Household Impact Study II (HIS2)

The impact of a carbon fee and dividend policy on the finances of U.S. households

Kevin Ummel

President, Greenspace Analytics

Research Affiliate, University of Pennsylvania

ummel@sas.upenn.edu

Prepared for Citizens' Climate Lobby

August, 2020

Working Paper v1.1

Abstract

This paper describes a microsimulation study of a "carbon fee and dividend" policy like that proposed in the Energy Innovation and Carbon Dividend Act of 2019 (H.R. 763). The short-run financial burden of a national carbon tax and allocation of all resulting revenue as per-capita "dividend" payments is simulated for a representative sample of 1.3 million U.S. households. Household-level financial impacts are estimated across demographic groups, states, and congressional districts. Compared to an earlier analysis (Ummel 2016), the current study uses an updated household sample, new consumption data, and improved emission factors, while also incorporating policy administrative costs and current tax law. Overall results are highly progressive. Under a conventional assumption of full pass-through of the carbon tax into consumer prices, the dividend payment exceeds the additional tax burden for 54% of all households and 92% of those in the lowest consumption quintile. Those figures rise to 67% and 98%, respectively, in a scenario that assumes 30% of the tax burden is shifted from consumers to owners of capital. An intermediate scenario with 85% pass-through delivers results of 61% and 96%, respectively. In general, larger families are benefited more than smaller ones. Younger and older households are benefited more than the middle-aged. Minority households are benefited more than white households. Households at similar income levels experience different impacts depending on where they live, largely due to geographic variation in climate, electricity supply, and population density. Tailored allocation of dividend payments could address geographic variation more effectively than other revenue recycling options, but this would require moving away from the simplicity of per-capita dividends.

Table of Contents

Overview

Citizens' Climate Lobby (CCL) commissioned a study in 2016 to assess the impact of a "carbon fee and dividend" policy on the finances of U.S. households. The resulting Household Impact Study (Ummel 2016; hereafter "HIS1") provided a uniquely high degree of socioeconomic and spatial detail, allowing impacts to be estimated for specific types of households and places.

Such detail is achieved via "microsimulation", whereby the policy impacts – given certain assumptions and limitations – are calculated for each household in a large, nationally-representative sample. The approach has three basic steps:

1. Estimate how much each household spends on different categories of goods and services (electricity, clothing, child care, etc.).

2. Estimate the carbon-intensity of different categories (e.g. CO_2 per dollar of clothing) and convert each household's spending into emissions (i.e. a "carbon footprint").

3. Simulate the tax burden and "dividend" payment for each household. The net financial effect is the difference between the additional tax burden and additional income from the dividend.

HIS1 relied on a household consumption dataset dating from 2008-2012 (Ummel 2014), along with other inputs of similar vintage. Since then, the U.S. economy emerged from the "Great Recession" and experienced a "fracking" boom that led to significant changes in the power sector. In addition, a number of new datasets became available and others were updated to the post-Recession period. In light of these changes, an update of the Household Impact Study seems prudent.

While the new study described here (HIS2) shares the same general approach, there are enough changes to technique, data, and assumptions to make it effectively new. Of the many changes, the following are particularly noteworthy:

1. HIS2 simulates two policy scenarios. Scenario 1 assumes that businesses pass 100% of the carbon tax "forward" in the form of higher prices (same as in HIS1). Scenario 2 assumes that businesses pass 70% of the tax forward and 30% "backward" onto owners of capital.

2. A new household consumption dataset underpins HIS2. The representative sample of 1.3 million households reflects U.S. demographics circa 2018 and contains household-level consumption estimates across 57 categories using data from 2015-2018.

3. The consumption dataset is calibrated to match national- and state-level totals for 2018. Results of a novel study by economists at the Federal Reserve Board (Feiveson and Sabelhaus 2019) are used to further correct for income- and age-specific underreporting of consumption.

4. Emission factors used to convert household-level consumption into greenhouse gas (GHG) footprints come from the "USEEIO" life cycle assessment model produced by the

U.S. EPA (Yang et al. 2017). It provides state-of-the-art results for nearly 400 commodities using emissions data from 2016.[1]

5. Taxation of the dividend payment is now simulated explicitly for each of the 1.3 million households and reflects relevant changes introduced by the Tax Cuts and Jobs Act of 2017.

As in HIS1, only the policy's immediate ("overnight") household financial effects of the policy are considered. The microsimulation model is "static" and does not consider "dynamic" effects on economic growth, consumption, employment, wages, trade, technology, or pollution over time. For example, a carbon tax might suppress employment or wages in particular industries (e.g. coal mining) (Ho, Morgenstern, and Shih 2008). This study ignores such details, effectively assuming that the policy only affects households through 1) higher consumer prices; 2) a reduction in business profits (Scenario 2); and 3) the dividend payment. Readers interested in other impacts are directed to Caron et al. (2018), Goulder et al. (2019), Kaufman et al. (2019), and the references within. The advantage of a static microsimulation model is that it allows a level of socioeconomic and spatial detail not found in dynamic models.

A carbon tax would likely put upward pressure on the general price level (inflation). This would trigger automatic increases in social security and food stamp benefits as well as federal income tax brackets. Cronin et al. (2019) estimate these "statutory" effects could cost the federal government ~$25 billion, limiting the amount of money available for dividend payments under a revenue neutral policy. However, the increase in spending confers benefits (higher transfer payments or lower income tax liability) on certain households. HIS2 does *not* model these considerations, but Cronin et al. find that the net effect of including them is generally progressive because transfers disproportionately benefit poorer households.

The USEEIO emission factors reflect domestic production of goods and services. The implicit assumption (same in HIS1) is that imports consumed by U.S. households have the same GHG-intensity as domestically-produced counterparts. In reality, the intensity varies depending on country of origin. However, the "domestic production" assumption is common in such analyses and could well reflect real-world border tax adjustments under unilateral carbon pricing (Metcalf and Weisbach 2009).[2]

Both Scenario 1 (100% pass-through) and Scenario 2 (70% pass-through) assume that the carbon tax revenue is returned to individuals as proposed in H.R. 763. Individuals age 19 and older receive a "full share" dividend; those age 18 and younger receive a "half-share". The dividend is subject to federal payroll and income taxes. In Scenario 2, the 30% of the tax burden passed "backwards" is assigned to households in proportion to investable assets (i.e. a proxy for household exposure to a reduction in business profits). CCL estimates that the policy initially generates $78.1 billion in revenue (net of export rebates) and that $73 billion is dispersed as

[1] HIS1 was limited to CO_2 emissions from fossil fuel combustion. USEEIO emissions factors include all GHG's (i.e. CO_2-equivalent). It is supplemented with the latest data from EPA eGRID such that residential electricity emissions reflect the state of the power sector in 2018.
[2] H.R. 763 applies a border adjustment only to imports of "carbon-intensive products". The analysis here effectively assumes the border adjustment is applied to *all* imports using the domestic production assumption. However, the effect of this discrepancy on distributional outcomes is likely small.

dividend payments after administrative costs. All model results are scaled to match this assumption.

Scenario 2 is not a prediction of how businesses will respond to the carbon tax. The comparatively scant carbon cost pass-through literature (mostly from Europe) generally finds high pass-through rates for direct energy goods like electricity and gasoline but otherwise highly heterogeneous results across industries (Neuhoff and Ritz 2019). A recent analysis of six U.S. industries estimates pass-through rates in the range of 70% (Ganapati et al. 2020). This is the primary basis for the 70% overall pass-through assumed in Scenario 2.[3] The additional assumption that the remaining tax is borne exclusively by owners of capital (as opposed to labor) is the most progressive of possible outcomes. Consequently, Scenarios 1 and 2 generally reflect the most regressive and progressive plausible allocation of the tax burden, respectively.

A third scenario assuming 85% overall pass-through (halfway between Scenarios 1 and 2) was also generated. The results section below discusses only Scenarios 1 and 2, as these constitute the plausible range of impacts. Results for Scenario 3 (85% pass-through) can be found in Annex D.

Additional information and technical details are provided in Annexes A through C.

Results

Throughout this section, a household is said to be "benefited" by the policy if it produces a positive net financial impact – i.e. the after-tax dividend exceeds the additional tax burden – given the assumptions and limitations noted above. A household experiences a "minor loss" if the net impact is negative but does not exceed 0.2% of pre-tax household income.

Both household consumption and income are used in the presentation of results. Consumption refers to total household spending, with the caveat that housing cost is either actual rent paid or (for owner-occupied homes) the imputed rental value. Consumption is generally better than income for measuring how "well-off" a household is, and it is used in a number of recent studies to describe distributional outcomes (Cronin et al. 2019; Goulder et al. 2019). The income measure used here (e.g. to define a minor loss) is a measure of "normal income" derived from the Survey of Consumer Finances; see Annex C for additional details. The terms "poorer" and "richer" are used generally to describe households that are relatively lower or higher on either of these measures.

A note regarding interpretation. The percent of households outright benefited by the policy is largely insensitive to assumptions about total revenue, because the total tax burden and dividend move up or down together. But this is not true of results that measure impact relative to income (e.g. minor losses). If revenue is higher (lower) than that assumed by CCL, the proportion of households incurring minor losses will decrease (increase) relative to the results shown here.

[3] Specifically, the pass-through rate is assumed to be 95% for direct energy goods and ~55% for all other goods and services, which results in an economy-wide pass-through rate of 70% given the current balance of direct and indirect emissions.

Impact across the consumption distribution

Figure 1 shows how the percentage of households benefited and/or incurring a minor loss varies across the consumption distribution in both scenarios. The x-axis is a percentile measure constructed using household consumption as a percentage of the Federal Poverty Level (FPL). Since the FPL varies with the number of people in a household, this measure accounts for economies of scale and provides a more realistic ranking of the economic well-being of different households.

Figure 1: Percent of household benefited across the consumption distribution

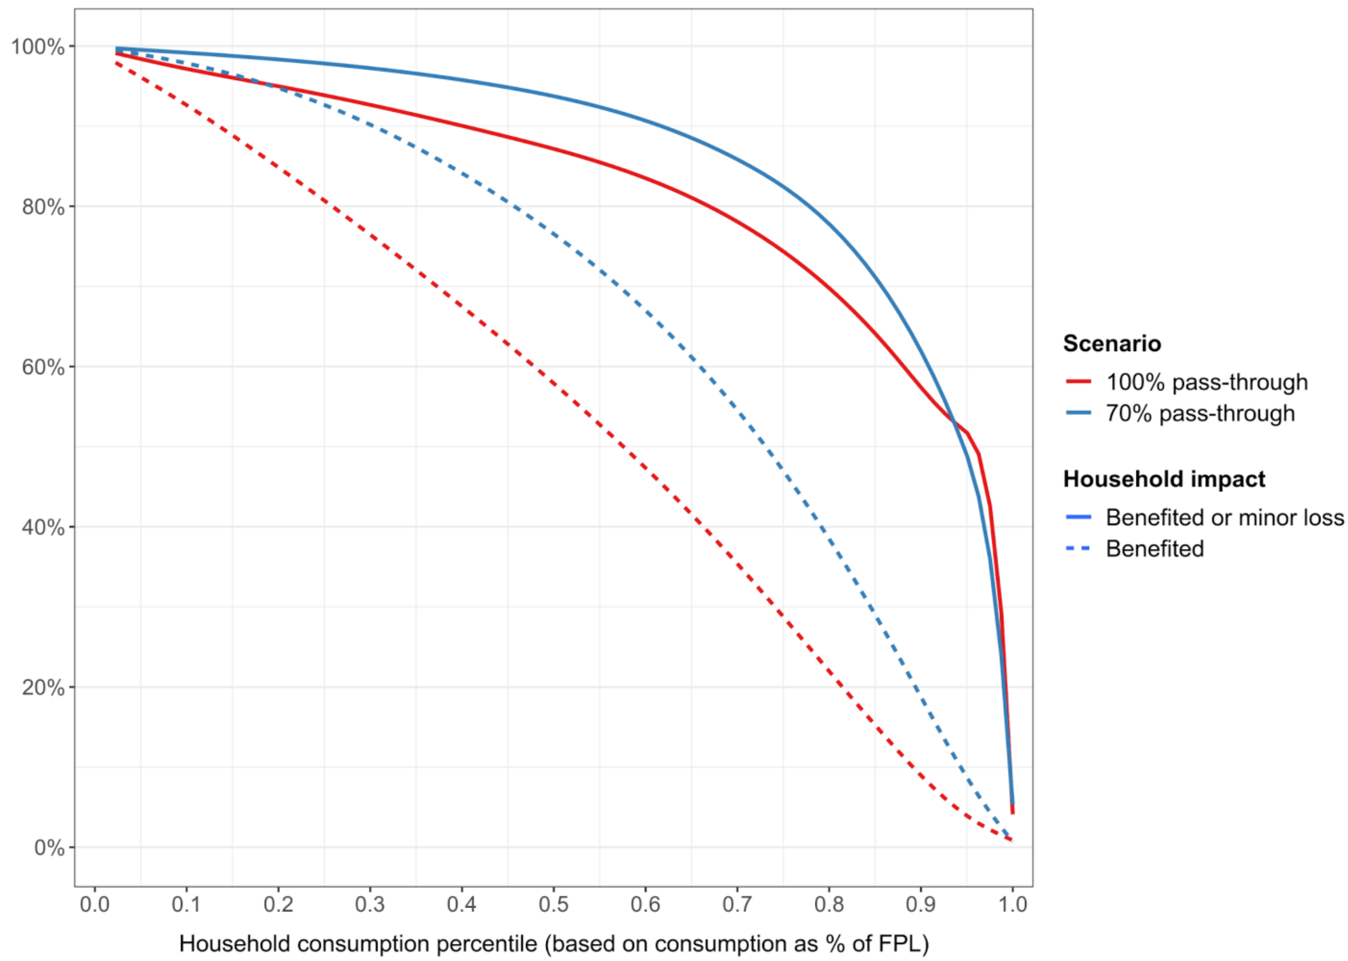

Both scenarios are progressive in the sense that poorer households are benefited more than richer households. Scenario 2 (70% pass-through) is more progressive, typically benefitting 10-15% more households than Scenario 1 through most of the income distribution. This is as expected: Scenario 2 assumes that 30% of the tax burden falls on owners of capital, which is highly-concentrated among richer (high-consumption) households. The percent of households either benefited or incurring a minor loss is relatively stable and high (>80%) through about the 65th percentile in both scenarios.

Figure 2 shows the net dollar amounts flowing in or out of each consumption decile. In Scenario 1, for example, the lowest-consuming 10% of households receive a net benefit of $3.3 billion while the highest-consuming 10% lose a total of $9.6 billion. The difference between the scenarios is particularly apparent in Scenario 2's large negative impact in the richest decile, reflecting the concentration of investable assets among those households. The sum of net impacts across all deciles is actually slightly negative, reflecting a net transfer from households to government due to the total tax burden exceeding after-tax dividends.[4]

Figure 2: Net financial impact by consumption decile

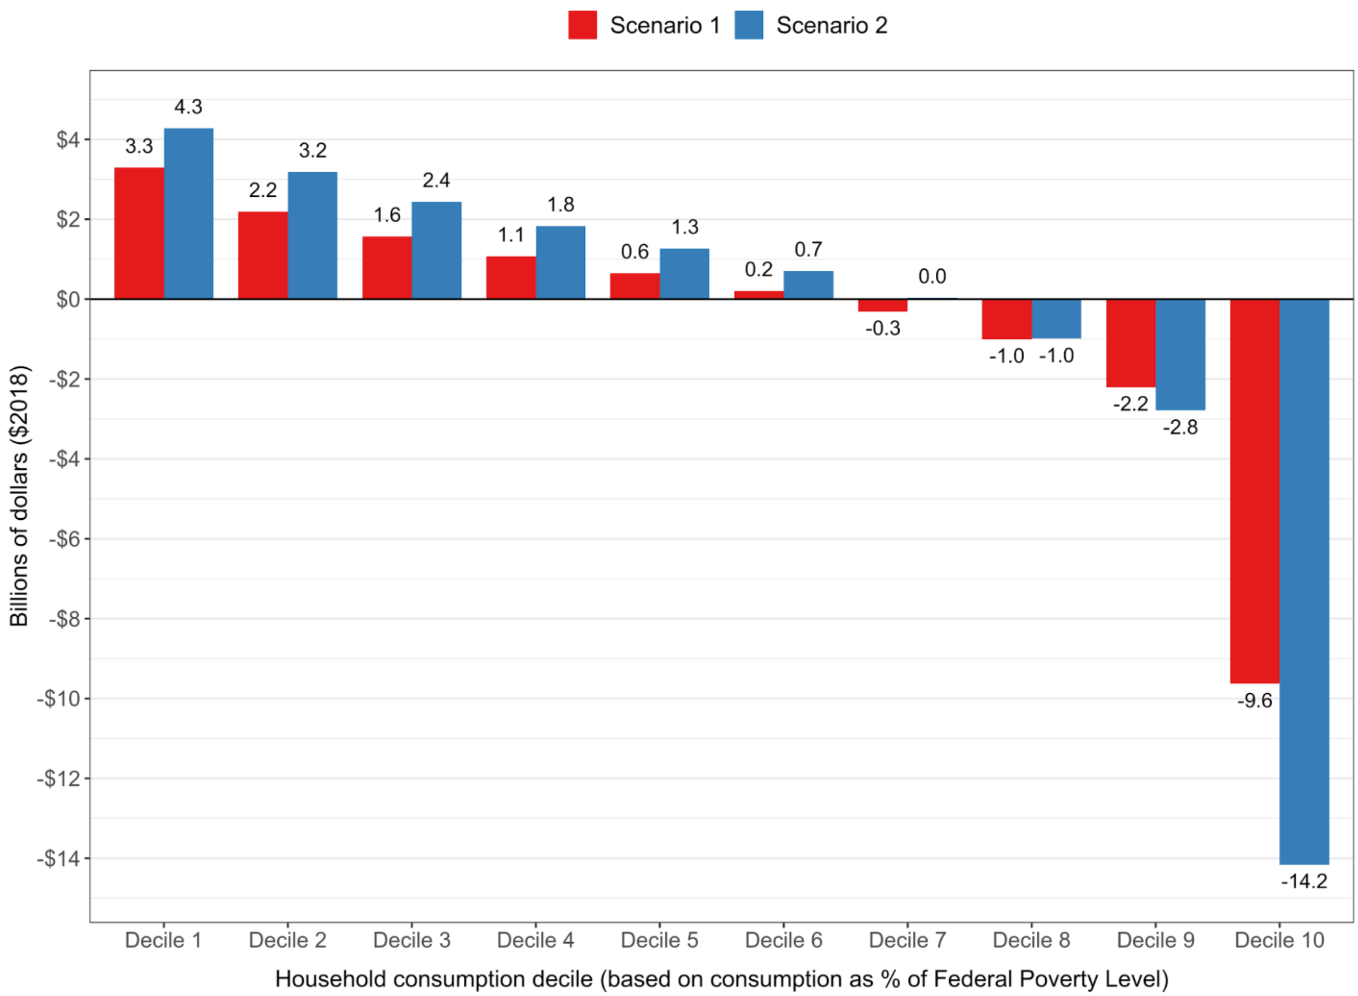

Figures 3 and 4 re-package the information in Figure 1 to display results by consumption quintile for Scenarios 1 and 2, respectively. The table at the bottom of each figure reports typical impacts in absolute dollar amounts and relative to household income. In Scenario 1, for example, 54% of all households are benefitted and another 28% incur a minor loss. Among households benefited, the typical (median) gain is $118 per household or 0.25% of income. The average net impact across all households is slightly negative (-$31), again reflecting the net transfer from households to government.

