constraint.

The analysis also assumes no income effects so that consumer surplus is the correct welfare measure.

4.1 Basic setup

Consider a market for good Q . Let $D(Q)$ denote inverse demand and $S(Q)$ denote inverse supply. The setup admits the possibility for a preexisting excise tax on good Q equal to $\tau \geq 0$. We can write the equilibrium quantity as a function of the tax such that $Q(\tau)$ satisfies the market clearing condition $D(Q) = \tau + S(Q)$. The exchange of Q generates an externality, where $m(Q)$ denotes the marginal external costs with $m'(Q) \geq 0$. It follows that the efficient quantity, denoted Q^* , will satisfy $D(Q^*) = m(Q^*) + S(Q^*)$.

Now consider the possibility for a new tax λ on each unit of Q such that the overall tax is $\tau + \lambda$. The new equilibrium quantity is $Q(\tau + \lambda)$ that satisfies $D(Q) = \tau + \lambda + S(Q)$. The efficient new tax is $\lambda^* = m(Q^*) - \tau$, where $\lambda^* + \tau$ is the familiar Pigouvian tax rate. Note that, in general, λ^* could be positive or negative (i.e., a subsidy) depending on the size of the preexisting tax.

4.2 The ratio

We can now define the change in welfare and tax revenue given implementation of a new tax. For any given tax rate λ , the change in welfare can be written as

$$\Delta W(\lambda) = \int_{Q(\tau+\lambda)}^{Q(\tau)} m(Q) - [D(Q) - S(Q)] dQ. \quad (4)$$

The first term is the change in total external costs, and the second term in brackets is the net foregone private surplus. Notice that while the expression is fully general for any new tax rate λ , the level that maximizes the expression is the Pigouvian level λ^* .¹⁴ The change in tax revenue from any given λ is

$$\Delta TR(\lambda) = (\tau + \lambda)Q(\tau + \lambda) - \tau Q(\tau), \quad (5)$$

which consists of the total revenue raised at the new equilibrium quantity and the new overall tax rate minus the revenue raised at the initial equilibrium quantity and the initial tax rate.

¹⁴ To see this formally, differentiate equation (4) using Leibniz Rule, which yields the first-order condition $[m(Q) - D(Q) + S(Q)] dQ(\tau + \lambda)/d\lambda = 0$. This is satisfied at Q^* , which sets the term in brackets equal to zero and is the equilibrium consistent with $\lambda^* = m(Q^*) - \tau$.

The measure of central interest is the ratio of the change in welfare over the change in tax revenue at any level of the new tax rate λ (i.e., not necessarily the efficient level):

$$R(\lambda) = \frac{\Delta W(\lambda)}{\Delta TR(\lambda)} = \frac{\int_{Q(\tau+\lambda)}^{Q(\tau)} m(Q) - [D(Q) - S(Q)]dQ}{\lambda Q(\tau + \lambda) - \tau[Q(\tau) - Q(\tau + \lambda)]}, \quad (6)$$

where the denominator has been rearranged a bit to facilitate interpretation. Two results follow immediately from equation (6) about the factors that contribute to having a larger or smaller ratio.

Result 1.—The ratio is increasing with a greater marginal externality, $m(Q)$. This increases the welfare gain of reducing the quantity for any given tax without affecting the equilibrium quantity or tax revenue. Note that the statement holds for a given tax λ and does not necessarily hold for taxes that adjust to the size of the marginal externality, which would be the case for an appropriately calibrated Pigouvian tax. If the tax rate is endogenous to the marginal externality, the change in the numerator of (6) is different and the denominator changes as well, resulting in an indeterminate sign overall.

Result 2.—The ratio is increasing with greater responsiveness of the equilibrium quantity to implementation of the new tax, $Q(\tau) - Q(\tau + \lambda)$. This result applies unambiguously provided that the new tax is not large enough to cause the combined tax to exceed the Pigouvian level, that is, when $\lambda \leq m(Q^*) - \tau$. If this condition holds, increasing the difference $Q(\tau) - Q(\tau + \lambda)$ increases the numerator and decreases the denominator, resulting in a higher ratio.¹⁵ Note that the market responsiveness is itself increasing in both the demand and supply elasticity, each of which themselves would contribute to a greater ratio.

4.3 The linear special case

A special case with linear functions, which is the standard textbook case, is helpful to build further intuition and establish some baseline results. Assume inverse demand takes the form $D(Q) = \alpha - \beta Q$, and inverse supply is $S(Q) = \gamma + \mu Q$. Assume further a constant marginal external cost $m(Q) = \phi$. With this setup and a new tax level λ , it is straightforward to

¹⁵ Although perhaps the less interesting case, if the new tax were to push the combined tax above the Pigouvian level, greater responsiveness would decrease the numerator, as the foregone private surplus would be growing at a faster rate than the reduction in external costs. In this case, the overall effect on the ratio would depend on whether the numerator is decreasing at a faster or slower rate than the denominator, which is generally indeterminate.

solve for the reduction in external costs and the foregone private surplus. The difference between these two changes is the overall change in welfare consistent with equation (4):

$$\begin{aligned}\Delta W(\lambda) &= \frac{\phi\lambda}{\mu + \beta} - \frac{\tau\lambda + \lambda^2 - \tau^2}{2(\mu + \beta)} \\ &= \frac{2\phi\lambda - \tau\lambda - \lambda^2}{2(\mu + \beta)}.\end{aligned}\quad (7)$$

Assuming the tax is set sufficiently low, this expression is the reduction in deadweight loss (DWL) associated with a given level of the new tax. Solving now for the change in tax revenue in equation (5) yields

$$\Delta TR(\lambda) = \frac{\lambda(\alpha - \gamma - 2\tau - \lambda)}{\mu + \beta}.\quad (8)$$

The ratio of interest is equation (7) over equation (8), which holds for any preexisting tax τ and any new tax λ . Because the ratio is cumbersome even for this linear, special case, it is helpful to narrow things even more, by assuming the most basic setup with no preexisting tax (i.e., $\tau = 0$). In this case, the ratio simplifies to

$$R(\lambda; \tau = 0) = \frac{2\phi - \lambda}{2(\alpha - \gamma - \lambda)}.\quad (9)$$

