Working Paper Series
Congressional Budget Office
Washington, D.C.

Distributional Effects of Reducing Carbon Dioxide Emissions With a Carbon Tax

Dorian Carloni
Congressional Budget Office
dorian.carloni@cbo.gov

Terry Dinan
Congressional Budget Office
terry.dinan@cbo.gov

Working Paper 2021-11

September 2021

To enhance the transparency of the work of the Congressional Budget Office and to encourage external review of that work, CBO's working paper series includes papers that provide technical descriptions of official CBO analyses as well as papers that represent independent research by CBO analysts. Papers in this series are available at http://go.usa.gov/xUzd7.

We are grateful to Nicholas Chase, Edward Harris, Joseph Kile, John Kitchen, Jeffrey Kling, John McClelland, Kevin Perese, Kerk Phillips, Joseph Rosenberg, and Julie Topoleski for helpful comments and suggestions. We also thank Lawrence Goulder of Stanford University, Marc Hafstead of Resources for the Future, and Gilbert Metcalf of Tufts University. Although those experts provided considerable assistance, they are not responsible for the contents of this paper. Rebecca Lanning was the editor.

Abstract

Putting a price on emissions of carbon dioxide, either by taxing them or by establishing a cap-and-trade program, is one policy that lawmakers could consider to address climate change. Although such a policy could encourage cost-effective reductions in emissions throughout the economy, lawmakers have expressed concern about whether it would disproportionately affect lower-income households.

Determining the distributional effects—that is, the effects on households at different income levels—of a policy that would price carbon emissions (referred to in this paper as a carbon tax) is challenging, and the results would vary substantially depending on how the effects were measured. Using a method that allocates the burden to households on the basis of their income rather than their consumption, the Congressional Budget Office estimates that the burden on households in the lowest income quintile, measured as a percentage of income before transfers and taxes, would be roughly twice as large as that imposed on households in the highest income quintile. The burden on households appears less regressive if measured as a percentage of income after transfers and taxes, largely because of the progressivity of the existing federal transfer and tax system.

This paper describes CBO's current method for measuring distributional effects and its rationale for choosing that method, while also comparing it to CBO's prior method and methods used by other researchers. Compared with CBO's prior method, the agency's updated method better reflects the tax burden on annual income and accounts for differences between consumption and income that are due to life-cycle patterns of households' spending, and it better facilitates a comparison between the burden of a carbon tax and the burden of other existing federal taxes.

We also describe the limitations of CBO's updated method relative to CBO's prior method and the methods used by some other researchers. Compared with CBO's updated method, the method the agency used in 2012 better captured the burden that the tax would impose on households consuming out of accumulated wealth and better aligned the revenues raised in a given year with the burden that households incur in that year.

In addition, CBO's updated method and prior method share limitations that are common to all methods that measure the tax burden using data on consumption and income from a single year. Specifically, they do not account for the effect of the tax on households at different points in their life cycle, some of the interactions between the carbon tax and the existing tax system, the additional tax burden on saving if the tax rate increases over time, the macroeconomic effects of the tax, and differential behavioral responses to the tax across income groups.

Keywords: Carbon taxes, cap-and-trade, distributional effects

JEL Classification: H00, H20, H23

Contents

Introduction

Emissions of carbon dioxide (CO_2) and other greenhouse gases (GHGs) accumulate in the atmosphere and contribute to climate change—a long-term and potentially very costly global problem. Policymakers could provide incentives for individuals and businesses to reduce their CO_2 emissions by requiring them to pay a price for emitting CO_2. Such a price could be created by levying a tax on CO_2 emissions or by enacting a cap-and-trade program that would establish rights to a limited number of emissions and allow those emission rights to be traded, thereby establishing a price for the emissions.[1] Those emission prices would increase the relative price of energy-intensive goods and services, providing businesses and households throughout the economy with incentives to produce and consume goods in a manner that led to fewer emissions. Alternatively, policymakers could reduce emissions through regulatory programs, such as clean energy standards. Regulations also affect prices in various ways, but they are generally less cost-effective than a carbon tax would be and achieve a given reduction in emissions at a higher overall cost.

Policymakers could collect a carbon tax "upstream" at the point where carbon enters the economy or, for the electricity sector, when emissions occur. For example, the tax could be collected from electricity generators on the basis of their CO_2 emissions and from producers of natural gas used in home heating and importers or refiners of oil used in transportation on the basis of the emissions that would result when those fuels were burned.

The economic incidence of a tax—that is, the estimation of who ultimately pays the tax—typically differs from who legally remits the payment.[2] Although a carbon tax is collected from fossil fuel producers and importers and from electricity generators, people ultimately bear the costs in the form of higher nominal prices and lower nominal returns on labor and capital. Specifically, the tax would introduce a wedge between prices that consumers pay for goods and services and the returns that workers and investors receive for producing them. Because that wedge would be larger for carbon-intensive goods and services, the carbon tax would increase the cost of producing and consuming carbon-intensive goods, such as gasoline and electricity, relative to other goods and services, such as clothing and food.

Higher relative prices would provide incentives for households to reduce their consumption of carbon-intensive goods and services, such as by buying a more fuel-efficient car. Likewise, decreases in profits and wages in carbon-intensive industries would cause workers and investors

---

[1] In the mid- to late-2000s, the House of Representatives passed, and the Senate considered but ultimately did not pass, legislation that would have imposed a price on CO_2 emissions through a cap-and-trade program. Most recent discussions have focused on imposing a price on CO_2 emissions through a carbon tax rather than a cap-and-trade program. This working paper focuses on a carbon tax.

[2] Similarly, the net burden of limiting emissions through a cap-and-trade program is typically not borne by those who own the right to emit CO_2.

to seek other opportunities. For example, a carbon tax could cause electricity to be generated from renewable sources rather than from fossil fuels. As investors and workers seek new opportunities outside of carbon-intensive industries, relative returns on capital and labor could also be altered.[3]

The changes in relative prices and factor returns would also affect households differently depending on the composition of their consumption or income. All else being equal, households that had relatively carbon-intensive consumption would bear a larger burden than households that spent a smaller share of their income on such goods. Likewise, if a carbon tax caused profits to fall relative to wages, households that received a relatively large share of their income from profits would incur a larger burden (measured relative to their total income) than those that received their income largely from wages.

Determining how the economic incidence of a carbon tax would vary by income is challenging, and many questions arise about the methods and data used to measure that incidence. For example, the annual effect of a tax on a household one year after it is implemented is likely to differ from its annual effect decades later. Similarly, the effect of the tax on a household in a single year is likely to differ from its effects over a household's lifetime.

This working paper outlines the ways in which a carbon tax would affect households and some of the methods researchers have recently used to estimate those effects. We describe the Congressional Budget Office's updated method for estimating how the effects of a carbon tax would vary across U.S. households, as well as the limitations of that method. Then, using the most recent data and CBO's updated distributional framework, we estimate the distributional effects of raising $100 billion in gross tax revenues.

The ultimate distributional effects of a policy that included a carbon tax would depend not just on the effects of the carbon tax but also on how the revenues generated by the tax were used. Lawmakers could offset the effects of the tax on some households by using the revenues to reduce existing federal taxes or to benefit those households in other ways. For example, lawmakers could direct some of the revenues to programs that benefited low-income households or workers in adversely affected industries. Distributional effects associated with the uses of those revenues are beyond the scope of this paper.

Effects of a Carbon Tax on Households

Although a carbon tax would help limit greenhouse gas emissions and reduce the effects of climate change, it would also cause real household income (that is, income adjusted to remove

[3] The effect on relative returns on capital and labor could also vary over time. For example, the demand for labor might be high during a transitional period, when new wind and solar plants were being constructed, but lower in the long run, after the new capacity was in place.

the effects of inflation) to fall, either through increases in the nominal price level or decreases in nominal profits and wages. Because prices and factor returns would not change uniformly, the relative burden that households would face could be affected by the composition of their consumption and their income.[4] Reductions in real household income that resulted from the carbon tax would be offset in two ways. First, income and payroll tax liabilities would be reduced as real household income fell. Second, real income of at least some households could rise depending on how the government used the revenues collected from the carbon tax. The revenues could be used to increase transfers to households, finance reductions in other tax collections, or reduce government debt.

Reduction in Average Real Household Income

A carbon tax would reduce aggregate real household income by the amount of revenues collected by the tax. Thus, before any utilizing of the revenues it generated, a tax that raised $100 billion annually would reduce aggregate household real income by that same amount each year.[5] That reduction could manifest itself as a percentage increase in the nominal price level or as a proportional decrease in nominal profits and wages by the same percentage. In either case, the average purchasing power of households' earnings would be lowered, and aggregate real income would fall.[6] Because the carbon tax is a form of consumption tax (designed to decrease the consumption of carbon-intensive goods and services), the tax-induced decline in real income that households experienced would vary by the fraction of their income that they saved.

---

[4] Throughout this paper, *price level* refers to overall consumer prices. For an in-depth discussion of the effects of consumption taxes on relative prices and income under the assumption of a fixed gross domestic product (GDP) price index and fixed nominal GDP, see Kitchen (2017).

[5] The total burden on households would exceed the revenues raised to the extent that households incurred costs to reduce their carbon consumption, for example, by installing insulation in their homes or purchasing a more fuel-efficient vehicle. As discussed below, because CBO does not have reliable information on how the behavioral responses to the tax would differ across households, the agency projects that, on average, households within each quintile would reduce their consumption of carbon-intensive goods and services by a similar percentage and that, on average, they would incur a similar burden for each unit of carbon consumption avoided (although the composition of those burdens, out-of-pocket versus inconvenience, might vary across quintiles). On the basis of that projection, those adjustments by households would not change the distribution of the burden across households. By contrast, if higher-income households were able to reduce their consumption of carbon-intensive goods and services at a lower cost than lower-income households, then our analysis would overestimate the burden on higher-income households relative to that on lower-income households. If lower-income households could reduce their consumption at a lower cost, then our analysis would overestimate their burden relative to that of higher-income households.