---

[4] This is for two reasons. 1) CCL assumes the policy incurs administrative costs of about $5 billion. 2) A portion of the dividend (~2%) goes to individuals living in non-correctional group quarters (e.g. nursing homes, dormitories) that are not part of the household sample analyzed here. Neither of these factors were considered in HIS1.

Figure 3: Impact by consumption quintile for Scenario 1 (100% pass-through)

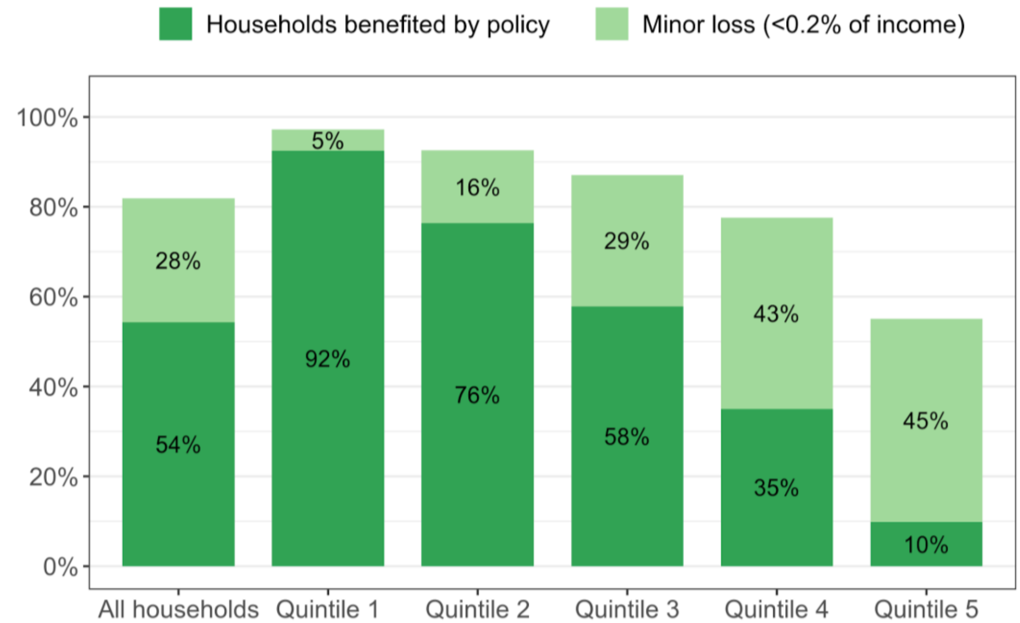

	All households	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Mean net benefit (% of avg. income)	$-31 (-0.04%)	$204 (0.54%)	$98 (0.17%)	$31 (0.04%)	$-49 (-0.05%)	$-441 (-0.24%)
Median gain (% of HH income)	$118 (0.25%)	$178 (0.61%)	$119 (0.24%)	$91 (0.15%)	$72 (0.10%)	$59 (0.06%)
Median loss (% of HH income)	$-118 (-0.15%)	$-34 (-0.13%)	$-46 (-0.11%)	$-66 (-0.12%)	$-103 (-0.14%)	$-261 (-0.20%)
Median HH income % of FPL	354%	154%	268%	361%	476%	815%

Figure 4: Impact by consumption quintile for Scenario 2 (70% pass-through)

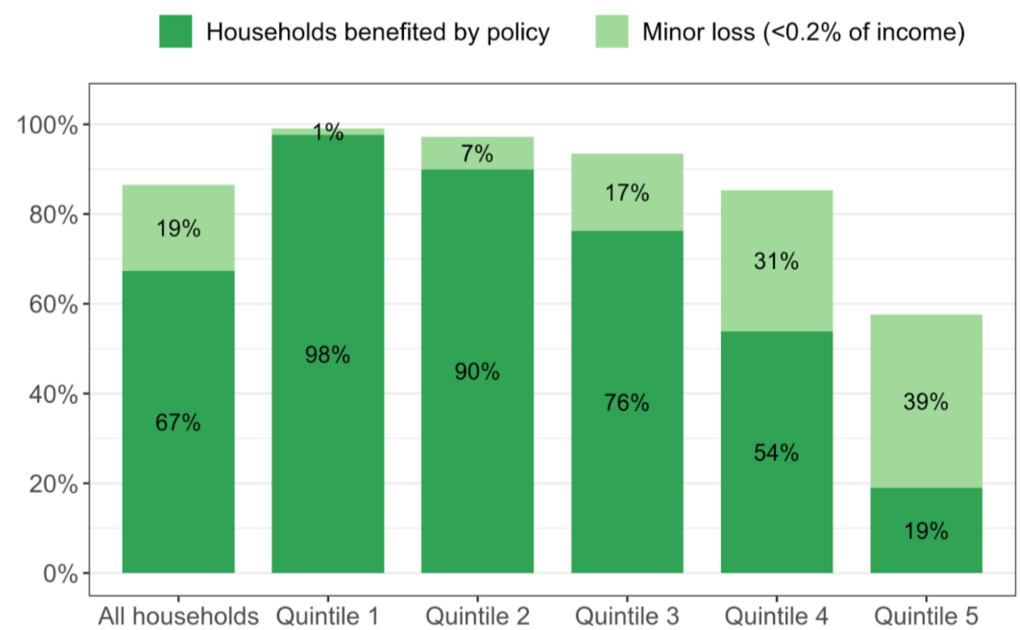

	All households	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Mean net benefit (% of avg. income)	$-31 (-0.04%)	$278 (0.74%)	$159 (0.28%)	$73 (0.10%)	$-36 (-0.04%)	$-631 (-0.34%)
Median gain (% of HH income)	$152 (0.33%)	$250 (0.84%)	$169 (0.36%)	$126 (0.22%)	$96 (0.14%)	$76 (0.08%)
Median loss (% of HH income)	$-142 (-0.16%)	$-37 (-0.13%)	$-45 (-0.10%)	$-62 (-0.10%)	$-96 (-0.12%)	$-282 (-0.21%)
Median HH income % of FPL	354%	154%	268%	361%	476%	815%

Impact by household type

The richness of the microsimulation data allow impacts to be estimated for almost any sub-population. Figures 5 and 6 provide results for different types of households. “Poverty” and “Low income” refer to households with income below 100% and 200% of FPL, respectively. “Older adults” refers to households with no children, no more than two adults, and an average age of at least 65.

The vast majority of households in poverty are benefited in both scenarios (83% and 95%), with the median gain for such households ranging from $133 to $166 per household (1.5-2% of income). A similar pattern is observed, though less pronounced, for minority and single parent households. Scenario 2 confers greater net benefits for all household types presented here – typically increasing the percent benefited by 10-15% over Scenario 1.

The most apparent change from HIS1 concerns older adult and “Family of four” households in Scenario 1, where the percent benefited drops in the former and increases in the latter by about 10%. CCL’s original (HIS1) proposal capped the number of minors receiving dividends at two per household; H.R. 763 has no such cap. This change benefits larger households (families) at the expense of smaller households that tend to be older. In addition, recent evidence suggests that older households may under-report consumption more than younger households. Correcting for this increases the relative tax burden of older households (see Annex B for details).

Figure 5: Impact by household type for Scenario 1 (100% pass-through)

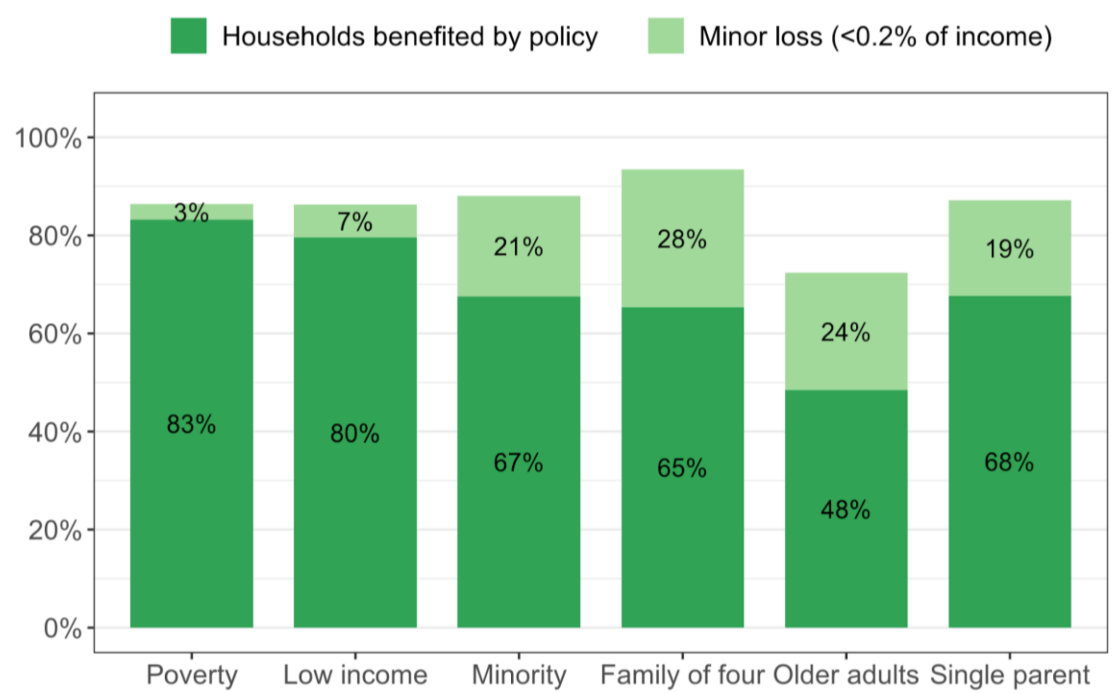

	Poverty	Low income	Minority	Family of four	Older adults	Single parent
Mean net benefit (% of avg. income)	$147 (1.39%)	$135 (0.64%)	$65 (0.09%)	$12 (0.01%)	$-99 (-0.15%)	$38 (0.08%)
Median gain (% of HH income)	$133 (1.49%)	$139 (0.76%)	$149 (0.35%)	$164 (0.21%)	$89 (0.28%)	$123 (0.50%)
Median loss (% of HH income)	$-47 (-0.66%)	$-53 (-0.34%)	$-96 (-0.14%)	$-146 (-0.09%)	$-122 (-0.22%)	$-94 (-0.16%)
Median HH income % of FPL	67%	121%	278%	389%	300%	171%

Figure 6: Impact by household type for Scenario 2 (70% pass-through)

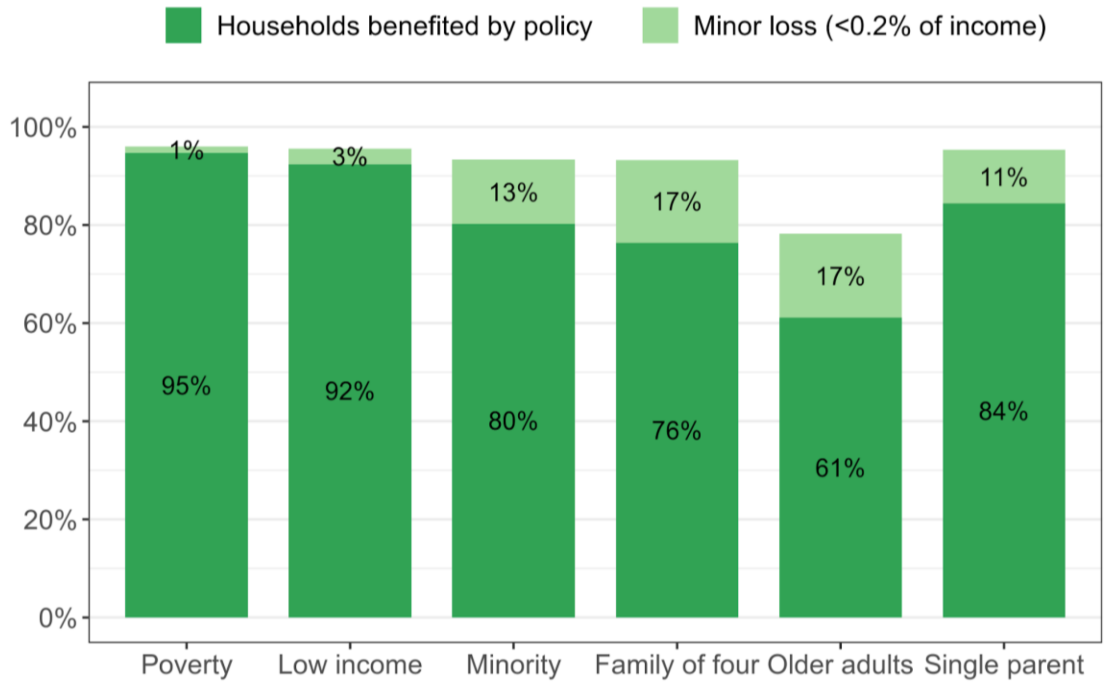

	Poverty	Low income	Minority	Family of four	Older adults	Single parent
Mean net benefit (% of avg. income)	$223 (2.11%)	$214 (1.02%)	$109 (0.15%)	$27 (0.02%)	$-143 (-0.21%)	$117 (0.25%)
Median gain (% of HH income)	$166 (1.96%)	$178 (1.01%)	$194 (0.46%)	$246 (0.31%)	$116 (0.38%)	$179 (0.64%)
Median loss (% of HH income)	$-37 (-0.54%)	$-42 (-0.27%)	$-108 (-0.13%)	$-215 (-0.11%)	$-158 (-0.24%)	$-95 (-0.12%)
Median HH income % of FPL	67%	121%	278%	389%	300%	171%

Impact by age group

Figures 7 and 8 report impacts by age group – based on the age of the household's primary earner – for Scenarios 1 and 2, respectively. Both scenarios show evidence of a U-shaped trend with age consistent with expected life cycle changes. Older households (on average) have less carbon-intensive lifestyles. Younger households tend to be larger – and disproportionately benefited by the dividend allotment mechanism – in addition to having less income/consumption in early career. Households in the "50 to 65" group, on the other hand, have higher income/consumption and, as children age and move out, receive fewer offsetting dividend payments.

That said, differences across age groups are muted and shifted relative to HIS1: younger households are comparatively better off in HIS2 and older households worse off. Again, this change is consistent with the idea that both the H.R. 763 dividend formula and higher consumption among the elderly than previously thought lead to a moderate shift of policy benefits from the old to the young.

Figure 7: Impact by age group for Scenario 1 (100% pass-through)

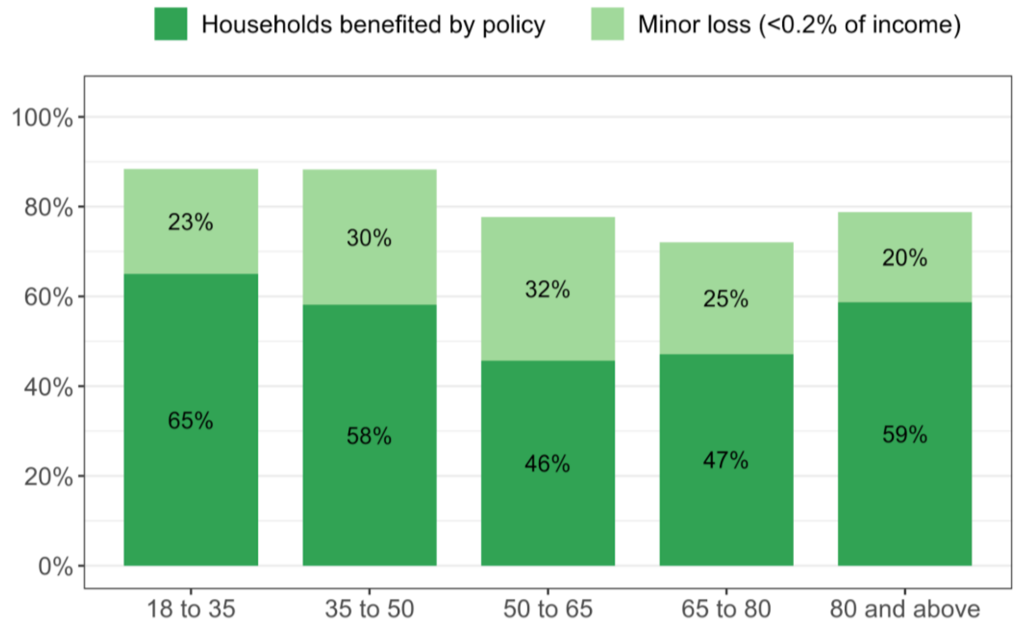

	18 to 35	35 to 50	50 to 65	65 to 80	80 and above
Mean net benefit (% of avg. income)	$61 (0.09%)	$15 (0.01%)	$-98 (-0.09%)	$-117 (-0.15%)	$-18 (-0.03%)
Median gain (% of HH income)	$123 (0.29%)	$143 (0.23%)	$109 (0.22%)	$97 (0.27%)	$100 (0.36%)
Median loss (% of HH income)	$-79 (-0.13%)	$-114 (-0.11%)	$-146 (-0.16%)	$-136 (-0.21%)	$-98 (-0.21%)
Median HH income % of FPL	288%	386%	433%	331%	237%

Figure 8: Impact by age group for Scenario 2 (70% pass-through)

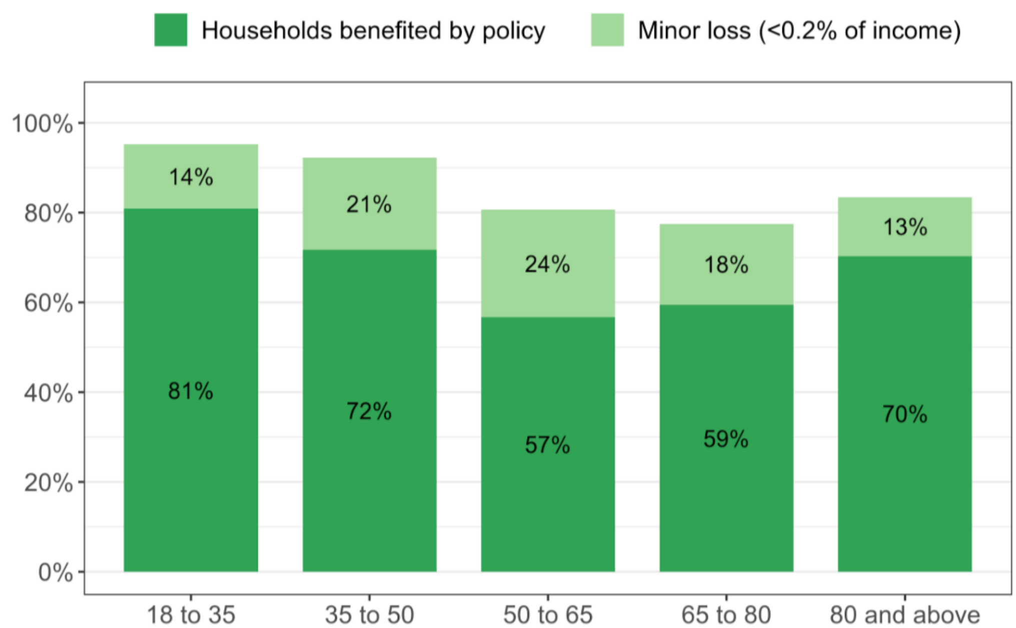

	18 to 35	35 to 50	50 to 65	65 to 80	80 and above
Mean net benefit (% of avg. income)	$123 (0.19%)	$46 (0.04%)	$-142 (-0.14%)	$-155 (-0.20%)	$-63 (-0.12%)
Median gain (% of HH income)	$161 (0.37%)	$192 (0.29%)	$138 (0.29%)	$126 (0.36%)	$127 (0.47%)
Median loss (% of HH income)	$-68 (-0.10%)	$-125 (-0.11%)	$-180 (-0.17%)	$-176 (-0.23%)	$-133 (-0.24%)
Median HH income % of FPL	288%	386%	433%	331%	237%

Impact by race

Figures 9 and 10 report impact by race, based on the self-identified race of the household's primary earner. The percent of households either benefited or incurring a minor loss is relatively equal across racial groups; about ~80-85% in Scenario 1 and increasing slightly in Scenario 2. All groups other than Asian households see a sizable increase (>10%) in percent benefitted moving from Scenario 1 to 2. The comparatively muted increase for Asian households presumably reflects higher rates of saving and business ownership, leading to greater exposure to the *de facto* tax on investable assets in Scenario 2.

