NBER WORKING PAPER SERIES

IS BROADER ALWAYS BETTER? PREEXISTING DISTORTIONS, EMISSIONS ELASTICITIES, AND THE SCOPE OF EMISSIONS PRICING

Lawrence H. Goulder
Marc A.C. Hafstead
Roberton C. Williams III

Working Paper 32915
http://www.nber.org/papers/w32915

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2024

We thank Dallas Burtraw and Don Fullerton for early conversations on this topic. We also thank participants at the 2023 National Tax Association Spring Symposium and 2024 Allied Social Science Associations Meetings for helpful comments and suggestions. Hafstead thanks RFF's Carbon Pricing Initiative for financial support. A portion of Williams's work on this project was supported by the USDA National Institute of Food and Agriculture (Hatch project 7006024) Roberton C. Williams III has served during calendar years 2019-2021 as Chief Economist for the Climate Leadership Council, which advocates for a US federal carbon fee and dividend. The views expressed herein are those of the authors 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Lawrence H. Goulder, Marc A.C. Hafstead, and Roberton C. Williams III. 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.

Is Broader Always Better? Preexisting Distortions, Emissions Elasticities, and the Scope of Emissions Pricing
Lawrence H. Goulder, Marc A.C. Hafstead, and Roberton C. Williams III
NBER Working Paper No. 32915
September 2024
JEL No. D58, D62, H21, H23, Q58

ABSTRACT

Economists often regard broad-based carbon pricing (whether in the form of a carbon tax or cap and trade) as the most efficient policy to reduce carbon dioxide emissions. Relative to a narrower policy that exempts some emissions sources, a broader policy is often favored because it can exploit more low-cost emissions reduction opportunities and cause less emissions leakage to uncovered sources. Yet narrower approaches have gained considerable political support, partly because they avoid price increases for outputs (such as gasoline) regarded as especially critical to household budgets. Some analysts might lament any departure from broad carbon pricing, citing efficiency costs. This paper offers theory and numerical simulations revealing that such a shift need not sacrifice efficiency. This result reflects differences across sectors in distortions from preexisting taxes and in the elasticity of emissions with respect to the carbon price. Our analytical model reveals that a narrower policy that exploits these differences can be more cost-effective than a policy with a broad, economy-wide tax base. Our numerical model of the US economy compares quantitatively the effects of an economy-wide carbon price with those of several narrower policies, including one that applies only to the power sector, one that exempts gasoline, and one that exempts energy-intensive trade-exposed industries. We compare policies under alternative specifications for policy stringency and find that the broader policy always becomes more cost-effective at sufficiently high stringency.

Lawrence H. Goulder
Department of Economics
Landau Economics Building 328
Stanford University
Stanford, CA 94305
and NBER
goulder@stanford.edu

Marc A.C. Hafstead
Resources for the Future
1616 P St NW
Washington, DC 20036
Hafstead@rff.org

Roberton C. Williams III
Department of Agricultural
and Resource Economics
University of Maryland
Symons Hall
College Park, MD 20742
and Resources for the Future
and also NBER
roberton@umd.edu

A data appendix is available at http://www.nber.org/data-appendix/w32915

1 Introduction

Decisionmakers frequently debate alternative policies to curb emissions of carbon dioxide (CO_2) and other greenhouse gases at both state and federal levels. Among the policy options for reducing CO_2 emissions, economists tend to support a broad-based price on carbon that spans all major sources of CO_2 emissions. Such a policy receives support from basic economic theory: a narrow policy that excludes significant emissions sources could pass up some low-cost emissions reduction opportunities. It also can bring about shifts in production or consumption toward products from uncovered sources of emissions; such shifts are viewed as economic distortions that come with costs to the economy.[1]

At the same time, narrower carbon price approaches continue to gain considerable attention. Their appeal partly reflects political considerations, including concerns about potential political resistance to any policy that raises prices of certain goods. For example, Democratic Party negotiations on a climate-related reconciliation bill in the fall of 2021 considered a carbon tax that would exempt gasoline.[2] This pattern also has emerged from subnational US policies and policies outside the United States. For example, the Regional Greenhouse Gas Initiative (RGGI) in the Northeast covers only the power sector, and the European Union's carbon-trading system excludes major sectors, such as transportation and agriculture, and covers only about half of EU carbon emissions.

[1] Metcalf and Weisbach (2009) articulate this view: "Absent administrative, enforcement, and political costs, an ideal tax system would include all activities that produce climate externalities" (p. 521). Aissa Assia (2011), in an OECD report, provides a full explanation, noting, "Homogenous taxes encourage abatement at the lowest-cost source, helping to ensure that environmental goals are achieved at the lowest social cost. A tax applied on a uniform basis also minimises the costs of compliance for taxpayers and the costs of administration for government, and reduces the opportunities for tax evasion;" the passage concludes, "Governments should therefore try to implement environmental taxes as broadly as possible, with few or no exemptions" (p. 5).

[2] See, for example, "Read Sen. Wyden's lips: No new gas taxes" Washington Post, October 5, 2021. (`https://www.washingtonpost.com/politics/2021/10/05/read-sen-wyden-lips-no-new-gas-taxes`).

Some analysts might lament any shift in focus away from a broad-based carbon-pricing policy toward a narrower one, citing the potential efficiency costs. However, this paper offers theory and numerical simulations that reveal that such a shift need not involve an efficiency sacrifice. It identifies plausible circumstances under which a narrower carbon tax costs less than a broad one that achieves the same overall reduction in CO_2 emissions.[3]

This potentially surprising result arises because of preexisting tax distortions (labor and capital taxes in particular) that cause the marginal cost of a carbon tax to differ significantly across sectors. As discussed in prior literature, preexisting factor taxes interact with a carbon tax (or other price on carbon), changing the cost of achieving any given level of emissions abatement.[4] In the absence of interactions with tax-related or other distortions, a broader policy will always have (weakly) lower costs than a narrower policy: exempting some sectors passes up potential low direct-cost reductions in uncovered sectors and thus requires additional high direct-cost reductions in its covered sectors to achieve a given emissions target.[5] Interactions with pre-existing taxes affect the cost in each sector, and we find that the distortions from these interactions can vary significantly across sectors (varying with capital/labor intensities, effective marginal tax rates, and other differences). If those differentials are sufficiently large, excluding sectors with relatively high tax-related distortions can lower the cost of policy.

[3]For the remainder of this paper, we refer to carbon tax policies, but the conclusions also apply to cap-and-trade programs and other emissions-trading programs in which the market prices of limited supplies of emissions allowances serve as the carbon price.

[4]See, for example, Bovenberg and de Mooji (1994), Parry (1995), Fullerton and Metcalf (2001), West and Williams (2007), Goulder (2013), and Barrage (2020). This literature distinguishes two effects —the tax-interaction effect (which typically increases the cost of environmental policy) and the revenue-recycling effect (which typically decreases cost) —that arise in a model with preexisting tax distortions. We use the term "interactions" to embrace both effects, though we will separate them later in the paper. To our knowledge, the only paper in this literature that compared broad and narrower environmental taxes is Parry and Williams (1999)), which compared a wide range of policy alternatives.

[5]In this paper, we focus on preexisting tax distortions, but as discussed later, the point generalizes to other preexisting distortions. The key is that interactions with preexisting distortions —tax-related or of other forms —are necessary for a narrower carbon tax to be more cost-effective than a broad tax.

A second critical factor is the responsiveness of emissions to the carbon tax in the excluded sectors, which influences both the direct cost disadvantage of excluding sectors and the magnitude of interactions with preexisting distortions in those sectors. Let the "direct cost" of the carbon tax refer to its cost excluding the previously mentioned welfare effects of interactions with the tax system. A narrower tax will always have a higher direct cost than a broader tax because (as just noted) it passes up low direct-cost reductions in the exempted sectors. Importantly, the magnitude of that disadvantage—the opportunity cost of passing up the low direct-cost options—depends on the emissions elasticity in the excluded sectors. If a given sector's emissions are relatively unresponsive to the tax, excluding or "carving out" that sector will not pass up much in terms of emissions cuts, and the additional direct costs from needed increases in the tax on the other sectors will be small. But the greater the emissions elasticity of the carved-out sector, the greater the increase in direct cost. In addition to this direct effect, the sectoral responsiveness to the carbon price also influences the magnitude of the interactions with preexisting distortions: all else equal, the less responsive emissions are, the larger the magnitude of those interactions for the marginal cost of emissions reductions in that sector.

Finally, the relative importance of tax interactions depends on the stringency of the carbon tax policy. As indicated below, the direct cost disadvantage of a narrower tax is roughly proportional to the square of the reduction in emissions, while the effect of tax interactions increases more slowly. Thus, although a narrower tax can have a cost advantage when emissions reductions are small, the broader tax will gain that advantage as policy becomes sufficiently stringent.

Our paper's analytical model brings out those factors and conveys the links among them. Our numerical model then considers the quantitative implications of the analytical model's main findings. The general equilibrium simulations reveal large tax-interaction effects that vary across sectors within the US economy, as well as significant differences in the price responsiveness of emissions across sectors. We apply the model to compare the outcomes of a broad (economy-wide) carbon tax with those under a range of narrower policies that have received recent attention. These comparisons identify several real-world circumstances in which the narrower carbon tax is more cost-effective. In particular, a power-sector-only policy and policies that exempt motor vehicle fuels consumed by households and/or energy-intensive trade-exposed (EITE) sectors are more cost-effective than an economy-wide policy for modest reductions in emissions. In addition, in several cases involving more ambitious emissions reductions, the narrower policy is only slightly less cost-effective than the broader one. In such cases, attaching slight weight to other considerations (e.g., political feasibility or distributional equity) could tilt the balance toward the narrow alternative.

Those numerical results also reinforce the analytical model's finding that the cost-effectiveness advantage of breadth depends on policy stringency: the ratio of the broader tax's cost to the narrower one's cost declines with the ambitiousness of the CO2 reduction target and the magnitude of the tax rates needed to achieve the target. Correspondingly, even in cases where the narrower tax is more cost-effective for small to moderate reductions in emissions, the broader tax always becomes more cost-effective for a sufficiently large emissions reduction. This suggests that if policymakers are committed to increasing policy stringency over time, it could be most cost-effective to start with a narrower tax and subsequently broaden the tax as stringency increases.[6]

For the cases where a narrower tax is more cost-effective, our results should not be interpreted as suggesting that it is the economic ideal. When the narrower option is more cost-effective, its advantage rests on a preexisting inefficiency of the tax system on nonenvironmental grounds. Tax reforms that reduce these preexisting inefficiencies would alter the differences in cost between the broad carbon tax and

[6] The EU emissions trading system has followed this pattern (albeit perhaps not for cost-effectiveness reasons): revisions to the system over time have both broadened coverage and increased stringency.

various narrower ones. Indeed, they generally would tilt the balance toward the broad-tax option. However, if such tax reforms are not politically feasible, the narrower carbon tax could emerge as the most attractive, feasible environmental policy.