Both of the results highlighted above can be seen or derived from this expression. The first is that the ratio is increasing in the size of the marginal externality ϕ (Result 1). The second is that the ratio is increasing in the equilibrium responsiveness (Result 2), but this one requires an additional step. In order to change the slope of supply or demand, we must also change the intercepts to maintain the initial equilibrium at $Q(0)$ as the pivot point.¹⁶ The conditions for shifting the slopes of supply, demand, or both yet maintaining the initial equilibrium are $d\alpha = d\beta Q(0)$ and $d\gamma = -d\mu Q(0)$. Substituting these expressions for $d\alpha$ and $d\gamma$ into (9), we have the following that accounts for the possibility of changing the slopes of both supply and demand:

$$R(\lambda; \tau = 0) = \frac{2\phi - \lambda}{2(\alpha + d\beta Q(0) - [\gamma - d\mu Q(0)] - \lambda)}.\quad (10)$$

¹⁶ In particular, we solve for the necessary intercept adjustments by totally differentiating the demand or supply curve, holding price and quantity constant, and solving for the change in the intercept given a change in the slope.

This expression shows that making the market more responsive—by decreasing the slope parameters β , μ , or both—makes the denominator smaller and therefore the ratio larger, conforming with Result 2.

A further special case of (10) is one where the tax is set not at any arbitrary level, but at the optimal (Pigouvian) level such that $\lambda = \lambda^* = \phi$. In this case, equation (10) simplifies a bit further such that the numerator becomes simply ϕ , and λ changes to ϕ in the denominator. The basic insight that follows is that the same results apply in this special case even when taxes are set optimally: the ratio is increasing in the marginal externality and the market responsiveness.

5 Carbon tax model and assumptions

Let us now turn from theory to application with implementation of a carbon tax in the United States. The analysis considers the markets for four different fossil fuels (natural gas, coal, gasoline, and diesel) separately and combined. In particular, the approach is to consider carbon taxes at different levels, along with a range of marginal external costs, to derive ratios of the welfare gain to tax revenue for each fuel individually and all four fuels combined. The focus on each market separately is intended to illustrate the range of potential results pertaining to the ratio of welfare gains to tax revenue; for as we will see, the differences between fuels highlight how results may differ in markets for other goods. The model and assumptions build on Kotchen (2021) and are briefly summarized in this section.¹⁷

5.1 The model

Supply and demand are of the constant elasticity functional form. Supply is $Q = Ap^\eta$ where $\eta > 0$, and demand is $Q = Bp^\varepsilon$ where $\varepsilon < 0$. The initial market price satisfies $p = p_b = p_s + \tau$, where p_b is the price consumers pay, and p_s is the price suppliers receive net of the preexisting tax. Market clearing implies $Q(\tau) = Bp_b^\varepsilon = Ap_s^\eta$, and this establishes prices as a function of the tax, written as $p_b(\tau)$ and $p_s(\tau)$. The initial equilibrium enables estimation of the

¹⁷ Readers are referred to the original source for further details, especially regarding background on the data and parameter assumptions. While the setup here is identical, the research questions are different. Kotchen (2021) focuses on the incidence between producers and consumers of correcting for implicit fossil fuel subsidies, and there was no consideration of the relationship between the welfare gains and tax revenue.

scale parameters $B = Q(\tau)/p^\varepsilon$ and $A = Q(\tau)/(p - \tau)^\eta$, which fully specifies both supply and demand.

With this simple and transparent setup, it is straightforward to solve for the way that implementing a new tax λ alters the equilibrium quantity. Market clearing establishes the new prices $p_b(\tau + \lambda)$ and $p_s(\tau + \lambda)$, as described above, and the new equilibrium quantity will satisfy

$$Q(\tau + \lambda) = \left(\frac{p_b(\tau + \lambda)}{p_b(\tau)} \right)^\varepsilon Q(\tau) = \left(\frac{p_s(\tau + \lambda)}{p_s(\tau)} \right)^\eta Q(\tau). \quad (12)$$

This is the standard result that with constant elasticity, changes in quantity scale in the same proportion as changes in price, and this applies to both demand and supply. Using the second equality and substituting out the supply price, we have

$$\left(\frac{p_b(\tau + \lambda)}{p_b(\tau)} \right)^\varepsilon = \left(\frac{p_b(\tau + \lambda) - \tau - \lambda}{p_b(\tau) - \tau} \right)^\eta \quad (13)$$

This equation leaves only one unknown to solve for, $p_b(\tau + \lambda)$, to fully define the equilibrium, as both $p_s(\tau + \lambda)$ and $Q(\tau + \lambda)$ can then be solved for directly. Finally, these variables, combined with the calibrated scale parameters B and A , and an assumed constant marginal external cost ϕ , are all that is needed to solve for $\Delta W(\lambda)$, $\Delta TR(\lambda)$, and $R(\lambda)$ according to equations (4), (5), and (6).

5.2 Data and parameter assumptions

Price and quantity data are needed for establishing initial conditions of the model. These were obtained from the U.S. Energy Information Administration (EIA) for the year 2018 and are reported in Table 1.¹⁸ The quantities of all four fuels—natural gas, coal, gasoline, and diesel—are based on annual domestic consumption in the United States. The price of natural gas is the annual average Henry Hub spot price. The price for coal is the annual average cost of coal delivered to electric power plants, which accounts for 93 percent of all domestic coal consumption in 2018. The prices for gasoline and diesel are the average retail prices in each year.

¹⁸ For the broad comparisons being made in this paper, variation from year-to-year over the last decade is relatively inconsequential, except for avoiding years when COVID-19 hit and had substantial effects on energy consumption.

Table 2 reports the preexisting tax rates and marginal external costs for each fuel. The first column reports the excise equivalent, preexisting taxes taken from Kotchen (2021). These account for existing royalties, severance taxes, and excise taxes, including state and federal taxes on gasoline and diesel. The second column reports marginal external costs due to local pollution, and these health-based estimates are taken from Parry et al. (2014). The third column reports the carbon dioxide (CO₂) intensity of each fuel that is multiplied by the social cost of carbon (SCC) to estimate the marginal external costs due to climate damages. The combined marginal external costs (i.e., local pollution and climate damages) are reported in the remaining columns for different estimates of the SCC that range from \$50 to \$125 per ton.¹⁹ The \$25 increments of the SCC are intended to span the preferred, interim estimate of the U.S. Interagency Working Group on the Social Cost of Greenhouse Gases (2021) and recent estimates of updated values (Carlton and Greenstone 2021).