Group differences in previous sections largely correlated with underlying differences in income and consumption; i.e. poorer households are benefited more than richer households. The results by race muddy this apparent relationship. For example, Latino households are benefited more than Black households – especially in Scenario 1 – despite similar income levels. Asian households are benefited more than White households despite the former, on average, having higher income. Part of the explanation is that groups with larger households (e.g. Latino) receive larger dividend payments. But an additional factor is the role of geography: Latino and Asian households are more likely to be located in coastal states that benefit from the policy for reasons explored in the following section.

Figure 9: Impact by race for Scenario 1 (100% pass-through)

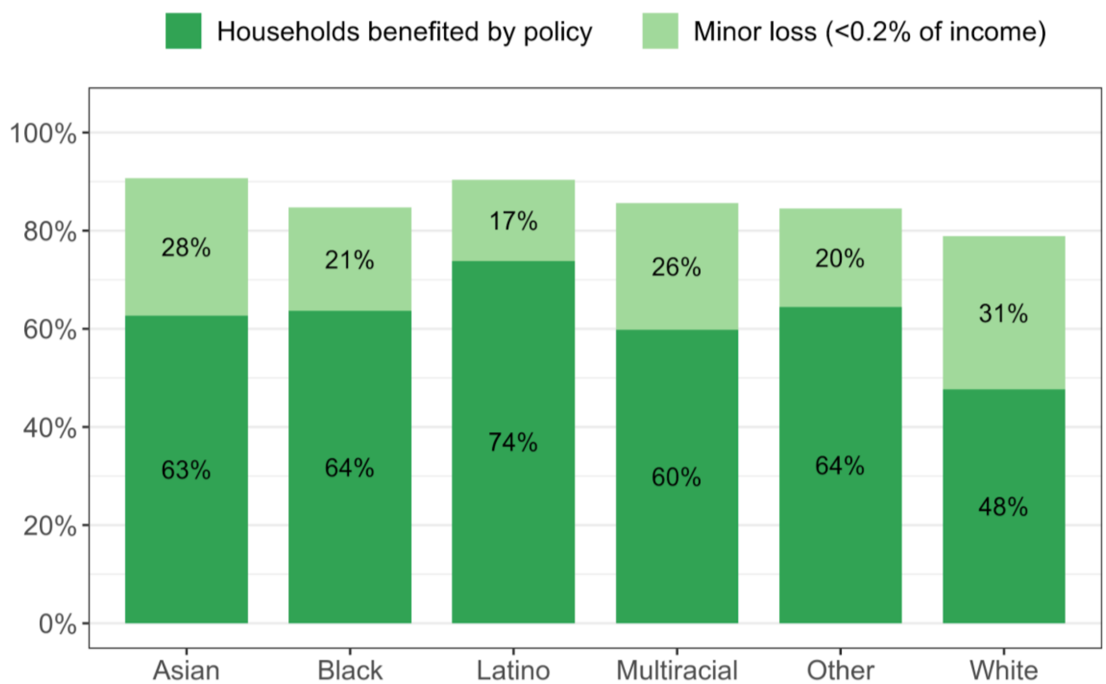

	Asian	Black	Latino	Multiracial	Other	White
Mean net benefit (% of avg. income)	$4 (0.00%)	$41 (0.07%)	$117 (0.17%)	$12 (0.01%)	$51 (0.08%)	$-80 (-0.08%)
Median gain (% of HH income)	$147 (0.24%)	$121 (0.35%)	$182 (0.41%)	$118 (0.26%)	$146 (0.36%)	$101 (0.20%)
Median loss (% of HH income)	$-122 (-0.10%)	$-91 (-0.16%)	$-90 (-0.14%)	$-102 (-0.14%)	$-102 (-0.17%)	$-126 (-0.16%)
Median HH income % of FPL	456%	255%	249%	331%	269%	396%

Figure 10: Impact by race for Scenario 2 (70% pass-through)

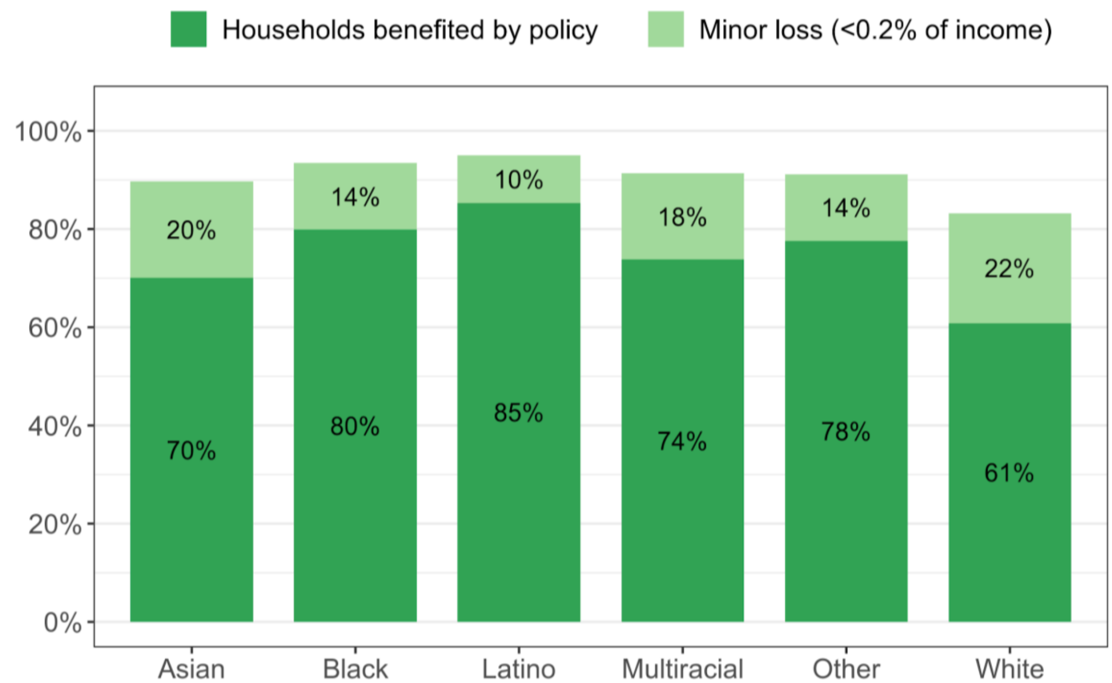

	Asian	Black	Latino	Multiracial	Other	White
Mean net benefit (% of avg. income)	$-56 (-0.05%)	$105 (0.17%)	$182 (0.27%)	$38 (0.05%)	$109 (0.16%)	$-102 (-0.11%)
Median gain (% of HH income)	$182 (0.31%)	$159 (0.45%)	$242 (0.53%)	$153 (0.34%)	$185 (0.46%)	$132 (0.27%)
Median loss (% of HH income)	$-184 (-0.13%)	$-85 (-0.12%)	$-101 (-0.13%)	$-112 (-0.13%)	$-110 (-0.15%)	$-152 (-0.16%)
Median HH income % of FPL	456%	255%	249%	331%	269%	396%

Impact by location

Figures 11 and 12 show the percent of households either benefited or incurring a minor loss, by the type of community a household lives in: rural, suburb/town, or urban. Households in urban areas are slightly more likely to be benefited than rural or suburban areas, but the difference is small (4%) in both scenarios. This differs somewhat from HIS1, which found slightly more variation (7% difference) and is likely due to technical differences. Regardless, the results do not suggest that rural areas (on average) are disproportionately harmed relative to urban areas of the country.

Figure 13 shows the percent of households either benefited or incurring a minor loss, by state and scenario. The x-axis arranges states by median household income. Given the highly-progessive national outcomes observed so far, the results may look surprising. Scenario 1 shows a definite pattern: richer states are benefited more than poor ones. Scenario 2 mitigates the trend to some degree by shifting policy benefits from richer states (with more investable assets) to poorer states.

The overall national results are still highly progressive, because the poorest households in all places tend to do reasonably well under a per-capita-type dividend allocation. In addition, states like California, New York, and Florida have large populations and comparatively low exposure to carbon pricing. These states effectively “pull up” the percent of households benefited nationally. Consequently, national results can mask considerable diversity among states.

Figure 11: Impact by community type for Scenario 1 (100% pass-through)

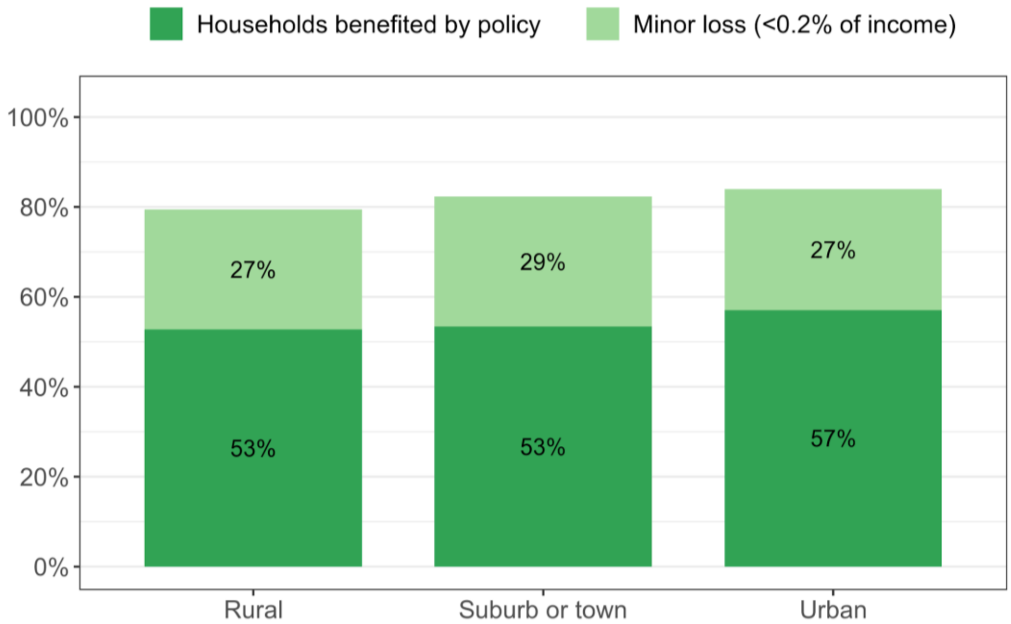

	Rural	Suburb or town	Urban
Mean net benefit (% of avg. income)	$-31 (-0.04%)	$-40 (-0.04%)	$-21 (-0.02%)
Median gain (% of HH income)	$113 (0.25%)	$120 (0.24%)	$119 (0.29%)
Median loss (% of HH income)	$-117 (-0.17%)	$-123 (-0.15%)	$-113 (-0.14%)
Median HH income % of FPL	332%	379%	344%

Figure 12: Impact by community type for Scenario 2 (70% pass-through)

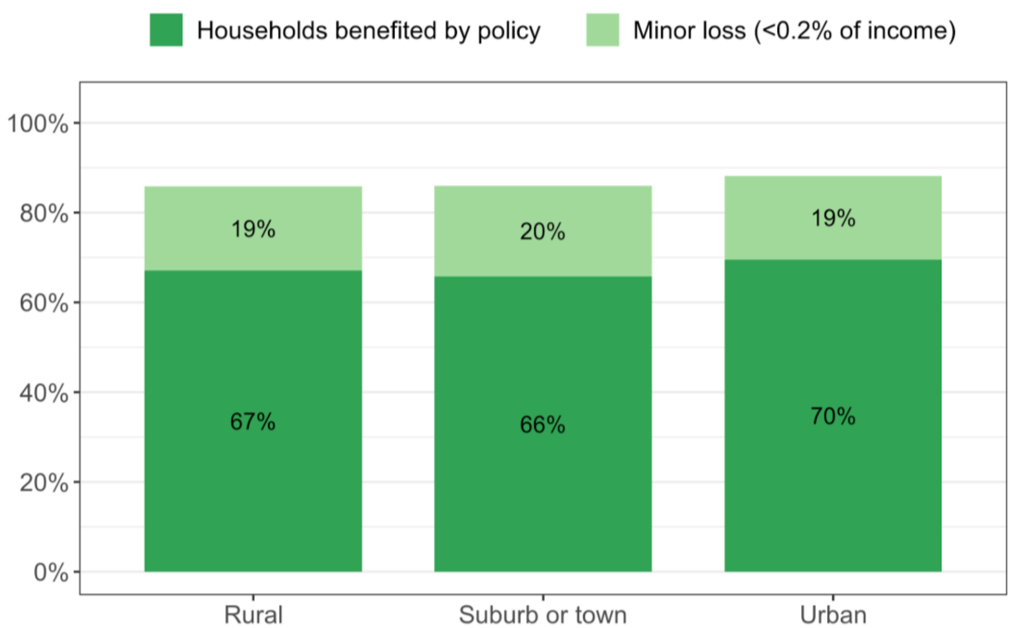

	Rural	Suburb or town	Urban
Mean net benefit (% of avg. income)	$-13 (-0.02%)	$-48 (-0.05%)	$-29 (-0.03%)
Median gain (% of HH income)	$150 (0.33%)	$155 (0.31%)	$151 (0.36%)
Median loss (% of HH income)	$-134 (-0.17%)	$-152 (-0.16%)	$-137 (-0.15%)
Median HH income % of FPL	332%	379%	344%

Figure 13: Impact by state

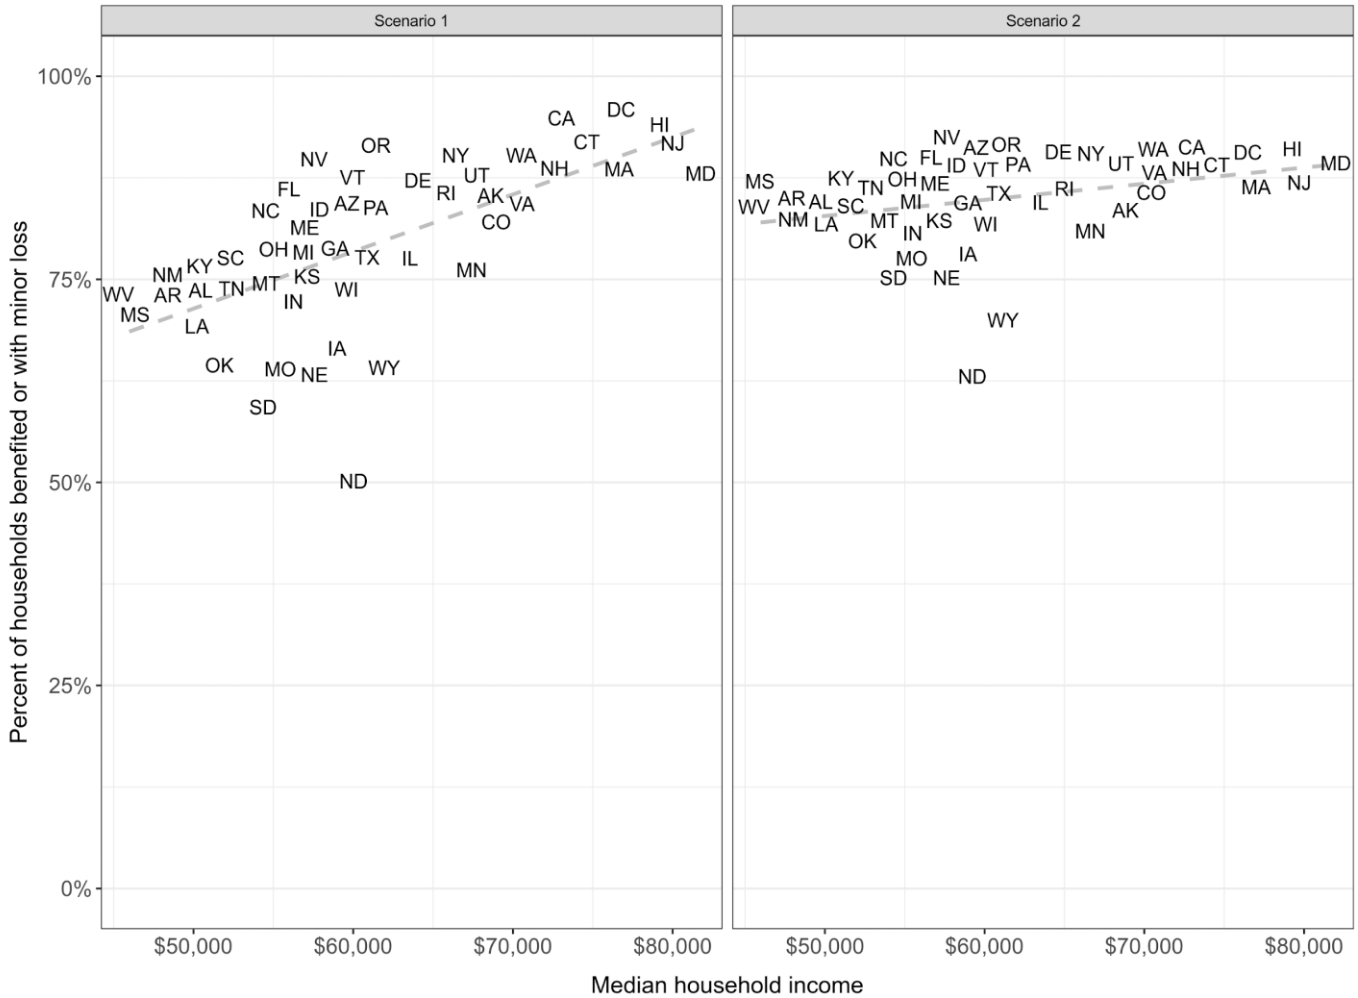

This requires some explanation. Why do poorer states benefit less, especially in Scenario 1? Households in these states have comparatively high carbon footprints – even at lower levels of income – making them more exposed to a carbon tax. High-income states with low footprints typically enjoy some combination of milder climate, cleaner electricity, and/or greater population density. The opposite tends to be true in low-income states.