Although this paper focuses on preexisting tax distortions, the major points are more general and apply to other preexisting distortions.[7] Considering additional distortions could change specific numerical results. For example, it could change the magnitude of cost differences and potentially change which specific narrow policies are less costly than broad policies. But the core qualitative points would remain unchanged. First, to the extent that the interactions between carbon prices and preexisting distortions differ across sectors, a narrower carbon price focused on sectors where those interactions are most beneficial or least harmful will gain in relative cost-effectiveness. Second, emissions elasticities matter in determining the influence of distortions: low elasticities will tend to magnify the relative importance of preexisting distortions. Third, policy stringency will remain important: the significance of interactions with preexisting distortions (relative to the direct cost of policy) generally falls as the carbon price rises. And finally, directly correcting the preexisting distortions would shift the balance toward the broader option.

The rest of the paper is organized as follows. The next section provides an analytical general equilibrium model that identifies the conditions that determine the relative cost-effectiveness of a broader versus a narrower tax. Section 3 presents the structure of our numerical general equilibrium model. Section 4 describes the data and parameters for the numerical model and conveys the nature of the numerical model's reference (business-as-usual) case. Section 5 describes the policy alternatives considered, and Section 6 compares and interprets the simulation results. Section 7 presents a sensitivity analysis. The final section offers conclusions.

[7] Other preexisting distortions that could be particularly relevant for climate policy include imperfect competition and innovation spillovers.

2 An Analytical Model

We develop an analytical model to illustrate the main effects influencing the relative cost-effectiveness of broader versus narrower emissions taxes, with the goal of developing a simple model that can capture the main effects.[8] The model shows that in the absence of other preexisting distortions, a broader tax is more cost-effective than a narrower one, with the magnitude of the difference depending on characteristics of the sectors that are taxed or untaxed. Relevant characteristics include the elasticity of emissions in taxed vs. untaxed sectors with respect to the tax rate, and the extent of leakage to untaxed sectors (the effect of the tax on untaxed emissions through changes in fuel prices, demand levels, etc.).

The model then shows that preexisting tax distortions introduce additional effects—the tax-interaction and revenue-recycling effects—and demonstrates that the combined influence of these effects can either increase or reduce costs of the narrower policy relative to the broader one. The model also reveals an interaction between these effects and the emissions elasticity: all else equal, the lower the elasticity, the larger the magnitude of the gain or loss from the combined tax-interaction and revenue-recycling effects in a given sector. When these combined effects favor the narrower policy, that narrower policy will be more cost-effective for a sufficiently small emissions reduction.

2.1 The Model

A representative agent's utility depends on consumption of n private goods, given by the vector $\boldsymbol{X}$, the quantity of a public good G, and its supply of labor L. The

[8]For brevity, this section refers to carbon taxes, but the effects discussed here also apply to other carbon-pricing policies (e.g., a cap-and-trade program).

utility function is given by

$$U(\boldsymbol{X}, G, L), \tag{1}$$

where U is continuous, quasi-concave, and twice-differentiable. This function is increasing in consumption of each of the private goods and the public good and decreasing in the amount of labor supplied. We assume that the quantity of the public good is fixed.

The agent's budget constraint is given by

$$\sum_i p_i X_i = (1 - \tau_L)L + T, \tag{2}$$

where p_i is the price of the private good, τ_L is the tax rate on labor income, and T is government transfers to the household. Without loss of generality, the (pretax) wage is normalized to one (i.e., labor is the numéraire).

The agent maximizes utility (1) subject to the budget constraint (2), taking prices, the labor tax rate, the government transfer, and the quantity of the public good as given. This yields the first-order conditions:

$$U_{X_i} = p_i \lambda \ ; \quad -U_L = (1 - \tau_L)\lambda, \tag{3}$$

where λ is the marginal utility of income.

For simplicity, we assume that labor is the only factor of production and production exhibits constant returns to scale. Thus, production of each good is proportional to labor used to produce that good. Units are normalized such that one unit of labor produces one unit of any of the produced goods:

$$X_i = L_{X_i} \ ; \quad G = L_G, \tag{4}$$

where L_{X_i} and L_G are the quantities of labor used to produce X_i and G, respectively. Total labor used in production must equal labor supply (i.e., the labor market clears):

$$L = L_g + \sum_i L_i. \tag{5}$$

Pollution emissions associated with each good are proportional to the quantity of that good.[9] That is,

$$Z_i = z_i X_i, \tag{6}$$

where Z_i is pollution emissions from production and/or consumption of good i, and z_i is the (exogenous) pollution intensity of that good.[10]

The government imposes an emissions tax at the rate τ_{Z_J} on emissions from a subset J of the N industries; the tax rate is zero for industries not in J. Total emissions covered by the tax are then

$$Z^C = \sum_{i \in J} z_i X_i. \tag{7}$$

Production is competitive, so the output price equals the marginal cost of production (equal to 1, given the earlier normalization assumptions) plus the emissions taxes paid per unit of output:

$$p_i = 1 + \tau_{Z_J} z_i \; \forall \; i \in J \quad \text{and} \quad p_i = 1 \; \forall \; i \notin J. \tag{8}$$

The government uses tax revenue to finance the transfer to the household and

[9]The model does not capture any harmful effects of pollution emissions. If carbon is the only pollutant affected, this omission has no effect on comparisons between broader and narrower taxes because in all cases we compare policies that yield the same reductions in carbon emissions (i.e., we focus on cost-effectiveness). However, to the extent that they differentially affect other pollutants —such as local copollutants —the policies could differ in their overall environmental consequences.

[10]We assume that production of the public good is nonpolluting. To facilitate welfare comparisons, we hold the level of the public good fixed throughout the analysis. Hence the assumption of a nonpolluting public good has no effect on the results.

provision of the public good. The government's budget constraint is therefore

$$\tau_L + \tau_{Z_J} Z^C = G + T. \tag{9}$$

Taken together, equations (1) through (9) implicitly define utility and all prices and quantities as functions of the set of goods subject to the pollution tax, the tax rates on pollution and labor, and the government transfer.

2.2 Effects of the Emissions Tax

Consider the effect of a marginal increase in the emissions tax rate τ_{Z_J}, with revenue returned via either a reduction in the labor tax rate τ_L or an increase in the government transfer T. For now, we assume that the emissions tax does not cause emissions leakage; that is, we assume no effect on emissions associated with industries not covered by the tax. We will relax this assumption later. The total derivatives here (i.e., the $d/d\tau_{Z_J}$ terms) represent the combined effect of the change in the emissions tax and the effects of the associated change in the labor tax or government transfer to make the overall policy revenue-neutral. The marginal cost per ton of emissions reductions can be expressed as (see Appendix A for derivation)

$$\begin{aligned} MC \equiv \frac{1}{\lambda}\frac{dU}{d\tau_{Z_J}} / \frac{dZ^C}{d\tau_{Z_J}} = \\ \underbrace{\tau_{Z_J}}_{\text{direct}} \underbrace{- (\eta_R - 1)\left(\frac{Z^C + \tau_Z dZ^C/d\tau_{Z_J}}{-dZ^C/d\tau_{Z_J}}\right)}_{\text{revenue-recycling}} + \underbrace{\eta_R \mu_{IJ} \frac{Z^C}{-dZ^C/d\tau_{Z_J}}}_{\text{tax-interaction}}. \end{aligned} \tag{10}$$

The first term on the right-hand side of (10) is the direct cost of the policy: the firm's marginal cost of emissions abatement. At the margin, this is equal to the emissions price. In a model without any distortions other than the emissions externality, this would be the entire marginal cost.

The second term is the *revenue-recycling* (RR) effect: the welfare effect (measured per unit of emissions reduction) attributable to recycling the revenue through either labor tax cuts or an increase in the government transfer. This effect is a function of η_R, the marginal cost of public funds (MCPF) for whatever instrument is used to recycle the revenues. If revenues are recycled via a reduction in the labor tax rate, then η_R is the MCPF of the labor tax; if revenues are recycled via an increase in the government transfer, then η_R is the MCPF of the transfer.[11] The RR term is then equal to the distortionary cost per dollar of marginal revenue (i.e., the MCPF minus one) times the marginal revenue from the environmental policy per unit of emissions reduced (the term in large brackets).

The third term on the right-hand side of (10) is the *tax-interaction* (TI) effect: the welfare effect (per unit of emissions reduction) resulting from the effect of the emissions tax on markets with preexisting tax distortions (in this case, the labor market). This term depends on μ_{IJ}, which captures the extent to which a tax on emissions from the industries in set J interacts with the preexisting tax distortion:

$$\mu_{IJ} \equiv \frac{\eta_{IJ} - 1}{\eta_{IJ}}, \tag{11}$$

where η_{IJ} is the MCPF relevant for tax interactions for the emissions tax, given by

$$\eta_{IJ} = \frac{Z^C}{Z^C + \tau_L \frac{\partial L}{\partial \tau_{Z_J}}}. \tag{12}$$

The tax interaction effect, then, is the product of η_R (which appears because the interactions affect revenue, thus influencing how much can be recycled), μ_{IJ} (marginal tax interactions per dollar of burden from the emissions tax), Z^C (the marginal burden of the emissions tax per dollar of the tax rate) and $1/\left(dZ^C/d\tau_{Z_J}\right)$

[11]More generally, η_R should represent the MCPF for whatever recycling method is used. In this analytical model, we consider only two methods: cutting the labor tax and increasing the transfer. The numerical simulations below consider additional forms of recycling.

(the marginal tax rate increase per unit of emissions reductions).

At times we shall refer below to the *combined* tax-interaction effect, representing the overall impact of the TI and RR effects above.

2.2.1 Narrow vs. Broad without Preexisting Tax Distortions

First, consider the differences in outcomes of narrow vs. broad emissions taxes in the absence of preexisting tax distortions. In this case, the second and third terms in (10) are zero and the marginal cost per unit of emissions reduction is simply the first term, the tax rate τ_{Z_J}.