A literature review informs the point estimates chosen for the supply and demand elasticities. The aim in each case is to select an estimate that represents a long-run elasticity applicable to the U.S. domestic market.²⁰ The estimates used here are reported in Table 3. The U.S. Department of the Interior’s (2015) MarketSim model of U.S. energy markets provides a useful resource for some estimates, including those for coal and natural gas. Whenever possible, the estimates are based on the empirical, peer-reviewed literature, and several are the same as those used in the EIA’s National Energy Modeling System (NEMS). In order to account for uncertainty in the estimates and carryout sensitivity analysis, scenarios are considered with low- and high-elasticity estimates, where these cases are consistent with a simultaneous 50-percent decrease or increase in both elasticity estimates. These ranges are reported in brackets in Table 3.

6. Results

This section reports the carbon tax results for each of the four fuels separately, all four fuels combined, and an alternative where the carbon taxes only take effect when they exceed

¹⁹ In particular, the combined marginal external cost is equal to $\phi_l + \phi_c$, where ϕ_l is the marginal external costs due to local pollution and ϕ_c is due to climate damages. The climate damages are calculated based on $\theta_k \times SCC$, where θ_k is the carbon intensity of fuel k , and SCC is the estimate of the SCC.

²⁰ In the cases of gasoline and diesel, the supply elasticity estimates take account of the upstream and global market for oil. In contrast, the markets for coal and natural gas involve fewer steps on the supply chain and are primarily domestic. See Kotchen (2021) for details.

preexisting taxes. Additionally, comparisons are made to other model estimates to provide a sense of external validity of the magnitudes.

6.1 Fuel-specific results

Figure 1 reports the change in welfare, $\Delta W(\lambda)$, across all four fuels at different levels of the carbon tax. The figure also shows how the results vary substantially with different assumptions about the magnitude of the SCC. This is not surprising because the reduction in the total external cost, which is one component of the welfare change, scales linearly with the SCC. While the results are also sensitive to changes in the elasticities, this variation is less than that for the SCC. Appendix Figure A1 illustrates the results across the range of elasticity scenarios assuming a SCC of \$50 per tonne.

The welfare gains from the carbon tax on coal are far greater than that for the other fuels, owing to its large external costs (more on this below). The gains for coal are increasing rapidly with a higher carbon tax up to about \$50 per ton, where they start to level off.²¹ While the welfare gain for coal at a \$50 carbon tax ranges between \$110 and \$192 billion, those for natural gas range between \$24 and \$60 billion. The estimates for gasoline and diesel are quite different. For gasoline, the change in social welfare is negative at all tax rates when the SCC is \$50. The reason is that the preexisting excise tax on gasoline already exceeds the marginal external cost at this level of the SCC (56 vs. 53 cents per gallon). See Borenstein and Bushnell (2022) for a detailed analysis on this topic with county-specific estimates. This means that imposing the carbon tax over and above the preexisting tax results in a welfare loss.²² The same holds for low levels of the tax even with a SCC of \$75. It is only when the SCC is greater than \$100 that the welfare gain is always positive.²³ Yet the magnitude remains small, reaching a maximum of \$4.7 billion with an SCC of \$125 and a carbon tax of \$70 per ton. The results for diesel follow a

²¹ Note that the maximum of any curve does not occur where the carbon tax is equal to the SCC. The reason is twofold: preexisting taxes and the externalities associated with local pollution. It does, however, equal its maximum value when the tax is equal to $\lambda^* = \phi_l + \phi_c - \tau$.

²² Later in the paper, we will examine how the results differ if the carbon tax is only applied when it exceeds the preexisting tax.

²³ An interesting observation in Appendix Figure A1 is that greater elasticity leads to greater welfare gains when the gains are positive, but it also leads to greater losses when they are negative. This is apparent with a comparison of the results between gasoline and diesel, where the order of the curves is reversed.

similar pattern, though generally remain positive. The reason is that diesel has a significantly larger externality due to local pollution.²⁴

Figure 2 reports the change in total revenue, $\Delta TR(\lambda)$, at different tax rates, across fuels, and for the three elasticity scenarios (central, low, and high). Total revenue does not depend on the SCC because tax rates are independently set. The low- and high-elasticity scenarios correspond with a simultaneous 50-percent decrease or increase in the central estimate of both the supply and demand elasticities.²⁵ The pattern of results illustrates how lower elasticity is associated with greater tax revenue. The magnitudes for the central estimates are roughly similar for both natural gas and gasoline, followed by diesel. The results for coal follow a different pattern. In the low-elasticity scenario, revenue continues to increase with the carbon tax rate. Yet in the central and high-elasticity scenarios, the results are consistent with a Laffer curve: revenues are initially increasing, reach a maximum, and then begin the fall. In both scenarios, revenue is falling in the tax rate above carbon taxes exceeding \$30 per tonne.

Figure 3 reports the measure of central interest—the ratio of the change in welfare over the change in tax revenue—for all four fuels and at different tax rates. To limit the proliferation of scenarios to consider, only those for the central estimates of the elasticities are considered, along with the range of SCC estimates. Other scenarios based on alternative assumptions are easily constructed from the available data. The four fuels illustrate the range of possibilities about the relative magnitudes captured by the ratio. Those for natural gas range between approximately 0.2 and 1.2. The ratios for coal are substantially larger, beginning above 6 for all levels of the SCC and reaching as high as 43 for a SCC of \$125 and a carbon tax of the same value. Coal also differs from the other fuels because the ratio is increasing in the level of the carbon tax, owing to the fact that revenue (the denominator) is increasing at a slower rate (or even decreasing) than the welfare gain (the numerator). The ratios for gasoline and diesel, in contrast, are exceeding low or even negative.

²⁴ The results for transportation fuels exclude other externalities due to congestion and accidents. While an argument can be made that in the short-run these externalities should be included (Parry et al. 2014; Kotchen 2021), the focus here is on pollution. It is also the case that congestion and accident externalities would remain even as transportation shifts from internal combustion engines to electric vehicles.