[6] Taxes paid by producers of goods and services can be passed forward to consumers in the form of higher prices or passed back to factor income in the form of lower wages or lower returns on capital. However, the economic incidence of a tax measures changes in real, not nominal, consumer prices and factor income and therefore is mostly unaffected by a change in the overall price level. One exception is returns from debt that are fixed in nominal terms and whose purchasing power changes only with a change in the average price level. For a discussion of changes in the real value of income with or without changes in the average price level, see Toder et al. (2011).

Change in Relative Prices of Household Purchases

The decline in real income that each household experienced because of the tax would also vary by the composition of their consumption. Because the tax would raise the relative price of carbon-intensive goods and services (such as gasoline and electricity consumption) more than the price of less carbon-intensive goods and services (such as food and entertainment), the fall in purchasing power that a household experienced would depend on the goods and services that they consumed. A household whose consumption consisted of a relatively large share of carbon-intensive goods would experience a larger decline in real income than a household whose consumption comprised a smaller share of such goods. The change in relative prices would induce households to shift their consumption to less carbon-intensive goods.

Change in Relative Returns That Households Receive From Investing and Working

The decline in real income that each household experienced because of the tax could also vary by the composition of their income. That variation would occur if the carbon tax caused firms to change the mix of inputs they used for production or to produce goods in a manner that was more capital intensive, or more labor intensive, than it would be otherwise. For example, if firms' efforts to use less fossil fuel caused them to increase their reliance on capital, rather than on labor, then investment income, the factor return on capital, would rise relative to wages, the factor return on labor. In that case, the carbon tax would reduce the real income of a household whose income consisted solely of wages by a greater amount than an otherwise similar household whose income consisted of a mix of investment income and wages.

In addition, the tax could affect returns on capital and wages differently across sectors. For example, returns on capital and wages could fall more in carbon-intensive sectors than in other less carbon-intensive sectors. Such differences across sectors would probably be more important in the short run, before the economy adjusted, than in the longer run, when factors of production could move across sectors to better equalize changes in factor returns.

Reduction in Income and Payroll Tax Liabilities

The decline in households' real income would also reduce the real burden of income and payroll tax liabilities on those households, somewhat offsetting the effect of the carbon tax on their real after-tax income. The offsetting effect associated with the decline in the real value of income and payroll tax liabilities would occur regardless of whether the average nominal price level increased or average nominal profits and wages fell. If the overall price levels remained unchanged and factor income fell, households' nominal payments of income taxes and payroll taxes would fall because their nominal income and wages would decline. The ultimate effect on households would be similar if the nominal price level rose, although the manner in which it occurred would be different. In that case, the carbon tax would not lower their nominal income and wages, but the real burden of income and payroll taxes would be lower because of the higher nominal price level.

Change in Households' Current Real Income Due to Use of Revenues

Households' current real income would rise if the government spent the revenues that it collected from the carbon tax or used those revenues to finance reductions in other tax collections. Whereas the distributional effects of the carbon tax would depend on the carbon intensity of households' consumption and the extent to which their income declined because of the tax, the distributional effects associated with the use of the additional revenues would be determined by policymakers' choices about how to use those revenues. Although beyond the scope of this paper, some uses of the revenues (such as sending each household an equal share in the form of a rebate) could more than offset the economic incidence that a subset of households experienced as a result of the carbon tax.

Existing Estimates of Distributional Effects

Previous studies have measured the distributional effects of a carbon tax in different ways and, correspondingly, have found a wide array of results. This section presents a selection of estimates, describing how different methodologies affect distributional measures. Figure 1 shows the distributional effects estimated by those studies.

In 2012, CBO estimated the distributional effects of a carbon tax by allocating the tax to each household on the basis of its reported annual consumption and measuring that burden relative to the household's reported annual income.[7] That comparison captured differences in the amount of consumption and income reported for households in each income quintile as well as differences in the mix of goods they consumed. The 2012 estimate did not account for any potential changes in relative profits and wages.

Using that method, CBO estimated that the carbon tax would be regressive, meaning that lower-income households would face a larger reduction in income than higher-income households because of the tax (see Figure 1). The main cause of that regressivity was a large gap between reported consumption and reported income for households at both ends of the income spectrum, with consumption exceeding income by more than 200 percent for households in the lowest income quintile and amounting to less than half of income for households in the highest income quintile. Such large differences could be due both to reporting errors and to life-cycle effects: In early and late stages of a household's life, it may consume more than its income by borrowing from future earnings or by spending out of savings. By contrast, at the peak of its earning years, a household may consume less than it earns to fund consumption after its members retire.

Rausch et al. (2011) also estimated distributional effects by allocating the tax to households on the basis of their annual reported consumption and measuring that burden relative to households' annual income. One main difference between that study and CBO's 2012 results is that Rausch et al. imputed additional income to households if their consumption exceeded their income. That

[7] See Dinan (2012).

imputation caused consumption to more closely align with income, making the method used by Rausch et al. closer to the method used by Horowitz et al. (2017) and Rosenberg et al. (2018), described below, both of which assigned the burden of the tax to households according to their income rather than their consumption.

Other choices not linked to allocating the tax burden on the basis of households' income or consumption or to using annual or lifetime measures also affected the distributional outcomes that Rausch et al. obtained. For example, they estimated that the tax would cause returns on capital to fall relative to wages.[8] Because returns on capital make up a larger share of total income in higher-income households than in lower-income households, such a change imposes a larger burden on higher-income households than on lower-income households. In addition, Rausch et al. assumed that all price-indexed transfer payments, which make up a larger share of income for lower-income household than for higher-income households, were not affected because they are indexed to price changes.

Goulder et al. (2019) also allocated the tax to households on the basis of their consumption but used a life-cycle analysis, which ensured that lifetime income matched lifetime consumption. Their results were therefore not driven by gaps in reported annual consumption and income. Constraining consumption to more closely align with income brought their analysis closer to the studies that allocated the carbon tax burden to households according to their income (the analyses by Horowitz et al. and Rosenberg et al. described below).

Like the results from Rausch et al., those from Goulder et al. were also determined by factors that were not necessarily linked to the decision to allocate the tax to households on the basis of their consumption or income or by the use of annual versus lifetime measures. For example, Goulder et al. estimated that returns on capital fell relative to labor in the near term but that such returns evened out over time as the supply of capital adjusted. In addition, Goulder et al. assumed that all price-indexed transfer payments, such as Supplemental Nutrition Assistance Program (SNAP) payments, were protected from an average loss in purchasing power. In their estimate, that protection offset the negative effect of more carbon-intensive consumption enough for the households in the lowest two income quintiles so that, on average, their lifetime burden was less than the lifetime burden of households in the middle income quintile.

Horowitz et al. allocated the reduction in purchasing power caused by the carbon tax to the time when households earned their income rather than when they spent their income to consume goods and services. Allocating the tax to households on the basis of their reported annual income tended to reduce the estimated regressivity of the tax by eliminating the effects of gaps between reported annual consumption and reported annual income for households in different quintiles.

---

[8] Rausch et al. (2011) noted, however, that the effect could be a result of their closed-economy assumption and may not hold in an open economy.

Specifically, Horowitz et al. allocated the loss in purchasing power to households according to the profits and wages that they earned, implicitly holding constant the purchasing power of returns on savings and of transfer income, such as Social Security benefits or SNAP payments. The analysis by Horowitz et al. accounted for differences in the carbon intensity of households' consumption; unlike Rausch et al. and Goulder et al., however, Horowitz et al. did not account for the potential effect of the tax on changes in relative returns on capital and labor.

Like Horowitz et al., Rosenberg et al. allocated the tax to households on the basis of their reported annual income—specifically, allocating the reduction in purchasing power to households at the time they earned their income.[9] Those analyses accounted for differences in the carbon intensity of households' consumption and did not account for potential changes in relative returns on capital and labor.

Although Horowitz et al. and Rosenberg et al. both allocated the tax burden to households according to their annual income, other choices caused Rosenberg et al. to estimate a more regressive outcome. Most importantly, unlike Horowitz et al., Rosenberg et al. assumed that all cash transfer payments, including Social Security benefits and SNAP payments, are burdened by a carbon tax. They made that assumption because they estimated that, over the long run, such payments are linked to wages. Because transfer payments tend to benefit low-income households more than high-income households, that assumption increased the estimated regressivity of the tax.

CBO's Updated Method

CBO's updated method allocates the burden of the carbon tax to households on the basis of their annual income, whereas the agency's 2012 method allocated it according to their consumption. The new method accounts for the effect of the carbon tax on households in two distinct ways: It accounts for how changes in relative prices would affect households on the basis of the carbon intensity of their consumption, and it accounts for how changes in relative returns on capital and labor would affect households according to the sources of their income. In addition, the new method accounts for the effect that the tax would have on price-indexed means-tested transfers.

Allocating the Tax Burden on the Basis of Annual Income

CBO's updated method allocates the burden of a carbon tax on the basis of households' pretax annual income. That approach is similar to the method used by Horowitz et al. and Rosenberg et al., but it differs from CBO's prior method and other methods that allocated the burden of the tax

---

[9] Both studies exempted the normal return on capital from the tax when allocating changes in households' average purchasing power. That is because, although a carbon tax considered in isolation would increase the price of companies' initial investments (which reflects the expected future stream of income generated by those investments), that effect would be offset because the value of the deductions those companies could claim would also increase.

according to annual consumption. CBO’s updated method, which allocates the burden of the tax to households on the basis of their income it, has several advantages:

- It better aligns the tax burden with the measure of ability to pay used to rank households,
- It better accounts for differences between consumption and income that are due to life-cycle patterns of households’ spending, and
- It better facilitates a comparison between the burden of carbon taxes and that of other existing federal taxes.

Allocating the tax according to households’ annual income rather than annual consumption better aligns the estimated tax burden allocated to each household to the base on which the tax is imposed. Because annual consumption can be financed from past savings or borrowing against future income, allocating the burden to households from their current consumption and expressing it as a fraction of current-year income can lead to a misleading measure of the tax rate on current-year income.