This is apparent in Figure 14, which shows how the percent of households either benefited or incurring a minor loss varies across the income distribution in seven select states. There is significant variation in impacts at equivalent levels of income. For example, a household at 400% of FPL (roughly the national median) experiences different outcomes depending on location. In California, nearly all such households are benefited or negligibly-affected; in Minnesota about three-quarters are; in North Dakota about half are. Scenario 2 "closes the gap" among states to some degree but significant variation remains.

Figure 14: Impact across income distribution for select states

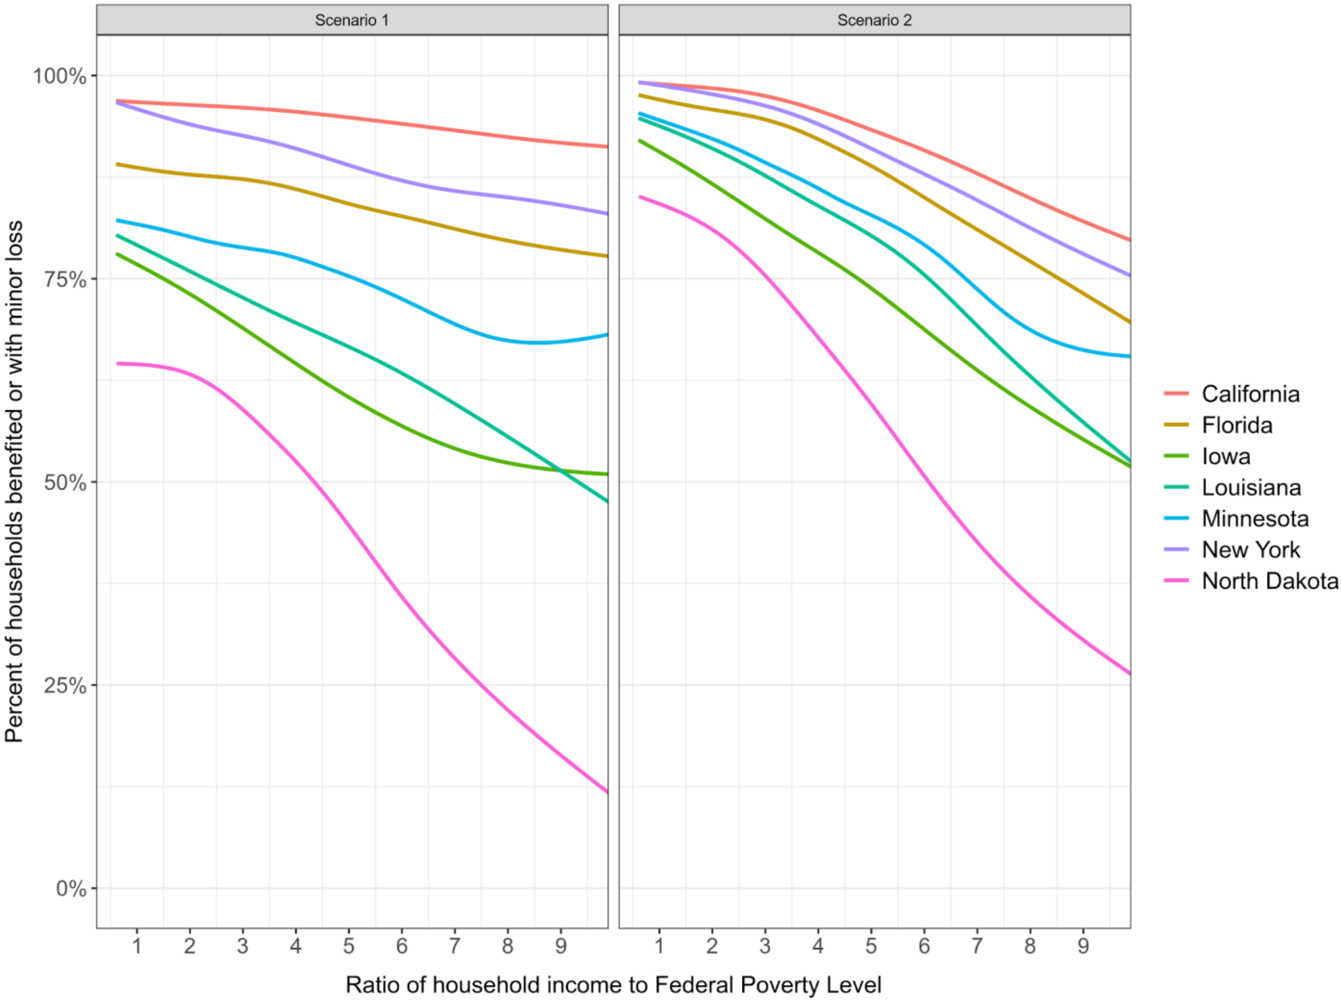

Figure 14 is unique in that it shows how household impacts vary across *both* a household-level characteristic (income) and state of residence (location). National-level results ignore location and effectively "compress" the various curves into one. Social policies often use income relative to the FPL to determine eligibility for public benefits (e.g. Medicaid), ostensibly providing some notion of fairness. Figure 14 suggests that carbon pricing policies may struggle to equalize impacts across households that public policy otherwise deems "equally well-off". This issue is returned to in the conclusion.

Finally, it is possible to estimate impacts for individual congressional districts. Figure 15 shows how the percent of households either benefited or incurring a minor loss varies with the political leaning of each congressional district. The x-axis displays the Cook Partisan Voting Index (PVI), a measure of how strongly Democratic or Republican a district is based on recent election results.

There is a slight trend in Scenario 1 indicating better (worse) outcomes for more left (right) leaning districts. However, the overall mean difference between the parties is not exceptionally large: About 86% of households in Democratic districts are benefited or negligibly-affected; 78% in Republican districts. The difference is nearly erased in Scenario 2.

Figure 15: Impact by congressional district

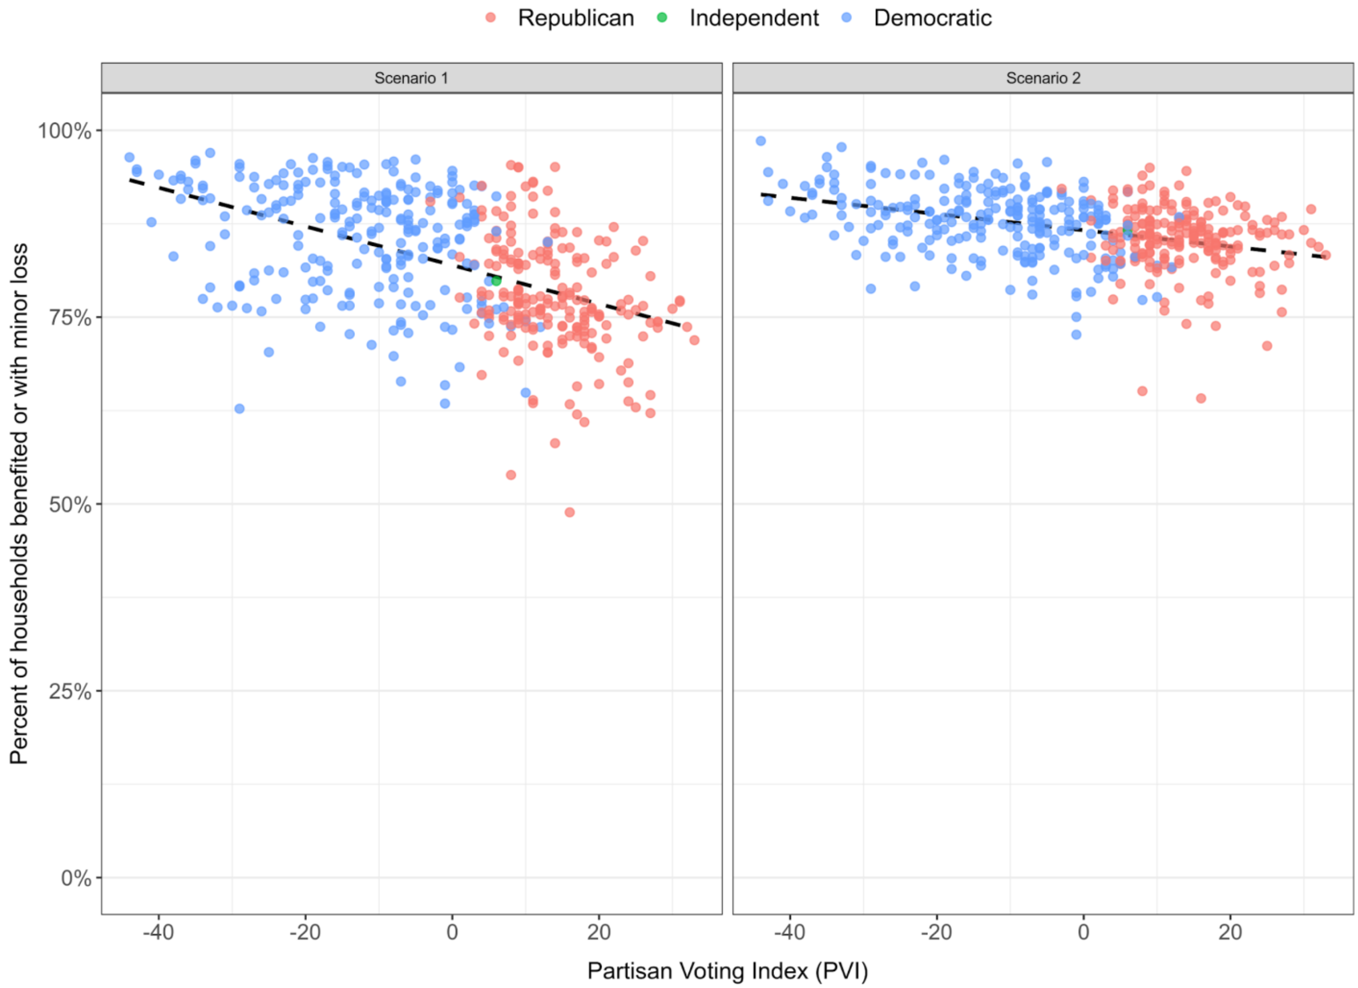

Conclusion

Microsimulation of a “carbon tax and dividend” policy like that envisioned in H.R. 763 indicates that the short-term financial impact for U.S. households is highly progressive. The after-tax dividend payment exceeds the additional tax burden for about 80% of low-income households, assuming full pass-through of the carbon tax into consumer prices. The percent benefited increases to over 90% in a scenario assuming partial (70%) pass-through.

The results are even more progressive when consumption (rather than income) is used to sort households. The percent of households benefited or negligibly-affected exceeds 90% among the lowest-consuming 40% and 60% of households in Scenarios 1 and 2, respectively. To the extent that consumption reflects the relative economic well-being of households, the results imply that few (<10%) of the worse-off half of households are likely to be harmed by the policy – at least when ignoring employment and macroeconomic effects.

These results are consistent with general equilibrium studies that find per-capita-style dividends to be progressive, albeit comparatively inefficient (from a macroeconomic perspective) compared to other revenue recycling options (Caron et al. 2018; Goulder et al. 2019).

Analysis of population subgroups reveals how policy design, demographics, and geography combine to influence household-level impacts. Larger families are benefited more than smaller ones. Younger and older households are benefited more than the middle-aged. Latino households are benefited more than black households, and Asian households see greater net benefits than white households. Multiple phenomena drive these differences, but between-group variation in income, household size, and location are central. That said, differential impacts by age and race are small relative to differences across the income distribution. The policy's progressivity with respect to income (or consumption) is its predominant feature.

A partial exception may apply in the case of geography. Impacts vary across for both states and congressional districts. In general, the comparatively mild-climate, "blue" coastal states benefit more under a policy like H.R. 763 than "red" states in the south and interior of the continent – particularly in Scenario 1. Climate and attendant heating and cooling needs is an important factor in explaining this pattern, as are differences in population density and the carbon-intensity of the electricity supply.

Spatial variation in household exposure to carbon pricing is a potential challenge, both politically and in terms of fairness. It is not immediately obvious if equally well-off families in North Dakota and California should face significantly different financial impacts – or if such differences are politically tenable.[5] However, this is a challenge for *any* carbon pricing policy. A dividend-based policy is more easily tailored to address geographic concerns than other revenue recycling options, but taking advantage of this potential would require moving away from the simplicity of per-capita dividends.

New to this study is the inclusion of a partial pass-through scenario (Scenario 2), which is effectively a hybrid consumption and capital tax. Scenario 2 may constitute something like a "best case" scenario from a progressivity standpoint (though not in terms of macroeconomic efficiency). The microsimulation results reflect this, as Scenario 2 consistently increases the share of households benefited by 10-15% over Scenario 1.

The problem, of course, is that policymakers cannot control how the tax will be allocated among consumers, labor, and capital. Those decisions will be made by businesses. As noted earlier, empirical research on this topic is hazy, suggesting that the allocation of the tax burden is not merely uncontrollable but also unpredictable. Ultimately, how the dividend is allocated among households (i.e. the "revenue recycling" strategy) is the principal lever that can be pulled to shape the policy's distributional effects.

As mentioned earlier, this study does not include automatic (statutory) increases in transfer payments like social security and food stamps that result from an increase in the general price level. If revenue neutrality is enforced, then inclusion of statutory effects shifts a portion of the

[5] One could argue that California residents are entitled to some initial advantage due to early adoption of GHG mitigation policies that other states chose not to pursue. But such logic surely does not extend to inequities that result from geographic variation in climate (i.e. heating and cooling needs).

carbon tax revenue to households receiving transfers and reduces the size of per-capita dividend payments. The available evidence suggests that inclusion would not significantly alter the overall progressivity of the policy, since poorer households benefit disproportionately from transfer payments. But impacts would be altered for certain types of households. For example, older households would be benefited more (and younger households less) than suggested in this study, as a result of automatic increases in social security benefits.

A final complication concerns taxation of the dividend payment. Consistent with H.R. 763, this study subjects each household's dividend to both federal payroll and income taxation. However, the real-world mechanics of this approach are decidedly messy.[6] In principle, the Treasury Department could simply retain a percentage of the gross carbon tax revenue to offset expected additional costs to government – thereby ensuring approximate revenue neutrality – and the remaining funds disbursed to households as a tax-free payment (as was recently done for coronavirus "recovery rebates"). Assuming a per-capita-style dividend, this approach would result in somewhat less progressive impacts than reported here, since it eliminates progressive taxation of the dividend at the household level.

[6] Payroll and income taxes are typically remitted to the IRS by employers. Subjecting the dividend payment to such taxation would require wage earners to notify employers of the additional income or, for the ~50% of households without wage income, to report the dividend as self-employment income and pay appropriate taxes when returns are filed. In addition to hassle and annoyance (e.g. an unexpected tax bill), this process seems likely to result in a lower effective tax rate on the dividend, threatening the policy's intended revenue neutrality. Since the dividend payments originate from Treasury, it seems unnecessary to burden households and businesses with the task of recouping a portion of the payments.

References

Buuren, S. V., & Groothuis-Oudshoorn, K. (2011). Mice: Multivariate Imputation by Chained Equations in R. *Journal of Statistical Software, 45*(3). doi:10.18637/jss.v045.i03

Caron, J., Metcalf, G. E., & Reilly, J. (2017). The CO_2 Content of Consumption Across U.S. Regions: A Multi-Regional Input-Output (MRIO) Approach. *The Energy Journal, 38*(1). doi:10.5547/01956574.38.1.jcar

Caron, J., Cohen, S. M., Brown, M., & Reilly, J. M. (2018). Exploring The Impacts Of A National U.S. CO_2 Tax And Revenue Recycling Options With A Coupled Electricity-Economy Model. *Climate Change Economics, 09*(01), 1840015. doi:10.1142/s2010007818400158

Cook, J.D. (2010). Determining distribution parameters from quantiles. Working paper. Houston, TX. https://www.johndcook.com/quantiles_parameters.pdf

Cronin, J. A., Fullerton, D., & Sexton, S. E. (2019). Vertical and Horizontal Redistributions from a Carbon Tax and Rebate. Journal of the Association of Environmental and Resource Economists, 6(S1), S169-S208. doi:10.1086/701191

Feiveson, L., & Sabelhaus, J. (2019). Lifecycle Patterns of Saving and Wealth Accumulation. *Finance and Economics Discussion Series, 2019*(010). doi:10.17016/feds.2019.010

Ganapati, S., Shapiro, J. S., & Walker, R. (2020). Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes. *American Economic Journal: Applied Economics, 12*(2), 303-342. doi:10.1257/app.20180474

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. doi:10.1016/j.jpubeco.2019.04.002

Graetz, N., Ummel, K., & Cohen, D. A. (2020). Small-Area Analyses Using Public American Community Survey Data: A Tree-Based Spatial Microsimulation Technique. *SSRN Electronic Journal*. doi:10.2139/ssrn.3574679

Hannagan, A., & Morduch, J. (2015). Income Gains and Month-to-Month Income Volatility: Household Evidence from the Us Financial Diaries. *SSRN Electronic Journal*. doi:10.2139/ssrn.2659883

Ho, M. S., Morgenstern, R., & Shih, J. S. (2008). Impact of Carbon Price Policies on U.S. Industry. *SSRN Electronic Journal*. doi:10.2139/ssrn.1320201

Iman, R. L., & Conover, W. J. (1982). A distribution-free approach to inducing rank correlation among input variables. *Communications in Statistics - Simulation and Computation, 11*(3), 311-334. doi:10.1080/03610918208812265

Institute of Medicine (2005). *Dietary reference intakes: For energy carbohydrate, fiber, fat, fatty acids, cholesterol, protein, and amino acids*. National Academy Press.

Jorda, V., Sarabia, J.M., & Jantii, M. (2018). Estimation of income inequality from grouped data. arXiv: 1808.09831

Kaufman, N., Larsen, J., Marsters, P., Kolus, H., & Mohan, S. (2019). An Assessment of the Energy Innovation and Carbon Dividend Act. Center on Global Energy Policy, Columbia University. https://energypolicy.columbia.edu/research/report/assessment-energy-innovation-and-carbon-dividend-act

McDonald, J. B., & Ransom, M. (2008). The Generalized Beta Distribution as a Model for the Distribution of Income: Estimation of Related Measures of Inequality. *Modeling Income Distributions and Lorenz Curves,* 147-166. doi:10.1007/978-0-387-72796-7_8

Millard, S. P. (2013). *EnvStats an R package for environmental statistics*. Springer.