²⁵ It is useful to note that with constant elasticity of supply and demand the same proportional changes in both elasticities will affect equilibrium quantities but not prices.

The basic conceptual framework in Section 4 provides a systematic way of understanding why the results differ across fuels, and these differences provide a useful illustration of how ratios might differ in other setting where externality correcting taxes might be applied. Following the general derivation of $R(\lambda)$ in equation (6), two factors are identified as affecting the magnitude: the marginal external costs and the market responsiveness.

Table 4 illustrates the differences across the four fuels. To facilitate comparison, the units of each fuel are converted to British thermal units (Btus).²⁶ There are substantial differences in the marginal external costs across fuels. The first row shows the full marginal external costs (assuming a SCC of \$50), and perhaps more relevant for the analysis here, the second row shows the marginal external costs net of the preexisting tax, that is, the uninternalized portion of the externality. Per unit of energy, coal has the highest uninternalized externality, followed by natural gas, diesel, and a negative value for gasoline. All else equal, increases in the uninternalized externality will increase ratios (Result 1). Finally, to facilitate comparison, the market responsiveness is reported as an elasticity, showing the percentage change in Btus given a percentage change in the carbon tax. Note that this elasticity does not equal the supply or demand elasticity, nor is it constant. Hence an illustrative value is reported for a discrete change from a carbon tax of \$45 to \$50 for all four fuels. The elasticity is the greatest for coal, by a wide margin, and this further explains why coal has a substantially higher ratio (Result 2). Natural gas is also relative elastic compared to gasoline and diesel, and this contributes to its relatively higher ratios.

6.2 Aggregate results

While the fuel-specific analysis illustrates the potential variability of results for externality motivated taxes, we now turn to the aggregated results for a carbon tax applied uniformly across all four fuels. This is, of course, the most efficient way to apply a carbon tax. The results for the change in welfare and the change in tax revenue are simply sums of those shown previously.²⁷ Graphical summaries of the aggregated results, in parallel with the fuel-

²⁶ Using data from the EIA, these conversions are based on 1,036 Btu per cf of natural gas, 9,592 Btu per pound of coal, 120,286 Btu per gallon of gasoline, and 137,381 Btu per gallon of diesel.

²⁷ When aggregating the results across fuels, it is worth noting that the analysis does not account for cross-price effects. While other studies also employ this simple and transparent approach (e.g., Parry et al. 2014; Davis 2017; Kotchen 2021), cross-price effects could be important in cases where fuels are substitutes (as would hold for coal

specific results for $\Delta W(\lambda)$ and $\Delta TR(\lambda)$ in Figures 1 and 2, are shown in Appendix Figure A2. With a carbon tax of \$50 per tonne, the welfare gains range from \$132 to \$261 billion for SCCs of \$50 and \$125, respectively. The revenue at the same carbon tax level is \$152 billion for the central-elasticity estimate, with a range of \$129 and \$188 billion for the high- and low-elasticity scenarios.

The left-hand side of Figure 4 shows the aggregate ratios at different levels of the SCC and the central-elasticity scenario. The aggregate ratios indicate more of a relative balance between the welfare gains and the tax revenue. In fact, the ratios are greater than 1 at all levels of the SCC and carbon taxes less than \$40. For low levels of the carbon tax, the ratios are more centered around 2. These results suggest, rather strikingly, that for low levels of a carbon tax, there is robust evidence that the policy would pass a benefit-cost test even if all of the tax revenue were “burned.” Yet it is worth emphasizing that the tax revenue is not in fact a cost, even though some may view it that way. Returning to the focal point of a \$50 carbon tax, the ratio ranges between 0.9 and 1.7 for a SCC of \$50 and \$125, respectively. Nevertheless, as illustrated above, these aggregated results mask very different results across fuels.

Let us now consider the question of how much the results change if rather than impose carbon taxes on top of preexisting taxes, the carbon tax only takes effect if it exceeds the level of preexisting taxes. In particular, with a carbon tax of λ per tonne, the approach considered thus far has been to impose a tax of $\lambda_k = \theta_k \lambda$ for each fuel k , where the fuel-specific carbon intensity is θ_k . The alternative approach now being considered is to apply an effective new tax per fuel of

$$\hat{\lambda}_k = \max\{0, \theta_k \lambda - \tau_k\}, \quad (14)$$

where τ_k is the preexisting tax for fuel k . This means that when the carbon tax rate falls below the preexisting tax rate, no additional tax is applied. Moreover, when the carbon tax exceeds the preexisting tax rate, only the difference between the two is applied as a new tax.

Before turning to the impact of implementing (9) on ratios, it is worth noting the possibility for indeterminate effects, in general, on the numerator and denominator. When the

and natural gas in electricity generation, and gasoline and diesel for transportation), but incorporating these into the analysis in a complete way, even with estimates, is not a simple matter. One would need estimates on the rate of substitution away from fuels and not just between them. Additionally, long-run supply-side responses are another factor not included but that matter. Fully accounting for these different effects would require a general equilibrium model that simultaneously considers interactions among the four different markets for each fuel. These types of considerations are included in models that are discussed in the next subsection.

combined tax rate exceeds the marginal external costs (i.e., $\tau_k + \theta_k \lambda > \phi$), the alternative approach will increase welfare by not over taxing.²⁸ We saw this condition in the case of gasoline for sufficiently low carbon taxes and values of the SCC. In contrast, when the combined tax is less than the marginal external costs, the alternative approach will decrease welfare by exacerbating the under taxation from an efficiency perspective. As for the tax revenue denominator, whether a decrease in the effective tax rate (for a given carbon tax) increases or decreases tax revenue depends on the side of Laffer curve, and the previous results for coal illustrate the possibility for both. Appendix Figure 3A shows the aggregated results for $\Delta W(\lambda)$ and $\Delta TR(\lambda)$ and the alternative approach. Comparing these to the baseline results in Figure 2A reveals little differences in aggregate welfare, but meaningful changes in tax revenue. For example, the tax revenue for a carbon tax of \$50 with the central-elasticity estimate declines from \$152 billion in the baseline scenario to \$72 billion with the alternative approach.

