Taxing Externalities: Revenue *vs*. Welfare Gains with an Application to U.S. Carbon Taxes

Matthew J. Kotchen*†
Yale University and NBER
January 2024

Abstract

This paper illustrates how the ratio of welfare gains to tax revenue plays a central role in a political economy and efficiency framing of externality correcting taxes. Yet this comparison plays virtually no role when economists recommend Pigouvian type taxes. Intuitive results are shown about how the ratio is increasing in the marginal external costs and the equilibrium responsiveness to the tax. The second part of the paper illustrates the wide range of potential results with carbon taxes applied to different fossil fuels in the United States. For example, assuming a social cost of carbon (SCC) and a carbon tax equal to $50 per ton, 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 range between 0.9 and 1.8. 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.

JEL Codes: H23, P18, Q48, Q54

---

* Yale School of the Environment, Kroon Hall, Yale University, 195 Prospect Street, New Haven, CT 06511 USA; email: matthew.kotchen@yale.edu.

† 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 appreciate generous comments on an earlier draft from Joe Aldy, Don Fullerton, Robert Hahn, Arik Levinson, Gib Metcalf, Kyle Meng, Andrew Plantinga, Kathy Segerson, and two anonymous referees. My thinking was also refined based on helpful discussions while presenting this work at the University of California Santa Barbara, Claremont McKenna College, University of Illinois Urbana-Champaign, the University of Wyoming, and the 2023 AERE Summer Conference. Any remaining errors or misstatements are my own.

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 to 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? One explanation put succinctly by political scientist Barry Rabe (2018) is that "…compelling ideas from economics do not necessarily suspend the laws of politics."

There is a wealth of evidence in support of the notion that many people—and therefore political leaders—simply dislike taxes, and the reasons are often ideological.[1] Indeed, a pledge against increasing or implementing any new taxes is a highly influential litmus test for candidates in one of the two major political parties in the United States.[2] When it comes to correcting externalities, 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 questions and concerns about

[1] See Eisenstein (2010) and Chait (2011) for general discussions, and Carattini et al. (2018) and Andreson et al. (2023) for discussions in relation to carbon taxes in particular.

[2] See the Taxpayer Protection Pledge sponsored by the Americans for Tax Reform at https://www.atr.org/take-the-pledge.

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.[3] 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.

In practice, economists rarely make comparisons between these two measures—welfare gains and tax revenue—when analyzing and recommending externality correcting taxes. 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 typically focus on the welfare gains of imposing a tax, whereas those categorically opposed to new taxes and raising government revenue often mount the opposition. The latter group rejects the basic assumption that tax revenue is welfare neutral. Instead, they view it more as a cost or mechanism for unwelcome transfers, often focusing on government inefficiency.

[3] 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.

In recognition of these tradeoffs, the fundamental assertion in this article is that economists should pay more attention to revenue when analyzing externality correcting taxes. Specifically, I argue that the ratio of the welfare gain to the tax revenue should be an object of interest when evaluating the potential 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.[4]

What follows in this article is organized into four main parts. First, I describe a simple conceptual framework to illustrate why the ratio of welfare gain to tax revenue might matter in both a political economy and efficiency framing of Pigouvian type taxes. Second, I summarize selected examples of Pigouvian taxes from the literature to illustrate some of the key ideas. Third, I return to the framework to identify general features of markets that affect the size of ratios, including the magnitude of the marginal external costs and the market responsiveness to the tax. Finally, I provide an empirical application based on implementation of a carbon tax in the United States to further illustrate results based on the conceptual framework.

[4] One might argue that the literature on "recycling" revenue from environmental taxes provides an exception because of its focus on the use of tax revenue (see, for example, Goulder 1995; Parry et al. 1999; Parry and Bento 2000). These papers consider the potential welfare gains from using externality correcting tax revenue to reduce 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 itself is rarely an object of interest on its own.

The empirical analysis considers four fossil fuels—natural gas, coal, gasoline and diesel—where the tax is considered to apply separately and combined. Assuming a social cost of carbon (SCC) and a carbon tax equal to $50 per ton, for example, 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 range between 0.9 and 1.8. Despite the fact that these are not benefit-cost ratios, they do have a useful interpretation when greater then unity: imposing the 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 analysis is primarily intended to illustrate heterogeneity of results across fuels, but the overall magnitudes exhibit a high degree of externality validity when compared to more detailed energy-economy models of U.S. carbon taxes. While the ratios themselves do not explain why carbon taxes have not been implemented in the United States, nor do they solely indicate future prospects, they do provide insight about the ways in which some may view carbon taxes (for good or bad) through the lens of raising revenue or addressing climate change. I conclude the paper with a discussion of the key findings and an appeal to economists to pay more attention to tax revenue vis-à-vis welfare gains when analyzing and advocating for externality correcting taxes.

Ratio of the Welfare Gain to Tax Revenue

My aim in this section is to show in the simplest possible way how the ratio of the welfare to gain to tax revenue can play a central role in a political economy and efficiency framing of externality correcting taxes. 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 $W(\tau)$, where τ is the level of the tax, and it is assumed that the associated tax revenue, denoted $TR(\tau)$, is returned lump sum. If the market for the good subject to the tax is perfectly competitive, and there are no other distortions, welfare is maximized setting the tax equal to zero. However, if the good is associated with a negative externality, then welfare is maximized at some tax greater than zero, set to internalize the social marginal costs at the efficient quantity (Pigou 1932).

Figure 1 illustrates the standard, partial equilibrium framework included in most introductory textbooks. The unregulated equilibrium quantity is Q^0 at the intersection of demand and supply. The efficient quantity is Q^1 at the intersection of demand and the social marginal costs, denoted SMC. The Pigouvian tax is τ, and it is associated with a welfare gain, $W(\tau)$, equal to area of triangle ABC. The tax revenue, $TR(\tau)$, equal to the rectangle CDEF.

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 generating government revenue through implementation of new taxes. We can characterize the relative influence of these two groups with a weight $\kappa \geq 0$ that reflects the extent to which a monetary unit of tax revenue is perceived in the economy as a cost compared to the same unit's worth of welfare gain. Note that the level of tax aversion here is intended to reflect an ideological opposition to the tax rather than a cost-benefit calculation based on individuals' having to pay the tax themselves.

This simple way of characterizing the level of tax aversion and its influence necessarily overlooks context-specific features that could influence its magnitude. For example, the level of tax aversion may depend on the issue being addressed, the intended purpose of the revenue, or both.[5] And the level of influence will depend on the degree of political power of voters and lobbies. The simple parameter κ nevertheless enables us to see how the extent of tax aversion interacts in various ways with the ratio of welfare again to tax revenue. Indeed, the weighting is consistent with standard approaches in political economy for analyzing how different groups and interests affect outcomes based on their relative influence.[6]

With this setup, the question to consider is whether a particular policy—defined as having a set tax rate—is a winner or loser from a political economy perspective. The answer is yes or no depending on whether $W(\tau) - \kappa TR(\tau)$ is positive or negative, respectively. Using this condition and rearranging the expression, we find that a tax policy is a political economy winner if and only if $W(\tau)/TR(\tau) \geq \kappa$. The key insight is how the expression depends on whether the ratio of the welfare gain to the tax revenue exceeds a critical threshold, defined by the level of political influence to tax aversion. The condition makes clear the pivotal role that the ratio plays when evaluating the political feasibility of an externality correcting tax in the presence of political opposition to raising tax revenue. With full information, one could calculate the magnitude of the ratio compared to κ. But more realistically, leaders are only likely to have a rough sense of κ (and perhaps the ratio), in which case the threshold inequality still illustrates a key qualitative result: a policy with a greater ratio of the welfare gain to tax revenue is more

[5] There is evidence suggesting that the way tax revenue is spent (i.e., earmarking) can affect ideological opposition and support. For studies examining this question with respect to carbon taxes, see Kotchen et al. (2017), Klenert et al. (2018), Carattini et al. (2018), and Ewald et al. (2022). It is assumed here that this feature is taken into account in the determination of κ.