In addition, allocating the tax burden on the basis of households’ annual income rather than annual consumption reduces the effect of life-cycle differences on estimated tax burdens. Because households tend to borrow early in their life cycle, save during their peak earning years, and spend out of savings after retirement, a method allocating the tax burden according to annual consumption would allocate very different burdens to two households that had the same lifetime earnings and consumption but were at different stages of life at the time the tax burden was measured. Younger or retired households would be allocated higher burdens than otherwise similar households that were in their peak earning years. Because younger and older households are more likely to be ranked as lower-income households on the basis of their annual income, allocating the tax burden according to households’ consumption would lead to a more regressive outcome than allocating the burden according to households’ income.

Finally, allocating the burden of the carbon tax to households on the basis of their annual income makes it easier to compare the distributional effects of the carbon tax (or other consumption-based taxes) with the distributional burden of current federal taxes. Furthermore, to the extent that policymakers direct the revenues raised by a carbon tax back to households (for example, by reductions in income taxes or as household rebates), allocating the tax according to household income better facilitates a consistent evaluation of both the distributional effects of the carbon tax itself and the recycling of the revenues that it raises. Finally, it ensures that taxes that are economically equivalent are allocated the same distributional effects.[10]

[10] For a broader discussion of tax equivalences, see Auerbach (2019).

Accounting for Changes in Relative Consumer Prices

CBO's updated method accounts for the effect of changes in relative prices that would be associated with a carbon tax. Such an effect is also accounted for in all the other analyses discussed in this paper. Those relative price changes, which are essential to the success of the tax in reducing emissions, have distributional effects because of differences in the composition of consumption among households with different income. Previous research and CBO's own analyses have found that lower-income households spend a greater fraction of their income on carbon-intensive goods than higher-income households, which suggests that lower-income households would face a larger increase in the price of their consumption bundles if a carbon tax was introduced.

Accounting for Changes in Relative Returns on Capital and Labor

CBO's updated method accounts for changes in relative returns on capital and labor that result from the tax. That approach is consistent with the approaches used by Rausch et al. and Goulder et al., which estimate that the carbon tax would reduce returns on capital relative to returns on labor, largely because carbon-intensive industries are also capital-intensive.[11] That effect, which reduces returns on capital relative to wages, has a progressive effect on the overall distributional effect of the tax because returns on capital make up a larger share of income for higher-income households than for lower-income ones.

CBO's updated method incorporates available estimates on changes in the relative returns on capital and labor.[12] While doing so, it also holds the distribution of wages and other factor returns constant across the income distribution. Such differential changes in wages and returns on capital across sectors are more likely to occur in the short run than in the long run and depend on the extent of labor and capital mobility across sectors.

Holding the Average Purchasing Power of Means-Tested Transfers Constant

Like all the other analyses described above, CBO's updated method holds constant the purchasing power of means-tested transfers, such as Medicaid. That purchasing power remains constant because the level of those benefits, a large fraction of which are in-kind transfers, is indexed to changes in prices and does not depend on changes in nominal factor income.

---

[11] The exact magnitude of the changes in relative factor returns depends on model assumptions about the mobility of labor and capital. For example, models that assume relatively greater labor mobility tend to find smaller effects on wages.

[12] As shown in Goulder et al. (2019), the effects on relative factor returns vary over time. Specifically, they found that the reduction in interest rates relative to wages is greatest in the initial years after the tax is put in place and decreases over time.

Limitations of CBO's Updated Method

The distributional framework defined in CBO's updated method also has limitations, most of which apply to any cross-sectional distributional analysis that relies on consumption and income in a single year:

- Because it focuses on average effects by income quintile, it ignores variation in tax burdens within those income quintiles.
- By focusing on the long-run effects of a carbon tax, it ignores the transitional wealth effects of the tax, which can be large.
- It does not align the revenues raised in a given year with households' aggregate burden in that year (as measured by the loss in purchasing power of their income) if the tax rate is expected to increase over time.
- It does not distinguish between households at different stages of their life cycle.
- It does not account for the effects of a carbon tax on the size of the economy.
- To measure the tax burden with respect to prepolicy outcomes, it is based on consumption bundles and production technologies in place before the introduction of the tax.

Variation of Effects Within Quintiles

The results presented here are the average effects among households in each quintile. However, like all analyses that report effects at the quintile level, the effects that we report can differ substantially among households within a given quintile. For example, households that live in different regions of the country depend more or less on cars for transportation and rely on electricity sources that differ with respect to their carbon emissions.[13]

Transitional Effects on Wealth of Imposing a New Tax

CBO's updated method focuses on how to allocate the economic burden of a carbon tax after the tax has been in place for some time. In doing so, it does not capture the transitional effects of the tax, which can be large. Specifically, it fails to capture the burden on wealth accumulated before the tax went into effect, a transitional burden that applies to consumption taxes more generally. A carbon tax, like other consumption taxes, would reduce the purchasing power of that wealth, either because of an increase in the overall price level or because, absent changes in the overall price level, the reduction in nominal income that would result from the tax would be capitalized in the value of that wealth. For example, the reduction in future real returns on capital would lower the current real value of equity assets because the real values of those assets reflect their

[13] For more discussion of heterogeneous effects within income groups, see Rausch et al. (2011) and Cronin et al. (2017).

expected future stream of real income.[14] Methods that allocate the burden of the tax on the basis of annual consumption capture some of that transitional burden on wealth because some households finance consumption with past savings.

Alignment Between Revenues Raised and Burden If Tax Rate Increases Over Time

CBO's updated method for allocating the burden of a carbon tax assigns the aggregate revenues collected from the tax to households according to their current income. That method accurately captures the reduction in the present value of consuming out of current income only if the tax rate is constant over time. Methods that allocate the burden of the tax on the basis of annual consumption help overcome that limitation because the tax revenues raised more closely align with actual households' consumption than with changes in the purchasing power of households' income.

Most tax proposals to reduce carbon emissions, however, contemplate a rising tax rate, which implies that the expected future cost of consuming goods and services is higher than the current consumption cost of consuming those same items. In such a case, the rising tax rate is equivalent to a tax on savings, reducing the purchasing power of income that is saved for future consumption.

In CBO's updated method, that additional tax on future consumption relative to current consumption would be equivalent to a tax on the normal return on capital for that future consumption. Accounting for that would increase the burden on high-income households relative to low-income households because business and capital income is a larger share of total income for those high-income households. Therefore, it would reduce the estimated regressivity of the tax.

Effects at Different Points in Households' Life Cycle

The distributional framework discussed above does not capture the differential effect of the tax on households at different points in their life cycle. Households are compared within an annual framework, and the distributional framework used does not account for age differences of the individuals within those households.[15] That limitation is sometimes addressed by focusing on a longer-term or lifetime measure of tax burden.[16] Although distributing the burden of the tax over a longer time horizon is probably preferable from a conceptual point of view, it involves empirical challenges. For example, studies focusing on lifetime incidence have to rely on strong modeling assumptions about household behavior.

[14] For a discussion of transitional effects on wealth in the context of consumption taxes such as a value-added tax, see Toder et al. (2011).

[15] Estimating distributional effects for specific age cohorts within an annual framework would be possible under our method but is outside the scope of this paper.

[16] See Joint Committee on Taxation (1993) and Goulder et al. (2019).

Macroeconomic Effects on the Overall Size of the Economy

The estimates presented in this paper do not account for changes in the overall size of the economy that might result from the tax. The distributional effects of any such macroeconomic changes are uncertain and would depend on how consumers and producers reacted to the tax change.[17]

Our analysis also does not capture the welfare effects associated with additional distortions created by the interaction of the carbon tax with income and payroll taxes, referred to as the tax interaction effect. Existing income and payroll taxes create a wedge between the pre- and posttax benefits of working and investing, thus reducing the amount that individuals would work and invest in the absence of such taxes. By further reducing the purchasing power of returns on working and investing, a carbon tax would accentuate those existing distortions, creating both further reductions in economic efficiency and additional distributional effects not captured in our analysis.

Fixed Consumption Bundles and Production Technologies

Like all analyses that estimate the burden of a tax at a single point in time, our analysis allocates the tax burden to households on the basis of their prepolicy consumption bundles and production technologies. To the extent that higher- or lower-income households could more easily reduce their consumption of carbon-intensive goods and services, our results would under- or overestimate the regressivity of a carbon tax. To the extent that changes in production technologies would alter relative prices in ways that affected different income groups differently, our analysis would underestimate or overestimate the regressivity of the tax.[18]

CBO's method does, however, account for aggregate responses by households and businesses to a specific tax—that is, for the aggregate change in the mix of goods consumed and in the ways they are produced. The aggregate revenues raised by the tax and allocated to households reflect those aggregate responses and their effects on aggregate emissions.

Results Using Recent Data and CBO's Distributional Framework

CBO's updated methodology allocates a potential carbon tax to households on basis of the composition of their income (allocating a loss in purchasing power to profits and wages and to unemployment insurance and workers' compensation, both of which are indexed to wages) and on the composition of their consumption (allocating larger burdens to households whose consumption is more carbon intensive than that of the average U.S. household and smaller

[17] Metcalf and Stock (2020) and Bernard and Kichian (2021) showed that the carbon tax would have limited macroeconomic effects, but the carbon tax considered in those papers was smaller than the one considered in this paper.

[18] For example, if a carbon tax motivated research and development that lowered the cost and increased the use of renewables in producing electricity, the extent to which the average relative price of electricity increased as a result of the tax would be overestimated under CBO's method.

burdens to those whose consumption is less carbon intensive than that of the average household). The agency's analysis is based on households' 2019 characteristics (see the appendixes for details).