The right-hand side of Figure 4 shows aggregate ratios for the alternative approach at different levels of the SCC and for the central-elasticity scenario. The panel is directly comparable to left-hand side for the baseline approach. The use of equation (14) to adjust the carbon tax rates results in substantially higher ratios. With a carbon tax of \$50, the ratio ranges between 1.8 and 3.5 for a SCC of \$50 and \$125, which is up from 0.9 to 1.7 for the baseline approach. And for lower carbon taxes, many of the ratios imply that the welfare gain is between 3 and 5 times the size of the tax revenue. The figure also shows the kink at a carbon tax of \$65, which is the level at which the effective tax rate for gasoline and diesel become greater than the preexisting tax for both fuels.

6.3 External validity

The final question to consider in this section is the extent to which the simple and transparent model used here has external validity in terms of the magnitudes it produces. There are, of course, more complicated models that analyze various topics related to carbon taxes. The comparisons made here are to a meta-analysis across 11 different energy and economy models included in the Stanford Energy Modeling Forum exercise 32 (EMF 32). Barron et al. (2018) provide a set of policy insights across models for implementation of different carbon tax

²⁸ Recall that the setup here assumes a constant marginal external cost, and the analysis admits the possibility for a carbon tax set at any level, only one of which is the efficient Pigouvian tax.

scenarios. Two sets of results they report are the CO₂ emission reductions and tax revenues for different carbon tax policies, and these can be compared to the results here as a check of external validity.

A summary of the results across all 11 models is reported on the left-hand side of the two panels in Figure 5. For example, in the top panel, the average emission reduction across EMF 32 models is 18 and 29 percent below the baseline for carbon taxes of \$25 and \$50, respectively.²⁹ Also reported are the minimum and maximum values across models. The right-hand side shows comparable results of emissions reductions for the simpler model used in the present paper, and the magnitudes are quite similar. For the central-elasticity estimate, the reductions are 25 and 35 percent for carbon taxes of \$25 and \$50, respectively.³⁰ The magnitudes of the low- and high-elasticity scenarios also correspond well with the min and max results of the EMF models. The bottom panel of Figure 5 compares results for annual tax revenue.³¹ Here again the magnitudes averaged across the 11 EMF models and those produced here are quite similar. For the \$25 carbon tax scenario, the EMF average tax revenue of \$103 billion compares to \$95 billion here. And for the \$50 carbon tax scenario, the respective magnitudes are \$170 billion and \$158 billion.

The overall conclusion, based on these comparisons, is that the very simple and transparent model employed in this paper, which can readily accommodate alternative assumptions, produces similar results on the impacts of implementing carbon taxes in the U.S. economy. The primary aim has been to illustrate the different possibilities for the ratio of welfare gains to tax revenue, both across fuels and in aggregate; yet these comparative results to the EMF meta-analysis provide a degree of external validity for interpreting the magnitudes as well.

²⁹ Specifically, the carbon tax scenarios in Barron et al. (2018) that are referenced here are for \$25 and \$50 that increase 1 percent per year. Moreover, the percentage decrease in emissions from the baseline is the cumulative effects over the years 2020 through 2030.

³⁰ To produce results for taxes of similar magnitude to those in the EMF meta-analysis, the \$25 and \$50 scenarios used here are the average results based on an assumed carbon tax of \$25 and \$30 and of \$50 and \$55. The averaging of results at these different prices is intended to come closer to the 1% increase over 10 years that was assumed in the EMF analyses.

³¹ The revenue results, which are consistent with lump-sum recycling, are taken from Table 2 in Barron et al. (2018) with two adjustments. The first is conversion to 2018 dollars using the GDP deflator. The second is to adjust the gross tax revenue to reflect an estimate of the net revenue, which accounts for reductions in tax revenue from other sources. The adjustment is made using the average “haircut” across models for each scenario as reported in Table 5 of McFarland et al. (2018). In particular, the haircut is 24 percent for the \$25 scenario and 27 percent for the \$50 scenario. The EMF revenue results are reported for 2021 through 2030, and the annual estimate is obtained by simply taking the average.

Indeed, the results are well within the bounds established in the literature by more detailed and computationally intensive approaches.

7. Discussion and conclusion

There are two central aims of this paper: The first is to argue that reporting of the ratio of welfare gains to tax revenue should be standard protocol in economic analyses of externality correcting taxes. The second is to illustrate the range of potential results with an application to carbon taxes on fossil fuels in the United States.

7.1 The case for ratios

Economists rarely hesitate to recommend Pigouvian taxes as a preferred approach for addressing externalities. Even when taxing at the optimal level is not feasible, using taxes to internalize some portion of the external costs is considered advantageous because of the associated efficiency gain. But these recommendations rely on a critical assumption that tax revenue is welfare neutral. Is this assumption always reasonable? Because political debates revolve around this question, shouldn't it matter if the tax revenue is orders of magnitude larger than the welfare gain or perhaps only a small fraction? The fundamental assertion of this paper is that the ratio of the welfare gain to tax revenue should be an object of interest when analyzing and advocating for or against an externality correcting tax.

There are several reasons why consideration of the proportional magnitudes of welfare gains to tax revenue is compelling for externality correcting taxes. The first is straightforward positive political economy. While economists advance Pigouvian type taxes with the aim of correcting a market failure, concerns about government failure often drive opposition. In other words, economists typically focus on the efficiency gains, whereas the hurdle to implementation is often the raising of tax revenue. The ratio of these two magnitudes provides a sense of proportionality between these two sets of interests and may therefore help to evaluate potential feasibility, which will also depend on whether there is an interest or need to raise revenue (Barthold 1994).

A second reason relies only on the efficiency criterion when there exists administrative, compliance, or corruption costs. In these cases, whether a policy promotes efficiency might hinge on whether the ratio of the welfare gain to the tax revenue is something like 20:1 or 1:20.

Nevertheless, the information to make such a calculation is typically not even reported in economic analyses, nor does the way that economists recommend Pigouvian type taxes acknowledge that such proportions might matter.

A third reason is because of increasing concern about the distributional consequences of policy interventions. As noted previously, economists have begun paying more attention to the different ways in which revenue generated from externality correcting taxes can be redistributed, especially in the context of economy-wide carbon taxes. Without needing to take a stand on how revenue should be redistributed, the ratio of the welfare gain to tax revenue clarifies how important decisions about the aims of tax revenue might be. When the ratios are lower, an argument can be made that its increasingly important to make the aims of the tax revenue an explicit part of the policy proposal, rather than assuming this aspect away.