This first term will always favor a broad policy. We can consider emissions tax rates over a range starting from zero. To achieve any level of emissions reduction greater than zero, the tax rate must be weakly higher under the narrower tax (and must be strictly higher under the narrower tax unless emissions are completely insensitive to the tax in the industries covered by the broader tax but not by the narrower tax). The magnitude of the difference will depend on emissions elasticities: the difference in required tax rates between broad and narrow policies will be relatively small if the narrow policy excludes relatively inelastic sectors, but relatively large if it excludes relatively elastic sectors.

2.2.2 Narrow vs. Broad with Preexisting Tax Distortions

Now consider how the outcomes change in the presence of preexisting tax distortions. In this case, the second and third terms in (10) no longer are zero. To focus sharply on the difference from the previous case, it is useful to rearrange (10) to yield

$$MC = \underbrace{\tau_{Z_J}}_{\text{direct}} + \underbrace{(\eta_R - 1)\,\tau_{Z_J} + \eta_R\,(\mu_{IJ} - \mu_R)\,\frac{1}{\varepsilon_J}}_{\text{combined TI and RR}}, \tag{13}$$

where

$$\mu_R \equiv \frac{\eta_R - 1}{\eta_R}, \tag{14}$$

where ϵ_J is the semi-elasticity of emissions with respect to the emissions tax applied to the set of industries J, given by

$$\epsilon_J = -\frac{1}{Z^C}\frac{dZ^C}{d\tau_{Z_J}}. \tag{15}$$

The first term in (13) is the direct effect, which is the same as in equation (10). Here it exhibits the same pattern as what emerged in the absence of preexisting distortions: it is lower for a broader tax than for a narrower tax, with the magnitude of the difference depending on the sensitivity of emissions to the tax in the sectors not covered by the narrower tax.

The second and third terms combine the TI and RR effects. The second term is part of the RR effect: reductions in emissions lower the revenue from the emissions tax, thus reducing revenue available for recycling, which produces this term. It is equal to τ_{Z_J} (the marginal cost in the absence of preexisting distortions) times the MCPF for the RR effect minus one.

The third term is the rest of the RR effect, along with the TI effect. This term is the efficiency gain or loss from raising revenue via the pollution tax and then recycling it (either to cut labor taxes or to provide a lump-sum transfer). This depends on the relative magnitudes of μ_{IJ} and μ_R (which in turn depend on η_{IJ} and η_R), or in other words, on the relative efficiency of the pollution tax (as a revenue raiser) versus the efficiency gain from recycling the revenue. If $\eta_{IJ} > \eta_R$ (and therefore if $\mu_{IJ} > \mu_R$), this term will be positive and will increase the marginal cost of emissions reductions. This can occur if the industries covered by the pollution tax have larger-than-average tax interactions (i.e., η_{IJ} is relatively high) or if the revenue is recycled in a way that generates a relatively small efficiency gain (i.e., η_R

is relatively low).[12] In the opposite case, where $\eta_{IJ} < \eta_R$, this term will be negative and will decrease the marginal cost of emissions reductions. Note that the magnitude of this term depends on the semi-elasticity: the less elastic emissions are, the larger this tax-shift effect will be for a given reduction in emissions.

Now focus on the intercept: the marginal cost for the initial marginal increment of emissions reduction stemming from an increment from zero in the tax rate. The first and second terms in (13) will start from zero (i.e., they are zero when the tax rate is zero). The third term, however, will not start from zero in general, and it could favor either the broader or the narrower tax, depending on how the level of tax interactions (measured by μ_{IJ}) and the emissions semi-elasticity (ϵ_J) differ. Put differently, the marginal direct cost (and the corresponding part of the marginal RR effect) starts from zero, but marginal revenue from the pollution tax and marginal effects on the prices of polluting goods do not, which means that the RR and TI effects (and therefore the third term) are nonzero initially.

Because the first two terms are intially zero, the initial marginal cost —the intercept of the marginal cost curve —will equal the third term. This term is proportional to the difference between μ_{IJ} and μ_R: if μ_{IJ} is larger (i.e., the TI effect is larger than the RR effect) then the intercept is positive, and if μ_R is larger (i.e., the RR effect is larger than the TI effect), then the intercept is negative.

The intercept is inversely proportional to the semi-elasticity of emissions with respect to the tax. This arises because we're looking at the cost per ton of emissions reduction, and neither the TI nor the RR effect (each evaluated for a marginal

[12] In this simplified analytical model, η_{LI} depends on the weighted-average cross-elasticity between the goods covered by the pollution tax and leisure. In a more general model with multiple factors of production —such as the numerical model later in this paper —η_{LI} will depend on the extent to which a pollution tax on the covered sectors interacts with taxes on different factors of production (being higher to the extent that the burden of the pollution tax falls on factors of production with relatively large tax distortions). Similarly, in the analytical model, η_R is relatively large when revenue is recycled to cut the tax on labor and relatively small (indeed, negative) when used to increase the lump-sum transfer. In a more general model, η_R depends on how distortionary the tax is that is cut using the revenue from the pollution tax.

increase in the tax from zero) depends on the emissions response to the tax, which implies that the more responsive emissions are to the tax, the smaller the magnitude (absolute value, whether positive or negative) of this term.

Equation 13 helps identify the determinants of the relative cost-effectiveness of the broad and narrower policy alternatives. As was noted, one major factor is policy stringency—the required emissions reduction (or extent of emissions abatement). The first two terms in (13) are zero for an infinitesimal reduction in emissions. Thus, for policies involving sufficiently low stringency, the cost difference between the broader and narrower taxes will be determined by the third term. But as stringency increases (and thus the tax rate rises), the relative importance of the first two terms increases: these terms will increase roughly in proportion to the tax rate, whereas the third term will remain roughly constant as the tax rate increases.[13] For a sufficiently large reduction in emissions, the first two terms will dominate, and the broader policy will be more cost-effective.[14]

The implications of the third term for the relative cost-effectiveness of the narrower and broader policies depend on two additional factors. The first is a difference in the extent of tax interactions associated with the policy options. If the sectors included in the broader tax but excluded from the narrower tax are sectors with relatively high tax interactions (i.e., including them increases μ_{IJ}), then that will favor the narrower tax.

The second is a difference in the semi-elasticity of emissions with respect to the emissions price. If $(\mu_R - \mu_{IJ})$ is positive (i.e., if the gain from the RR effect outweighs the cost of the TI effect), then the third term will lower the cost of the

[13]If η_R, η_{IJ}, and ϵ_J are all constant as the tax rate increases, then the first two terms are proportional to the tax rate and the third term remains constant. In practice, they are likely to change somewhat as the tax rate changes but will change much more slowly than the tax rate.

[14]Equation (13) shows the marginal cost. One can also show that the total cost of the policy must be higher under a narrower tax when the policy is sufficiently stringent. Consider the cost of a level of emissions reduction slightly greater than the total emissions of the sectors covered by the narrower tax. This will be infeasible under the narrower tax but feasible under the broader tax (unless the additional sectors included in the broader tax are completely insensitive to the tax).

tax (and the intercept will be negative). A lower semi-elasticity will magnify that effect. In that case, if the sectors included in the broader tax but excluded from the narrower tax are high-elasticity sectors (implying that including them increases ϵ_J), then the third term will favor the narrower tax. Put differently, in this case, excluding high-elasticity sectors will make the initial cost of the tax more negative, favoring the narrower tax. If $(\mu_R - \mu_{IJ})$ is negative, however, the effect is reversed: the third term raises the cost of the tax and the intercept is positive. Again, a lower semi-elasticity will magnify that effect. So if the sectors included in the broader tax but excluded from the narrower tax are low-elasticity sectors (and thus including them decreases ϵ_J), then the third term will favor the narrower tax. Put differently, in this case, excluding low-elasticity sectors will make the initial cost of the tax less positive, favoring the narrower tax.

Thus, this third term could make a narrower tax more cost-effective for a sufficiently small reduction in emissions. This depends on both the level of tax interactions and the semi-elasticity of emissions for the sectors included in the broader tax but excluded from the narrower tax.

2.2.3 Allowing for Leakage

Leakage occurs to the extent that emissions in the uncovered sectors are influenced by the covered sectors' emissions tax. Up to this point, we have assumed no leakage. The presence of leakage would not change the marginal cost of a given increase in the pollution tax, but it can affect the overall emissions reduction caused by that tax increase. Allowing for leakage, the expression for the marginal cost per unit of emissions reductions becomes (see Appendix A for details)

$$\begin{aligned} MC &\equiv \frac{1}{\lambda}\frac{dU}{d\tau_{Z_J}} \Big/ \frac{dZ^C}{d\tau_{Z_J}} \\ &= \left(\frac{1}{1-\theta}\right)\left[\tau_{Z_J} + (\eta_R - 1)\tau_{Z_J} + \eta_R(\mu_{IJ} - \mu_R)\frac{1}{\epsilon_J}\right], \end{aligned} \tag{16}$$

where θ indicates emissions leakage to uncovered sectors as a proportion of covered-sector reductions. θ is given by

$$\theta = 1 - \frac{dZ/d\tau_{Z_J}}{dZ^C/d\tau_{Z_J}}. \tag{17}$$

If the tax increases emissions in uncovered sectors, $\theta > 0$.

Note that allowing for leakage simply multiplies the marginal cost of emissions reductions in equation (13) by $1/(1-\theta)$. For a given tax rate, positive (negative) leakage increases (reduces) the cost of emissions reductions. The greater the extent of leakage, the higher the tax rate must be to achieve a given reduction in emissions. That latter effect doesn't appear as a term in (16), but its influence stems from the fact that with higher leakage, a higher τ_{Z_J} will be needed to achieve any given level of emissions reduction.

Leakage provides another channel that can influence the relative cost-effectiveness of broader versus narrower emissions prices. One might expect that broadening the tax would reduce leakage, since uncovered sectors would then represent a smaller share of the economy. To the extent that broadening the tax reduces leakage, it will tend to improve the relative cost-effectiveness of the broader tax.