Neuhoff, K. & Ritz, R. (2019). "Carbon cost pass-through in industrial sectors," Cambridge Working Papers in Economics 1988, Faculty of Economics, University of Cambridge.

Pya, N., & Wood, S. N. (2014). Shape constrained additive models. *Statistics and Computing, 25*(3), 543-559. doi:10.1007/s11222-013-9448-7

Ummel, K. (2014). Who Pollutes? A Household-Level Database of America's Greenhouse Gas Footprint. *SSRN Electronic Journal*. doi:10.2139/ssrn.2622751

Ummel, K. (2016). Impact of CCL’s proposed carbon fee and dividend policy: A high-resolution analysis of the financial effect on U.S. households. Working Paper produced for Citizens’ Climate Lobby. https://citizensclimatelobby.org/household-impact-study/

Weisbach, D. A., & Metcalf, G. E. (2009). The Design of a Carbon Tax. *SSRN Electronic Journal*. doi:10.2139/ssrn.1327260

Whitworth, A., Carter, E., Ballas, D., & Moon, G. (2017). Estimating uncertainty in spatial microsimulation approaches to small area estimation: A new approach to solving an old problem. *Computers, Environment and Urban Systems, 63*, 50-57. doi:10.1016/j.compenvurbsys.2016.06.004

Yang, Y., Ingwersen, W. W., Hawkins, T. R., Srocka, M., & Meyer, D. E. (2017). USEEIO: A new and transparent United States environmentally-extended input-output model. *Journal of Cleaner Production, 158*, 308-318. doi:10.1016/j.jclepro.2017.04.150

Annex A: Modeling of household consumption

The core household survey inputs are the American Community Survey Public Use Microdata Sample (PUMS) and the Consumer Expenditure Survey (CE). The PUMS solicits demographic, income, employment, and housing expenditure information; the single-year (2018) sample contains 1.3 million households and is representative for individual states and Public Use Microdata Areas (PUMA's).

The CE solicits detailed household expenditures along with demographics, income, and employment; pooling interview survey waves for 2015-2018 yields annual expenditures for ~35,000 households after imputation.[7] The CE's comparatively small sample size means it is only representative for large Census regions. Household-level expenditures are aggregated into 57 custom categories (Table A1) that allow good alignment with NIPA PCE categories.

The CE collects expenditure data for individual "consumer units", while the PUMS solicits information for the "household". In most cases, the two concepts are identical. The exceptions are largely cases where unrelated roommates are economically independent. To better align the two data sources, PUMS households consisting of unrelated roommates are disaggregated into individual consumer units. PUMS housing expenditure variables are allocated equally across resulting consumer units. The term "household" is used throughout for convenience, but the analysis is actually performed on consumer units.

Ummel (2014) introduced a technique to simulate CE expenditure variables for PUMS households. In general, this can be thought of as "fusing" variables from a "donor" survey onto "recipient" microdata. Using the large-sample PUMS as the recipient allows for analysis of fused variables at levels of geo-demographic detail typically impossible in the donor. The fused variables are fully "synthetic" in the sense that they are randomly drawn from modeled conditional distributions for each recipient household (details below). Conventional "sample matching" techniques are problematic when the donor sample is much smaller than the recipient, as is the case in almost any situation involving the PUMS.[8]

Survey fusion requires that the donor and recipient datasets be "aligned" to create a set of identically-defined predictor variables observable in both; these variables provide the statistical "link" between donor and recipient households that the modeling process can exploit. Given the extensive subject matter overlap between the CE and PUMS, the set of aligned variables is quite large. In addition, household location is partially/approximately known for both the CE and PUMS, allowing additional "spatial predictors" to be added to both the donor and recipient.

The goal of the household consumption modeling is to simulate detailed expenditures for PUMS households such that the univariate distributions and correlation structure of the original CE

[7] Multivariate imputation via predictive mean matching (van Buuren & Groothuis-Oudshoorn 2011) is used to generate complete annual expenditure profiles for all respondent households, regardless of the length of time the household stays in the sampling frame. This addresses a known problem when using CE data for cross-sectional analysis, as a result of higher sample drop-out among poorer households.

[8] One advantage of data fusion over sample matching is the former's ability to simulate rare "tail" behavior. While sample matching necessarily replicates donor observations (perhaps many times), data fusion generates a greater variety of outcomes in the recipient -- including outcomes drawn from the tails of the conditional distributions. This is particularly relevant for consumption modeling, since high-consumption households are under-represented in surveys like the CE.

data are generally replicated. This is a non-trivial task given the relatively large number of expenditure categories. The technique described below builds upon that introduced by Ummel (2014).

Given a CE-derived "donor" dataset with *N* household observations, non-negative expenditures for each of *P* categories, and a set of household-level predictors *X* shared with a PUMS-derived "recipient" dataset, the algorithm proceeds as follows:

1) Logistic regression is used to model the probability of non-zero (binary) expenditures as a function of *X*. Models are fit sequentially (*P* total models), allowing binarized variables to become predictors in subsequent models. This "chaining" of the models induces more realistic correlation among simulated outcomes. Since the potential predictors in *X* are numerous and sometimes highly correlated, LASSO regression is used to identify a suitable set.

2) Gradient boosting machine (GBM) regression is used to model the distribution of non-zero expenditures. A mean-response model is fit along with quantile models for the conditional median, 15th percentile, and 85th percentile (4*P* total models). Models are again chained; fast testing of linear models is used to determine a pseudo-optimal order. LASSO regression is used to reduce the set of potential predictors available to the GBM models. The preferred interaction depth and number of trees is determined via out-of-sample testing on random validation subsets.

3) The GBM models are used to predict conditional means and quantiles for the original donor observations, to which parametric conditional distributions are then fit. This entails *NP* distributions, making explicit optimization prohibitive. Instead, I leverage the fact that plausible candidate parameters can be quickly calculated for the location-scale family of distributions (Cook 2010).[9] The candidate distribution that best matches the conditional quantiles is selected.

4) Each non-zero expenditure observation is converted to a conditional percentile using the conditional distribution derived in (3) and then a rank correlation matrix is calculated across all households. This matrix reveals the extent to which expenditure categories are *conditionally* correlated (i.e. correlation that the chained models do not explicitly capture).[10]

5) A binary expenditure matrix is simulated for the recipient dataset using the models from (1).

[9] The algorithm includes the normal, logistic, log-normal, Weibull, and gamma distributions, which collectively provide a wide range of distribution shapes.
[10] Pairwise conditional correlations are generally low (mean absolute value = 0.03), as expected if the chained models capture most of the correlation structure. However, moderate correlations are observed for expenditure categories that are particularly linked within households. For example, "Jewelry and handbags" category exhibits conditional correlation of 0.13 and 0.14 with "Clothing and footwear" and "Food away from home", respectively.

6) A matrix of correlated random percentiles is generated for non-zero observations simulated in (5). The matrix uses the conditional rank correlation matrix from (4) and the algorithm of Iman and Conover (1982) as implemented by Millard (2013).

7) Step 3 is performed using the recipient dataset, and nonzero expenditure values are drawn from the resulting conditional distributions using the random percentiles from (6).

The resulting dataset contains simulated values for each of the 57 consumption categories and 1.3 million PUMS households.[11] The simulated or "synthetic" consumption variables re-create the basic statistical structure of the donor CE data. Variable means, proportion of nonzero values, the shape of univariate distributions, between-variable correlations, and the shape of non-linear relationships with, say, household income are all reasonable (though not exact) facsimiles of the original CE data. In effect, the fusion process provides estimates of how we expect PUMS respondent households might have responded had they completed the CE survey questionnaire.

This begs the question: Why not just use the pooled CE data as-is? If the CE was a large sample representative at the level of individual states or (ideally) even smaller geographic entities, fusion would have little additional value. Absent that, fusion to the PUMS allows subsequent analysis to operate with a degree of spatial resolution otherwise impossible, since PUMS households are identified at the level of individual PUMA's.[12]

A known problem with CE data (and, hence, the PUMS synthetic variables) is significant under-reporting of household spending. There is good reason to believe that households (in aggregate) self-report only about 75% of actual consumption. The degree of underreporting varies across types of spending and, presumably, types of households. HIS1 (Ummel 2016) provides a thorough explanation of the problem and attempts to mitigate it by scaling CE-based consumption estimates to match known national totals for similar categories in the NIPA Personal Consumption Expenditures (PCE) tables.

This approach is taken further here by incorporating two additional sources of data. First, the BEA now publishes state-level PCE tables that reflect inter-state variation in consumption patterns, albeit with less categorical resolution than in the national tables. Second, a recent study by Feiveson and Sabelhaus (2019; hereafter "FS") provides unique information regarding the distribution of household consumption by age and income.

FS employs a "pseudo-panel" approach combining two decades of microdata from the Survey of Consumer Finances (SCF) with aggregate macro data from the Financial Accounts of the United States and the National Product and Income Accounts. The technique allows them to construct estimates of how income, wealth, saving, and consumption vary over the lifecycle by

[11] A handful of the 57 categories (rent, property tax, home insurance, other owner costs, and total mortgage payment) are not simulated, because they are observed directly in the PUMS and are, in fact, part of the joint set of predictor variables *X*.

[12] There are conceptual overlaps here with "spatial microsimulation" techniques, whereby microdata observations (e.g. the CE) are re-weighted to generate representative samples at some smaller geographic scale (e.g. PUMA's) (Graetz, Ummel, and Cohen 2020; Whitworth et al. 2017). The difference is that the fusion process does not perform re-weighting (thereby avoiding observation replication and attendant problems) but rather simulates entirely new values for households in the PUMS.

"agent type". SCF respondents are assigned to agent types on the basis of age and pre-tax income. The latter is used to assign individuals to the "Bottom 50 Percent", "50th to 90th Percentiles", or "Top 10 Percent" income group within their age group. There are six age groups ranging from "Under 35" to "75 and over", yielding a total of 18 income-age "cohorts".

FS deduce mean household consumption by cohort by imposing an intertemporal household budget identity, such that after-tax income less consumption is equal to the change in wealth over some period less capital gains and net inter-family transfers. In short, they "back out" consumption from observed changes in wealth and income over time. This provides novel estimates of how total household consumption varies across the life cycle, while avoiding the use of CE self-reported data.

Consistent with the hypothesis that consumption under-reporting increases with income, FS "estimate consumption for high earners that is well above estimates based on household spending data such as the Consumer Expenditure Survey". However, they note that their results likely underestimate consumption of younger cohorts and overestimate consumption of older cohorts, leading to exaggerated consumption inequality across age cohorts.

The exact nature of this bias is not known. However, it is possible to compare FS cohort consumption shares with analogous figures derived from self-reported CE data. The comparison indicates that self-reported consumption *exceeds* FS consumption among the young, especially those in the top-10% of the income distribution. This is consistent with the idea that the FS results are "missing" consumption among younger households (since we don't expect households to *over*-report actual consumption). Shifting 4-5% of total consumption from older households to the youngest households produces an arbitrary but more reasonable pattern of implied CE under-reporting – i.e. one consistent with the idea that households do not generally over-report consumption and that under-reporting increases with consumption.[13]

The adjusted FS results are used to estimate the full consumption distribution for each age group by fitting a generalized beta distribution of the second kind. The GB2 distribution is a flexible, four-parameter distribution known to provide good fits with income data (McDonald and Ransom 2008; Jorda et al. 2018). The fitted distributions are then used to derive the percentile points in the income distribution that define twenty equal-consumption groups within each age cohort.

[13] FS provides two explanations for why the pseudo-panel approach might underestimate consumption of the young. First, SCF data likely underestimate *inter vivos* transfers from parents to offspring, especially among rich families. Second, young households may enjoy greater capital gains than assumed by SF, possibly due to greater home value appreciation as a result of "sweat equity". The proposed 4-5% reallocation equates to $400-500 billion per year. Such a large sum suggests that both mechanisms are likely involved.

Table A1: Consumption categories (57)

Group	Category	Description
Apparel	CLOFTW	Clothing and footwear
Apparel	JWLBG	Jewelry and handbags
Education	EDUC	Education services
Education	STDINT	Student loan interest
Entertainment	CABLE	Cable, satellite, and streaming services
Entertainment	GAMBL	Gambling and lotteries
Entertainment	HOTEL	Hotels and motels
Entertainment	OEPRD	Other entertainment products
Entertainment	OESRV	Other entertainment services
Entertainment	RVHPRD	Recreational vehicles and products
Entertainment	TKTMEM	Admission fees and memberships
Entertainment	TOBAC	Tobacco and smoking supplies
Entertainment	TVCOMP	Televisions and computers
Food and drink	AAWAY	Alcohol away from home
Food and drink	AHOME	Alcohol at home
Food and drink	FAWAY	Food away from home
Food and drink	FHOME	Food at home
Health care	HCOOP	Out-of-pocket health care
Health care	HLHINS	Health insurance premiums
Household operation	FURAPP	Furniture, appliances, and housewares
Household operation	HSRVCC	Household services and child care
Household operation	HSUPL	Household supplies
Housing	HINSP	Home insurance, primary
Housing	HMTIMP	Home maintenance and improvement
Housing	MRTGIP	Mortgage interest, primary
Housing	MRTGPP	Mortgage principal, primary
Housing	OWNOP	Other owner costs, primary
Housing	OWNOS	Other owner costs, secondary
Housing	OWNRNT	Owner-occupied rental value
Housing	PTAXP	Property tax, primary
Housing	RENT	Rent
Miscellaneous	MSCFIN	Miscellaneous financial and insurance
Miscellaneous	PROSRV	Professional services and organization dues
Personal care	PERPRD	Personal care products
Personal care	PERSRV	Personal care services
Transfers	AACSP	Alimony and child support payments
Transfers	CHRTY	Charitable contributions
Transfers	OCASH	Other cash transfers
Transportation	AIRSHP	Air and ship travel
Transportation	GAS	Gasoline and other motor fuel
Transportation	NEWVEH	Gross value of new vehicle purchases
Transportation	PUBTRN	Public transportation
Transportation	VEHINS	Vehicle insurance
Transportation	VEHINT	Vehicle loan interest
Transportation	VEHPRD	Vehicle parts, accessories, and supplies
Transportation	VEHREG	Vehicle licensing, registration, and inspection
Transportation	VEHRNT	Vehicle rental, leasing, and parking
Transportation	VEHSRV	Vehicle maintenance and repair services
Utilities and phone	ELEC	Electricity
Utilities and phone	FOIL	Fuel oil
Utilities and phone	INTNET	Internet
Utilities and phone	LPG	LPG and propane
Utilities and phone	NGAS	Natural gas
Utilities and phone	OFUEL	Other fuels
Utilities and phone	PHONE	Phone services and products
Utilities and phone	TRASH	Trash and recycling
Utilities and phone	WATER	Water and sewer

The final step consists of adjusting the synthetic PUMS consumption variables to produce a "calibrated" set of values that match the following aggregate constraints: 1) national PCE totals; 2) state PCE totals; 3) state EIA totals for residential fuels; and 4) the FS-inspired consumption shares across income and age groups. The original values are adjusted iteratively (i.e. "raked") until the new values are stable. It is assumed that households never over-report consumption; i.e. calibration can increase a household's self-reported consumption but never decrease it.

Figure A1 shows how the distribution of total household consumption changes as a result of the calibration step. There is a decided shift to the right, and the tail of the post-calibration distribution is noticeably "fatter". Figure A2 shows the GAM-smoothed mean degree of under-reporting across the consumption distribution, by age group. It implies that lower-consumption (poorer) households tend to self-report consumption relatively well, and that reporting becomes less reliable the more a household earns/consumes.

It is not clear how to interpret the differences across age groups. If this difference is real, it suggests that higher levels of wealth reduce reporting accuracy. Older households are generally wealthier, and this may lead (on average) to less accurate recall regardless of one's position in the consumption distribution. Cognitive decline could also play a role. But the significant gap between younger and older cohorts suggests that this is more likely a data artifact. It is possible that the FS results – even after a plausible adjustment – still over/under-estimate consumption of older/younger households, leading to an exaggerated gap in spending recall.[14]

Estimating the distribution of household consumption remains surprisingly difficult, especially in light of its centrality to not only consumption tax policy but also more systemic issues of welfare and inequality. The dataset generated here offers a clear step in the right direction. Certainly, analyses that rely on self-reported consumption alone risk significant under-estimation of both overall levels of consumption and the extent of inequality.

[14] It is also possible that survey respondents from families with significant in-kind *inter vivos* transfers (e.g. parents pay child's rent) could incorrectly classify such transfers as the child's own expenditure, leading to upward bias in self-reported expenditures of the young and vice versa. But this effect would have to be systematic and quite large to explain the patterns seen here.