Our analysis uses information on households' expenditures from the Bureau of Labor Statistics' Consumer Expenditures Survey (CE) and information on households' income from the Internal Revenue Service's Statistics of Income–Public Use File. The carbon tax considered in this paper would be imposed on natural gas, oil, and coal and would raise $100 billion in gross tax revenues. The revenues raised from each source of energy would reflect the 2019 aggregate level of emissions from the household consumption of each source. We estimated the consumer price changes that would result from the tax first by quantifying the effects of the tax on producers' prices using an input-output model described in Appendix A and then by mapping those changes to CE price changes. Because of the different carbon content of different commodities, a carbon tax would change the relative prices of those commodities, with carbon-intensive expenditure items such as natural gas, electricity, and gasoline facing the largest price increase (see Figure 2).

Measures of Income Used to Rank Households and Scale the Tax Burden

The measures of income used to rank households and to scale the tax burden on each household when distributing a carbon tax are consistent with the measure CBO has previously used to allocate other taxes.[19] Households are ranked using a measure of their income, before accounting for means-tested transfers (such as Medicaid) and federal taxes. Furthermore, households are ranked using a version of that income measure that adjusts for differences in household size. The burden of the tax on each household is expressed as a percentage of income before means-tested transfers and taxes, which is referred to as income before transfers and taxes. Income before transfers and taxes consists of market income and social insurance benefits (such as benefits from Social Security and Medicare) and excludes means-tested transfers and federal taxes.[20] Means-tested transfers are cash payments and in-kind benefits from federal, state, and local governments that are designed to assist individuals and families who have low income and few assets. They include benefits from government programs such as Medicaid and the Children's Health Insurance Program, the Supplemental Nutrition Assistance Program, and Supplemental Security Income. Federal taxes consist of individual income taxes net of refundable tax credits, such as the earned income tax credit.

Distributional results can be sensitive to the choices of the income measure used to rank households and to scale the burden of the tax. The appropriate measures may depend on the question the analysis is trying to answer. For example, it may be preferable to use income

---

[19] For details about the measures of income used, see Congressional Budget Office (2020).

[20] Market income comprises labor income (including cash wages, employers' contributions for health insurance premiums, and payroll taxes paid by employers), business income, capital income (including realized capital gains), and income from other nongovernmental sources.

measures after transfers and taxes to approximate the incremental effects of a carbon tax when layering it on the existing tax and transfer system. As shown below, ranking households and scaling their tax burden with a postpolicy measure of income results in a measure of the burden that is less regressive than with a prepolicy measure.

Allocation of the Tax Burden

Our analysis allocates the burden caused by the revenues raised by a carbon tax. As discussed earlier, we allocate that tax burden on the basis of income earned by households in any given year and account for households' composition of consumption and income. The allocation of the burden on income earned by households reflects CBO's estimates of the split between labor and investment income, which is uncertain for some income categories, and estimates of the supernormal returns on investment, which come from the earlier literature. The allocation of the burden when accounting for households' composition of consumption reflects differences in households' expenditure bundles as measured in the Consumer Expenditure Survey.

Average Real Income Effect. A carbon tax would create a wedge between producers' and consumers' prices, which would reduce households' real income, either through a reduction of nominal income or through an increase in the overall price level. We define that effect as the average real income effect. The tax would reduce the purchasing power of households' wages and investment income, but, because it is a tax on consumption, it would exempt households' savings from the tax. Exempting households' savings is equivalent to exempting the normal return on those savings, as previously discussed. After reviewing previous literature, CBO estimates that 33 percent of investment income constitutes normal return on investment and is therefore not taxed under a carbon tax.[21] The remaining portion of the return on investment is defined as supernormal return on investment and is burdened by the tax.

Because the values of wage-indexed transfer payments, such as unemployment benefits and workers' compensation, are closely tied to wages, CBO also allocates burden to those forms of income linked to wages. By contrast, no burden is imposed on Medicare benefits (measured by the average cost to the government of providing those benefits), means-tested transfers, and, in CBO's main analysis, Social Security benefits.

The effect of the carbon tax on households' purchasing power reflects an increase in the overall price level or a proportional reduction of nominal income, both of which would decrease households' real income in a similar way. The burden associated with the bulk of the revenues

[21] CBO's estimate of the percentage of investment income that constitutes a normal return is an average of estimates across four prior studies: Cronin et al. (2013), Gentry and Hubbard (1997), Power and Frerick (2016), and Toder and Rueben (2007). Those studies estimated that supernormal returns range between 60 percent and 75 percent and average 67 percent of total returns on investment. As a result, normal returns average 33 percent of total returns on investment.

raised by the tax is allocated to households proportionally to the amount of each household's income facing a decrease in purchasing power because of the tax. Households in the two bottom income quintiles receive a larger share of income as Social Security and Medicare benefits than other income groups, which reduces the decline in real income that they would experience because of the tax. Specifically, Social Security benefits account for about 12 percent of income for households in the bottom quintile and 17 percent of income in the second quintile, whereas imputed Medicare benefits account for roughly 11 percent of income for each of those two quintiles. By contrast, a large fraction of income of higher-income households is made up of wages and investment income, which would be subject to the tax. As a result, a larger fraction of high-income households' income would be subject to the tax, and the real income effect would be progressive: In our estimate, the average household in the bottom quintile would face a 0.59 percent reduction in income before transfers and taxes, compared with 0.53 percent for the average household in the second income quintile and 0.64 percent for average households in the fourth and fifth quintiles.

Relative Consumer Price Effect. Although the average real income effect is estimated as the average loss in the purchasing power of households' wages and investment income, some households would have greater losses in purchasing power than the average loss, and others would have smaller losses. One reason is that a carbon tax would change the relative prices of goods and services, and households differ in the composition of their consumption.

The relative price effect measures changes in households' purchasing power that is driven by differences in households' composition of consumption. By changing the relative prices of different expenditure items, a carbon tax would cause the loss in purchasing power of each dollar of profits and wages earned by households to differ according to the carbon intensity of their consumption. The loss for households with low carbon-intensive consumption (high-income households in our analysis) would be less than that for households with high carbon-intensive consumption (low-income households in our analysis)—that is, households whose utilities and gasoline expenditures are relatively high when measured as a share of their total expenditures.

As shown in Table 1, the average household in the lowest income quintile has more carbon-intensive consumption than a U.S. household with average income, when measured both as a fraction of income and as a fraction of total expenditures. The relative price effect would increase its loss in purchasing power, by an amount equal to 0.23 percent of income, beyond the reduction it would experience if it only experienced an average real income effect. By contrast, because the average household in the highest income quintile has less carbon-intensive consumption than a U.S. household with average income, the relative price effect would partly offset the loss in purchasing power it would experience through the average real income effect, by an amount equal to 0.07 percent of income.

We estimate that relative price effect by focusing on households' consumption bundles before a carbon tax is introduced because changes in households' composition of consumption, which are reflected in the aggregate revenues collected from the tax, are challenging to estimate for each income quintile. In our analysis, we project that, on average, households in each quintile would reduce their consumption of carbon-intensive goods and services by a similar percentage and that, on average, they would incur a similar burden for each unit of carbon consumption avoided (although the composition of those burdens, out-of-pocket versus inconvenience, might vary across quintiles). To the extent that higher-income households could more easily reduce their consumption of carbon-intensive goods and services, our results would overestimate the burden on those households and understate the regressivity of a carbon tax. If the converse was true, then our analysis would overestimate the burden on lower-income households and overestimate the regressivity of the tax.

Relative Factor Returns Effect

In addition to changing relative consumer prices, a carbon tax could also change relative factor returns—that is, returns on capital and labor. That effect on relative factor returns would be uncertain and would depend on a series of factors, including the time horizon considered, changes in households' and businesses' behavior as relative consumer and input prices changed, and the capital intensity of carbon-intensive industries compared with the rest of the economy.

In our analysis, we use estimates of changes in factor returns based on Goulder et al. (2019). They estimated that the return on capital would initially decrease more than the return on labor but that the effect would be reduced over time. As a result, because high-income households earn more income from capital (as a share of their income before transfers and taxes) than other households, the relative factor returns effect would be progressive. Figure 2 shows that the relative factor returns effect would increase the purchasing power of income for households in the bottom quintile by 0.05 percent and would decrease it by 0.06 percent for households in the top income quintile.

Burden on Social Security Benefits

Social Security benefits are based on past earnings. As a result, a carbon tax would not have any effect on the average purchasing power of people who were already receiving Social Security benefits when the tax was implemented.[22] By contrast, the carbon tax would reduce the purchasing power of Social Security benefits for people who began receiving benefits after the carbon tax was in place. The share of Social Security benefits that were affected by the tax would increase over time, and the carbon tax would fully burden Social Security benefits once all current beneficiaries had exited the system.

[22] If the overall price level increased as a result of the tax, that reduction in purchasing power would be offset by cost-of-living adjustments that apply to Social Security benefits each year.

When reflected in CBO's baseline projections, lower Social Security benefits would decrease the income of households receiving Social Security benefits. Therefore, the burden of the tax on those benefits would already be accounted for through lower income, and no additional tax burden would need to be allocated to Social Security benefits. Figure 3 shows the distributional effect of the tax without explicitly allocating any burden to Social Security benefits.

By contrast, the burden on Social Security benefits would need to be accounted for if not already reflected in CBO's baseline projections, and the fraction of Social Security benefits received and bearing the tax burden would vary depending on the time horizon. Figure 4 includes that additional long-run burden on Social Security benefits, which would reflect the burden once that tax had been in effect long enough to reduce the purchasing power of all beneficiaries.[23]

Income and Payroll Tax Offsets

Because a carbon tax would reduce households' real profits and wages, it would also reduce their income and payroll tax burdens, which would partially offset the burden imposed by the carbon tax. That reduction of income and payroll taxes would be larger for high-income households than for low-income households because high-income households pay more in income and payroll taxes. As discussed above, the effects of such tax offsets on households' income would not depend on whether the average price level was increasing.