[6] See Grossman and Helpman (1994, 1995) for seminal contributions, and Fredriksson (1997) and Aidt (1998) for applications to environmental policy.

likely to be perceived as a net winner from a political economy perspective when there exists aversion to taxes.[7]

Efficiency

The ratio also has importance when the focus is solely on efficiency. Pigou (1947) recognized the potential role of administrative and compliance costs, and others have more recently considered their potential implications (e.g., Polinsky and Shavell 1982; Kampas and Horan 2016). In such cases, assume the administrative costs are some function of the tax revenue. Then a policy is more likely to pass a benefit-cost test by simply having a higher welfare gain, which is equivalent to having a higher ratio of the welfare gain to tax revenue because the latter is pinned down by the specific policy. Indeed, in the special case where administrative costs increase in proportion to the tax revenue, the preceding discussion applies exactly with an efficiency interpretation, where κ need only be reinterpreted as the administrative or compliance cost per unit of tax revenue. The threshold condition is therefore a benefit-cost condition, based on a comparison of whether the welfare gains per unit of tax revenue are greater than (less than) the per unit costs of administration, compliance, or both. The presence of corruption in a way that imposes an efficiency cost on tax revenue (see, for example, Besley and Persson 2014; Conway and Hermann 2021) is another potential interpretation that can give rise to the same conditions. The important observation here is identification of additional circumstances—ones based on standard economic efficiency—where the ratio of welfare gains to tax revenue is an important object of interest.

[7] One can also arrive at the same condition using a median voter approach, where each voting agent has their own preference κ_i for opposition to tax policy. An individual will support the policy if and only if $W(\tau)/TR(\tau) \geq \kappa_i$, and whether a policy passes will then depend on the distribution of κ_i in the population.

Another efficiency related argument for the importance of ratios is rent seeking. Tullock (1967) argues that tax revenues should not be treated a welfare neutral because agents will spend resources trying to obtain a share of the revenue for themselves. This is an example of rent seeking behavior (Kruger 1974; Dal Bó 2006; Congleton et al. 2008), and several studies have considered it in the context of emissions taxation and environmental policy instruments more generally (Dijkstra 1998; MacKenzie and Ohndorf 2012; MacKenzie 2017). Drawing on the preceding framework, we can once again capture the phenomenon with a reinterpretation of κ. If κ represents the fraction of the (transferred) tax revenue benefit that is dissipated through rent seeking, then $\kappa TR(\tau)$ represents a true cost, and the threshold condition is again an efficiency-based, benefit-cost comparison.

Finally, the ratio of the welfare gain to tax revenue can be important in the context of setting priorities among different externality correcting taxes with the goal of maximizing efficiency. 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. Which tax policies should be selected? It turns out that the key input for answering this question is the ratio of the welfare gain to tax revenue, as choosing those with higher ratios produces greater efficiency gains per unit of revenue. The logic here is similar to that of using the marginal value of public funds (MVPF) to set priorities among government expenditure policies.[8]

[8] A series of recent papers (Hendren 2016; Hendren and Sprung-Keyser 2020; Finkelstein and Hendren 2020) advance the notion of the 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 polices 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." Lumping all policies that pay for themselves into the same category can, however, over overlook important information.

Summary

We have illustrated circumstances where the ratio of welfare gain to tax revenue can inform how one might evaluate the political feasibility and desirability of an externality correcting tax. In the presence of politically motivated tax resistance, a bigger ratio may help political feasibility. When there is an efficiency cost to raising tax revenue (e.g., administrative, compliance, corruption), a greater ratio is more likely to ensure the policy is efficient. Moreover, ratios can help when setting priorities among potential policies in order to promote overall efficiency. There are, of course, other criteria that matter when it comes to choosing environmental policy instruments (Fullerton 2001; Goulder and Parry 2008; Stavins 2022), but the aim here has been to show how the ratio of welfare gain to tax revenue can be among them when considering externality correcting taxes. In doing so, I have assumed a constant tax rate with no exemptions, and the tax revenue is redistributed in a welfare neutral way. Relaxing either of these assumptions would change the ratios and may well serve as areas where policymakers can fine tune policy proposals to produce more favorable ratios and therefore potential support.

Selected Examples

The selected examples summarized in this section are intended to illustrate ways in which policy evaluation can be aided with additional information about the ratio of the welfare gain to tax revenue. An 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.

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.[9] 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 that yield a present value of tax revenue equal to $12.15 billion.[10] 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 certainty play a role in

---

[9] 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.

[10] 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.

debate about whether the tax policy is desirable. 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, quantity-based regulation 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.

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 dioxide emissions consistent with proposed legislation at the time. The analysis considers three different tax rates (10, 15, and 20 cents per pound of emissions) and a range of estimated control costs across nine 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 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 a low of 0.25 to a high of 9.96.[11] 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 are 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.

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.

[11] 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.

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 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.

Ratios Large or Small?

We have considered circumstances, from both a political economy and efficiency perspective, where a larger ratio of the welfare gain to tax revenue might be a desirable feature of an externality correcting tax. This raises a further question: Under what conditions would we expect an externality correcting tax to be associated with a large or small ratio? To provide some basic insight on this question, I focus on the canonical case taught in most undergraduate classes: linear demand and supply and constant marginal external costs. I take a partial equilibrium perspective, where the market has no distortion other than the externality of interest, and the analysis assumes no income effects so that consumer surplus is the correct welfare measure. I first show the results when the tax is set optimally and then discuss how the results apply even if the tax is not set optimally.

A Change in Marginal External Costs

Figure 2 expands upon the initial setup in Figure 1 to consider a change in the marginal external costs. A larger marginal external cost shifts the social marginal costs up to the left from SMC to SMC'. Assuming the policy changes to the newly calibrated Pigouvian tax at τ' and quantity Q', the new ratio is area IBH over area HJKL. It is then a matter of relatively straightforward geometry to see that the ratio must now be greater with the greater external costs, as the numerator must increase by proportionally more than the increase in tax rate, and the denominator must increase by proportionally less. We have thus established, in the linear case, that when setting optimal Pigouvian taxes, the ratio of the welfare gain to tax revenue is increasing in the size of the marginal external costs.

A Change in the Market Responsiveness

Figure 3 illustrates what happens when there is greater market responsiveness to the tax, which can be captured with greater elasticity of supply, demand, or both. The initial setup is identical to Figure 1, but the alternative scenario is now one with greater elasticity of demand. Greater market responsiveness is clear because Q' is less than Q^1, and it is also clear that the ratio must be greater. Because the tax rate remains the same, but the quantity is lower, tax revenue must decline from area CDEF to GHIJ. At the same time, the welfare gain is greater from triangle ABC to ABG. Although not shown here, similar graphs can show the same qualitative results for the cases of greater elasticity of supply and both supply and demand. We have therefore established a second result: when setting optimal Pigouvian taxes, the ratio of the welfare gain to tax revenue is increasing in the market responsiveness to the tax.

Generalization to Non-optimal Taxes

The preceding analysis assumed that taxes were set at the optimal Pigouvian level. In the case of a change in the marginal external costs, this means that the tax rate changes. In many real-world settings, however, taxes are not likely to be set optimally. So we might wonder how the results change for any given tax that need not be set optimally. We might also wonder whether the results hold even in cases where demand and supply are not linear. It turns out that both of the preceding results hold. For the linear model, it is easy to see in Figure 2 how increasing the marginal external costs, without adjusting the tax rate, increases the welfare gain with no change to the tax revenue. Moreover, in Figure 3, the tax rate does not change, so the same analysis goes through, again with the linear model. The overall conclusion is that the ratio of the welfare gain to tax revenue is increasing in the marginal external costs and the market responsiveness, regardless of whether the tax is set optimally or not.[12]

Carbon Taxes on U.S. Fossil Fuel Consumption

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 across markets.