3 The Numerical Model

We obtain quantitative results by employing a computable general equilibrium (CGE) model similar to the Goulder-Hafstead E3 model documented in Goulder and Hafstead (2017).[15] The model comprises 22 domestic production sectors, a single representative household, a representative government, and a simple

[15]The model used here differs from that in Goulder and Hafstead (2017) in containing new consumption and production nests, a simplified treatment of international trade, and an alternative aggregation of producer and consumer goods. Detailed documentation of the current model is given in Appendix B.

treatment of foreign trade. It captures the interactions of supplies and demands for goods and factors of production and solves for market-clearing prices. The representative consumer and managers of firms have perfect foresight. The model is solved annually, beginning in 2019.

Several features of the model make it well suited to compare the cost-effectiveness of carbon-pricing policies that differ in their sectoral coverage. First, it includes a detailed treatment of the US tax system, enabling it to measure the size of the tax-interaction effect across sectors. This also allows it to measure how welfare costs change under different forms of revenue recycling. Second, it recognizes the adjustment costs associated with the introduction or removal of physical capital; this affects the rate at which capital stocks will adjust through time and introduces effects on capital incomes that may differ by the sector (or sectors) covered. Third, it allows for a range of emissions reduction options by producers, including fuel switching across fossil fuels and from fossil fuels to electricity and other inputs. It also incorporates emissions reductions stemming from consumers' shifts in demand from relatively high carbon intensity products to products with lower carbon intensity. The ability of producers or consumers to mitigate their emissions through changes in patterns of demand will depend on both elasticities of substitution and the initial relative fuel demands for each agent.[16] Finally, the model can address a wide range of tax policies in terms of sectors covered and the ways that the revenues are recycled to the private sector.

3.1 Households

A single representative household chooses between work and leisure, savings and consumption, and consumption expenditure across various consumer goods and

[16] A producer that obtains a relatively high share of its energy from electricity will have lower opportunities for mitigation than a producer that uses both gas and coal as fuel inputs and has the opportunity to switch between the two fuels (through its elasticity of substitution).

services to maximize its utility subject to a budget constraint. The treatment of household behavior closely follows Goulder and Hafstead (2017). One difference here is the use of a nested constant-elasticity-of-substitution (CES) utility structure, in contrast with the Cobb-Douglas specification in Goulder and Hafstead (2017). At each level of the nest, households choose consumption intensities to achieve the least-cost combination of goods. The nested CES utility function is described in B-1 in Appendix B.

3.2 Production

The model's production sectors have been chosen to give focus to industries that supply fossil fuels, use these fuels intensively, or supply alternatives to fossil fuels. [17] A representative firm in each industry chooses variable inputs (energy, nonenergy, and labor) and investment (subject to capital adjustment costs) to maximize the value of the firm, which is the discounted flow of after-tax profits net of new share issues. Output from each sector stems from a nested structure of CES production functions.[18]

The model introduces technological change through labor-augmenting Harrod-neutral technological change. We assume a constant rate of change that is the same across all industries. Producers face output taxes, payroll taxes, property

[17] Sectors include fossil fuel producers (crude oil, natural gas, coal), an electricity sector with a transmission and distribution sector and multiple generators (coal, gas or petroleum, nuclear, wind, solar, hydro or other), secondary energy producers (natural gas distribution and petroleum refining), nonmanufacturing industrial sectors (agriculture and forestry, other mining, mining services, water utilities and construction), manufacturing (EITE and non-EITE manufacturing), transportation, services, and owner-occupied housing.

[18] The structure includes a subenergy nest that adds greater (potential) flexibility in the substitution between electricity and other energy inputs, potential "fixed" intermediate inputs, and natural resources to the nesting structure in Goulder and Hafstead (2017). Fixed intermediate inputs are the inputs required per unit of capital. For example, the input of crude oil into petroleum refiners is a fixed input, to prevent refiners in the model from reducing how much oil is needed to produce petroleum products. Natural resources are used for hydro or other generation and nuclear generation to substantially reduce the elasticity of supply for these generators in response to changes in relative fuel prices. The nest is described in Figure B-2 in Appendix B.

taxes on existing capital, and corporate income taxes on operating profits (net of property tax payments, depreciated capital deductions, and interest payments on bonds).[19]

To investigate how exempting certain fuel uses changes policy costs, we implement a downstream tax—that is, a tax on the agent directly responsible for the combustion of the carbon-based fuel.[20] Although the tax is imposed downstream, the resulting emissions depend on changes in production (and fossil fuel combustion) all along the supply chain, reflecting general equilibrium connections across industries. Emissions are based on emissions factors applied to the use and combustion of fossil fuels throughout the economy.

3.3 Government

A representative government incorporates public expenditure and taxation at the federal, state, and local levels. The government levies taxes and issues debt to finance spending on goods and services, labor, and transfers. Government expenditures are fixed in real terms, and input intensities of goods, services, and labor are held fixed in both the reference case and the policy experiments. Transfers are fixed in real terms by indexing initial transfers to the consumer price index. The government also makes interest payments to households on outstanding public nominal debt. The real level of debt is assumed to grow exogenously over time, and nominal debt is equal to the real debt level times the price level.

Revenues from the various taxes on households and producers are collected; a

[19] Following Goulder and Hafstead (2017), firms finance their activities through both debt and equity. Debt is held fixed relative to the capital stock, and payments to households through dividends are a constant fraction of after-tax profits. A cash-flow identity links the sources (profits and new borrowing) and uses of revenues (dividends, investment expenditures, and share repurchases).

[20] The model allows for introducing carbon taxes upstream (on producers of fossil fuels) or midstream (on the first purchaser of fossil fuels as inputs into production), but the downstream approach taken here yields the most effective contrast of policies in terms of sectoral coverage.

budget constraint equates endogenous government revenues to exogenous government expenditure (net of new debt issued). Lump-sum taxes on households adjust to satisfy this constraint. In policy simulations, revenues from the carbon tax are an additional source of revenue. These revenues can be rebated to households through lump-sum dividends or reductions in any of the preexisting tax rates.

3.4 Foreign Trade

The model adopts a simplified approach to foreign trade, one that resembles the approach in Rutherford and Schreiber (2019) and does not require an explicit foreign economy. Domestic agents can purchase goods produced domestically or abroad; Armington elasticities determine relative demand. Domestic producers use a constant elasticity of transformation (CET) function to determine the shares of their output for the domestic market and for export. Foreign prices are held fixed, and an endogenous exchange rate balances total trade each period.

4 Data, Parameters, and the Baseline

To obtain a consistent data set and baseline time path for the model, we adopt a multistep process. First, we collect raw data for the benchmark year, 2019. The data include intermediate inputs, final demands (consumption, investment, and government spending), factor inputs and payments, and taxes paid from the Industry Accounts of the US Bureau of Economic Analysis (BEA). Subsequently, we scale components of the raw data to produce an internally consistent social accounting matrix (SAM). The SAM is then combined with production and utility parameters to produce the benchmark data set. Parameters are identified by the requirement that the outputs from the model match the values in the benchmark data set.

This benchmark data set creates a time path of prices and outputs with balanced growth: relative inputs and prices do not change over time. However, most forecasts of energy use and emissions suggest that energy per unit of output will decline over time and that these changes will differ across sectors. Further, common forecasts also suggest changes in the relative prices of various energy goods over time. To create a more realistic baseline comparison path, we calibrate the model to approximate the reference case forecast by the US Energy Information Administration's (EIA) Annual Energy Outlook (AEO) 2022.

4.1 Primary Data

The primary data come from BEA's annual Input-Output Accounts, augmented by data on energy expenditures to be consistent with information from EIA's State Energy Data System (aggregated to the national level) and capital stocks from BEA's Industry Fixed Assets data. We consolidate this information and perform consistency procedures to ensure that the ultimate data set satisfies the household budget constraint and firms' zero-profit conditions.

4.2 Primary Parameters

Primary parameters for the model are independent of the input-output table described above. These parameters are chosen based primarily on empirical estimates in the literature.[21]

The model assumes Harrod-neutral technological change, which implies that all quantity variables grow at a fixed rate over time on the balanced growth path. We

[21] Secondary parameters, such as share and scale parameters in the CES production and consumption functions, are calibrated to be consistent with primary parameters and the consistent SAM, described above. The critical primary parameters are for the rate of technological change and for the elasticities of substitution in production and utility.

assume a value of 1.9 percent, which is close to the average economic growth from 2020 to 2050 predicted by AEO 2022. We assume that nominal prices grow at 2 percent per year (a fair assumption prior to more recent inflation) and a real interest rate of 4 percent.

The elasticities of substitution in production for intermediate inputs, labor, and capital are derived from estimates from Jorgenson and Wilcoxen (1993) and Jorgenson et al. (2013).[22]

In the household utility function, we calibrate the elasticity of substitution between consumption and leisure to yield a compensated elasticity of labor supply of 0.3.[23] The representative household's constant relative risk aversion parameter is set to 2, implying an intertemporal elasticity of substitution in consumption of 0.5, a value between time-series estimates (Hall (2016)) and cross-sectional studies (Lawrance (1991). Finally, we use elasticities of substitution across consumer goods in the consumption nest, based on the values used by the DIEM model (Ross (2014).

4.3 The Reference Case Path

As indicated above, projections for US energy prices and energy demand forecast changes in the relative prices of fuels and other inputs, as well as reductions in the energy intensity of the economy due to changes in technologies and state and federal policies. Accordingly, we adjust the time-paths of specific share and scale parameters in various sectors so that the model will approximate projections for fossil fuel prices, generation levels by type of generator, and emissions and electricity demand per unit of GDP, by fuel sector, from AEO 2022.

[22]For the nuclear and hydro generators, we set the elasticity of substitution between variable inputs, capital, and natural resources such that the elasticity of supply is 0.05 for each generator.

[23]As indicated in Goulder and Hafstead (2017), this value is at the high end of estimates for married men and single women and in the middle of the range of estimates for married women.

5 Policies Considered

We consider carbon tax policies that differ in the breadth of their coverage. Specifically, the policies differ according to whether the tax is:

- Economy-wide,
- Covering only the electricity sector,
- Economy-wide with an exemption for household motor fuel purchases,[24] or
- Economy-wide with an exemption for EITE industries.[25]

We examine each policy under a wide range of carbon tax rates and associated emissions reductions. We also explore how outcomes differ depending on the method of revenue recycling—a lump-sum rebate or a reduction in individual and corporate tax rates. We generally measure the welfare effects over the infinite horizon considered in the model. In all cases, the carbon tax increases at a constant rate of 4 percent (the prevailing annual interest rate in the model) above inflation from the date of implementation in 2024 through 2050. We vary the initial carbon tax rate in order to consider policies of varying stringency. For comparability, we contrast outcomes from policies achieving the same cumulative discounted CO2 emissions reductions.