Income and payroll tax offsets would reduce the burden of the overall tax system on households but are not reflected in the carbon tax burdens allocated to households. Therefore, we quantify the distributional effect of such offsets and include them in our distributional analysis. Accounting for such effects would not be necessary once income changes resulting from a carbon tax were accounted for in CBO's baseline projections.[24]

As shown in Figure 3, those offsets would reduce the burden for high-income households more than for low-income households. We estimate that the loss of purchasing power in households' income for an average household in the top income quintile would be reduced by 0.20 percent of income before transfers and taxes. By contrast, an average household in the bottom income quintile would experience no reduction in the loss of their purchasing power as a result of the offset because income and payroll tax liabilities are close to zero for those households.

[23] That effect on Social Security benefits is not included in the estimated average real income effect because Social Security benefits are classified as government social benefits rather than as income in the National Income and Product Accounts, which justifies their exclusion from the real income effect.

[24] In that case, reductions in income and payroll tax revenues would already be accounted for when distributing the burden of income and payroll taxes on households and could therefore be ignored when focusing on the distributional effects of the overall tax system.

Total Burden of the Tax

The total burden of the tax includes an average real income effect, a relative price effect, and an income and payroll tax offset. In our main estimate, households are ranked on the basis of their income before means-tested transfers and federal taxes (see Figure 3). The total burden of the carbon tax on the average household in each income quintile is expressed as a percentage of their average income before transfers and taxes. We estimate that the average household in the bottom quintile would face a decline in purchasing power equal to 0.76 percent of their income before transfers and taxes. Households in the top quintile would have their purchasing power reduced by 0.42 percent of their income before transfers and taxes.

Those reported effects are sensitive to the income measures used to rank households and to scale the tax burden.[25] Figure 5 shows the same estimated tax burdens, but it ranks households and scales tax burdens by income after transfers and taxes. Shown that way, the tax would still impose larger burdens on lower-income households but would not appear as regressive as the results shown in Figure 3. We estimate that the average household in the bottom quintile would face a decline in purchasing power equal to 0.61 percent of their income after transfers and taxes and that households in the top quintile would have their purchasing power reduced by 0.55 percent of their income after transfers and taxes That less regressive measure reflects the progressivity of the existing tax system, on which a new carbon tax would be layered.

[25] In 2012, CBO ranked households by a measure of income (including means-tested government transfers) before taxes and adjusted by household size and scaled the burden by income after transfers and taxes. For a discussion of how choices about ranking households and scaling the burden affect distributional outcomes, see Perese (2017).

Appendix A: Estimating Changes in Relative Prices

Changes in consumer price for expenditure categories included in the Bureau of Labor Statistics' Consumer Expenditures Survey (CE) require two separate steps. First, we use an input-output model to estimate changes in producers' prices as a tax on emissions from natural gas, coal, and oil is imposed. Second, those changes in producers' prices are mapped into changes in personal consumption expenditures (PCE) categories and CE categories, as discussed below.

The Input-Output Model

The input-output model used to estimate price changes of expenditure items follows the methodology developed by the Congressional Budget Office and described in Perese (2010). The model is estimated by fully passing the tax through to prices, and it uses information from the 2019 Make and Use tables from the Bureau of Economic Analysis (BEA) to estimate the effect of the revenues collected by taxing emissions from natural gas, coal, and oil on the producers' price of commodities.[26] A carbon tax on natural gas, coal, and oil is set to raise 33 percent of the total revenues from natural gas, 21 percent from coal, and 46 percent from oil. That split is based on information published by the Energy Information Administration on carbon dioxide emissions from household consumption of fossil fuels in 2019.[27]

The tax rates on natural gas, coal, and oil are set to raise revenues on domestically combusted inputs. Domestically produced inputs are determined by subtracting BEA's 2019 Import Make and Use tables from BEA's 2019 total Make and Use tables. The fraction of inputs that gets combusted in each industry is based on data from the Energy Information Administration's Manufacturing Energy Consumption Survey, which reports the manufacturing purchases of energy and the fraction of those purchases that gets combusted.[28] Currently, 2018 is the most recent year for which that information is available.

The annual BEA tables do not provide disaggregated information on the utilities and mining commodities. Instead, the 2012 BEA benchmark Make and Use tables are used to disaggregate the utilities commodity into electricity, natural gas, and water, and the mining commodity into coal mining and other mining.

Converting Producer Price Changes to Consumer Price Changes

The price changes estimated from the input-output analysis reflect changes for domestically produced commodities expressed in producers' values rather than for purchases made by households. Therefore, estimated producers' price changes need to be converted into changes in prices faced by households in the CE. That conversion requires two steps:

---

[26] For a detailed description of the content of BEA's Make and Use tables, see Perese (2010).

[27] See Energy Information Administration (2020).

[28] The data underlying our analysis are available at www.eia.gov/consumption/data.php#mfg.

- The first step is to map producers' price changes estimated through the input-output tables to consumers' price changes using PCE categories. That conversion is done by computing the fraction of each PCE that is domestically produced, to which price changes calculated in the input-output model apply (expenditures from imported commodities are held constant), and by using the input-output PCE bridge available from BEA.[29] The input-output PCE bridge also accounts for transportation costs as well as wholesale and retail margins, all of which create a wedge between producers' and purchasers' prices and can also change as a result of the carbon tax.[30]

- The second step is to map price changes faced by households as measured in the PCE price index to price changes faced by households as measured in the CE because the set of expenditure categories differs in the two cases. That mapping relies on a concordance file between PCE categories and Universal Classification Codes included in the CE.[31]

[29] The input-output PCE bridge is available from the BEA website: www.bea.gov/products/industry-economic-accounts/underlying-estimates.

[30] Specifically, changes in transportation costs are calculated from average changes predicted by the input-output model for air transportation (481), rail transportation (482), water transportation (483), truck transportation (484), transit and ground passenger transportation (485), and pipeline transportation (486). Changes in wholesale margins are calculated from the change predicted for wholesale trade (42), whereas changes in retail margins are calculated from average predicted changes for motor vehicle and parts dealers (441), food and beverage stores (445), general merchandise stores (452), and other retail (4A0).

[31] That mapping relies on the crosswalk file that the Bureau of Labor Statistics publishes on its website: www.bls.gov/cex/cepceconcordance.htm.

Appendix B: Statistical Matching of Data Files

The Internal Revenue Service's (IRS's) Statistics of Income–Public Use File (SOI-PUF) includes information on households' income but does not include detailed information on household characteristics and expenditures. Therefore, as described below, statistical matches with the Current Population Survey's (CPS's) Annual Social and Economic Supplement (ASEC) and the Bureau of Labor Statistics' (BLS's) Consumer Expenditures Survey (CE) are necessary. In addition, because our analysis focuses on 2019, aging methods are used to account for demographic changes over time and to impute income and expenditures for each household in 2019.

Creating the CE Extract

The CE is a panel household survey that collects information on households' expenditures on a quarterly basis. Each household is interviewed in five consecutive quarters. The first interview collects basic household demographic data, household income information, and expenditures information for the previous month. The subsequent four interviews collect detailed expenditure information incurred by the household over the previous three months.

Even if expenditure information is collected for a period of 12 months, the CE is not designed to produce annual spending estimates for each consumer unit. Instead, BLS calculates annual spending estimates by adding up the total spending in each quarter. Because our analysis focuses on annual aggregates, we convert CE quarterly files into annual files, so that each household is represented only once in any given calendar year. The calendar year weight allocated to each household accounts for the number of months that household is interviewed during that calendar year and for the quarterly population weight allocated to that household. Households' income and demographics are based on the information collected during the first quarter in which a household is surveyed, and households' estimated annual expenditures are based on an annualized average of households' expenditures over 12 months, which takes into account both the quarterly weights BLS allocates to each household in any given quarter and changes in the consumer price index over time.

Our analysis relies primarily on the CE interview survey files, which include information on households' expenditures at the Universal Classification Code level, but we also use information from CE diary survey files. Those detailed expenditures are aggregated to 56 expenditure categories, with the following adjustments:

- First, we adjust expenditure aggregates in the interview surveys when they differ substantially from those of the diary surveys. That includes adding the average expenditure amount reported in diary surveys to all households in interview surveys for expenditures that are not included in the interview surveys (for example, housekeeping services); it also includes adjusting aggregates in interview surveys for expenditures that are included in both surveys but are probably misreported in interview surveys (for example, food at home).

- Second, we impute utilities expenditures for renter households reporting zero utilities expenditures. The imputation procedure uses information on utilities-to-rent ratios observed for renter households with positive utilities expenditures, divides those households into seven groups with different compositions of utilities expenditures, and imputes utilities for renter households with no utilities expenditures by randomly allocating those to one of the seven groups.

Statistically Matching the SOI-PUF With the CPS ASEC File

The primary source we use to measure household income and federal tax liabilities is the IRS's 2013 Statistics of Income–Public Use File. The SOI-PUF contains a sample of tax returns submitted to the IRS, with an oversampling of high-income tax filing units, and contains detailed income and tax data. Because the data do not include information on certain demographic characteristics, income from transfer programs, and information on households that do not file an income tax return, the SOI-PUF is statistically matched to the Current Population Survey's Annual Social and Economic Supplement for the same year. (For example, the 2014 CPS ASEC is matched to the 2013 SOI-PUF.) The two sources are combined into a SOI-PUF/CPS file by statistically matching each SOI record to a corresponding CPS record on the basis of demographic characteristics and income. Each pairing results in a new record that takes on some characteristics of the CPS record and some characteristics of the SOI record. More details on the statistical matching process are included in CBO (2020).

Aging the CE Extract and the SOI-PUF/CPS Files

Our analysis focuses on the distributional effects of carbon taxes on households on the basis of their 2019 characteristics. The most recent data from BEA are for 2019, the most recent year for which we can construct an annual CE file from quarterly CE files is 2018, and the most recent SOI-OUF data are for 2013. To bring those three sources into alignment, we age both the CE and the SOI-PUF/CPS files to 2019.

Aging of both files accounts for income growth by income source observed in BEA's macroeconomic aggregates. For example, the growth rate applied to income from wages and salaries is different from the growth rate applied to dividend income. In addition, household weights are adjusted using growth rates in Social Security area population aggregates by age, gender, and marital status of households' main respondents. For single households, the CE weight is adjusted according to the characteristics of the main respondent; for married households, it is adjusted by the average population growth when taking into account the characteristics of both the main respondent and the spouse.