Figure A1: Effect of calibration step on distribution of household consumption

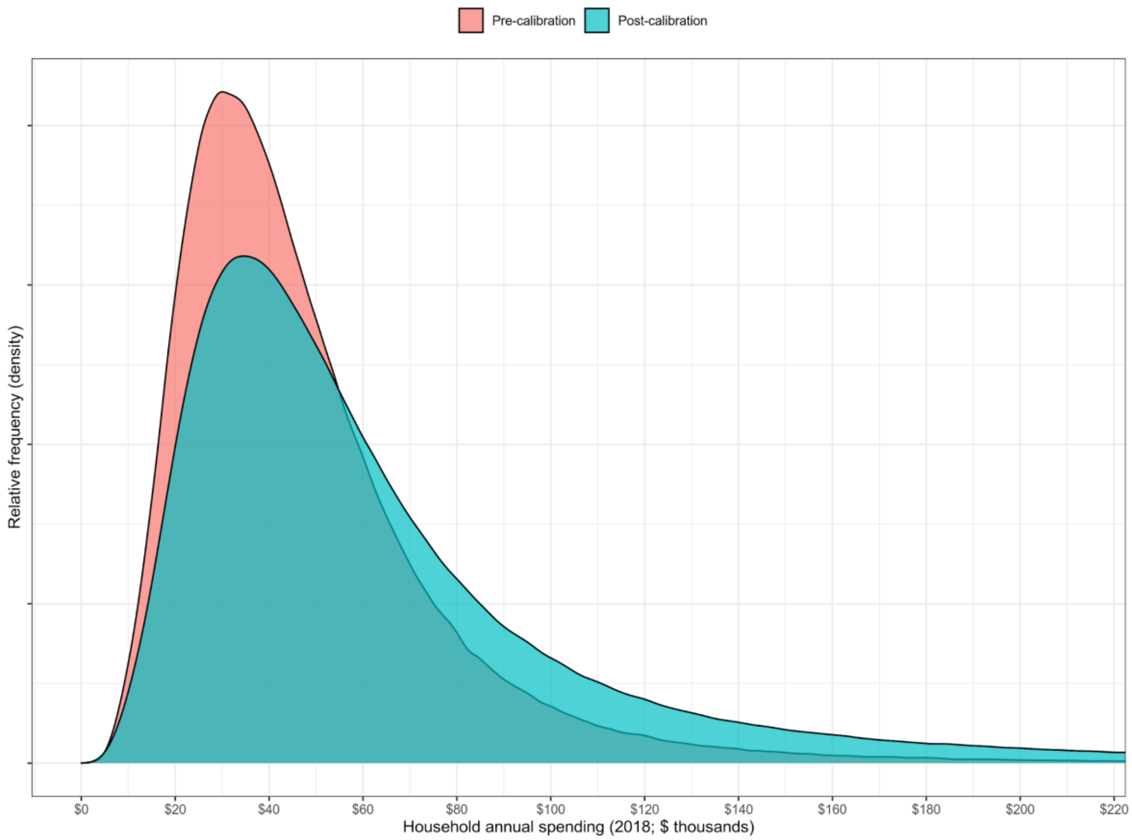

Figure A2: Implied degree of under-reporting in the Consumer Expenditure Survey

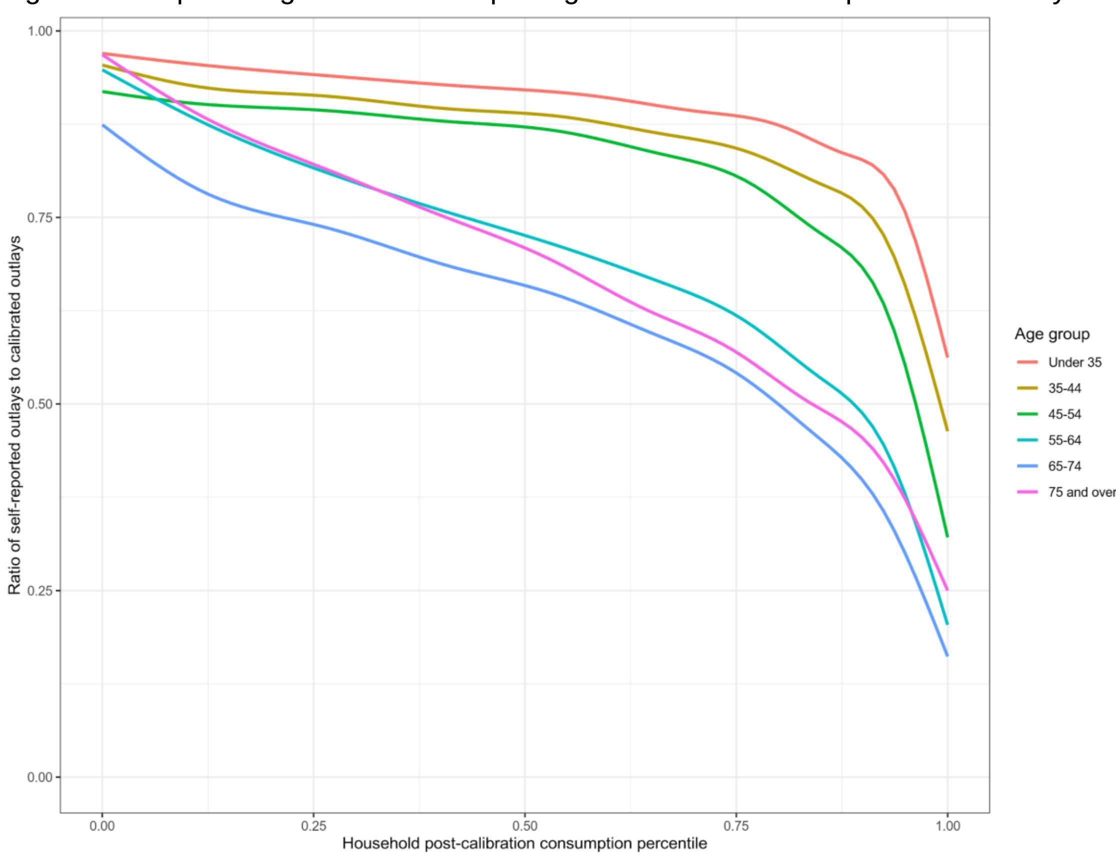

Annex B: Calculation of household GHG footprints

The EPA USEEIO model (v2.0) provides estimates of the GHG intensity of expenditure (i.e. $kgCO_2$-equivalent per dollar "emission factors") for nearly 400 commodities. The model is based on BEA input-output (I-O) tables from the 2012 benchmark paired with 2016 emissions data and updated to 2018 price levels. The BEA 2012 PCE "bridge" is used to estimate emission factors for individual PCE categories (called "series"), based on the amount of each commodity used to meet final household demand. Ummel (2016) provides a description of environmentally-extended input-output techniques and relevant references.

The 57 consumption categories (Table A1) used in this analysis are specifically constructed to allow good alignment with PCE series at the both and state and national level. This is made easier by the existence of a CE-PCE crosswalk constructed by BLS that links CE expenditure line items (UCC codes) to conceptually similar PCE series. For example, the "Clothing and footwear" category consists of 73 CE UCC codes that are collectively comparable to the sum of national PCE series DGARRC ("Garments"), DCSMRC ("Clothing materials"), and DSHURC ("Shoes and other footwear"), which are in turn all components of the DCLORC series ("Clothing and footwear") available at the state level. Consequently, it is possible to directly calculate a category-specific emission factor (EF) in most cases by combining EF's of the underlying PCE series. For expenditure categories lacking directly-comparable PCE series, an EF is assigned using the most logical PCE analog.

The USEEIO default EF's for fuel commodities do not include combustion emissions by the final end user; i.e. the EF for natural gas captures the upstream, indirect emissions of producing and delivering a therm of natural gas but not the emissions as a result of combustion itself. For natural gas, gasoline, fuel oil, and LPG, combustion EF's are calculated using EPA assumptions and added to the default upstream EF's – making adjustments as necessary so that the EF denominators reflect national average residential prices.

For electricity, the 2018 national GHG intensity of supply from the EPA eGRID database in conjunction with EIA residential electricity sales and price data are used to calculate a "downstream" EF for residential electricity consumption. It is then adjusted for grid losses and an estimate of power sector upstream emissions based on data from the NREL Life Cycle Assessment Harmonization project. The final result is a national average emission factor of 3.786 $kgCO_2e$ per dollar in 2018 for residential electricity.

The emission factors discussed to this point reflect *national average* $kgCO_2e$ per dollar of expenditure, by category. However, we know that emission factors for a given location and/or household often vary from the national average. In theory, this variation has four potential sources:

1) Variation due to difference between national and local average price level; e.g. a can of Coke may be more/less expense locally (the "Manhattan effect" described in HIS1).

2) Variation due to difference between local and household-specific average price level; e.g. some households may consume higher/lower-priced versions of otherwise-similar goods or services (the "Gucci effect" described in HIS1).

3) Variation due to difference between national and local average carbon content; e.g. goods or services consumed locally may contain more/less embedded carbon.

4) Variation due to difference between local and household-specific average carbon content; e.g. some households may consume higher/lower-carbon versions of otherwise-similar goods or services.

(1) and (2) – the Manhattan and Gucci effects – imply that national average EF's tend to overstate the carbon footprint of rich households, at least to the extent that richer households live in metropolitan areas with higher average price levels and (additionally) prefer upscale, higher-priced options within a consumption category (e.g. designer clothing instead of Walmart). (3) and (4) are more ambiguous with respect to distributional effects.

The Manhattan effect is estimated by incorporating BEA regional price parities that reflect differences in average price levels across metropolitan areas (and non-metropolitan areas within each state) for three broad consumption classes: Goods, Services, and Housing. Each of the 57 consumption categories is assigned to one of the three. For direct energy goods (gasoline, electricity, natural gas, fuel oil, and LPG), average prices by state are available from EIA SEDS. This allows adjustment factors to be calculated for each locale (metropolitan area or state) and consumption category, reflecting the ratio of local average price to national average price.

Estimating the Gucci effect requires information about the prices paid by different households for otherwise similar goods and services. Such information is typically only available in proprietary marketing datasets. However, it is possible to extract or otherwise deduce price variation across households for a limited set of goods using non-proprietary data that can serve as proxies for general consumer behavior.

As a rule, the CE collects expenditures and not prices or quantities. However, there are certain types of consumer durables that are purchased infrequently, increasing confidence that recorded transactions reflect the price paid for a single item. The CE also (separately) asks respondents for detailed information about vehicles, including the purchase year and price, even if the vehicle was financed and/or purchased used.

Figures B1 and B2 show how the typical purchase price varies for consumer durables and vehicles, respectively, *after* adjusting for local average price levels.[15] The curves are fitted to the data using generalized additive models (GAM) constrained to be monotone and convex (Pya and Wood 2015). The x-axis is a size-adjusted measure that ranks households according to total spending (not just on durables or vehicles) relative to household FPL. The two figures show similar patterns. Low-consumption households, on average, pay prices ~40% below the national average. High-consumption households typically purchase items at prices ~40-80% above the national average. The relationship is non-linear, with average price increasing sharply beyond the 90th percentile.

[15] The Gucci effect is estimated after adjustment for geographic variation in average price levels, so that adjustment factors for the Manhattan and Gucci effects are multiplicative when applied to national emission factors.

Figure B1: Gucci effect for consumer durables

Dashed black line is the average curve across all goods

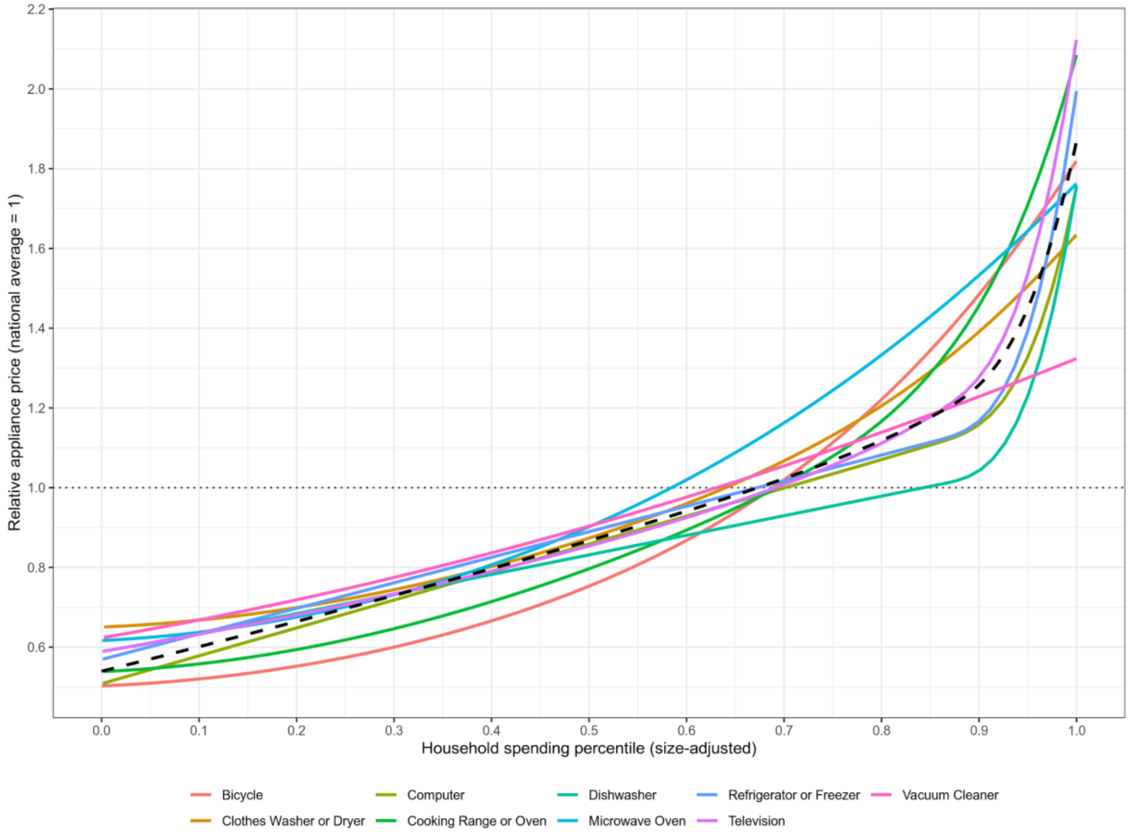

Figure B2: Gucci effect for vehicles

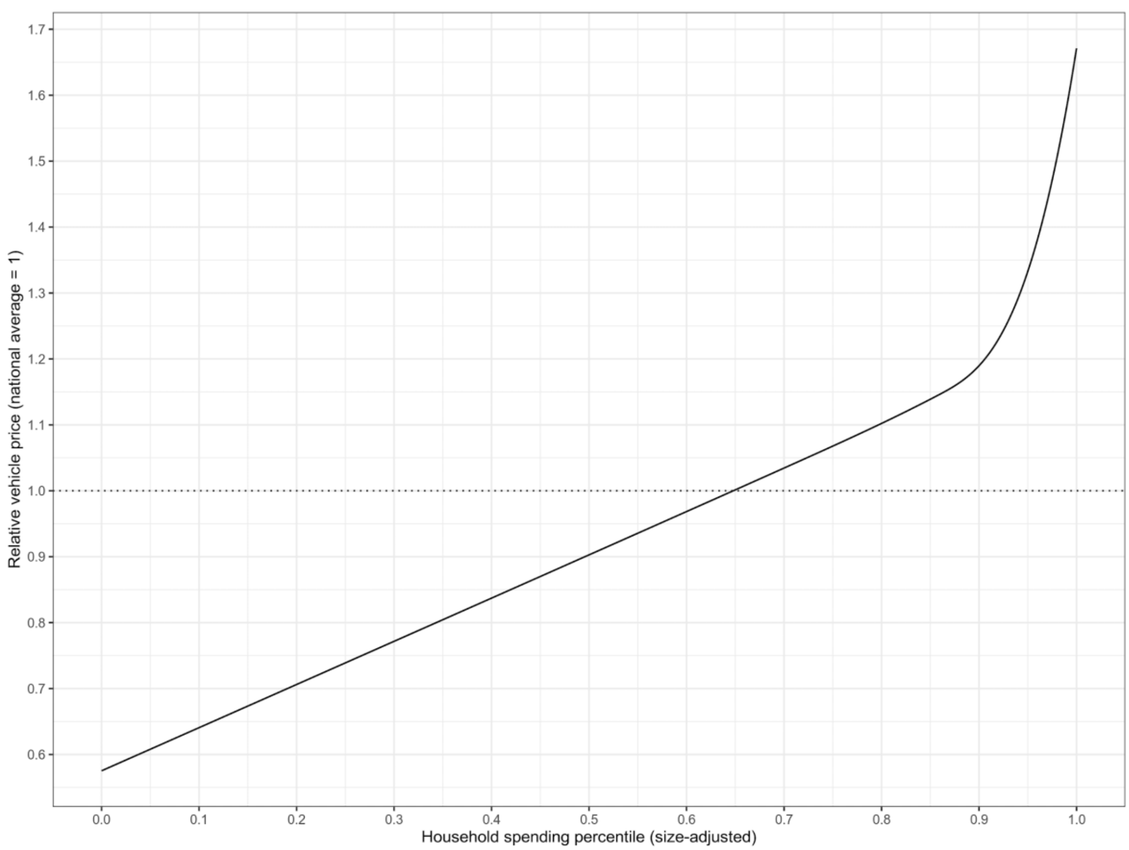

For the “Gasoline and other motor fuel” category, I use EIA data on national gasoline prices and consumption by fuel grade to estimate a plausible Gucci effect. Assuming that 1) consumption of higher-priced grades is strongly correlated with size-adjusted household spending and 2) the relationship is approximated by a power function leads to the assumed Gucci effect in Figure B3. The curve has a substantial flat portion, reflecting exclusive use of “regular” grade fuel by lower-consumption households. The curve plateaus for the highest spending percentiles, reflecting exclusive consumption of “premium” grade at prices ~18% above the national average.

Figure B3: Gucci effect for gasoline

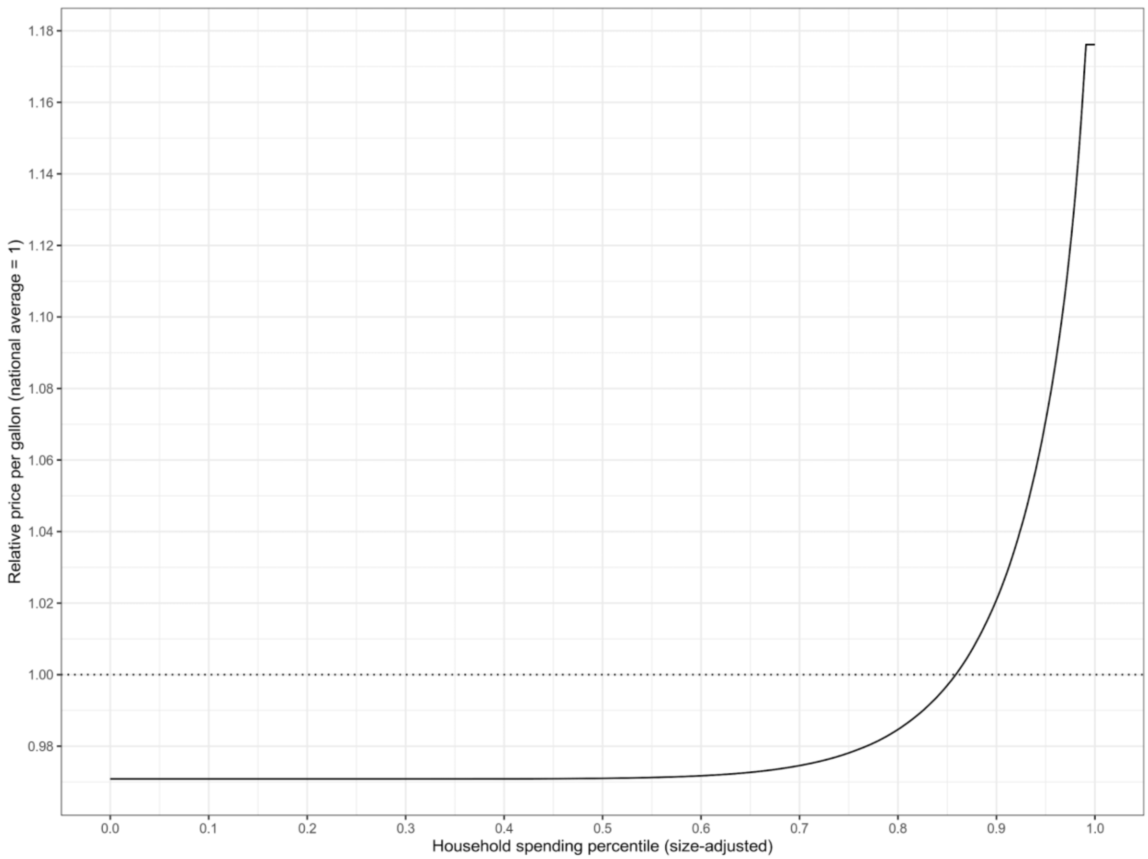

A unique case is food and drink. It is possible to compare PUMS (synthetic) food and drink expenditures to the number of calories consumed by each household, revealing an estimate of the cost-per-calorie. HIS1 estimated the Gucci effect for different food categories. HIS2 does not model food categories in detail, but a similar calculation is performed using total calibrated food and drink expenditures (including eating out and alcohol). The denominator (calories) is the household’s Estimated Energy Requirement (EER) (Institute of Medicine 2005). It is derived by applying EER equations to individuals in the 2017-2018 National Health and Nutrition Examination Survey (NHANES) with measured height and weight. A GBM model is trained to predict EER as a function of an individual’s age, sex, race, educational attainment, and marital status. This model is applied to individuals in the PUMS and EER then summed at the household level. The resulting Gucci effect for food and drink is shown in Figure B4.