Finally, the ratio of the welfare gain to tax revenue for an externality correcting tax is closely related to other concepts that are well established in the public economics literature. This paper has shown how, in general, it is equivalent to the average excess burden. And the ways in which it differs from the average cost of funds highlights the role of the ratio for evaluating the political feasibility, efficiency, or both of an externality correcting tax policy. The partial equilibrium approach adopted here, along with the assumption of no income effects, abstracts from much of the complexity that comes along with variants of the marginal cost of funds and excess burden measures. Yet in much the same way that recent work on the marginal value of public funds is helping bridge the gap to empirical work and applied policy evaluation of expenditure policies, empirically estimated ratios of the welfare gain to tax revenue can play a similar role for externality correcting tax policies.

7.2 Application to carbon taxes

The simple theory developed in this paper shows how the ratio of the welfare gain to tax revenue will depend on two fundamentals: the size of the marginal external costs and the market responsiveness to the tax. The range of potential results are illustrated across fossil fuels when considering implementation of a carbon tax in the United States. The ratios are high for coal because of its large external costs and relatively large market elasticity to implementation of a tax. The ratios are low (and even negative) for gasoline and diesel because of the substantial

preexisting tax rates and relative inelasticity. Natural gas provides an intermediate case with ratios indicating a more proportional balance between welfare gains and tax revenue.

The most efficient policy is, of course, to apply a uniform carbon tax across fuels and to account for preexisting taxes such that an overall tax rate is equal to the marginal external costs at the efficient quantity. Doing this implies ratios that range between 1.8 and 1.4 for the central-elasticity estimates and estimates of the SCC that range between \$50 and \$125 per tonne (see the right panel of Figure 4). These ratios imply that implementing a carbon tax is predicted to generate between \$1.80 and \$1.40 of welfare gains for each dollar of tax revenue raised. While decidedly not a benefit-cost ratio, these results do suggest that imposing an efficient tax would not only pass a benefit-cost test, it would do so even if all the tax revenue were “burned” and not used to generate additional social benefits. The high ratios also suggest that carbon taxes are likely to be efficient even if there are high administrative and compliance costs.

While the preceding analysis employs a simple and transparent model, the resulting estimates are very close to those of more detailed energy-economy models of the United States (Barron et al. 2018). The empirical results thus have a relatively high degree of externality validity, and the approach can readily accommodate alternative assumptions about external costs, demand and supply elasticities, and tax rates. Finally, it is worth noting that the comparisons for external validity were only possible because the other models reported information about tax revenue, which is far from standard in economic analyses of externality correcting taxes. Hence this paper concludes with a general appeal for economists: reporting on the ratio of efficiency gains to tax revenue is useful when analyzing and advocating for externality correcting taxes, but even short of that, reporting the magnitude of tax revenue itself should be standard practice.

References

- Allcott, H., B. B. Lockwood, and D. Taubinsky (2019) “Regressive Sin Taxes, With an Application to the Optimal Soda Tax,” *Quarterly Journal of Economics*, 134(3):1557-1626.
- Atkinson, A. B. and N. H. Stern. (1974) “Pigou, Taxation, and Public Goods,” *Review of Economic Studies*, 4(1): 119-28.
- Auerbach A. J. and J. R. Hines Jr. (2002) “Taxation and Economic Efficiency,” in *Handbook of Public Economics*, Vol. 3, A. J. Auerbach and M. Feldstein (eds), 1347-1421. Amsterdam: Elsevier.
- Ballard, C. and D. Fullerton. (1992) “Distortionary Taxes and the Provision of Public Goods,” *Journal of Economic Perspectives*, 6(3): 117-131.
- Barthold, T. A. (1994) “Issues in the Design of Environmental Excise Taxes,” *Journal of Economic Perspectives*, 8(1): 133–151.
- Barron, A., A. Fawcett, M. Hafstead, J. McFarland, and A. Morris. (2018) “Policy Insights from the EMF 32 Study on U.S. Carbon Tax Scenarios,” *Climate Change Economics*, 9(1): 1840003 1-37.
- Borenstein, S. and J. Bushnell (2022) “Headwinds and Tailwinds: Implications of Inefficient Retail Energy Pricing for Energy Substitution,” *Environmental and Energy Policy and the Economy*, 3, 37-70.
- Bovenberg, A. L. and L. H. Goulder (2002) “Environmental Taxation and Regulation,” in *Handbook of Public Economics*, Vol. 3, A. J. Auerbach and M. Feldstein (eds), 1471-1545. Amsterdam: Elsevier.
- Carattini, S., M. Carvalho, and S. Fankhauser. (2018) “Overcoming Public Resistance to Carbon Taxes.,” *WIREs Climate Change*. 9:e531. <https://doi.org/10.1002/wcc.531>.
- Carlton, T. and Greenstone, M. (2021). Updating the United States Government’s Social Cost of Carbon. Working Paper 2021-04, Energy Policy Institute at the University of Chicago.
- Davis, L. (2017) “The Environmental Cost of Global Fuel Subsidies,” *The Energy Journal*, 38(1): 7-21.
- Finkelstein, A. and N. Hendren (2020) “Welfare Analysis Meets Causal Inference,” *Journal of Economic Perspectives*, 34(4): 146-167.
- Grainger, C. and C. Kolstad. (2010) “Who Pays a Price on Carbon?” *Environmental and Resource Economics*, 46: 359-376.