Statistically Matching the CE Extract With the SOI-PUF/CPS File

The SOI-PUF/CPS file does not include information on households' expenditures. A statistical match is therefore necessary to impute expenditure information from the CE to households included in the SOI-PUF/CPS file. We use a flexible statistical match routine to impute expenditures to the SOI-PUF/CPS file. The match routine uses four parameters: income, age of the household's head, region, and family type. The region and family type parameters are fixed in the match, but age and income are structured to expand if a suitable match is not found in the first attempt. A potential match is a record in the CE extract that has the exact same family type and region, that has a head of household who is within plus or minus 1 year of age for the head of household in the base file, and that has total household income that is plus or minus 2 percent of the total household income in the base file. If no suitable matches are found for a given household, then the age and income criteria are expanded to be plus or minus 2 years' difference in the head of household's age and plus or minus 4 percent of income. That expansion continues until at least one suitable match is found. Given a pool of suitable matches, each weighted by their CE weight, one household is randomly selected. Matched CE expenditures-to-income ratios are then used to impute expenditure amounts for the household in the SOI-PUF/CPS file.

In addition, because both income and expenditures are top-coded in the CE, SOI-PUF high-income households are generally not statistically matched. Therefore, additional imputations are required for those high-income households. Such imputations rely on coefficients estimated from CE restricted data when regressing households' expenditure on households' income and selected demographic characteristics. Specifically, expenditure amounts for SOI-PUF high-income households are predicted from regression coefficients estimated using CE high-income households (defined as single households with income above $75,000 and married households with income above $150,000) with top-coded information in the publicly available 2019 quarterly CE files.[32]

[32] We are grateful to BLS for running those regressions on the CE confidential data. More information about those regressions is available on request.

Appendix C: Distributing the Burden of the Tax to Households

The total burden of the carbon tax is calculated by keeping consumption bundles and production technologies fixed at their prepolicy levels, and it includes four components: the average real income effect, the relative price effect, the relative factor returns effect, and the income and payroll tax offsets. The first two components of the burden are estimated from consumer price changes estimated through the input-output analysis, which are applied to the components of income burdened under each effect. The sum of the burdens under those two effects is then scaled to raise the target amount of tax revenues, which is $100 billion in our model.[33] The relative factor returns effect is based on estimates of changes in factor returns from Goulder et al. (2019). The loss of revenues from payroll tax offsets is then subtracted from the revenues raised to compute the total burden of the tax. In this appendix, we discuss the estimation of each component in more detail.

Fixed Consumption Bundles and Production Technologies

As previously discussed, our analysis allocates the tax burden to households according to their prepolicy consumption bundles and production technologies. That method is shared by other analyses that estimate the burden of a tax at a single point in time.

Because the Congressional Budget Office does not have reliable information on how the behavioral responses to the tax would differ across households, the agency projects that, on average, households within each quintile would reduce their consumption of carbon-intensive goods and services by a similar percentage. Because CBO does not have reliable information on how specific production technologies would change as a result of the tax, it projects changes in relative prices of goods and services based on existing production technologies.

Average Real Income Effect

The average real income effect is a measure of households' decrease in purchasing power resulting from the tax and is estimated as a proportional decrease in all income sources that are burdened by the tax. Because households differ in the fraction of their income before transfers and taxes (using CBO's definition) burdened by the tax, the real income effect (as a fraction of income before transfers and taxes) varies across income quintiles.

The total amount is computed by first calculating the average percentage decline in the purchasing power of income burdened by the tax across all households (given the price changes estimated through the input-output analysis and each household's composition of consumption) and by then multiplying that factor by the amount of each household's income burdened by the

[33] On the basis of the total output for coal, oil and natural gas, and estimated emissions from each fossil fuel, a tax that raised $100 billion in revenues would impose a 80.5 percent tax rate on coal inputs, a 22.1 percent tax rate on oil, and a 185.9 percent tax rate on natural gas.

tax. The aggregate real income effect of the tax on households is close to the revenues collected ($100 billion in the results shown in this paper).[34]

The burden of the tax is allocated to components of income previously defined by CBO, and the allocation of that burden reflects modeling choices regarding the split between labor and investment income within some income categories:[35]

- A full tax burden on labor income, which includes wages, unemployment insurance and workers' compensation as labor income;

- No tax burden on interest income, because it is projected to fully reflect the normal return on investment and is therefore exempt from the tax;

- A burden on two-thirds of net business income, dividend income, positive rental income, positive realized capital gains, and corporate income taxes borne by capital owners, because one-third is estimated to reflect the normal return on investment;[36]

- A full tax burden on nongovernmental sources of income, which include gambling earnings, cancellation of debt, taxable distributions from health spending accounts, and other net income, and is computed net of net operating losses and the foreign earned income exclusion;

- A burden on 70 percent of pension income, because we project that 60 percent of pension income is labor income (following Cronin et al., 2013) and that the remaining 40 percent is investment income, of which three-quarters constitutes normal returns on income; and

- No tax burden on means-tested government transfers, because the level of benefits for existing beneficiaries is either indexed to prices (if prices change from the tax) or not affected by reductions in factor income such as wages (if factor income falls as a result of the tax).

[34] The aggregate real income effect would aggregate to the revenues collected if the components of income burdened under the average real income and the relative price effects were the same. In our analysis, components of income burdened under the relative price effect include components of income burdened under the average income effect, but also Supplemental Nutrition Assistance Program (SNAP) benefits and Supplemental Security Income (SSI), which are indexed to changes in the average price level but not to changes in the specific consumption bundle consumed by each household. In our analysis, the average income effect aggregates to $98 billion for $100 billion of revenues.

[35] For definitions of the different components of income, see Congressional Budget Office (2020).

[36] The fraction of investment income that constitutes supernormal return does not vary by income quintile, which could lead to underestimating the fraction of supernormal investment income for high-income households and overestimating it for low-income households. For example, recent studies show that rates of returns on dividends and capital gains increase with income; see Smith et al. (2020).

Our estimate that one-third of the total return on many types of investment is normal—and not allocated a burden from the tax—is based on estimates from the literature.[37] The remainder that is estimated to be income from supernormal returns on investment, to which we allocate some tax burden, is particularly uncertain and challenging to measure.

Relative Consumer Price Effect

The relative price effect measures the additional change in households' purchasing power given their composition of consumption. Each household faces a decrease in the purchasing power of income burdened by the tax, which reflects its composition of consumption. The percentage decrease in each household's purchasing power probably differs from the average percentage decline in purchasing power faced by all households. We compute the relative price effect as the difference between the change in purchasing power a household faces on the components of income burdened by the tax and the change in purchasing power that would result if the purchasing power of those components of income decreased by the average percent decline in purchasing power faced by all households.[38] In addition, not all the revenues raised from the tax are raised from households' personal consumption expenditures (PCE), but some revenues are raised from other components of gross domestic product (GDP), such as investment, exports, and state and local government, which also lower real household income.[39] The effect of the tax revenues raised from those other components of GDP on relative consumer prices is less clear. Therefore, our analysis, which maps changes in producers' prices estimated through the input-output model to price changes in PCE, includes only changes in relative consumer prices that result from the burden on PCE. Because PCE account for roughly 70 percent of GDP, roughly 70 percent of overall tax revenues produce changes in relative prices in our analysis. That relative price effect is negative (reducing the real income effect discussed above) for households with low carbon intensity of consumption, and it is positive (increasing the real income effect discussed above) for households with high carbon intensity of consumption.

Although the correct measurement of the relative price effect does not depend on households' consumption-to-income ratios, which have been shown to be mismeasured in the Bureau of Labor Statistics' Consumer Expenditures Survey (CE) at the bottom and the top of the income distribution, it requires that there be no differential underreporting across consumption

[37] Our estimates are based on four prior studies: Cronin et al. (2013), Gentry and Hubbard (1997), Power and Frerick (2016), and Toder and Rueben (2007).

[38] The relative price effect aggregates to zero if the components of income burdened under the relative price effect are the same as the components of income burdened under the average income effect. That is not the case in our analysis because SNAP benefits and SSI are burdened by changes in relative prices but not by changes in the average price level. Because low-income households have more carbon-intensive consumption bundles and are more likely to receive SNAP benefits and SSI, the aggregate relative price effect is positive. In our analysis, it aggregates to $1.4 billion for $100 billion of gross revenues raised.

[39] See Perese (2010) for a discussion of tax revenues raised from each component of GDP.

categories. Such different underreporting would bias the estimated relative price effect and overall distributional effects of the tax.[40]

Relative Factor Returns Effect

The relative factor returns effect measures the additional change in households' purchasing power that results from differential changes in the returns on capital and labor as the carbon tax is imposed. We rely on the estimates of Goulder et al. (2019) for changes in factor returns using their intermediate time horizon (roughly 15 years after the tax is introduced) and scale those changes to the average real income effect estimated in this paper on those sources of income. In our analysis, that results in the return on capital decreasing by 1.22 percent and the return on labor decreasing by 0.61 percent. (The average real income effect for those sources of income is 0.77 percent in our model.) We apply Goulder et al.'s estimated change in the interest rate to both the normal and supernormal returns on capital, which includes interest income, net business income, dividend income, positive rental income, positive realized capital gains, corporate income taxes borne by capital owners, and 40 percent of pension income. We apply the estimated change in the return on labor to wages, unemployment insurance, workers' compensation, and 60 percent of pension income. Because the average reduction in purchasing power for those sources of income is already captured by the average real income effect, the relative factor returns effect is then normalized to that effect so that it aggregates to zero.