[24] In addition to matching policy discussions in 2021, this policy coverage is similar to a policy modeled in Bistline et al. (2024), which specifically notes that the policy exempts gasoline "given the political sensitivity of gasoline prices). It is also similar to a policy proposed in a recent Hamilton Project proposal by two former Biden administration Treasury officials (`https://www.brookings.edu/wp-content/uploads/2023/09/20230927_THP_SarinClausing_FullPaper_Tax.pdf`). That proposal differs in that it also exempts household heating oil. Household heating oil is a tiny fraction of household expenditures (substantial in some regions, but not for the country as a whole) and thus our results would not change significantly if our policy also excluded it.

[25] We categorize the following industries as EITE: cement and concrete, iron and steel, nonferrous metal production (primarily aluminum and copper), pulp and paper, and chemical product manufacturing. These sectors have higher-than-average emissions intensities and are often colloquially referred to as "hard-to-abate" sectors. They also compete in international markets, and increases in energy prices may make them globally less competitive.

6 Policy Outcomes

6.1 Results under Lump-Sum Recycling

Figure 1 displays the welfare costs of each of the four policy alternatives at various levels of emissions abatement (expressed as the percentage reduction in discounted cumulative emissions). Figure 1a shows results when policy-generated revenues are recycled via lump-sum rebates to the household; figures 1b and 1c present results when revenues finance cuts to individual and corporate income tax rates, respectively. We measure welfare cost using the equivalent variation and report this cost per discounted ton of cumulative emissions abatement (i.e., average cost per ton). Both the equivalent variation and the change in emissions are defined relative to the reference (baseline) case.

6.1.1 Power-Sector-Only Policy vs. Economy-Wide Policy

We begin by comparing the economy-wide tax with a tax covering only the power sector, with lump-sum recycling. The latter tax policy is the narrowest case among the ones considered. Figure 1a shows that for emissions reductions up to about 17 percent of reference case emissions, the power-sector-only policy has lower costs than the equivalently stringent (in terms of discounted cumulative emissions reductions) economy-wide policy. Results from our analytical model offer a basis for these perhaps unexpected results.

Cost-curve Intercepts

Consider first the different vertical intercepts under these two policies. Equation (16) from the analytical model expresses the marginal cost per ton as the leakage factor $1/(1-\theta)$ times the sum of three terms. The first two are proportional to the

Figure 1: Welfare Costs by Sectoral Coverage And Recycling Options

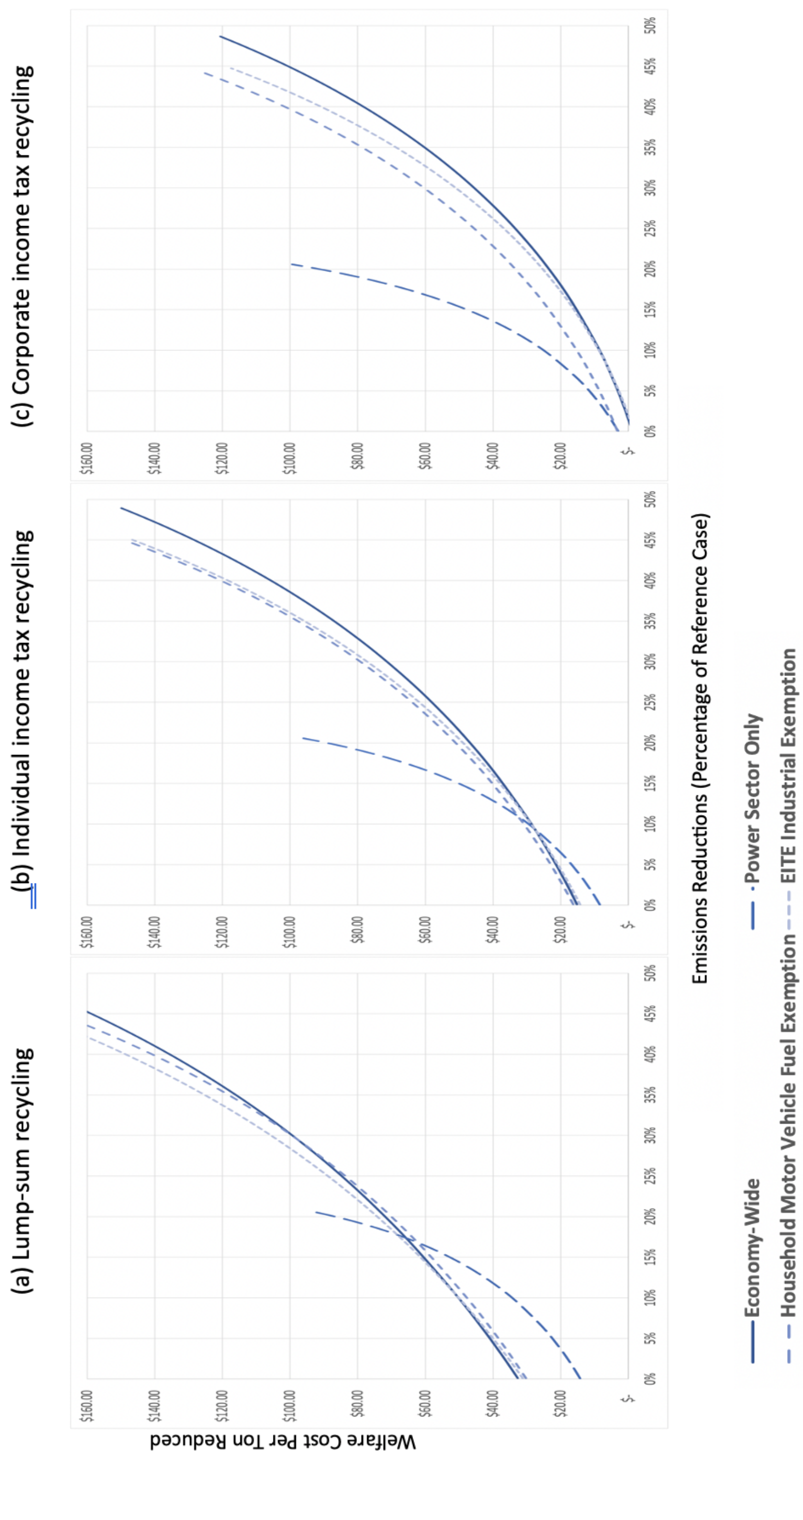

direct cost, and the third depends on the magnitudes of tax-interaction and revenue-recycling effects and on the semi-elasticity of emissions with respect to the imposed tax (adjusted for leakage from covered to noncovered sectors). Because the first two terms are zero at the intercept, the height of the intercept equals the third term.

As shown below, we can view the third term as ω_{IJR} divided by $\bar{\epsilon}_J$, where $\omega_{IJR} = \eta_R(\mu_{IJ} - \mu_R)$ and represents the combined TI and RR effect per dollar of gross carbon tax revenue. This is approximately the welfare cost per dollar of gross carbon revenue for a marginal increase in the carbon tax, starting from zero. $\bar{\epsilon}_J = (1-\theta)\epsilon_J$ represents the leakage-adjusted semi-elasticity: the percentage change in covered emissions for a marginal increase in the carbon tax from zero.[26] Hence the intercept equals $\omega_{IJR}/\bar{\epsilon}_J$, the combined TI-RR effect divided by the leakage-adjusted semi-elasticity.

Table 1 displays the derived values for the key elements of this intercept formula. As can be seen from the table, the estimated combined TI-RR effect ω_{IJR} is higher in the case of the power-sector-only carbon tax than for the economy-wide carbon tax. A main reason is that the power sector is particularly capital intensive, and preexisting taxes on capital are higher (and more distortionary) than preexisting labor taxes.[27] This causes the tax-interaction effect to be higher under the power-sector-only policy, and because both policies recycle revenue in the same way (lump-sum), they have the same revenue-recycling effect per dollar. Hence the combined effect is larger for the power-sector-only policy.

However, the estimated leakage-adjusted emissions elasticity is considerably higher for the power sector tax: within this sector, there are opportunities for emissions reductions at relatively low cost, even from very small carbon taxes. The

[26] We estimate $\bar{\epsilon}_J$ by introducing a carbon tax of \$0.05 and calculating $\bar{\epsilon}_J$ as the resulting percentage change in covered emissions divided by 0.05.

[27] See Goulder et al. (2016) for further discussion of this point.

Table 1: Marginal Costs per Ton Reduced at the Intercept

	Lump-Sum Recycling			Individual Income Tax Cuts			Corporate Income Tax Cuts		
	Combined TI-RR Effect per Dollar of Gross Revenue (ω_{IJR})	Leakage-Adjusted Semi-Elasticity ($\bar{\epsilon}_J$)	Marginal Cost per Ton Reduced at Intercept ($\omega_{IJR}/\bar{\epsilon}_J$)	Combined TI-RR Effect per Dollar of Gross Revenue (ω_{IJR})	Leakage-Adjusted Semi-Elasticity ($\bar{\epsilon}_J$)	Marginal Cost per Ton Reduced at Intercept ($\omega_{IJR}/\bar{\epsilon}_J$)	Combined TI-RR Effect per Dollar of Gross Revenue (ω_{IJR})	Leakage-Adjusted Semi-Elasticity ($\bar{\epsilon}_J$)	Marginal Cost per Ton Reduced at Intercept ($\omega_{IJR}/\bar{\epsilon}_J$)
Economy-Wide	$0.26	1.08%	$23.76	$0.08	1.06%	$7.17	-$0.10	1.05%	-$9.39
Power Sector Only	$0.31	2.60%	$11.89	$0.13	2.59%	$5.08	-$0.04	2.59%	-$1.55
Power Sector Exempt	$0.23	0.39%	$60.28	$0.05	0.38%	$13.58	-$0.12	0.36%	-$34.24
Motor Vehicle Fuel Exemption	$0.28	1.26%	$22.22	$0.12	1.25%	$9.32	-$0.04	1.24%	-$3.45
Motor Vehicle Fuel Only	$0.13	0.13%	$105.30	-$0.12	0.12%	-$97.48	-$0.37	0.12%	-$299.30
EITE Industry Exemption	$0.25	1.10%	$22.84	$0.07	1.08%	$6.48	-$0.10	1.07%	-$9.82
EITE Industry Only	$0.36	0.80%	$44.98	$0.18	0.80%	$22.93	$0.01	0.80%	$1.50

opportunities include fuel switching across fossil fuels (especially from coal to natural gas) and from fossil fuels to renewables. This higher elasticity implies an intercept closer to zero: even though the welfare cost per dollar of revenue is larger for a power-sector-only tax, the effect per ton of emissions reduction is smaller, and thus the intercept (welfare cost per ton) is lower.[28] Hence in this case, the narrow policy is more cost-effective than the broader policy, at least for modest carbon taxes; narrowing the tax to a relatively high-elasticity sector lowers the cost per ton of abatement.