Figure B4: Gucci effect for food and drink

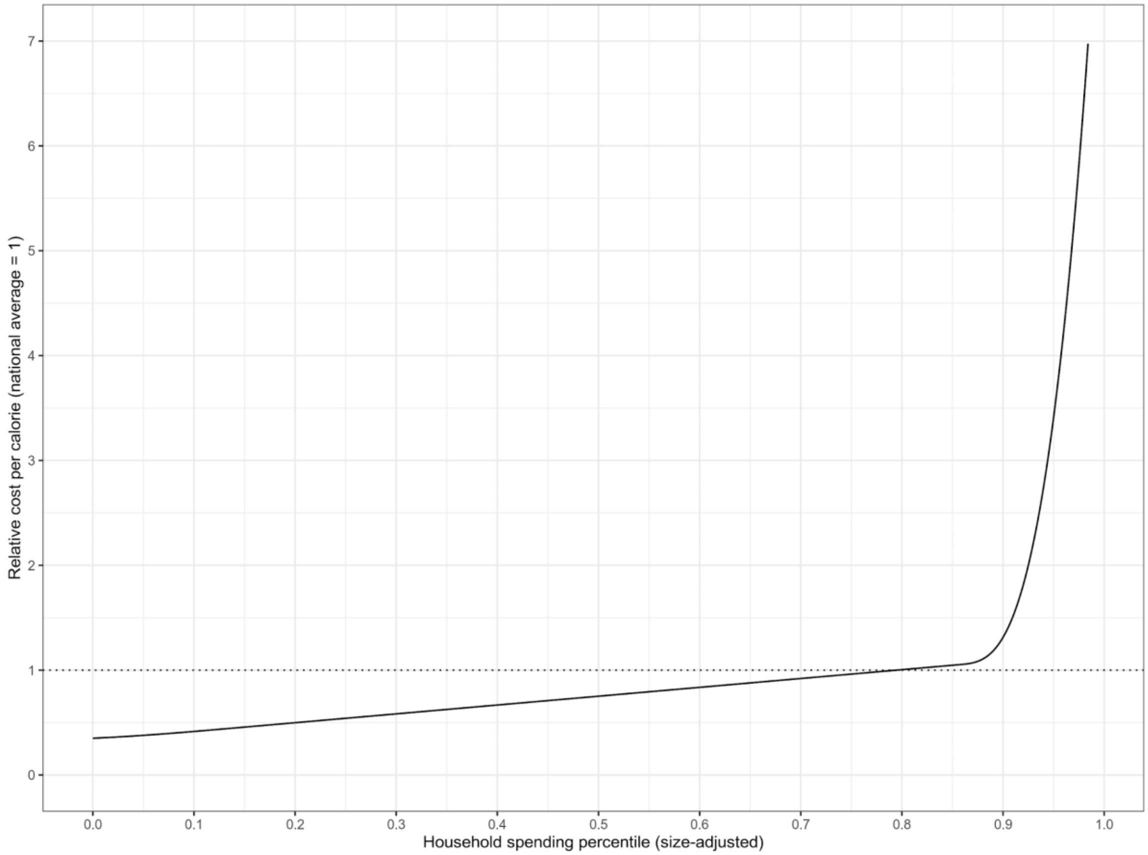

The Gucci effect for food and drink is similar to that observed for vehicles and appliances up until about the 90th percentile. Average cost-per-calorie increases dramatically beyond that point, with the highest-consuming households paying upwards of 7x the national average and 14x the price paid by the lowest-consumption households. This makes some sense, since food and drink is (uniquely) quantity-constrained for even the richest households. When large amounts of disposable income are directed to this category, it necessarily manifests as a steep rise in per-calorie prices.[16]

Finally, it is possible to estimate Gucci effects for electricity and natural gas for each of 27 regions, using the 2009 Residential Energy Consumption Survey (RECS) (Figures B5 and B6). Unlike the cases above, per-unit price variation for utilities reflects fixed costs in billing and tiered pricing structures (not consumer preferences) that vary considerably across the country. The same survey data were used for HIS1. The primary changes here are 1) to model the relationship using constrained GAM's for a smoother effect across the expenditure distribution and 2) calculation of a "generic" utilities Gucci effect -- derived from the relationship observed for natural gas -- that is applied to other billed utilities like water, sewer, trash, and recycling services (see Figure B6) where fixed costs are significant.

[16] The rich-to-poor price premium seems high but not improbable. The USDA estimates a U.S. family of four cooking all meals at home can provide a nutritionally-adequate diet at a cost of ~0.25 cents per calorie (Thrifty Food Plan calorie requirements at April 2020 food plan costs). A dinner including alcohol in a moderately up-scale restaurant might cost in excess of 10 cents per calorie.

Figure B5: Gucci effect for electricity

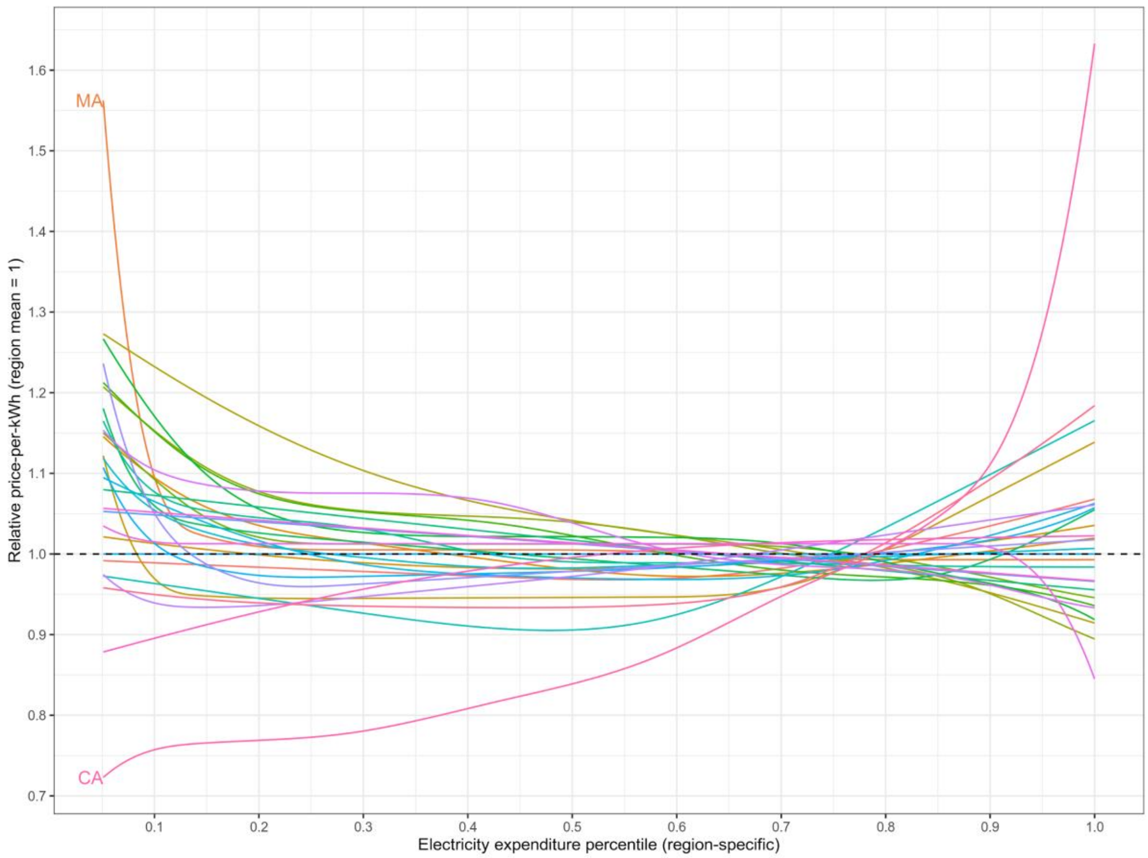

Households that consume less electricity in a given region typically pay more per-kWh (reflecting fixed monthly costs), as do the highest-expenditure households (reflecting tiered rate structures with increasing marginal cost). The lowest price-per-kWh is often paid by households in the middle of the distribution. A notable exception is California, where public-owned utilities pair low fixed monthly costs with aggressive tiered pricing. Consequently, the average price-per-kWh in California is considerably higher than the effective prices paid by most Californian households.

While Gucci effects are directly estimated for the expenditure categories above (consumer durables, vehicles, food and drink, gasoline, and utilities), other categories must be approximated. I use the “generic” Gucci effect derived for consumer durables in these cases (dotted line, Figure B1). This is clearly an over-simplification, since the magnitude of the Gucci effect must vary for different categories. But in the absence of better data, this approximation must suffice.

The one exception is the “Air and ship travel” category. Such spending is particularly important given its high average EF. Generally, both the carbon footprint and price of air travel correlate strongly with distance and cabin class. Air travel emission calculators typically assume that the footprint of business/first-class seats is about twice that of economy based on the additional space and weight of the former. This ratio does not seem obviously out-of-line with typical premiums paid for first-class seats. In light of this (and a lack of data proving otherwise), it seems reasonable to assume that airfare prices generally reflect the underlying carbon footprint (i.e. EF is more-or-less constant). Consequently, I assume that there is no Gucci effect for air and ship travel.

Figure B6: Gucci effect for natural gas

Dashed black line is the "generic" curve used for utilities other than electricity and natural gas

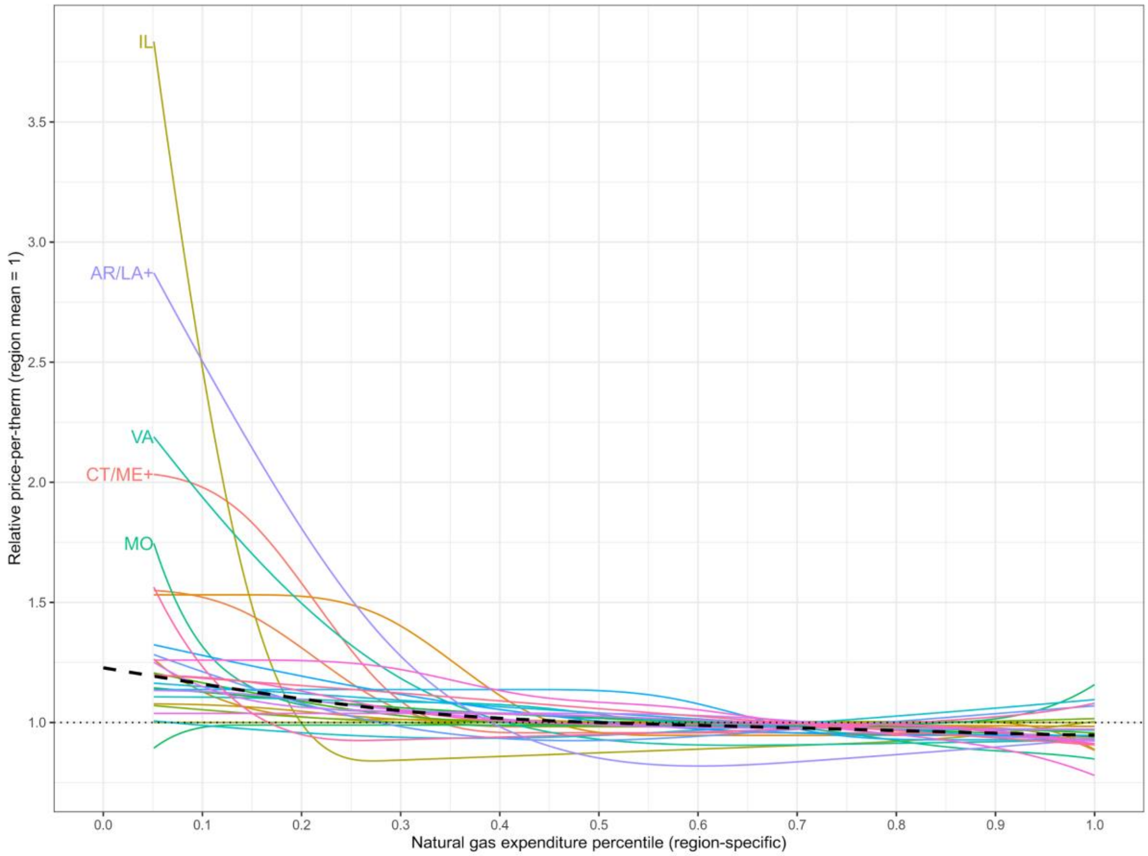

Both the Manhattan and Gucci effects concern how to adjust national average EF's for variation in *prices* across households – i.e. they adjust the EF denominator. The numerator also requires adjustment if the GHG content of particular goods or services varies across households, either due to geography or household preferences. An obvious case is electricity, for which there is clear geographic variation in GHG-per-kWh. Following the approach of HIS1, the 2018 EPA eGRID database is used to account for variation in GHG-per-kWh across 26 subregions of the electrical grid.

Beyond electricity, it is difficult to know how (or even whether) to make further adjustments for varying GHG content of goods and services. Caron et al. (2017) use a multi-regional input-output framework to estimate how the carbon intensity of expenditure varies across U.S. states and regions, accounting for trade flow between U.S. states and between states and foreign countries. They find substantial variation in carbon-per-dollar. At least part of that variation is surely due to price effects (i.e. the denominator), but the possibility remains that, for example, fruits and vegetables consumed in California exhibit significantly lower GHG-per-calorie than those consumed in Michigan. The Caron et al. analysis is based on I-O data dating from 2006, making the results too old to use here with confidence.

It is also possible that GHG intensity of otherwise similar goods could vary with household characteristics; if, say, rich households prefer "green" versions of products or (conversely) buy vehicles that are both more expensive *and* larger (i.e. more steel required). There is no convincing evidence on this front and, on its face, it seems unlikely that households make choices that have significant effects. But this remains an open question.

The adjustment factors estimated above are applied to the national, USEEIO-based EF's for each consumption category and household. This induces a considerable amount of variation across households. The most notable case is electricity, for which all three kinds of adjustments (Manhattan, Gucci, and spatial variation in GHG intensity) are applicable. Figure B7 shows the distribution of electricity emission factors across households. It makes clear that using a single, national emission factor in this kind of analysis would lead to significant misestimation of the distribution of GHG emissions.

Figure B7: Variation in household emission factors for electricity

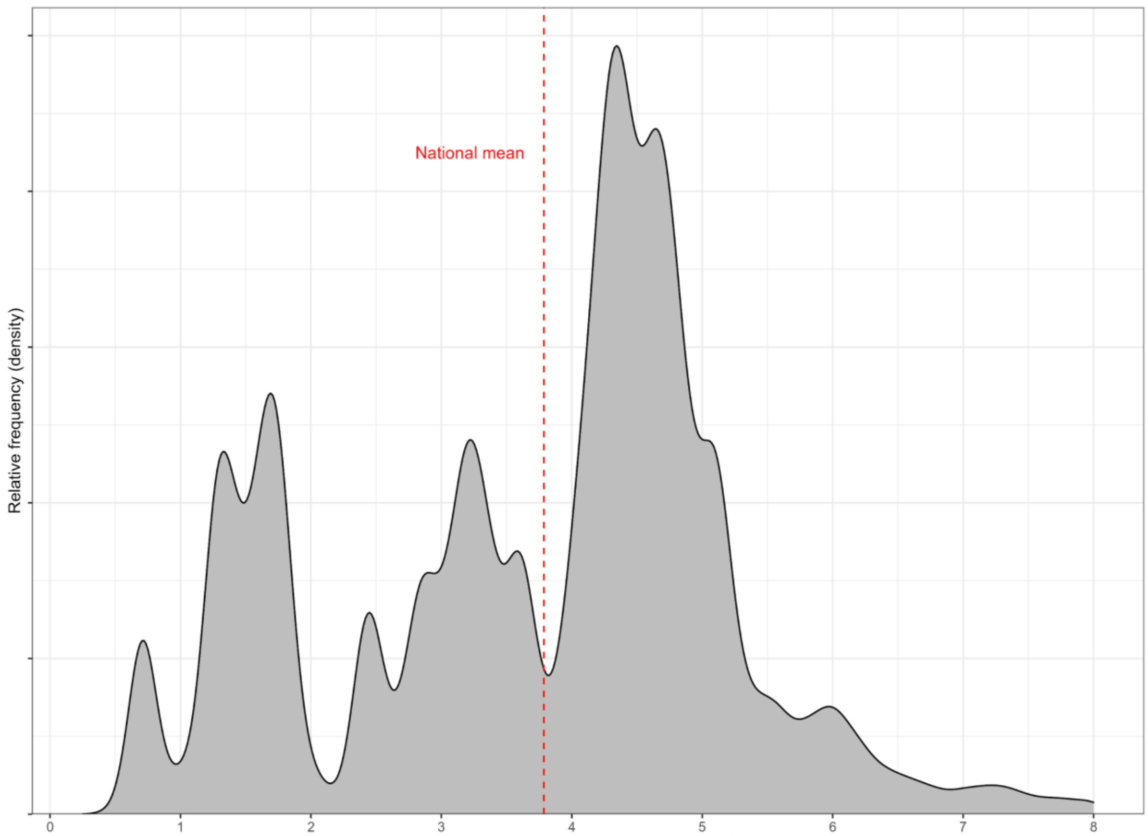

A household's GHG footprint for a given category is the product of expenditure and the fully-adjusted EF. The overall effect of the adjustments is to reduce the GHG footprint of richer households relative to what one would expect from looking at the distribution of household consumption (Figure A1). There are two trends at play here. As one moves up the consumption distribution, the share of consumption allocated to the most carbon-intensive goods (gasoline, electricity, etc.) declines. At the same time, greater purchasing power is increasingly directed to quality (or, at least, higher per-unit prices) over quantity. The combination of these trends leads to less inequality in household GHG footprints – though still substantial – than one might expect given the extent of inequality in income or consumption.

Annex C: Household simulation of tax burden and dividend payment

Microsimulation of the carbon fee and dividend policy entails estimating each household's additional indirect tax burden due to carbon pricing as well as the additional disposable income resulting from the dividend payment. The dividend side of the equation is identical for the two scenarios, but the nature of the tax burden differs. In Scenario 1, the additional tax burden is directly-proportional to the household's estimated GHG footprint. In Scenario 2, 70% of the overall tax burden is associated with GHG footprints, but the remaining 30% is assigned based on a household's relative exposure to a decline in business profits. Each of these components is described below.