[12] The result for market responsiveness and a non-optimal tax does require that the arbitrarily set tax not be too far above the marginal external costs, but this is generally not a case of central interest. Proofs of the additional results without linear supply and demand are available upon request.

Modeling Approach

The model and assumptions employed here are an extension to Kotchen (2021). Supply and demand are of the constant elasticity functional form. Price and quantity data are needed for establishing initial conditions, and these were obtained from the U.S. Energy Information Administration (EIA) for the year 2018 and are reported in Table 1.[13] The quantities of all four fuels 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 of coal is the annual average cost of coal delivered to electric power plants; and the prices for gasoline and diesel are the average annual retail prices.

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 (CO_2) 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.[14]

To inform the point estimates chosen for the supply and demand elasticities, I draw on the literature review contained in Kotchen (2021) and use the same estimates here. The aim in

[13] 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.

[14] Following convention in the literature and the most recent approach by the U.S. government (IWG-SCGHGs 2021), the global SCC is used to calculate welfare gains. See Kotchen (2018) for further discussion.

each case is to select an estimate that represents a long-run elasticity applicable to the U.S. domestic market that reflects the different substitution opportunities across sectors.[15] The estimates used here are reported in Table 3. 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 elasticities. These ranges are reported in brackets in Table 3.

Results

***Welfare gains* —** Figure 3 reports the welfare gain 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.[16] The welfare gains from the carbon tax on coal are far greater than those for the other fuels, owing to its large external costs. The gains for coal are increasing rapidly with a higher carbon tax up to about $50 per ton, where they start to level off.[17] 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 is 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). This means that imposing the carbon tax over and above the preexisting tax results in a welfare loss. The same holds for low levels of the tax even with a SCC of $75. It is only when the SCC is greater than

[15] 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.

[16] While the results are also sensitive to changes in the elasticities, this variation is less than that for the SCC. See online Appendix Figure A1 for results across the range of elasticity scenarios assuming a SCC of $50 per ton.

[17] 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.

$100 that the welfare gain is always positive. 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 similar pattern, though generally remain positive. The reason is that diesel has a significantly larger externality due to local pollution.[18]

***Tax revenue* —** Figure 4 reports the annual tax revenue 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. 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. 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 when carbon taxes exceed $30 per ton.

***Ratios by fuel* —** Figure 5 reports the measure of central interest—the ratio of the welfare gain to the tax revenue—for all four fuels and at different tax rates. To limit the proliferation of scenarios to consider, only those for the central elasticities are considered, along with the range of SCC estimates. 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

[18] 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.

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.

Relation to theory — The basic conceptual framework discussed previously provides a systematic way of understanding the results within and across fuels, and these differences provide a useful illustration of how ratios might differ in other setting where externality correcting taxes might be applied. Within fuels, Figure 5 illustrates how the ratios are increasing in the marginal external costs (i.e., the SCC), and although not illustrated in a figure, the ratios (when positive) are also increasing in the market responsiveness (i.e., from low- to high-elasticity scenarios). Across fuels it is instructive to compare the size of the externalities. Per unit of energy (i.e., normalized to British thermal units), coal has the highest uninternalized externality, followed by natural gas, diesel, and a negative value for gasoline. These differences help to explain why the ratios for each fuel follow in the same descending order.

Moreover, the elasticity estimates in Table 3 result in a market responsiveness across fuels that similarly helps to explain the results. For example, assuming a carbon tax increase from $45 to $50 and the central elasticity estimates, the market quantity elasticity is -1.15 for coal, -0.23 for natural gas, and approximately -0.065 for gasoline and diesel. The responsiveness is the greatest for coal, by a wide margin, and this further explains why coal has substantially higher ratios. Natural gas is also relatively elastic compared to gasoline and diesel, and this contributes to its relatively higher ratios.

Aggregate ratios — While the fuel-specific analysis illustrates a range of results for externality motivated taxes, we now turn to the aggregated results for a carbon tax applied uniformly across all four fuels. The results for the welfare gain and tax revenue are simply sums of those shown previously.[19] The left-hand side of Figure 6 shows the aggregate ratios at different levels of the SCC and the central-elasticity scenario.[20] 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.

Also shown in Figure 6 on the right-hand ride are results of an alternate scenario where the carbon tax is imposed on top of the preexisting tax only to the extent that it would exceed the preexisting tax. In other words, the carbon tax is imposed only when it is efficient to do so when taking account of climate damages. These results are directly comparable to those on the left-hand side for the baseline approach, and the ratios are substantially higher. 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

[19] 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 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 discussed shortly when considering external validity of the estimates.

[20] Graphical results of the aggregated welfare gains and tax revenue corresponding with the ratios in Figure 6 are shown in online Appendix Figures A2 and A3.

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.[21]

External Validity

The primary aim of the preceding carbon tax analysis is to illustrate different possibilities for the ratio of welfare gains to tax revenue, both across fuels and in aggregate. It is nevertheless worth considering how the magnitude of the estimates compare to those of more detailed energy and economy models. Barron et al. (2018) provide a set of policy insights across 11 models included in the Stanford Energy Modeling Forum exercise 32 (EMF 32). Their meta-analysis reports estimates of the CO_2 emission reductions and tax revenues for different carbon tax policies in the United States, 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 7. 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.[22] Also reported are the minimum and maximum values across models. The right-hand side shows comparable results of the model used here, 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.[23] The magnitudes of the low- and high-elasticity scenarios also correspond well

[21] Readers may also be interested in how the ratios differ considering only local pollution damages or only climate damages. While the different contributions to the marginal external costs can be seen in Table 3, online Appendix Figure A4 shows the aggregate ratios, in parallel with Figure 6, for the local and global pollutants separately.

[22] 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.

[23] 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

with the min and max results of the EMF models. The bottom panel of Figure 5 compares results for annual tax revenue.[24] Here again the magnitudes averaged across the 11 EMF models and those produced here are quite similar. 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.

Discussion

There are two central aims of this paper: The first is to argue that economists should pay more attention to tax revenue when analyzing and advocating for externality correcting taxes, and the ratio of welfare gains to tax revenue is a useful summary statistic for this purpose. The second is to illustrate the range of potential results with an application to carbon taxes on fossil fuels in the United States.

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

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.

[24] 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 deflater. 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.

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? A 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 set of reasons rely only on the efficiency criterion when, for example, there exists costs associated with administration, compliance, or corruption, along with the possibility for rent seeking. 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. The ratio is also useful for establishing efficiency-based priorities among candidate externality correcting taxes when revenue raising potential is limited. It therefore plays an analogous role to the MVPF for setting policy priorities with limits on expenditures.

A third reason is because of increasing concern about the distributional consequences of policy interventions. 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, especially when redistribution schemes are by design not welfare neutral. When the ratios are lower, for example, 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 when focusing efficiency.

Application to Carbon Taxes

The simple conceptual framework developed in this paper shows how the ratio of the welfare gain to tax revenue will depend on 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 and intermediate case with ratios indicating a more proportional balance between welfare gains and tax revenue.

When considering all fuels simultaneously, and accounting for existing taxes, the ratios range between 1.8 and 1.4 for the central-elasticity estimates and SCCs that range between $50 and $125 (see the right panel of Figure 6). These ratios imply the 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 in the presence of high administrative and compliance costs, and possibly rent seeking. While some may view these ratios a favorable, others may not because of ideological aversion to taxes and raising government revenue. The aim here is to show how the ratio of the welfare gain to tax revenue can help inform such debates, rather than provide a definitive assessment on the prospects of a carbon tax in the United States.