Table 1 also includes a scenario that is the "inverse" of the power-sector-only policy —a scenario in which the power sector is exempt from the tax but all other emissions are covered. Although this seems highly unlikely as actual policy, it is a useful illustrative example, since it helps confirm the significance of the factors we have identified. This policy has a slightly lower welfare cost per dollar of revenue and much lower semi-elasticity than the economy-wide tax. On each dimension, the difference is the opposite of the difference between the power-sector-only and economy-wide policies. Its much lower semi-elasticity leads to a much larger intercept than either the economy-wide policy or the power-sector-only policy. Even though the welfare cost per dollar of revenue is slightly lower for the

[28] Note also that this requires that the combined TI-RR effect be positive, and thus that the intercept be positive. If the combined TI-RR effect were negative (and thus the intercept were negative), then a higher elasticity would still imply an intercept closer to zero. In that case, a higher elasticity would raise costs (i.e., make costs less negative).

power-sector-exempt tax, the much lower elasticity means that the combined TI-RR effect per ton of emissions reduction is much higher. In this case, the narrow policy is initially much less efficient than the broader policy because it has exempted a highly elastic sector from the carbon tax. The results in this case offer an example of a general rule: if a narrow policy is initially less costly than an economy-wide policy, then the inverse of that narrow policy (i.e., a policy covering only the sectors excluded under the original narrow policy) will be initially more costly than the economy-wide policy (and vice versa: if a narrow policy is initially more costly, its inverse will be initially less costly).[29]

Relative cost as a function of stringency

As the stringency of the carbon tax increases, the power-sector-only policy's costs rise faster than the costs of the economy-wide policy, and eventually the power sector policy becomes more costly. Equation 16 from the analytical model demonstrates why the narrower policy's cost advantage declines with stringency. Note that the third term in that equation doesn't include the tax rate, whereas the first two terms are proportional to the tax rate and thus will increase as the policy becomes more stringent.[30] And those first two terms will increase faster for a narrower tax: the narrower the tax, the higher the tax rate needs to be to achieve any given reduction in emissions.[31] Narrowing the carbon tax base will make those terms increase faster with stringency and thus make the cost curve steeper. The

[29]The economy-wide policy is the combination of the narrow policy and its inverse, and thus the cost intercept for the economy-wide policy is a weighted average of the intercepts for the narrow policy and its inverse (with weights equal to the fraction of emissions each covers).

[30]Although the tax rate doesn't appear in the third term, the components of that term can potentially vary as the tax rate increases. Thus, the third term isn't necessarily constant with respect to stringency, and the slope of the cost curve isn't necessarily determined entirely by the first two terms. But the effect of those first two terms will tend to dominate in determining the slope, particularly at higher levels of stringency (and thus higher tax rates).

[31]For example, the crossover point (i.e., where the average cost per ton is the same under the economy-wide and power-sector-only policies) corresponds to a carbon price of $17.49 per ton under the economy-wide policy vs. $69.40 per ton under the much narrower power-sector-only policy. Appendix Table A-1 shows the carbon prices and level of emissions reduction at the crossover points for each of the narrower policies with the economy-wide policy.

magnitude of that difference in slopes, though, will depend on what sectors are excluded under the narrower tax policy. Excluding a high-elasticity sector will necessitate a much larger carbon tax rate to achieve a given emissions reduction and thus will result in a much steeper slope. In contrast, excluding a low-elasticity sector will make a much smaller difference. Indeed, if there were a sector in which emissions were completely nonresponsive to the carbon tax, then excluding that sector would not affect the tax rate needed to achieve any given level of emissions.

Moreover, the narrower policy's cost must eventually overtake that of the broader policy. The reason is simple: a narrower policy covering x percent of emissions cannot reduce total emissions by more than x percent, so the cost curves must intersect at (or before) x percent, the stringency level at which the narrower policy's costs go to infinity.[32]

6.1.2 Policies with Exemptions vs. Economy-Wide Policy

Figure 1a also displays, under lump-sum recycling, the welfare costs for policies that exempt either household motor vehicle fuels (MVF-exempt) or energy-intensive trade-exposed industries (EITE-exempt). For sufficiently low abatement levels (or carbon tax rates), the costs of these policies are nearly identical to those of the economy-wide alternative. The figure also shows that the economy-wide policy eventually becomes more cost-effective than either of the policies with exemptions once stringency reaches a certain level. Under the MVF-exempt policy, the crossover occurs at about 30 percent of reference case emissions; under the EITE-exempt policy, the crossover occurs at about 10 percent. In both cases, the combined TI-RR effect and the elasticity go in opposite directions.

[32]This argument implicitly assumes that leakage is nonnegative (i.e., the narrow carbon policy does not reduce emissions in noncovered sectors). The analysis can be extended to show that the main argument still holds when leakage is negative, though negative leakage will cause the crossover point to be at a greater level of emissions reduction.

If one were to focus solely on the combined TI-RR effect, the MVF-exempt policy would appear to be more costly than the economy-wide alternative because motor vehicle fuels have a relatively low tax-interaction effect. [33] However, household motor vehicle fuel demand is also highly inelastic: changes in the price of fuels have only small effects on drivers' behavior and fuel use. Therefore, as shown in Table 1, although the TI effect per dollar of revenue is low for a tax on MVF, the effect per ton of emissions reductions is very high. Correspondingly, the emissions elasticity for the MVF-exempt policy (which exempts an inelastic sector) is higher than for the economy-wide policy. The effect of that higher elasticity dominates, implying that the MVF-exempt policy is more cost-effective than the economy-wide policy for low to moderate carbon tax rates (under lump-sum recycling).

Table 1 also explains why, for low levels of abatement, costs are lower under the EITE-exempt policy than under the economy-wide alternative. As the table indicates, the EITE industries have both a relatively high combined TI-RR effect and a relatively low semi-elasticity of emissions. Each of these factors tends to increase the intercept. Hence, exempting the EITE industries lowers cost when carbon tax rates are sufficiently low.

As stringency increases, the costs of the MVF-exempt and EITE-exempt policies become higher than those of the economy-wide policy. But the differences in cost are relatively small. In both cases, the demand for output from the exempted sector is relatively inelastic, and thus the sector accounts for only a small share of emissions reductions under the economy-wide policy. Hence, excluding it makes relatively little difference for how fast the tax rate has to rise as stringency increases.[34] In policy discussions, these relatively small long-run cost differences

[33] Consider the inverse of the MVF-exempt policy: a carbon tax applied only to consumer purchases of motor vehicle fuels. This is effectively a narrow consumption tax (a relatively efficient tax), and thus it has a relatively low tax-interaction effect. Correspondingly, the MVF-exempt policy forgoes the opportunity to exploit a low tax-interaction effect. All else equal, this would cause the MVF-exempt policy to be more costly than the economy-wide policy.

[34] The crossover point for the MVF-exempt policy corresponds to an economy-wide carbon price of \$44.68 per ton versus an MVF-exempt carbon price of \$50.80. Similarly, the crossover for the

could be dominated by concerns about near-term equity or political concerns.

6.2 Results Under Individual Income Tax Recycling

Figure 1b displays the welfare costs of the four policy alternatives when policy-generated revenues are used to finance reductions in individual income tax rates. Unlike the case with lump-sum tax recycling, using revenues to reduce preexisting taxes works toward greater efficiency and lowers the cost of all the policies considered. Such revenue-recycling also affects the policies' relative costs.

6.2.1 Power-Sector-Only Policy vs. Economy-Wide Policy

Equation (16) from the analytical model again helps explain the differences in costs across policies. Using carbon revenues to reduce preexisting taxes implies a larger μ_R. This reduces the third term in that equation, which determines the intercept of the cost curve. Thus, for both the power-sector-only and the economy-wide policy, the intercept is lower under income tax recycling than under lump-sum recycling. But the intercept for the economy-wide policy shifts down by more than the intercept for the power-sector-only policy, thus narrowing the initial cost gap between them. The decomposition in Table 1 demonstrates why: the change lowers the combined TI-RR effect for both policies by approximately the same amount, but the higher elasticity under the power-sector-only policy implies a smaller reduction in the cost per ton of emissions reduction. A higher elasticity means less revenue generated for a given reduction in emissions, and thus a smaller gain from recycling that revenue to cut marginal tax rates.

EITE-exempt policy has an economy-wide carbon price of $7.58 per ton versus an EITE-exempt carbon price of $8.50. In each case, the exemption increases the carbon price necessary to achieve a given level of emissions reduction, but only modestly (increasing the carbon price needed by 12 to 14 percent).

The crossover point also moves substantially to the left: here the cost curves cross at emissions reductions of about 9 percent versus about 17 percent under lump-sum recycling. This is due primarily to the smaller initial gap in costs: the economy-wide policy has less ground to make up, so the crossover comes sooner.

6.2.2 Policies with Exemptions vs. Economy-Wide Policy

For the EITE-exempt policy, the change caused by going from lump-sum to income tax recycling is similar to the analogous change for the power-sector-only policy just discussed, and for the same reasons: the intercept falls, and because the emissions elasticity is higher for the EITE-exempt policy than for economy-wide one, the initial gap in costs between those two policies is smaller than in the lump-sum case.

But the effect is different for the MVF-exempt policy. In the earlier discussion of effects under lump-sum recycling, we showed that for modest levels of stringency, the MVF-exempt policy has a lower cost than an economy-wide carbon tax. However, under individual income tax recycling, the MVF-exempt policy is never less costly than the economy-wide carbon tax.