- Goulder, L. H. (1995) “Environmental Taxation and the Double Dividend: A Reader’s Guide,” *International Tax and Public Finance*, 2: 157-183.
- Goulder, L., M. Hafstead, G. Kim and X. Long. (2019) “Impacts of Carbon Taxes Across U.S. Household Income Groups: What are the Equity-Efficiency Tradeoffs?” *Journal of Public Economics*, 175: 44-64.
- Griffin, J. (1974) “An Econometric Evaluation of Sulfur Taxes,” *Journal of Political Economy*, 82(4): 669-688.
- Hahn, R. and J. Hird (1991) “The Costs and Benefits of Regulation: Review and Synthesis,” *Yale Journal on Regulation*, 6, 233-278.
- Harberger, A. C. (1964) “The Measurement of Waste,” *American Economic Review*, 54(3): 58-76.
- Hendren, N. (2016) “The Policy Elasticity,” *Tax Policy and the Economy*, 30(1): 51-90.
- Hendren, N. and B. Sprung-Keyser (2020) “A Unified Welfare Analysis of Government Policies,” *Quarterly Journal of Economics*, 135(3): 1209-1318.
- Interagency Working Group on Social Cost of Greenhouse Gases (2021). Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide. Interim estimates under executive order 13990, United States Government.
- Klenert, D., L. Mattauch, E. Combet, O. Edenhofer, C. Hepburn, R. Rafaty, and N. Stern. (2018) “Making Carbon Pricing Work for Citizens,” *Nature Climate Change*, 8, 669-677.
- Kotchen, M. (2021) “The Producer Benefits of Implicit Fossil Fuel Subsidies in the United States,” *Proceedings of the National Academy of Sciences*, 14(118). DOI: 10.1073/pnas.2011969118.
- Kotchen, M., Z. Turk, and A. Leiserowitz (2017) “Public Willingness to Pay for a U.S. Carbon Tax and Preferences for Spending the Revenue,” *Environmental Research Letters*, 12: 094012.
- Leape, J. (2006) “The London Congestion Charge,” *Journal of Economic Perspectives*, 20(4): 157-176.
- McFarland, J., A. Fawcett, A. Morris, J. Reilly, and P. Wilcoxon. (2018) “Overview of the EMF 32 Study on U.S. Carbon Tax Scenarios,” *Climate Change Economics*, 9(1): 1840002 1-37.
- Metcalf, G. (2019) “On the Economics of a Carbon Tax for the United States,” *Brookings Papers on Economic Activity*. 405-458.

- Morrison, S. A., C. Winston, and T. Watson (1999) “Fundamental Flaws of Social Regulation: The Case of Airplane Noise,” *Journal of Law and Economics*, 42(2): 723-743.
- Parry, I. and A. Bento (2000) “Tax Deductions, Environmental Policy, and the ‘Double Dividend’ Hypothesis,” *Journal of Environmental Economics and Management*, 39(1): 67-96.
- Parry, I., D. Heine, E. Lis and S. Li (2014). *Getting Energy Prices Right: From Principle to Practice*. International Monetary Fund, USA.
- Parry, I., R. Williams, and L. Goulder (1999) “When Can Carbon Abatement Policies Increase Welfare? The Fundamental Role of Distorted Factor Markets,” *Journal of Environmental Economics and Management*, 37(1): 52-84.
- Pigou, A. C. (1920) *The Economics of Welfare*. London: Macmillan & Company.
- Pigou, A. C. (1947) *A Study in Public Finance, Third Edition*. London: Macmillan.
- Polinsky, A. M. and S. Shavell (1982) “Pigouvian Taxation with Administrative Costs,” *Journal of Public Economics*, 19: 385-394.
- Rabe, B. (2018) “The Economics—and Politics—of Carbon Pricing,” Brookings FIXGOV Post, October 15, <https://www.brookings.edu/blog/fixgov/2018/10/25/the-economics-and-politics-of-carbon-pricing/>
- Stiglitz, J. E. and P. S. Dasgupta. (1971) “Differential Taxation, Public Goods, and Economic Efficiency,” *Review of Economic Studies*, 38: 151-74.
- U.S. Department of the Interior (2015) “Consumer surplus and energy substitutes for OCS oil and gas production: the 2015 revised market simulation model (MarketSim),” Bureau of Ocean Energy Management, BOEM 2015-054, Washington, DC.
- Williams, R., H. Gordon, D. Burtraw, J. Carbone, and R. Morgenstern (2014) “The Initial Incidence of a Carbon Tax Across U.S. States,” *National Tax Journal*, 67(4): 807-830.

Table 1: Price and quantity data for natural gas, coal, gasoline, and diesel

Fuel	Price	Quantity
Natural gas (Henry Hub spot)	\$3.26 per 1000 cf	30,075 billion cf
Coal (delivered for electric power)	\$39.63 per short ton	688 million short tons
Gasoline (retail)	\$2.73 per gallon	143 billion gallons
Diesel (retail)	\$3.18 per gallon	64 billion gallons

Notes: All data obtained from the EIA for 2018. See the main text for details.

Table 2: Preexisting taxes and marginal external costs by fuel

Fuel	Preexisting tax	Local pollution	CO ₂ intensity (mt)	Combined marginal external costs			
				SCC \$50	SCC \$75	SCC \$100	SCC \$125
Natural Gas (per 1000 cf)	\$0.16	\$1.42	0.0606	\$4.45	\$5.97	\$7.48	\$9.00
Coal (per short ton)	\$2.72	\$127.45	1.9041	\$222.66	\$270.26	\$317.86	\$365.46
Gasoline (per gallon)	\$0.56	\$0.09	0.0087	\$0.53	\$0.74	\$0.96	\$1.18
Diesel (per gallon)	\$0.63	\$0.70	0.0097	\$1.19	\$1.43	\$1.67	\$1.91

Notes: See the main text and Kotchen (2021) for details.

Table 3: Demand and supply elasticity estimates

Fuel	Demand	Supply
Natural Gas	-0.55 [-0.28, -0.83]	1.6 [0.80, 2.40]
Coal	-1.75 [-0.88, -2.63]	1.9 [0.95, 2.85]
Gasoline	-0.63 [-0.32, -0.85]	2.0 [1.00, 3.00]
Diesel	-0.58 [-0.29, -0.87]	2.0 [1.00, 3.00]

Notes: See the main text for details. Numbers in brackets are the range for a 50-percent decrease and increase in the estimate.

Table 4: The marginal externality and market responsiveness for each fuel

	Natural gas	Coal	Gasoline	Diesel
Marginal externality (\$s per million Btu)	\$4.30	\$11.61	\$4.41	\$8.66
Marginal externality net of preexisting tax (\$s per million Btu)	\$4.12	\$11.49	-\$0.26	\$3.98
Market elasticity to carbon tax (from a carbon tax of \$45 to \$50)	-0.23	-1.15	-0.07	-0.06

Notes: The marginal external costs are reported in term of Btus. The estimate is based on a SCC of \$50, and the market elasticity, as a measure of the market responsiveness, is based on the central-elasticity estimates and a change in the carbon tax from \$45 to \$50 per tonne. See the main text for details on unit conversions.