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

Aidt, T. 1998. Political internalization of economic externalities and environmental policy. *Journal of Public Economics* 69: 1-16.

Barthold, T. A. 1994. Issues in the design of environmental excise taxes. *Journal of Economic Perspectives* 8: 133–151.

Besley, T. and T. Persson. 2014. Why do developing countries tax so little? *Journal of Economic Perspectives* 28: 99-120.

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: 1840003 1-37.

Carattini, S., M. Carvalho, and S. Fankhauser. 2018. Overcoming public resistance to carbon taxes. *WIREs Climate Change*. 9:e531.

Chait, J. 2011. The triumph of taxophobia. *Democracy*, No. 20.

Congleton, R., A. Hillman and K. Konrad. 2008. Forty years of research on rent seeking: an overview. In *40 Years of Research on Rent Seeking 1: Theory of Rent Seeking*, Congleton, R., A. Hillman and K. Konrad (eds), Springer, Berlin/Heidelberg.

Conway, D. and B. Hermann. 2021. Dealing with corruption in adopting and implementing carbon taxes. U4 Brief 2021:5, Michelsen Institute (CMI), Norway.

Dal Bó, E. 2006. Regulatory capture: a review. *Oxford Review of Economic Policy* 22: 203-225.

Davis, L. 2017. The environmental cost of global fuel subsidies. *The Energy Journal* 38: 7-21.

Dijkstra, B. 1998. A two-stage rent-seeking contest for instrument choice and revenue division, applied to environmental policy. *European Journal of Political Economy* 14: 281-301.

Eisenstein, L. 2010. *The Ideologies of Taxation*, Cambridge, MA: Harvard University Press.

Ewald, J., T Sterner, and E. Sterner. 2022. Understanding the resistance to carbon taxes: drivers and barriers among the general public and fuel-tax protesters. *Resource and Energy Economics* 70: 101331.

Finkelstein, A. and N. Hendren. 2020. Welfare analysis meets causal inference. *Journal of Economic Perspectives* 34(4): 146-167.

Fredriksson, P. 1997. The political economy of pollution taxes in a small open economy. *Journal of Environmental Economics and Management* 33: 44-58.

Fullerton, D. 2001. A framework to compare environmental policies. *Southern Economic Journal* 68: 224-248.

Goulder, L. H. 1995. Environmental taxation and the double dividend: a reader's guide. *International Tax and Public Finance* 2: 157-183.

Goulder L. and I Parry. 2008. Instrument choice in environmental policy. *Review of Environmental Economics and Policy* 2: 152-174.

Griffin, J. 1974. An econometric evaluation of sulfur taxes," *Journal of Political Economy*, 82: 669-688.

Grossman, G. and E. Helpman. 1994. Protection for sale. *American Economic Review* 84: 833–850.

Grossman, G and E Helpman. 1995. The politics of free-trade agreements. *American Economic Review* 85: 667–690.

Hendren, N. 2016. The policy elasticity. *Tax Policy and the Economy* 30: 51-90.

Hendren, N. and B. Sprung-Keyser. 2020. A unified welfare analysis of government policies. *Quarterly Journal of Economics* 135: 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.

Kampas, A. and R. Horan. 2016. Second-best pollution taxes: revisited and revised. *Environmental Economics and Policy Studies* 18: 577-597.

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. 2018. Which social cost of carbon? A theoretical perspective. *Journal of the Association of Environmental and Resource Economists* 5: 673-694.

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.

Krueger, A. 1974. The political economy of the rent-seeking society. *American Economic Review* 64: 291-303.

Leape, J. 2006. The London congestion charge. *Journal of Economic Perspectives* 20: 157-176.

MacKenzie, I. 2017. Rent creation and rent seeking in environmental policy. *Public Choice* 171: 145-166.

MacKenzie, I. and M. Ohndorf. 2012. Cap-and-trade, taxes, and distributional conflict. *Journal of Environmental Economics and Management* 63: 51-65.

Anderson, S., I Marinescu, and B. Shor. 2023. Can Pigou at the polls stop us melting the poles? *Journal of the Association of Environmental and Resource Economists* 10: 903-945.

McFarland, J., A. Fawcett, A. Morris, J. Reilly, and P. Wilcoxen. 2018. Overview of the EMF 32 study on U.S. carbon tax scenarios. *Climate Change Economics* 9(1): 1840002 1-37.

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: 723-743.

Parry, I. and A. Bento. 2000. Tax deductions, environmental policy, and the ‘double dividend’ hypothesis. *Journal of Environmental Economics and Management* 39: 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.

Stavins, R. 2022. The relative merits of carbon pricing instruments: taxes versus trading. *Review of Environmental Economics and Policy* 16: 62-82.

Tullock, G. 1967. The welfare costs of tariffs, monopolies, and theft. *Western Economic Journal* 5: 224-232.

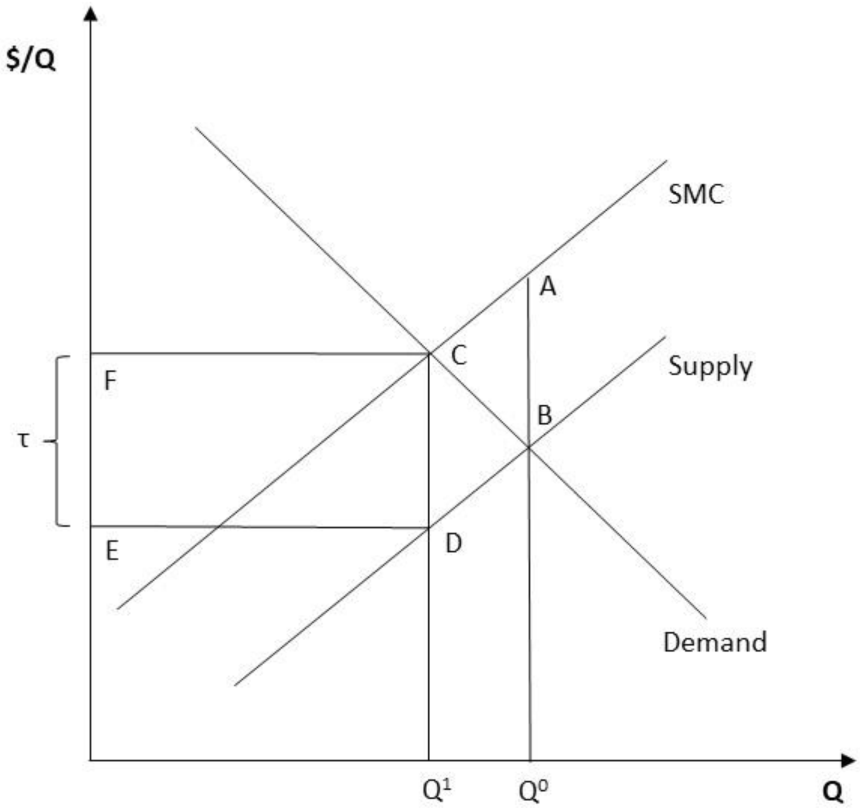

Figure 1: The textbook, linear setup of a Pigouvian tax, showing the welfare gain (ABC) and the tax revenue (CDEF).