This result is easiest to understand if one looks first at the inverse policy: the MVF-only tax. As discussed before, the MVF-only tax has a relatively small tax-interaction effect (it is effectively a consumption tax and thus is relatively efficient) but also a very low elasticity. Under lump-sum recycling, that combination led to a very high cost per ton: even though the cost per dollar of tax revenue was low, the very low elasticity implied a high cost per ton. Switching to more efficient recycling changes that result dramatically because the gain from recycling is now larger than the loss from tax interactions. So now, both the combined TI-RR effect and the initial marginal cost are negative. The low elasticity still implies the cost per ton is much larger in magnitude than the cost per dollar of revenue, but with the sign change, that low elasticity is magnifying a negative number. In this case, the estimated initial marginal cost for the MVF-only tax is negative $41.49 per ton.

The MVF-exempt policy thus leaves out a sector that is very efficient to tax, so its initial cost is higher than the cost of the economy-wide policy. And because it is narrower, its cost also rises faster as the policies become more stringent. As a result, the MVF-exempt policy always has a higher cost than the economy-wide carbon price. However, the low MVF elasticity implies that the difference between its cost and that of the economy-wide policy is relatively small: even though the initial cost per ton of taxing MVF is negative, only a very small share of emissions reductions under the economy-wide policy comes from reducing MVF emissions, which means that the negative cost per ton gets very little weight in determining the overall cost of the economy-wide policy.

6.3 Results Under Corporate Income Tax Recycling

Table 1 also shows the initial marginal cost when revenue is used to cut the corporate income tax. Because this tax has the highest marginal cost of public funds of all the taxes modeled (see Goulder and Hafstead (2017)), using revenue to cut this tax produces the largest revenue-recycling effect of any of the policies we consider. As the table shows, this larger recycling effect now exceeds the tax-interaction effect, and the combined TI-RR effect is negative (and the initial marginal cost is negative) for all the policies we model except the EITE-only policy. Under all policies and forms of recycling, the cost per ton is the product of the initial marginal cost per dollar of revenue and the inverse of the emissions elasticity. But in all but the EITE-only case, the initial marginal cost is negative, which implies that under corporate income tax recycling, in these cases the cost per ton is decreasing in the emissions elasticity. Correspondingly, if one were to rank the different coverage options by initial cost per ton, the ranking would be the opposite of the ranking when revenue is returned via lump-sum rebates.

7 Sensitivity Analysis

This sensitivity analysis helps reveal the robustness of our findings. It also further illustrates the major mechanisms that determine the relative cost-effectiveness of broader vs. narrower carbon pricing. We start by looking at simulations that vary the level of preexisting taxes. We consider several alternative cases. In the first three, we employ lower (and counterfactual) preexisting tax rates on labor income, capital income (personal and corporate), or both. In a fourth and final case, we assume (counter to fact) that the corporate tax rate cuts from the 2017 Tax Cuts and Jobs Act (the "Trump Tax Cuts") were repealed after the first year (2019) of the simulation time path. This offers a plausible scenario in which the US economy has higher preexisting taxes. Table 2A shows results from our original central case and these alternative cases.

In simulations with lower capital and/or labor taxes, the combined TI-RR effect terms are much lower, reflecting smaller preexisting tax distortions. These cases also reveal that capital taxes are responsible for a larger share of the combined TI-RR effect than labor taxes—the change in the intercept of the cost curve under low capital taxes (relative to the central case) is approximately three times larger than the corresponding change in the intercept under low labor taxes. Nonetheless, our core qualitative result—that narrow carbon-pricing policies can have lower costs than broad policies for sufficiently low levels of stringency—remains robust: even under very low capital and labor taxes, the initial marginal cost of pricing is highest under the economy-wide policy.[35]

[35]However, the range of stringency in which the narrower policies have lower cost is substantially reduced. Having a smaller difference in costs at the intercept means that the broader policy has less ground to make up before catching the initially lower-cost narrow policy. And if we were to reduce all preexisting taxes to zero, the initial marginal cost difference would go to zero.

Table 2A: Marginal Costs per Ton Reduced at Intercept - The Role of Preexisting Taxes

	Combined TI/RR Effect per Dollar of Gross Revenue (ω_{IJR})	Leakage-Adjusted Semi Elasticity ($\bar{\epsilon}_J$)	Marginal Cost per Ton Reduced at Intercept ($\omega_{IJR}/\bar{\epsilon}_J$)
Central Case			
Economy-Wide	\$0.26	1.08%	\$23.76
Power Sector Only	\$0.31	2.60%	\$11.89
Motor Vehicle Fuel Exemption	\$0.28	1.26%	\$22.22
EITE Industry Exemption	\$0.25	1.10%	\$22.84
Very Low Capital Taxation			
Economy-Wide	\$0.11	1.16%	\$9.15
Power Sector Only	\$0.13	2.65%	\$4.98
Motor Vehicle Fuel Exemption	\$0.10	1.33%	\$7.82
EITE Industry Exemption	\$0.11	1.16%	\$9.11
Very Low Labor Taxation			
Economy-Wide	\$0.21	1.08%	\$19.50
Power Sector Only	\$0.25	2.65%	\$9.79
Motor Vehicle Fuel Exemption	\$0.24	1.27%	\$18.96
EITE Industry Exemption	\$0.20	1.10%	\$18.57
Very Low Capital and Labor Taxation			
Economy-Wide	\$0.03	1.16%	\$2.47
Power Sector Only	\$0.04	2.65%	\$1.36
Motor Vehicle Fuel Exemption	\$0.03	1.33%	\$2.28
EITE Industry Exemption	\$0.03	1.17%	\$2.42
Pre-Trump Corporate Income Tax Rates			
Economy-Wide	\$0.28	1.07%	\$25.90
Power Sector Only	\$0.33	2.61%	\$12.70
Motor Vehicle Fuel Exemption	\$0.30	1.26%	\$24.20
EITE Industry Exemption	\$0.27	1.09%	\$24.82

Note: In the Very Low Taxation cases, corresponding taxes are 90 percent below the central case values, with the reductions introduced prior to the data consistency routines. In the "Pre-Trump Corporate Income Tax Rates" cases, corporate income tax rates are at the pre-Trump levels, with the change introduced after the data consistency routines.

The case with corporate tax rates at their levels prior to their Trump-era reductions illustrates the effect of a change in the opposite direction—that is, an increase in prior tax distortions. This increases the combined TI-RR effect and thus raises the cost curve intercept for each of the policies. But the magnitude here is much more modest than in the case of reduced capital taxes, since there is less than a 10 percent increase in cost at the intercept for each policy. Many observers considered the Trump-era cuts to corporate tax rates to be large, but they were still modest compared with the extreme change considered in the previous case—reducing all taxes on capital by 90 percent. Hence, it is not surprising that this tax change has a more modest effect. And again, the result that narrow policies can have lower costs than broad policies remains robust.

We next consider changes to elasticities. Whereas the previous set of changes affected the combined TI-RR effect, these changes affect the leakage-adjusted semi-elasticity. Here we look at changes to four elasticities, in each case varying the elasticity up and down and comparing the results with the central case. Table 2B presents the results of these simulations. Each of the chosen elasticities is particularly relevant for one of the narrow policies. When looking at the change in a given elasticity, we therefore compare results only for the economy-wide policy and the narrow policy for which that elasticity is particularly relevant.

First, we increase and decrease the elasticity of substitution among types of electric power generation (coal, natural gas, and nonfossil generators).[36] Increasing this elasticity makes it easier to reduce emissions in the power sector because it raises the elasticity of power sector emissions with respect to the carbon price. That higher elasticity implies a cost curve intercept closer to zero for the power-sector-only policy. And because some of the emissions reductions under the economy-wide policy come from the power sector, it also lowers the elasticity—and thus the cost-curve intercept—for the economy-wide policy. Lowering the elasticity has the opposite effect, raising the intercept of the cost curve for both policies.

[36]The benchmark elasticity of substitution across generators is 3, a value chosen to approximate emissions reductions from the power sector from detailed bottom-up electricity models at moderate carbon tax levels. In this exercise, we increase and decrease the substitution elasticity to 4 and 2, respectively.

Table 2B: Marginal Costs per Ton Reduced at Intercept - The Role of Key Elasticities

	Combined TI/RR Effect per Dollar of Gross Revenue (ω_{IJR})	Leakage-Adjusted Semi Elasticity ($\bar{\epsilon}_J$)	Marginal Cost per Ton Reduced at Intercept ($\omega_{IJR}/\bar{\epsilon}_J$)
Central Case			
Economy-Wide	\$0.26	1.08%	\$23.76
Power Sector Only	\$0.31	2.60%	\$11.89
Motor Vehicle Fuel Exemption	\$0.28	1.26%	\$22.22
EITE Industry Exemption	\$0.25	1.10%	\$22.84
Generation Elasticity - High			
Economy-Wide	\$0.26	1.15%	\$22.41
Power Sector Only	\$0.32	2.83%	\$11.12
Generation Elasticity - Low			
Economy-Wide	\$0.25	0.99%	\$25.51
Power Sector Only	\$0.30	2.32%	\$12.85
Gas vs. Electricity Elasticity - High			
Economy-Wide	\$0.26	1.08%	\$23.76
Motor Vehicle Fuel Exemption	\$0.28	1.26%	\$22.22
Gas vs. Electricity Elasticity - Low			
Economy-Wide	\$0.26	1.08%	\$23.76
Motor Vehicle Fuel Exemption	\$0.28	1.26%	\$22.22
EITE Energy Elasticity - High			
Economy-Wide	\$0.26	1.09%	\$23.62
EITE Industry Exemption	\$0.25	1.10%	\$22.78
EITE Energy Elasticity - Low			
Economy-Wide	\$0.26	1.07%	\$23.90
EITE Industry Exemption	\$0.25	1.09%	\$22.89
EITE Trade Elasticity - High			
Economy-Wide	\$0.26	1.08%	\$23.80
EITE Industry Exemption	\$0.25	1.10%	\$22.84
EITE Trade Elasticity - Low			
Economy-Wide	\$0.26	1.08%	\$23.72
EITE Industry Exemption	\$0.25	1.10%	\$22.84

Second, we increase and decrease the elasticity of substitution between gasoline and electricity as inputs to private vehicle transportation within the household consumption nest.[37] Increasing this elasticity is a way to represent easier consumer

[37]The benchmark elasticity between motor vehicle fuels and electricity is 0.3, which represents

substitution from gas to electric cars, which raises the elasticity of emissions with respect to motor vehicle fuel taxes. A higher elasticity lowers the cost of an economy-wide carbon price (which encourages gas-to-electric-car substitution) relative to a price that exempts motor vehicle fuel. But despite a relatively large change in this elasticity, the effects on the leakage-adjusted semi-elasticity and thus on the cost at the intercept are too small to be visible in the table. Gas-to-electricity switching accounts for such a small share of emissions reductions that even relatively large changes in the relevant elasticity have insignificant effects on overall costs.