The gross dividend payment is straightforward to simulate, since it is allocated to individuals based only on their age as prescribed in H.R. 763. Subsequent taxation of the payment is more complicated. To simulate taxation for each household, the PUMS sample is used to assign individuals to "tax units" and deduce the filing status of each (single, married, head of household).[17] The resulting overall number and share of tax units by filing status are comparable to IRS data.

The dividend payment is subject to federal income taxes and the employee portion of payroll (social security and Medicare) taxes. Payroll taxes are effectively "flat" taxes on individuals' wages; social security taxes are zero on individuals' wages beyond a specified threshold ($128,400 in 2018). For the purposes of this study, it is assumed that the dividend payment is treated as additional earned income for any individual with pre-existing wages or self-employment income. For individuals without such income (about half the population), the dividend is assigned to the tax unit's primary taxpayer as additional wages.

The tax simulation step assumes income tax brackets and rates for 2018 as introduced in the Tax Cuts and Jobs Act (TCJA) of 2017. It also includes household-level simulation of the Child Tax Credit (CTC) and Additional Child Tax Credit (ACTC) using TCJA credit amounts and limitations. The CTC and ACTC combine to function like a partially-refundable credit, so it is necessary to explicitly calculate how each respond to the additional income tax liability introduced by the dividend. The CTC often erases the dividend's income tax liability for lower-income households (but does not affect payroll tax liability). The other major tax credit affecting households – the Earned Income Tax Credit – is fully refundable and, therefore, unaffected by the dividend.

Simulated taxation of the dividend produces an effective tax rate of 18.2%. This is close to the share of the carbon tax burden (net of export rebates) falling on government as estimated in HIS1; indicating an approximately revenue neutral outcome. It is also nearly identical to the effective tax rate in HIS1, which used a stylized relationship between income and marginal tax rate based on results from the Urban-Brookings Tax Policy Center Microsimulation Model.

The total carbon tax burden allocated across government and households/individuals is $78.1 billion in accordance with CCL's own initial revenue assumption. The share assigned to

[17] The PUMS sample of 1.3 million households includes data for individuals within each household. Individuals' age, income, marital status, and relationship to the head of household are all used to create tax units.

government is 18.2%, reflecting the dividend's effective tax rate and an assumption of approximate revenue neutrality. About 2% is assigned to individuals in non-correctional group facilities (e.g. nursing homes, dormitories). The remainder is allocated among the PUMS household sample.

Each household's carbon tax burden in Scenario 1 is proportional to the estimated GHG footprint, and the same is true of the 70% of the tax burden that is passed forward into consumer prices in Scenario 2. The remaining tax burden in Scenario 2 is assigned in proportion to a household's "investable assets" – the value of all mutual funds, stocks, bonds, annuities, trusts, retirement accounts, pensions, and direct interests in businesses – i.e. a proxy for how exposed a household is to a general decline in business profits. The investable assets variable is constructed using microdata from the 2016 Survey of Consumer Finances (SCF) and is fused onto PUMS households using the same statistical technique described in Annex A.[18]

The distribution of investable assets is highly skewed: The 10% of households with the most investable assets own about 80% of the total. Those same households are responsible for about 20% of the national GHG footprint. Consequently, Scenario 2 entails a significant shift of the carbon tax burden onto wealthier households compared to Scenario 1.

Finally, the measure of household income used in the calculation and presentation of impacts (e.g. to define a "minor loss") is different from HIS1. HIS1 relied on household income as reported in the PUMS. An alternative measure of "normal income" from the SCF is fused to the PUMS at the same time as investable assets. The SCF normal income variable differs from PUMS income in two important ways: 1) It asks respondents to report their income in a "normal" year; 2) It includes income from capital gains. The two measures can differ significantly, especially for the self-employed and wealthy.[19]

[18] Investable assets include the NMMF, STOCKS, BOND, OTHMA, RETQLIQ, OTHFIN, and BUS components of the SCF Bulletin Asset and Debt Categories, which totaled about $54 trillion in 2016. This definition excludes transaction accounts and certificates of deposit but includes the value of directly-owned businesses.

[19] But not exclusively. In the SCF, about 17% of households report that 2016 calendar-year income (including capital gains) was at least 20% above or below income in a "normal" year. This phenomenon is surprisingly consistent across the income distribution, indicating that year-to-year variation in income is not limited to wealthier households, nor is it due solely to capital gains. This is consistent with the "financial diaries" research of Hannagan and Morduch (2015) that find considerable income variability over time among both poor and middle-class households.

Annex D: Scenario 3 (85% pass-through) results

Figure D1: Impact by consumption quintile for Scenario 3 (85% pass-through)

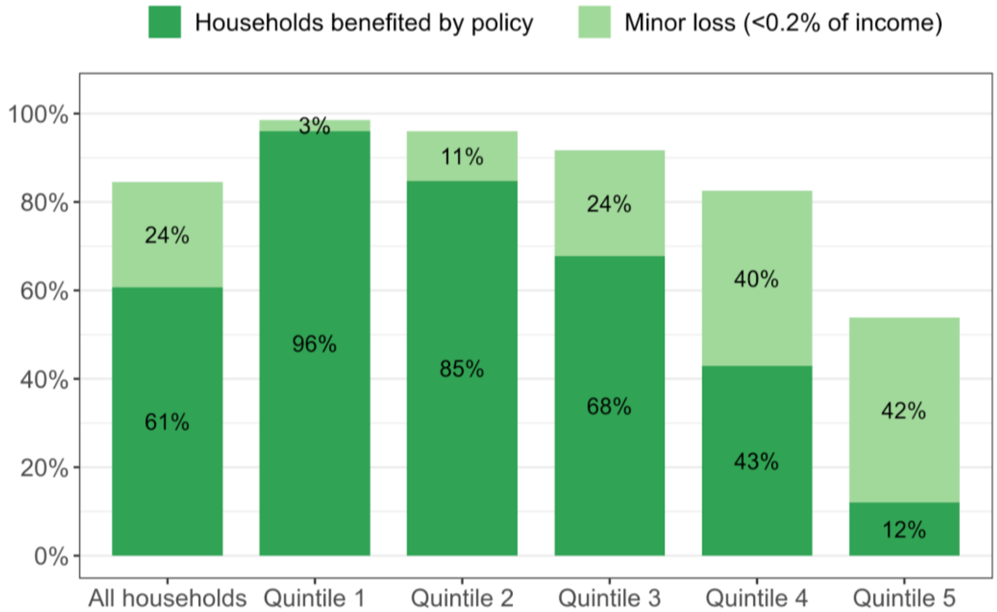

	All households	Quintile 1	Quintile 2	Quintile 3	Quintile 4	Quintile 5
Mean net benefit (% of avg. income)	$-31 (-0.04%)	$241 (0.64%)	$129 (0.23%)	$53 (0.07%)	$-42 (-0.05%)	$-538 (-0.29%)
Median gain (% of HH income)	$133 (0.29%)	$212 (0.72%)	$140 (0.29%)	$104 (0.18%)	$80 (0.11%)	$65 (0.07%)
Median loss (% of HH income)	$-126 (-0.15%)	$-32 (-0.12%)	$-40 (-0.10%)	$-59 (-0.10%)	$-95 (-0.12%)	$-273 (-0.21%)
Median HH income % of FPL	354%	154%	268%	361%	476%	815%

Figure D2: Impact by household type for Scenario 3 (85% pass-through)

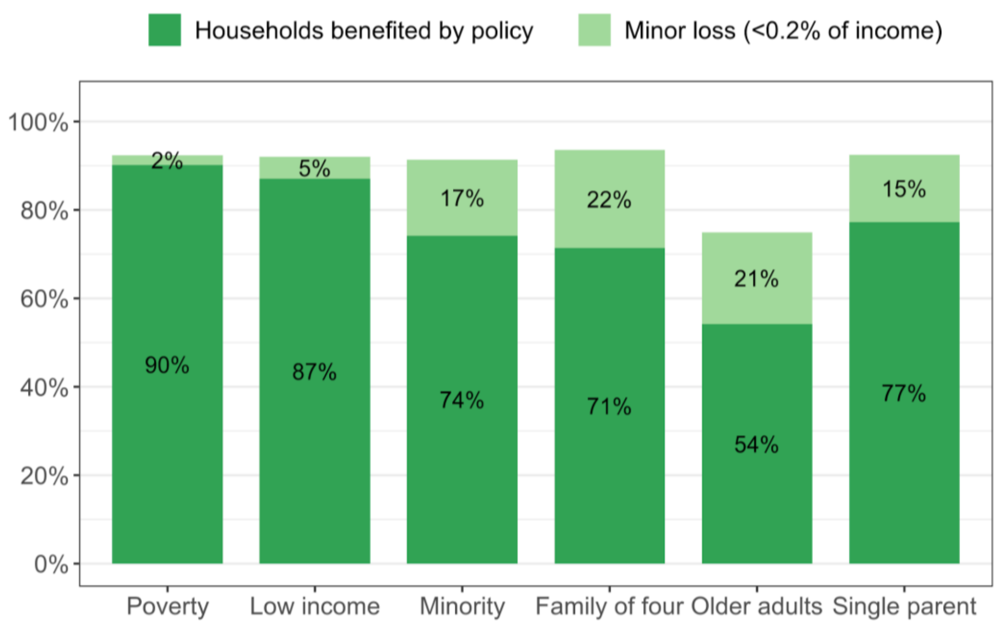

	Poverty	Low income	Minority	Family of four	Older adults	Single parent
Mean net benefit (% of avg. income)	$185 (1.75%)	$175 (0.83%)	$87 (0.12%)	$19 (0.01%)	$-121 (-0.18%)	$78 (0.16%)
Median gain (% of HH income)	$148 (1.70%)	$156 (0.87%)	$169 (0.41%)	$200 (0.26%)	$100 (0.33%)	$148 (0.56%)
Median loss (% of HH income)	$-40 (-0.58%)	$-44 (-0.29%)	$-98 (-0.13%)	$-174 (-0.09%)	$-134 (-0.22%)	$-94 (-0.13%)
Median HH income % of FPL	67%	121%	278%	389%	300%	171%

Figure D3: Impact by age group for Scenario 3 (85% pass-through)

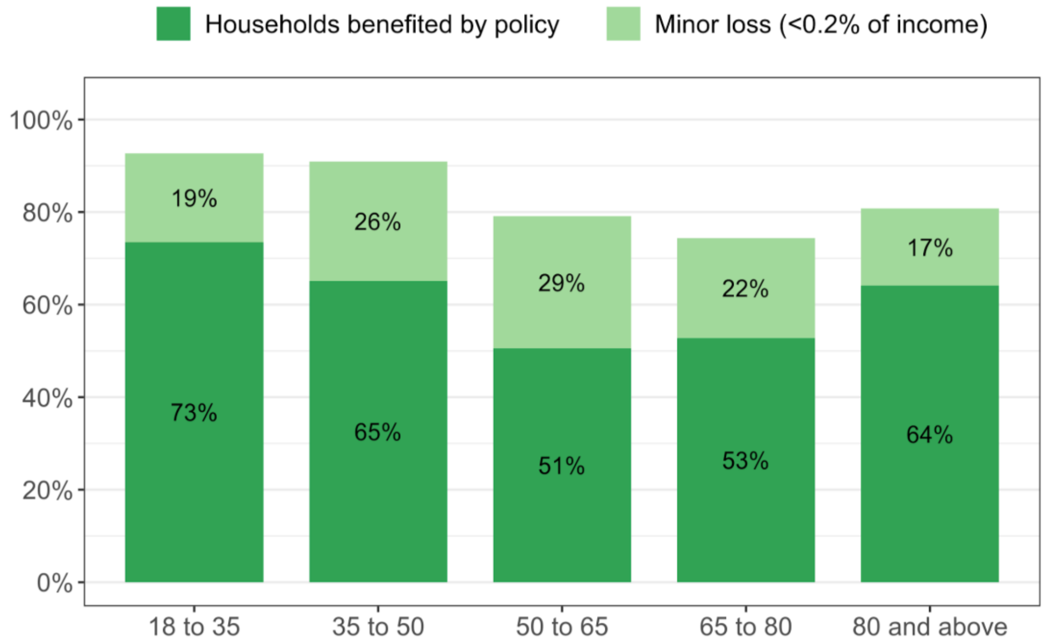

	18 to 35	35 to 50	50 to 65	65 to 80	80 and above
Mean net benefit (% of avg. income)	$92 (0.14%)	$30 (0.03%)	$-120 (-0.11%)	$-136 (-0.18%)	$-41 (-0.08%)
Median gain (% of HH income)	$140 (0.33%)	$164 (0.26%)	$122 (0.26%)	$109 (0.32%)	$111 (0.41%)
Median loss (% of HH income)	$-71 (-0.11%)	$-116 (-0.11%)	$-159 (-0.17%)	$-150 (-0.22%)	$-109 (-0.22%)
Median HH income % of FPL	288%	386%	433%	331%	237%

Figure D4: Impact by race for Scenario 3 (85% pass-through)

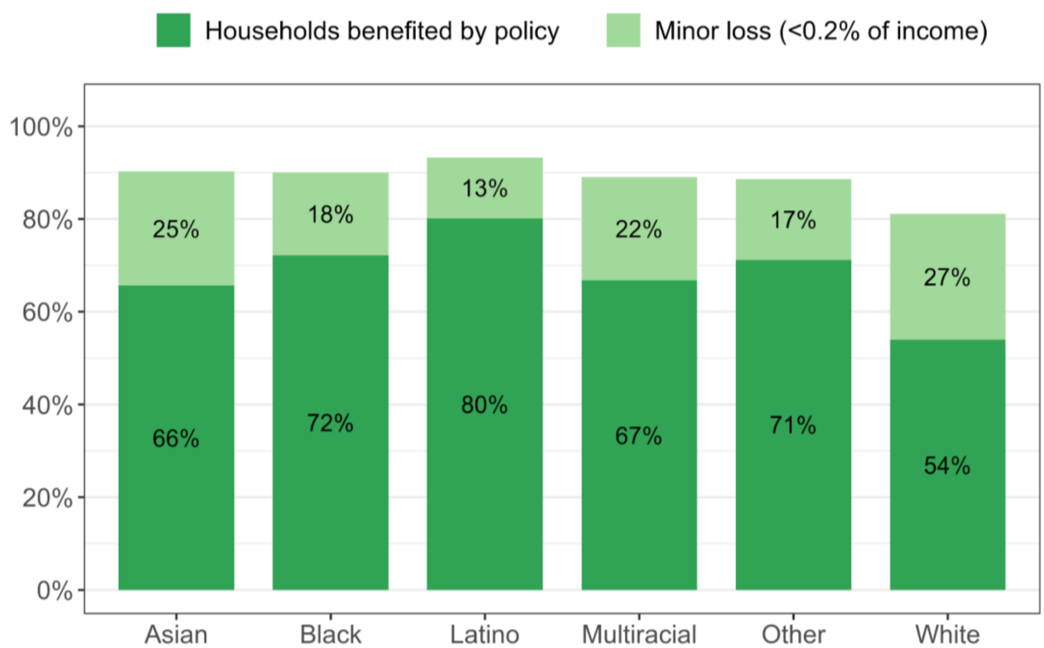

	Asian	Black	Latino	Multiracial	Other	White
Mean net benefit (% of avg. income)	$-27 (-0.02%)	$73 (0.12%)	$149 (0.22%)	$25 (0.03%)	$80 (0.12%)	$-91 (-0.09%)
Median gain (% of HH income)	$162 (0.28%)	$139 (0.40%)	$210 (0.47%)	$134 (0.30%)	$162 (0.41%)	$114 (0.24%)
Median loss (% of HH income)	$-147 (-0.11%)	$-84 (-0.14%)	$-92 (-0.13%)	$-102 (-0.13%)	$-101 (-0.15%)	$-135 (-0.16%)
Median HH income % of FPL	456%	255%	249%	331%	269%	396%

Figure D5: Impact by community type for Scenario 3 (85% pass-through)

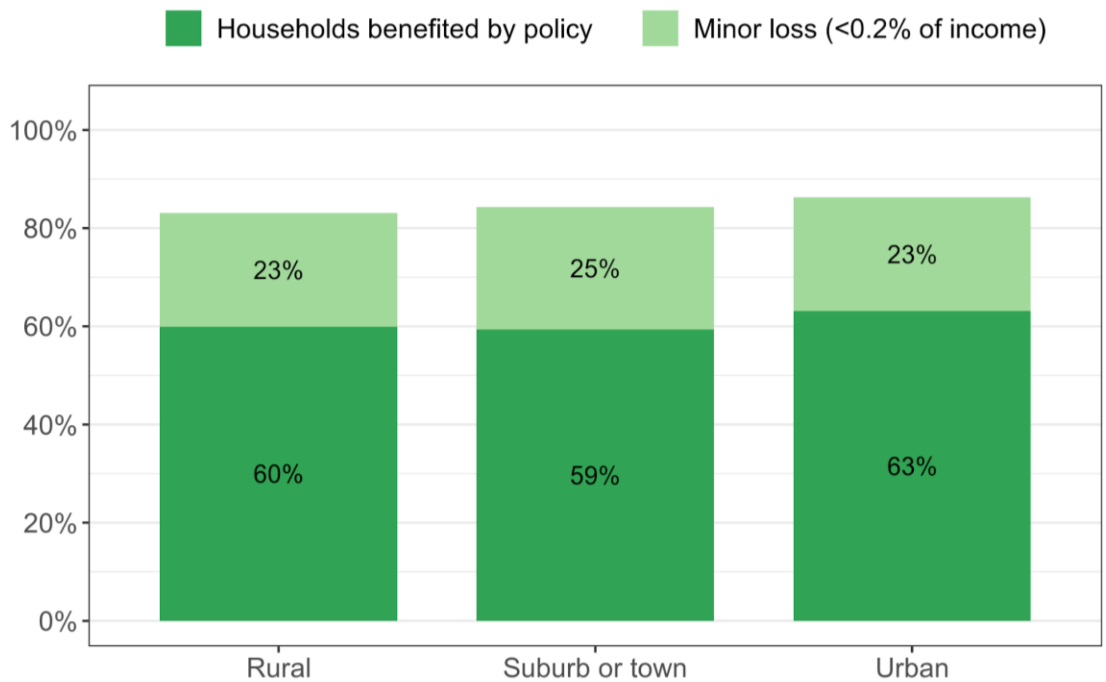

	Rural	Suburb or town	Urban
Mean net benefit (% of avg. income)	$-21 (-0.03%)	$-44 (-0.05%)	$-25 (-0.03%)
Median gain (% of HH income)	$129 (0.29%)	$135 (0.27%)	$133 (0.33%)
Median loss (% of HH income)	$-121 (-0.16%)	$-133 (-0.15%)	$-121 (-0.14%)
Median HH income % of FPL	332%	379%	344%