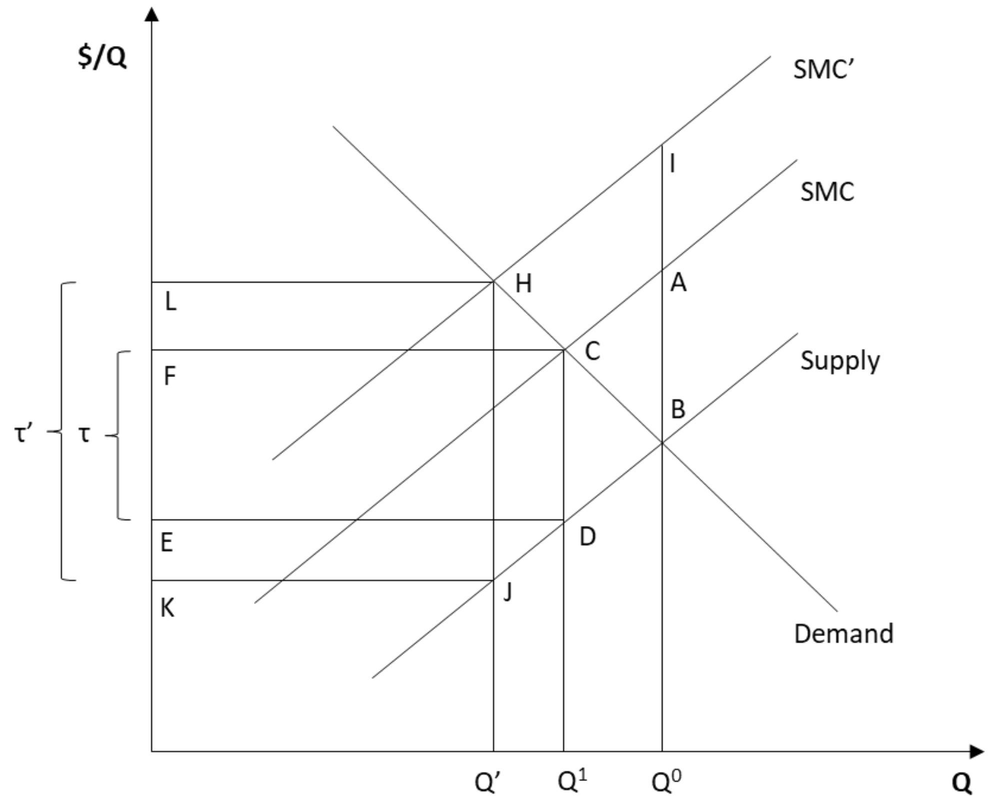

Figure 2: Illustration in a linear model of how an increase in the marginal external costs increases the ratio of the welfare gain to tax revenue for an optimal Pigouvian tax.

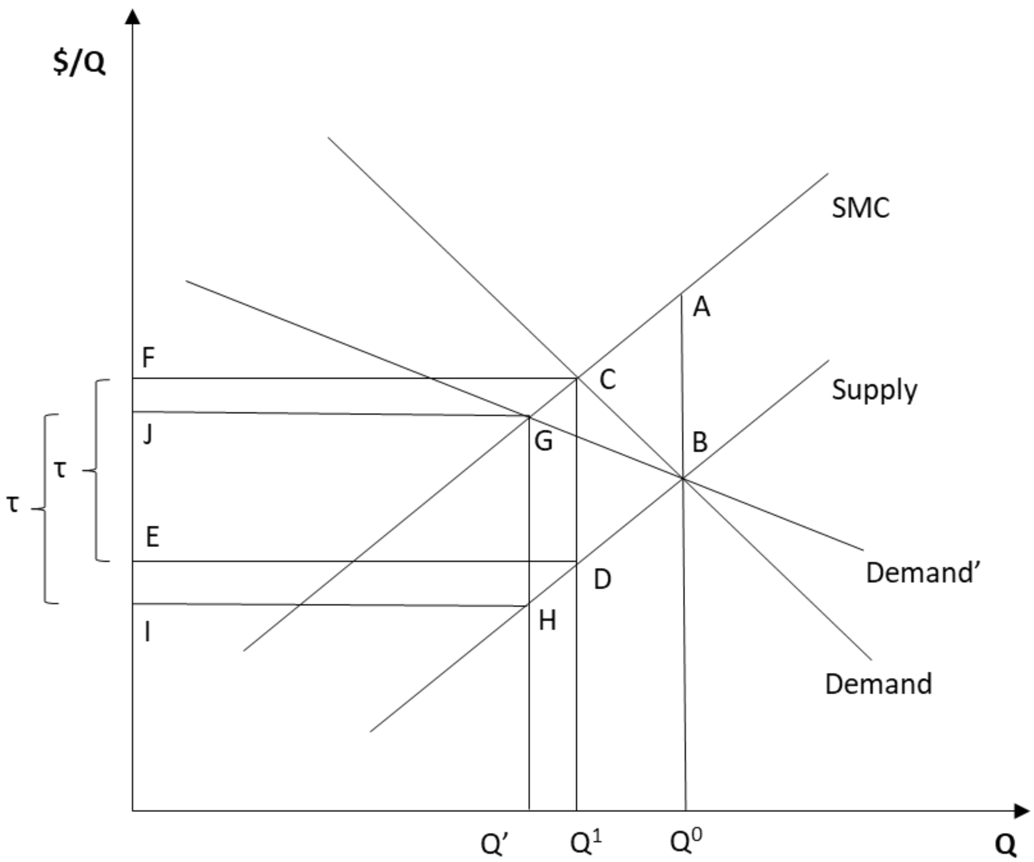

Figure 3: Illustration in a linear model of how an increase in the market responsiveness to the implementation of an optimal Pigouvian tax increases the ratio of the welfare gain to tax revenue. The case illustrated is for flatter demand, and similar results can be shown for flatter supply or flatter demand and supply simultaneously.

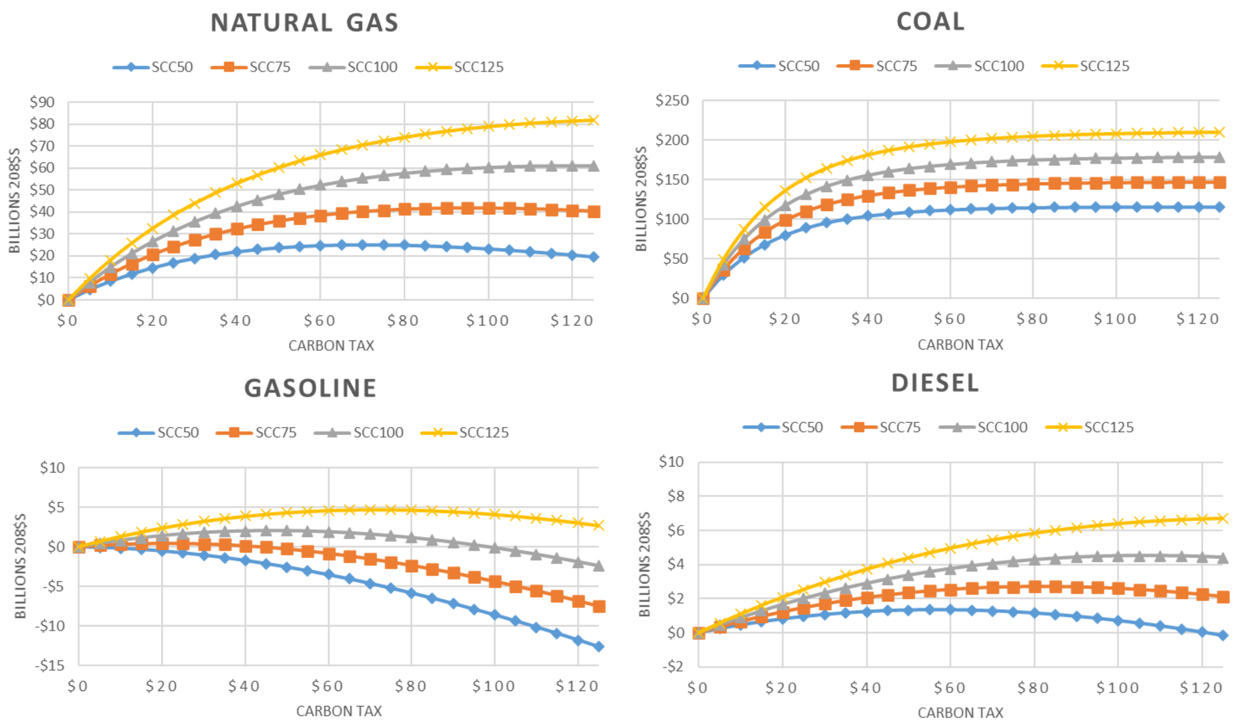

Figure 3: Annual change in welfare for all four fuels at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The scenarios of SCC50, SCC75, SCC100, and SCC125 correspond to the SCC at $50, $75, $100, and $125.