Figure 1: Annual change in welfare, $\Delta W(\lambda)$, for all four fuels at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The scenarios of SCC50, SCC75, SCC100, and SCC125 correspond to the SCC at \$50, \$75, \$100, and \$125.




Figure 2: Change in tax revenue, $\Delta TR(\lambda)$, for all four fuels at different carbon tax rates and assumptions about the market elasticities. The scenarios of central, low and high elasticity correspond to the central estimates, a 50-percent decrease in both the supply and demand elasticities, and a 50-percent increase in both elasticities. See Table 3 for the specific estimates employed.

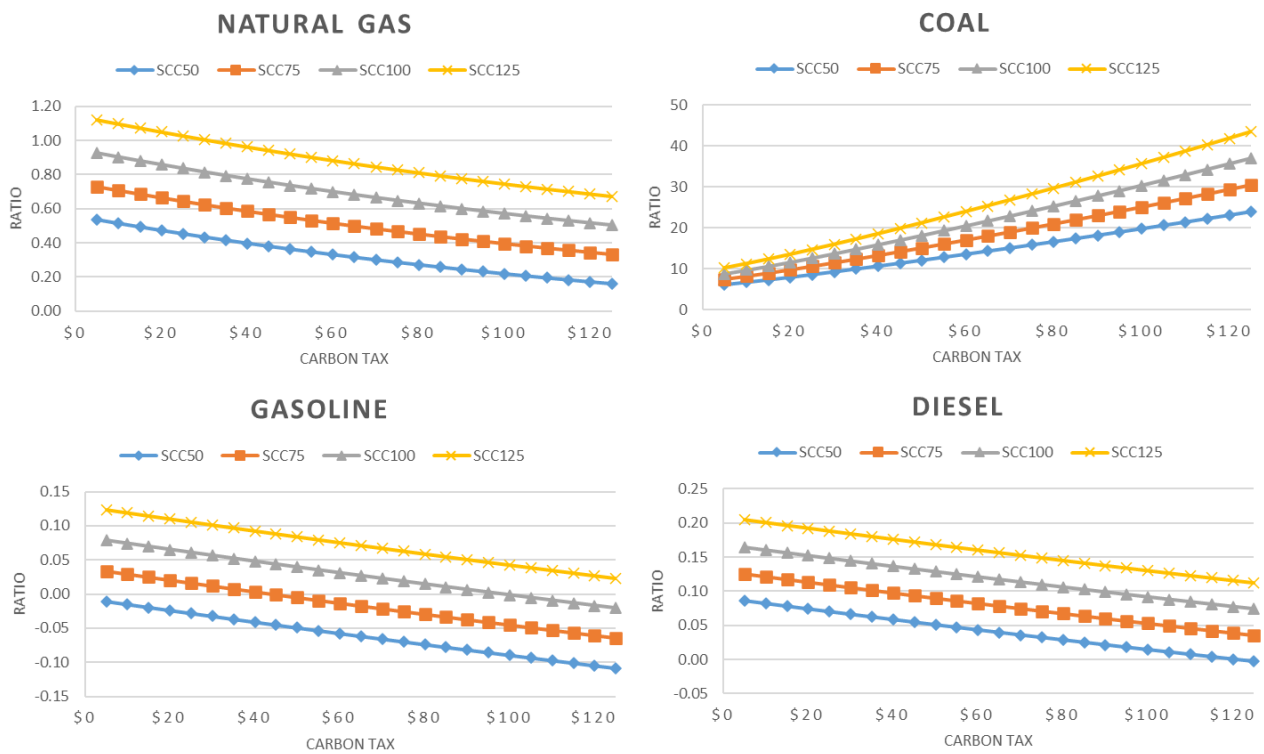


Figure 3: The ratio, $R(\lambda)$, for all four fuels at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The scenarios of SCC50, SCC75, SCC100, and SCC125 correspond to the SCC at \$50, \$75, \$100, and \$125.




Figure 4: The ratio, $R(\lambda)$, for all fuels combined at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The left-hand side is for the baseline case where the carbon tax is imposed over and above preexisting taxes. The right-hand side is for the alternative case where the carbon tax only takes effect if it exceeds the preexisting tax. The scenarios of SCC50, SCC75, SCC100, and SCC125 correspond to the SCC at \$50, \$75, \$100, and \$125.




Figure 5: Comparison of results generated here with those across models from the EMF 32 study on U.S. carbon tax scenarios. The top panel compares the percentage reduction in CO₂ emissions from a no tax baseline, and the bottom panel compare estimates of the annual tax revenue. Details on how comparisons were made within the \$25 and \$50 tax scenarios are reported in the main text.

Appendix figures




Figure A1: Annual change in welfare, $\Delta W(\lambda)$, for all four fuels at different carbon tax rates, a social cost of carbon equal to \$50, and assumptions about the market elasticity. The scenarios of central, low and high elasticity correspond to the central estimates, a 50-percent decrease in both the supply and demand elasticities, and a 50-percent increase in both elasticities. See Table 3 for the specific estimates employed.




Figure A2: The left panel shows the change in welfare, $\Delta W(\lambda)$, aggregated across fuels at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The scenarios of SCC50, SCC75, SCC100, and SCC125 correspond to the SCC at \$50, \$75, \$100, and \$125. The right panel shows tax revenue, $\Delta TR(\lambda)$, aggregated across fuels at different carbon tax rates and assumptions about the market elasticities. The scenarios of central, low and high elasticity correspond to the central estimates, a 50-percent decrease in both the supply and demand elasticities, and a 50-percent increase in both elasticities. See Table 3 for the specific estimates employed.




Figure A3: The graphs correspond with the alternative approach whereby carbon taxes are only applied for each fuel if the rate exceeds the preexisting tax, as specified in equation (14). As with Figure A2, the left panel shows the change in welfare, $\Delta W(\lambda)$, aggregated across fuels at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The right panel shows tax revenue, $\Delta TR(\lambda)$, aggregated across fuels at different carbon tax rates and assumptions about the market elasticities.