Income and Payroll Tax Offsets

A carbon tax, like other indirect taxes that place a wedge between compensation and spending, tends to reduce the amount of income paid to factors of production (such as profits and wages). Consequently, those lower-income payments reduce the revenues derived from existing income and payroll taxes.[41] The aggregate loss of revenues from income and payroll tax offsets is calculated from the Joint Committee on Taxation's (JCT's) excise tax offsets for 2030, the latest available year in its projections for excise tax offsets.[42] JCT estimates that $1 of excise taxes raises would reduce income and payroll tax revenues by $0.242 in 2030, the latest available year at the time the analysis was conducted. We then allocate that tax offset proportionally to the share of total income and payroll taxes paid by each household.

---

[40] For more details on the comparison between the CE and PCE aggregates by consumption category, see Garner et al. (2009). For tabulations of CE and PCE aggregates for recent years, see www.bls.gov/cex/cepceconcordance.htm.

[41] For a discussion of excise tax offsets, see Congressional Budget Office (2009).

[42] See Joint Committee on Taxation (2020).

Figures

Figure 1.
Comparison of Distributional Effects of a Carbon Tax Estimated From Prior Studies

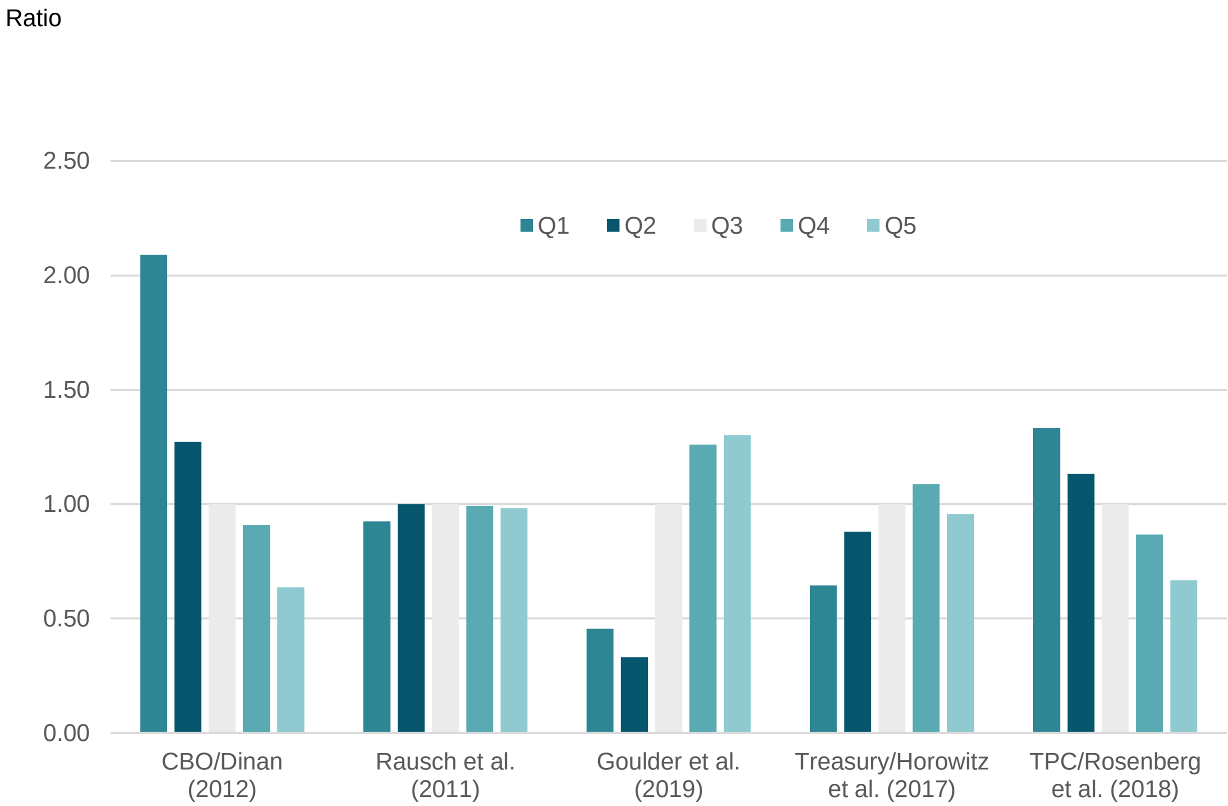

Data source: Congressional Budget Office.

The figure shows estimated tax burdens measured relative to the third quintile to compare across studies that examined taxes of different magnitudes. The results indicate the distributional effects of the carbon tax before the revenues it generated were used.

Figure 2.
Estimated Relative Price Changes for Expenditure Items Facing Largest Price Increase

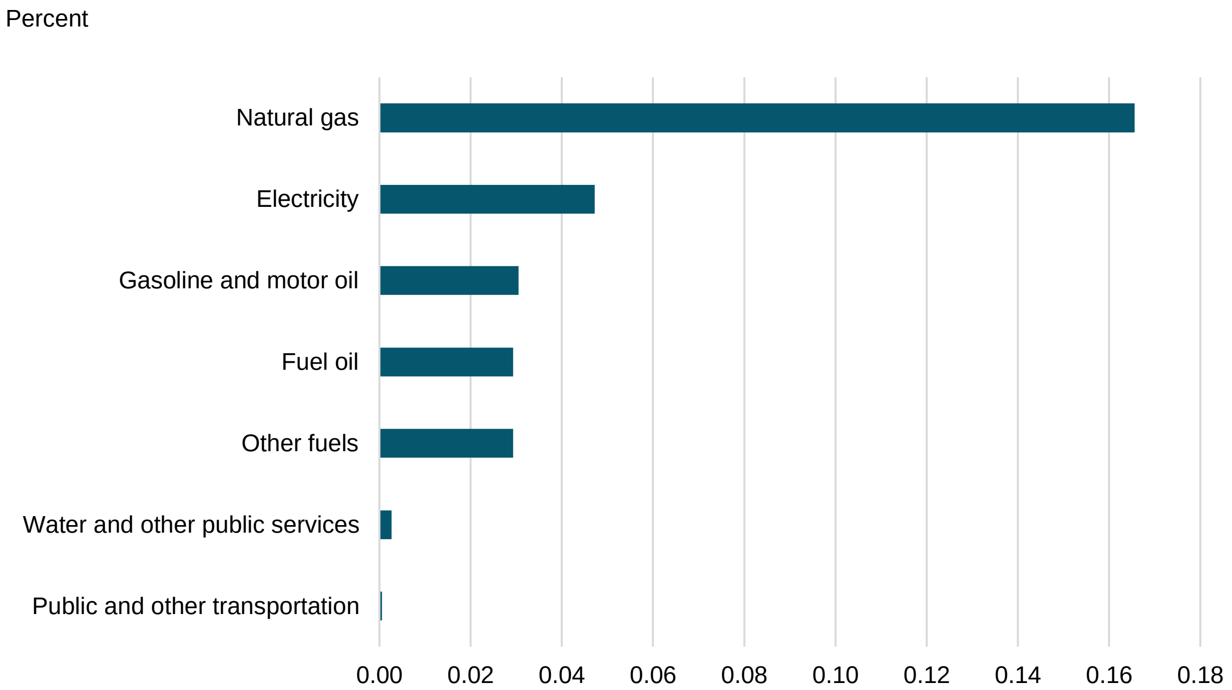

Data source: Congressional Budget Office.

The figure shows estimated price increases for expenditure items facing the largest price increase relative to the average price increase. The price increase for each commodity is estimated by imposing a tax on natural gas, coal, and oil in proportion to the 2019 emission levels from household consumption of those three energy sources, so that $100 billion is raised. The tax results in price changes of 73 commodities calculated with Make and Use tables from the Bureau of Economic Analysis (BEA). Those price changes are then mapped to changes in expenditure categories included in the Consumer Expenditure Survey (CE) through a crosswalk file that maps BEA's input-output categories into personal consumption expenditures (PCE) categories and a crosswalk file that maps PCE categories into CE categories.

Figure 3.
Estimated Distributional Effects of a Carbon Tax

Percent

Data source: Congressional Budget Office.

The figure shows the tax burden as a percentage of income before transfers and taxes, by quintiles of income before transfers and taxes and adjusted by household size. The real income effect is a measure of households' decline in purchasing power stemming from the revenues raised from the tax. The relative price effect accounts for differences in households' expenditure composition and aggregates to zero. The relative price effect is positive for households whose price of expenditures increases more than average, given those households' expenditure bundle, and negative for households whose price of expenditures increases less than average. The relative factor returns effect captures changes in relative returns on labor and capital estimated in Goulder et al. (2019) for 2035. Tax offsets indicate real reductions in income and payroll tax liabilities that result from the carbon tax. The total tax burden is equal to the sum of the real income effect, the relative price effect, and the tax offsets.