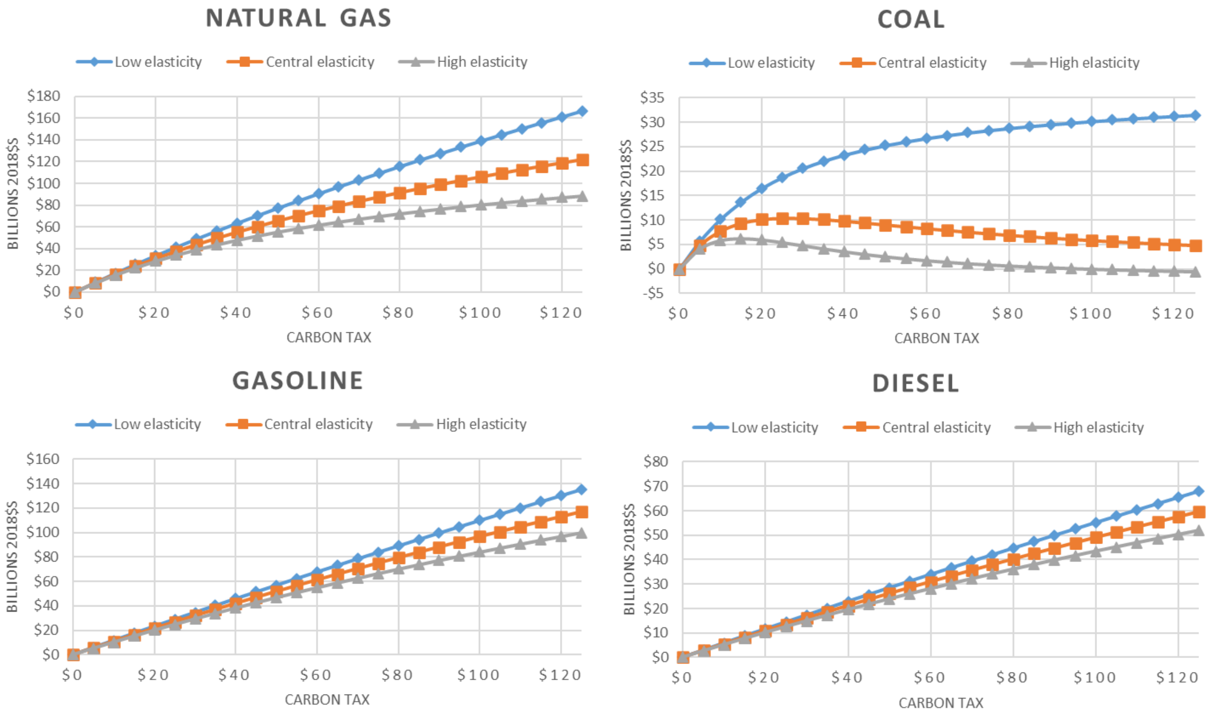

Figure 4: Annual tax revenue for all four fuels at different carbon tax rates and assumptions about the market elasticities. The scenarios of central, low, and high elasticity correspond to the central estimates, a 50-percent decrease in both the supply and demand elasticities, and a 50-percent increase in both elasticities. See Table 3 for the specific estimates employed.

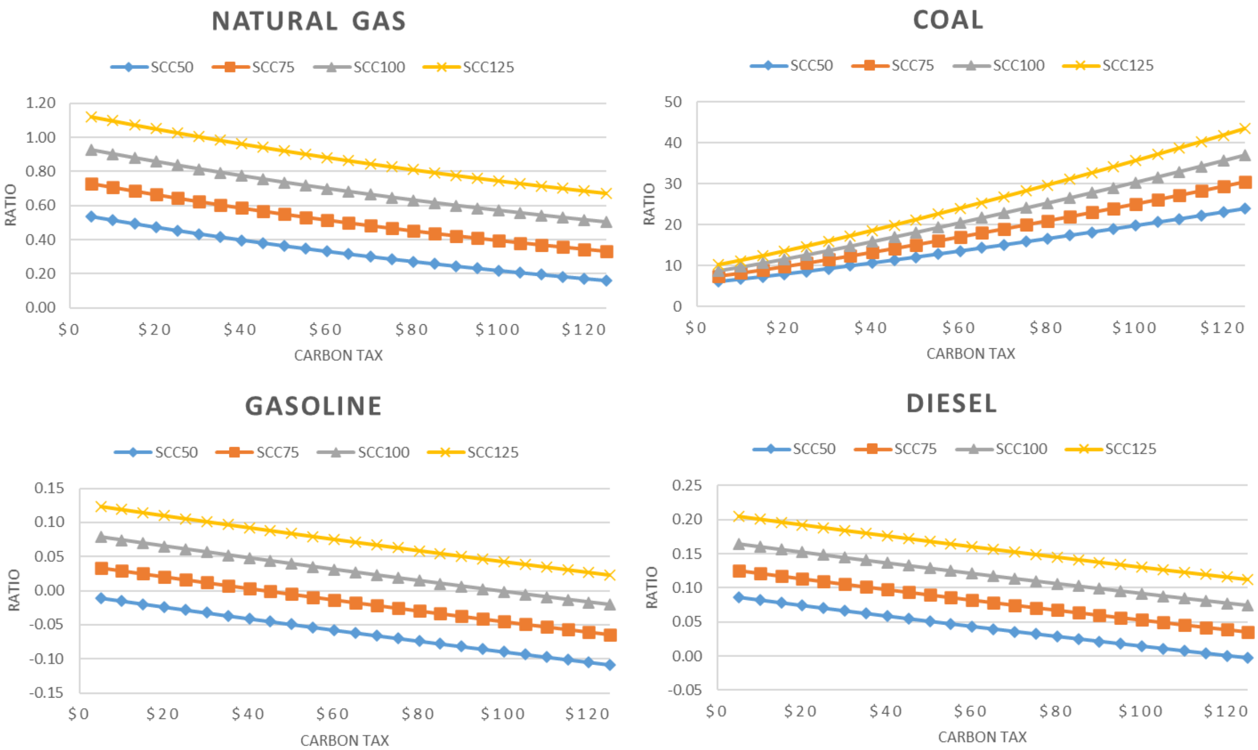

Figure 5: The ratio of the welfare gain to tax revenue for all four fuels at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates.