Third, we adjust the elasticity of substitution among energy inputs in production of EITE goods.[38] Increasing this elasticity makes it easier to reduce carbon emissions in the EITE sector, thus lowering the cost of the economy-wide policy relative to a policy that exempts EITE. This reduces the gap in intercepts between these two policies, but the narrower policy remains more cost-effective at sufficiently low stringency.

Finally, we adjust the Armington trade elasticity for EITE industries.[39] Raising this elasticity increases the cost of emissions reductions in the EITE sector: for a given carbon price applied to EITE, a higher elasticity implies a larger domestic welfare cost because a reduction in demand for exports and an increase in the demand for imports of these goods introduces larger terms of trade effects.[40] This affects the cost of the economy-wide policy but not that of the EITE-exempt policy. However,

the current relatively low level of adoption of electric vehicles and the lack of significant shifts between internal combustion engine and electric vehicles when gasoline prices change. We increase and decrease the elasticity by 50 percent, to 0.15 and 0.45, respectively.

[38]The benchmark elasticity of substitution within fuels (coal, natural gas, refined petroleum products) and between fuels and electricity is 0.71 (a weighted average of elasticities from various EITE sectors used in Goulder and Hafstead (2017)). The elasticity is increased and decreased by 50 percent, to 0.355 and 1.065, respectively.

[39]We generally use Armington elasticities for producer goods of 2; the alternative elasticities are 1.1 and 3.

[40]The model requires balanced trade in all periods. With larger EITE trade elasticities, the exchange rate must change by a larger amount given the larger changes in import and export demand for EITE goods, which implies a larger price increase for other traded goods (such as crude oil).

this effect is small.

The changes we have considered here affect the magnitudes of the cost differences across policies and the range of stringency under which each narrow policy has lower costs than the economy-wide policy. However, the core qualitative result is robust across the different cases: for sufficiently low stringency, narrow carbon-pricing policies can have lower costs than economy-wide carbon pricing.

8 Conclusions

Economists often tout emissions pricing as the most cost-effective way to bring about reductions in CO_2 emissions. And among the alternative emissions-pricing options, broad-based pricing approaches generally have been favored over narrower ones on the grounds that they have greater potential to capture low-cost abatement opportunities and lead to fewer distortionary shifts in production or consumption toward uncovered sectors.

While taking account of these potential attractions, the theory and numerical simulations offered in this paper reveal circumstances under which a narrower emissions-pricing policy yields emissions reductions at lower cost per ton than a broader one. These results have political significance, given that the political resistance to emissions pricing in certain sectors can be especially stiff. The results identify cases in which potential political constraints can be addressed without sacrifice of cost-effectiveness.

In the absence of prior tax distortions, greater breadth is an advantage. Under these conditions, carving out some sectors compromises cost-effectiveness: it forgoes potential low-cost abatement opportunities from the excluded sectors and thereby leads to higher economy-wide cost by requiring extra abatement and cost in the remaining, covered sectors.

Our analytical model shows that the situation is very different in realistic economies involving prior distortionary taxes. Such taxes imply that marginal costs of abatement can differ significantly across sectors, even for initial (infinitesimal) abatement in each sector. The model identifies the distortionary costs and the associated (and differing) marginal costs of abatement. It reveals circumstances in which policy costs can be reduced by exempting a sector with particularly high tax-interaction effects.

The model shows that the significance of prior distortionary taxes in a given sector depends on the extent of tax interactions (a function of the sector's factor tax rates and factor intensities): excluding a sector with relatively large preexisting distortions will tend to reduce costs. The significance of these interactions also depends on the elasticity of the sector's emissions with respect to the carbon tax rate: the less elastic emissions are, the larger (on a per ton basis) the influence of the combined interaction effects will be. If that combined effect is a net negative, this provides another reason to omit less elastic sectors. (Likewise, if the combined TI-RR effect is a net positive, it offers a reason to omit more elastic sectors.)

The analytical model further indicates that the relative cost of the narrower versus broader alternative depends on policy stringency. It finds that the potential advantage of a narrower policy is especially significant when the desired reduction in emissions (and corresponding emissions price) is "modest." When much higher abatement is required, the opportunity cost from exempting a given sector increases, since the marginal cost of the extra abatement in other sectors necessitated by the exemption of the given sector becomes much higher. Thus, for a sufficiently stringent abatement target, the broader policy is more cost-effective.

Our numerical general equilibrium model's results reinforce the analytical model's finding while offering quantitative assessments. Under central values for parameters, we find that a power-sector-only carbon tax is more cost-effective than an economy-wide one, for economy-wide emissions reductions of less than roughly 9 percent (when carbon tax revenue is recycled via cuts in the individual income tax) to 17 percent (for lump-sum recycling).

Exempting motor vehicle fuels, in contrast, increases costs for any level of emissions reductions when carbon tax revenue is used to cut income taxes, but under lump-sum recycling, it is more cost-effective than an economy-wide tax for emissions reductions less than roughly 30 percent. In both cases, the cost differential is small and might be outweighed by political or distributional considerations. Motor vehicle fuels are relatively inelastic, so excluding them has a relatively small effect on emissions at a given price. And taxing motor vehicle fuels is more efficient as a revenue-raiser than the income tax (so excluding them raises costs when tax revenue is used to cut the income tax) but less efficient than a lump-sum tax (so excluding them can lower costs in that case). These results complement the analytical model's results: they reveal that accounting for differences in marginal abatement costs and the responsiveness of emissions to a carbon tax has important implications for the choice among carbon-pricing policies.

Our results generalize in several ways. First, although we have considered policies that omit particular industrial sectors, our results would also apply to geographically narrow policies, such as subnational policies that cover only one state (e.g., California's cap-and-trade system) or a small group of states (e.g., the Northeast's Regional Greenhouse Gas Initiative). Our results suggest that exempting states with particularly large tax rates on labor and capital income could be more cost-effective than a federal carbon price for relatively modest emissions reduction targets. However, existing policy (at least in the United States) seems to go in the opposite direction: the states covered by existing subnational carbon prices tend to be those with higher preexisting tax rates, implying that interactions with tax distortions magnify the cost disadvantage of geographically narrow policies rather than overcoming it.

Second, although we have focused on preexisting tax distortions, our qualitative results apply to other preexisting distortions, such as imperfect competition or innovation spillovers. To the extent that such distortions are larger in some sectors than others, a narrow policy omitting sectors in which those distortions raise costs the most (or lower costs the least) would have lower costs than a broad policy (for sufficiently low stringency). And the potential of the narrow policy to have lower cost than a broader alternative would fall with stringency. Thus, even if the ultimate goal is deep reductions in carbon emissions, starting with a narrow policy at relatively low stringency and then gradually broadening and tightening could serve as a politically feasible and cost-effective path to a stringent broad-based policy in the longer run.

Climate change is an exceptionally serious global problem. But political considerations often make it difficult to address. Excluding certain economic sectors when implementing a carbon price can help overcome political roadblocks —and can do so while lowering —or at least not substantially increasing —the overall policy cost.

References

Aissa Assia, Z. 2011. Environmental taxation a guide for policy makers. OECD.

Barrage, L. 2020. Optimal dynamic carbon taxes in a climate-economy model with distortionary fiscal policy. *Review of Economic Studies*, 87(1): 1–39.

Bistline, J., K. A. Clausing, N. Mehrotra, J. H. Stock, and C. Wolfram. 2024. Climate policy reform options in 2025. Working Paper 32168. National Bureau of Economic Research, Boston, MA.

Bovenberg, A. L. and R. de Mooji. 1994. Environmental levies and distortionary taxation. *American Economic Review*, 84(4): 1085–89.

Fullerton, D. and G. E. Metcalf. 2001. Environmental controls, scarcity rents, and pre-existing distortions. *Journal of Public Economics*, 80(2): 249–67.

Goulder, L. H. 2013. Climate change policy's interactions with the tax system. *Energy Economics*, 40(Supplement 1): S3–S11.

Goulder, L. H. and M. A. Hafstead. 2017. *Confronting the Climate Challenge: US Policy Options*. Columbia University Press, New York, NY.

Goulder, L. H., M. A. Hafstead, and R. C. Williams III. 2016. General equilibrium impacts of a federal clean energy standard. *American Economic Journal: Economic Policy*, 8(2): 186–218.

Hall, R. E. 2016. Intertemporal substitution in consumption. *Journal of Political Economy*, 96(2): 339–57.

Jorgenson, D. W., R. J. Goettle, M. S. Ho, and P. J. Wilcoxen. 2013. *Double Dividend: Environmental Taxes and Fiscal Reform in the United States*. MIT Press, Cambridge, MA.

Jorgenson, D. W. and P. J. Wilcoxen. 1993. Reducing US carbon emissions: An econometric general equilibrium assessment. *Resource and Energy Economics*, 15(1): 7–25.

Lawrance, E. 1991. Poverty and the rate of time preference: Evidence from panel data. *Journal of Political Economy*, 99(1): 54–77.

Metcalf, G. E. and D. Weisbach. 2009. The design of a carbon tax. *Harvard Environmental Law Review*, 33: 499–566.

Parry, I. W. 1995. Pollution taxes and revenue recycling. *Journal of Environmental Economics and Management*, 29(3): S64–S67.

Parry, I. W. and R. C. Williams, III. 1999. A second-best evaluation of eight policy instruments to reduce carbon emissions. *Resource and Energy Economics*, 21(3-4): 347–73.

Ross, M. T. 2014. Structure of the dynamic integrated economy/energy/emissions model: Computable general equilibrium component, diem-cge. Working Paper 14-12. Nicholas Institute for Energy, the Environment, and Sustainability, Duke University, Durham, NC.

Rutherford, T. F. and A. Schreiber. 2019. Tools for open-source, subnational cge modeling with an illustrative analysis of carbon leakage. *Journal of Global Economic Analysis*, 4(2): 1–66.

West, S. E. and R. C. Williams, III. 2007. Optimal taxation and cross-price effects on labor supply: Estimates of the optimal gas tax. *Journal of Public Economics*, 91(3-4): 593–617.