Figure 4.
Estimated Distributional Effects When Accounting for the Burden of a Carbon Tax on Social Security Benefits

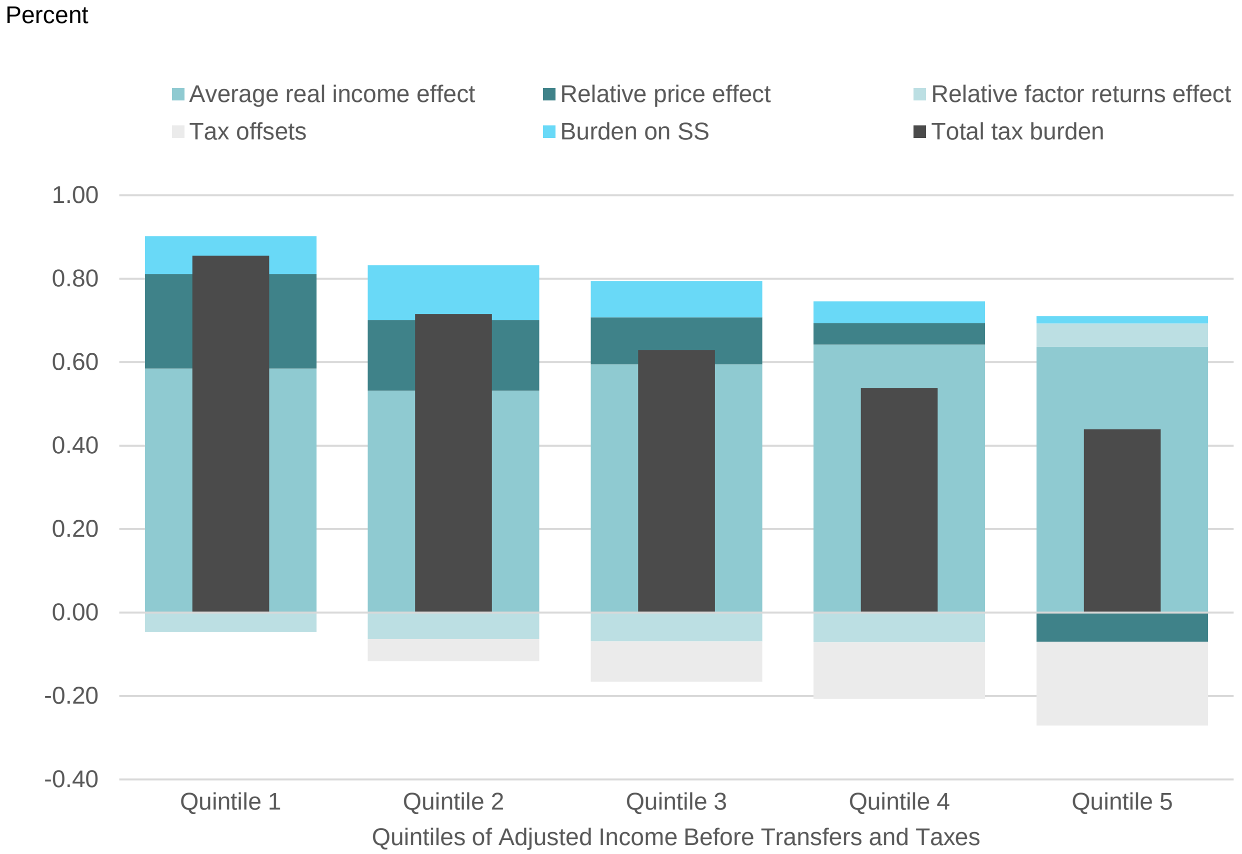

Data source: Congressional Budget Office.

The figure shows the tax burden as a percentage of income before transfers and taxes, by quintiles of income before transfers and taxes and adjusted by household size. The real income effect is a measure of households' decline in purchasing power stemming from the revenues raised from the tax. The relative price effect accounts for differences in households' expenditure composition and aggregates to zero. The relative price effect is positive for households whose price of expenditures increases more than average, given those households' expenditure bundle, and negative for households whose price of expenditures increases less than average. The relative factor returns effect captures changes in relative returns on labor and capital estimated in Goulder et al. (2019) for 2035. Tax offsets indicate real reductions in income and payroll tax liabilities that result from the carbon tax. The burden on Social Security is measured as the reduction in the purchasing power of households' Social Security benefits when the carbon tax is fully phased in. The total tax burden is equal to the sum of the real income effect, the relative price effect, the tax offsets, and the burden on Social Security benefits.

Figure 5.
Estimated Distributional Effects When Using Income After Means-Tested Transfers and Federal Taxes to Rank Households and to Scale the Tax Burden

Data source: Congressional Budget Office.

The figure shows the tax burden as a percentage of income after transfers and taxes, by quintiles of income after transfers and taxes and adjusted by household size. The real income effect is a measure of households' decline in purchasing power stemming from the revenues raised from the tax. The relative price effect accounts for differences in households' expenditure composition and aggregates to roughly zero. The relative price effect is positive for households whose price of expenditures increases more than average, given those households' expenditure bundle, and negative for households whose price of expenditures increases less than average. The relative factor returns effect captures changes in relative returns on labor and capital estimated in Goulder et al. (2019) for 2035. Tax offsets indicate real reductions in income and payroll tax liabilities that result from the carbon tax. The total tax burden is equal to the sum of the real income effect, the relative price offsets, and the tax offsets.

Table

Table 1.
Average Annual Household Expenditures by Income Quintile, 2019
Dollars

	Quintile					All Households
	Lowest	**Second**	**Third**	**Fourth**	**Fifth**	
Utilities Expenditures	1,724	1,970	2,155	2,354	2,852	2,206
Gasoline Expenditures	1,485	1,818	2,270	2,643	2,967	2,230
Total Expenditures on Utilities and Gasoline	3,209	3,788	4,425	4,997	5,819	4,436
Total as a Percentage of Income Before Transfers and Taxes	**13.4**	**7.6**	**5.6**	**4.1**	**1.8**	**6.5**
Total as a Percentage of Expenditures	**8.9**	**8.2**	**7.4**	**6.5**	**4.6**	**7.1**

Data source: Congressional Budget Office, using data from the Bureau of Labor Statistics, Consumer Expenditures Survey (CE) and Statistics of Income Public Use Files (SOI-PUF) aged to 2019.

The table shows average expenditures by quintiles of income before taxes transfers and taxes calculated by statistically matching a 2018 CE extract aged to 2019 and the 2013 SOI-PUF file aged to 2019. Utilities expenditures include natural gas, electricity, fuel oil, and other heating fuels. Gasoline expenditures include expenditures on both gasoline and motor oil.

References

Auerbach, A. J. (2019). Tax Equivalences and Their Implications. *Tax Policy and the Economy*, *33*(1), 81–107.

Bernard, J. T., & Kichian, M. (2021). The Impact of a Revenue-Neutral Carbon Tax on GDP Dynamics: The Case of British Columbia. *The Energy Journal*, *42*(3).

Bureau of Economic Analysis (2019). Input-Output Accounts Data (www.bea.gov/industry/input-output-accounts-data, accessed November 2, 2020).

Congressional Budget Office (2009). *The Role of the 25 Percent Revenue Offset in Estimating the Budgetary Effects of Legislation*. January 2009 (www.cbo.gov/publication/20110).

Congressional Budget Office (2020). *The Distribution of Household Income, 2017*. October 2020 (www.cbo.gov/publication/56575).

Cronin, J. A., Lin, E. Y., Power, L., & Cooper, M. (2013). Distributing the Corporate Income Tax: Revised Us Treasury Methodology. *National Tax Journal*, *66*(1), 239.

Cronin, J. A., Fullerton, D., and Sexton, S. E. (2017). *Vertical and Horizontal Redistributions From a Carbon Tax and Rebate*. National Bureau of Economic Research Working Paper 23250.

Dinan, T. (2012). *Offsetting a Carbon Tax's Costs on Low-Income Households*. Congressional Budget Office Working Paper 2012-16, November 2012 (www.cbo.gov/publication/43713).

Energy Information Administration (2020). *Monthly Energy Review*, Environment, July 2020 (www.eia.gov/totalenergy/data/monthly/archive/00352007.pdf).

Garner, T. I., McClelland, R., & Passero, W. (2009). Strengths and Weaknesses of the Consumer Expenditure Survey From a BLS Perspective. Draft presented at the National Bureau of Economic Research Summer Institute (www.bls.gov/cex/pce_compare_199207.pdf).

Gentry, W. M., & Hubbard, R. G. (1997). Distributional Implications of Introducing a Broad-Based Consumption Tax. *Tax Policy and the Economy*, *11*, 1–47.

Goulder, L. H., Hafstead, M. A., Kim, G., & Long, X. (2019). Impacts of a Carbon Tax Across US Household Income Groups: What Are the Equity-Efficiency Trade-offs? *Journal of Public Economics*, *175*, 44–64.

Horowitz, J., Cronin, J. A., Hawkins, H., Konda, L., & Yuskavage, A. (2017). *Methodology for Analyzing a Carbon Tax*. US Department of the Treasury, Washington, D.C.

Joint Committee on Taxation (1993). *Methodology and Issues in Measuring Changes in the Distribution of Tax Burdens*, JCS-7-93, June 1993 (www.jct.gov/publications/1993/jcs-7-93).

Joint Committee on Taxation (2020). *Updated Income and Payroll Tax Offsets to Changes in Excise Tax Revenues for 2020-2030*, JCX-20-20, August 2020 (www.jct.gov/publications/2020/jcx-20-20).

Kitchen, J. (2017). A Technical Note on Revenue Estimation for an Add-On VAT: The Fixed GDP Convention, Relative Prices, Income Offsets and Monetary Policy. Unpublished working paper, April 2017 (https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3047455).

Metcalf, G. E., & Stock, J. H. (2020). Measuring the Macroeconomic Impact of Carbon Taxes. *AEA Papers and Proceedings 110*, 101–106.

Perese, K. (2010). *Input-Output Model Analysis: Pricing Carbon Dioxide Emissions.* Congressional Budget Office Working Paper 2010-04, June 2010 (www.cbo.gov/publication/21538).

Perese, K. (2017). *CBO's New Framework for Analyzing the Effects of Means-Tested Transfers and Federal Taxes on the Distribution of Income.* Congressional Budget Office Working Paper 2017-09, December 2017 (www.cbo.gov/publication/53345).

Power, L., & Frerick, A. (2016). Have Excess Returns to Corporations Been Increasing Over Time? *National Tax Journal*, *69*(4), 831–846.

Rausch, S., Metcalf, G. E., & Reilly, J. M. (2011). Distributional Impacts of Carbon Pricing: A General Equilibrium Approach with Micro-data for Households. *Energy Economics*, *33*, S20–S33.

Rosenberg, J., Toder, E., & Lu, C. (2018). *Distributional Implications of a Carbon Tax*. Columbia Center on Global Energy Policy and Tax Policy Center.

Smith, M., Zidar, O., & Zwick, E. (2020). *Top Wealth in America: New Estimates and Implications for Taxing the Rich*, Working Paper 264, Princeton University, Department of Economics, Center for Economic Policy Studies.

Toder, E., & Rueben, K. (2007). Should We Eliminate Taxation of Capital Income? *Taxing Capital Income*, 89–141.

Toder, E. J., Nunns, J., & Rosenberg, J. (2011). *Methodology for Distributing a VAT*. Brookings Institution.