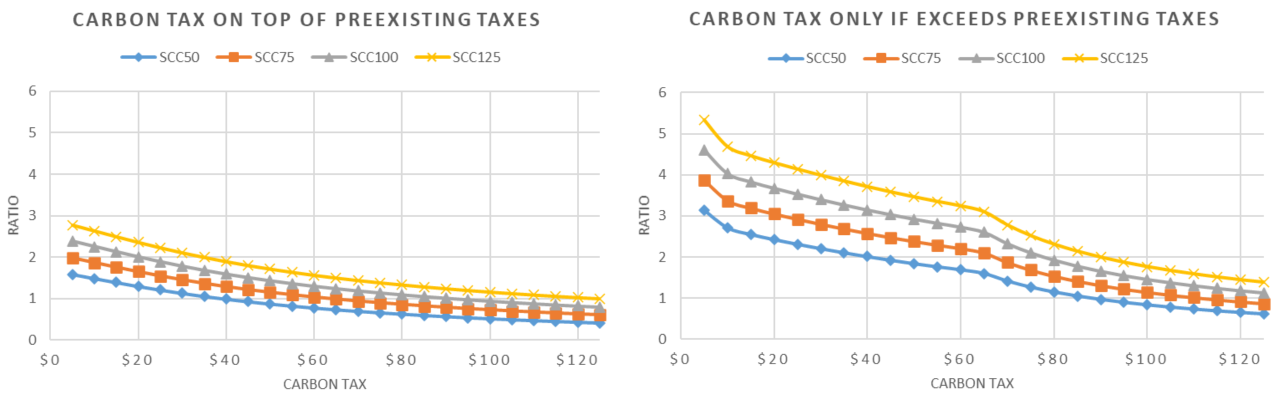

Figure 6: The ratio of the welfare gain to tax revenue for all fuels combined at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The left-hand side is for the baseline case where the carbon tax is imposed over and above preexisting taxes. The right-hand side is for the alternative case where the carbon tax only takes effect if it exceeds the preexisting tax.

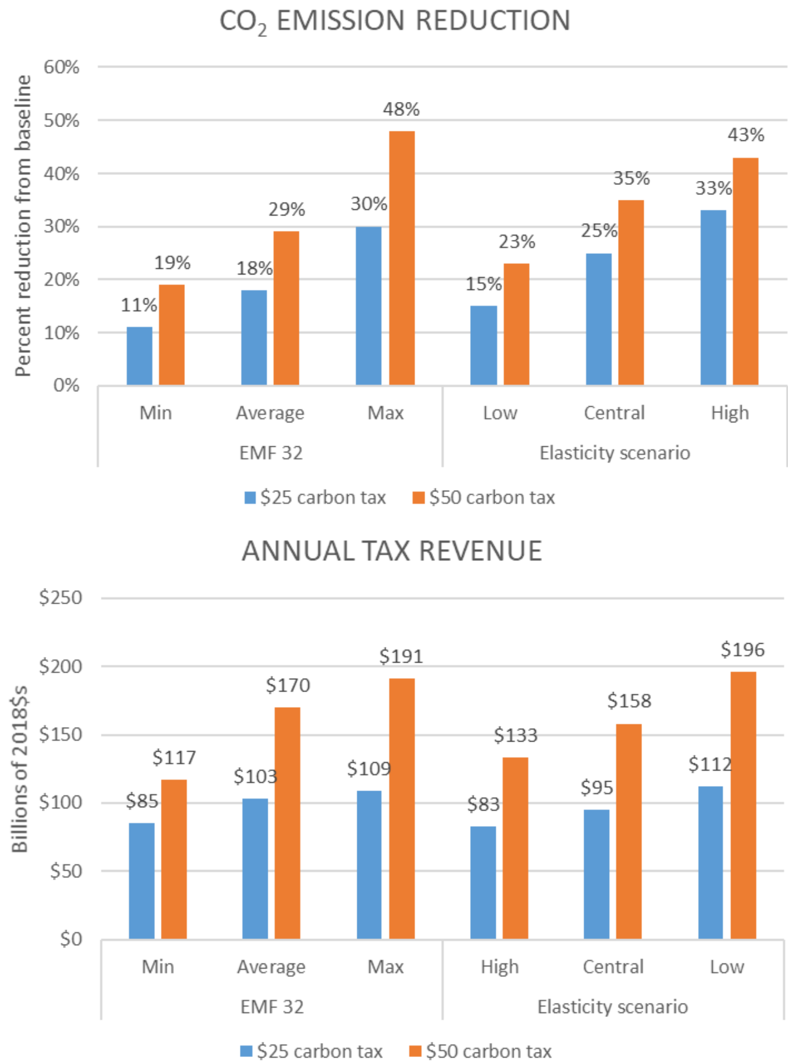

Figure 7: Comparison of results generated here with those across models from the EMF 32 study on U.S. carbon tax scenarios. The top panel compares the percentage reduction in CO_2 emissions from a no tax baseline, and the bottom panel compares estimates of the annual tax revenue. Details on how comparisons where made within the $25 and $50 tax scenarios are reported in the main text.

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	1.6
	[-0.28, -0.83]	[0.80, 2.40]
Coal	-1.75	1.9
	[-0.88, -2.63]	[0.95, 2.85]
Gasoline	-0.63	2.0
	[-0.32, -0.85]	[1.00, 3.00]
Diesel	-0.58	2.0
	[-0.29, -0.87]	[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.

Supplementary Appendix for
Online Publication for

Taxing Externalities: Revenue *vs.* Welfare Gains with an Application to U.S. Carbon Taxes

Matthew J. Kotchen

Yale University and NBER

September 2023

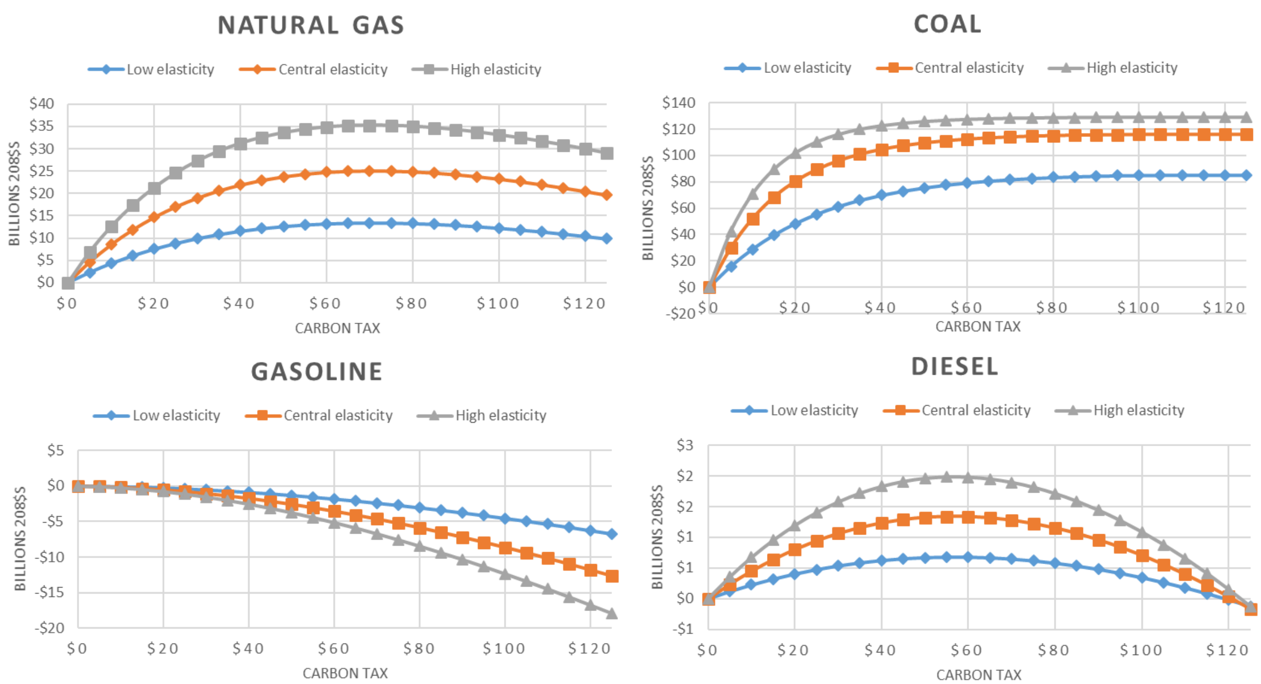

Figure A1: Annual change in welfare for all four fuels at different carbon tax rates, a social cost of carbon equal to $50, and assumptions about the market elasticity. The scenarios of central, low, and high elasticity correspond to the central estimates, a 50-percent decrease in both the supply and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

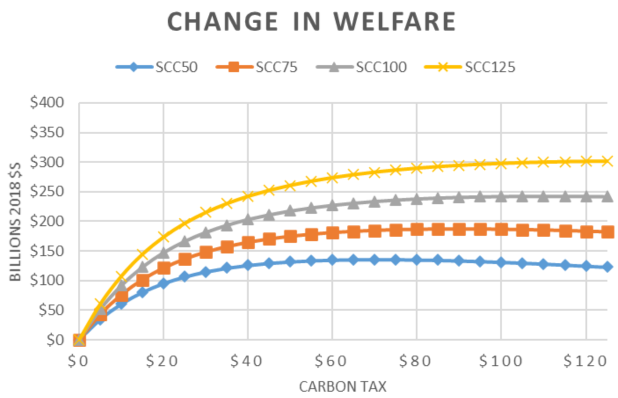

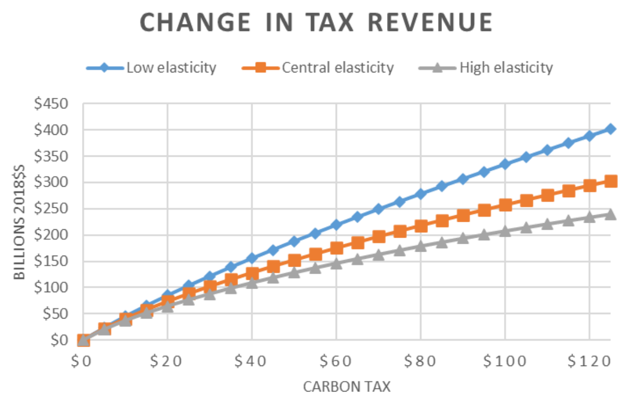

Figure A2: The left panel shows the change in welfare aggregated across fuels at different carbon tax rates, assumptions about the social cost of carbon, and the central-elasticity estimates. The scenarios of SCC50, SCC75, SCC100, and SCC125 correspond to the SCC at $50, $75, $100, and $125. The right panel shows tax revenue 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.

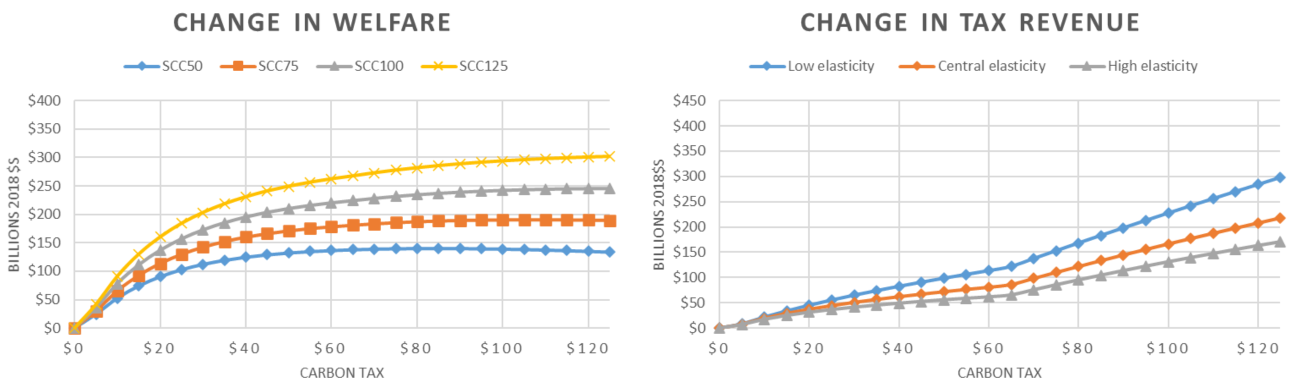

Figure A3: The graphs correspond with the alternative approach whereby carbon taxes are only applied for each fuel if the rate exceeds the preexisting tax. As with Figure A2, the left panel shows the change in welfare aggregated across fuels at different carbon tax rates, assumptions about the social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

Panel A: Local pollution only

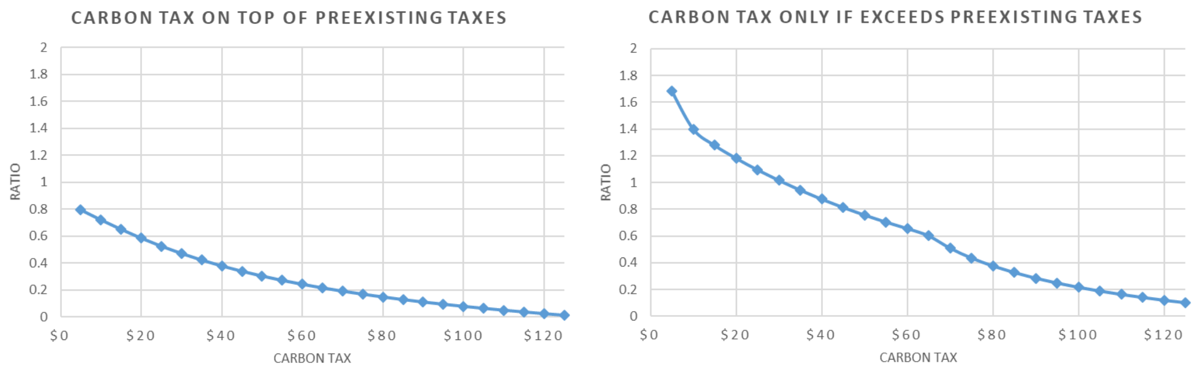

Panel B: Climate damages only

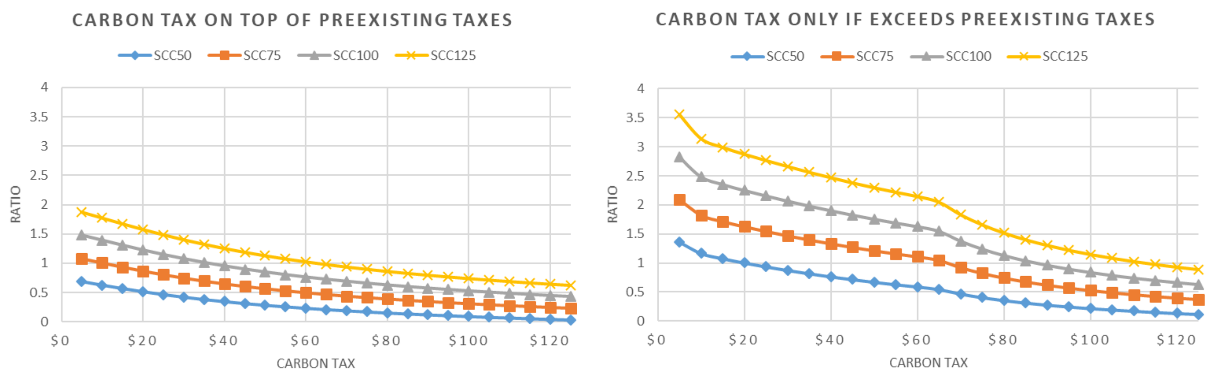

Figure A4: The ratio of the welfare gain to tax revenue for all four fuels combined at different carbon tax rates. Panel A includes local pollution damages only. In parallel with Figure 6, the left-hand side is for the baseline case where the carbon tax is imposed over and above preexisting taxes. The right-hand side is for the alternative case where the carbon tax only takes effect if it exceeds the preexisting tax. Panel B includes only climate damages, assuming different values for the social cost of carbon.