MIT CEEPR
MIT Center for Energy and Environmental Policy Research

Working Paper Series

The Roosevelt Project Special Series

Emission Inequities: The Distributional Effects of Climate Policy

TOMAS GREEN, CHRISTOPHER R. KNITTEL, SHEREEIN SARAF

September 2020 (revised April 2025)

CEEPR WP 2020-R3

HARVARD UNIVERSITY

Massachusetts Institute of Technology

Roosevelt Project Report Sponsor

The Roosevelt Project participants thank the Emerson Collective for sponsoring this report and for their continued leadership on issues at the intersection of social justice and environmental stewardship.

Emission Inequities:
The Distributional Effects of Climate Policy

Tomas Green, Christopher R. Knittel, Shereein Saraf *

Revised: April 28, 2025

Abstract

This paper applies machine learning methods to estimate household carbon footprints across the roughly 80,000 US census tracts, uncovering significant disparities based on income, geography, and urbanity. We evaluate the distributional impacts of seven climate policies, including carbon pricing, gasoline taxes, and intensity standards such as electricity clean energy standards and fuel economy standards. Policies vary considerably in their burden on low-income households, with the carbon tax policies being progressive while intensity standards are regressive. A simple carbon tax-and-dividend policy and the intensity standards also lead to much larger burdens on middle-American and rural households. We show that adjusting carbon tax dividends based on geography and urbanity can address these disparities and improve policy equity. However, this is not possible within intensity standards since they do not generate revenue.

JEL Classification Codes: Q54, Q52, Q58, D31, R14

Keywords: Climate Policy, Carbon Footprints, Distributional Impacts, Geographic Disparities, Rural-Urban Divide

* *Green:* Acting Chief of Staff of Sustainable Transportation and Fuels for the DOE's Office of Energy Efficiency and Renewable Energy, tomasgreen@alumn.mit.edu. *Knittel:* George P. Shultz Professor, Sloan School of Management, Director, MIT Center for Energy and Environmental Policy Research, Director, MIT Carbon Policy Center, and NBER, href hrefknittel@mit.eduknittel@mit.edu. *Saraf:* Center for Energy and Environmental Policy Research, MIT, ssaraf@mit.edu. This paper has benefited from conversations with David Hsu, Michael Kearney, Ernie Moniz, and the rest of the Roosevelt Project research team. We thank seminar participants at UC Berkeley, University of Chicago, Dallas Fed, Georgia Tech University, Mannheim University, MIT, and Yale University.

1 Introduction

Addressing climate change requires not only aggressive action but also a policy approach that balances environmental goals with economic and equity considerations. Policymakers, businesses, and local governments face complex decisions about which climate policies to implement, how to structure them, and how to mitigate their unintended distributional effects. Carbon emissions associated with household consumption vary considerably across several dimensions, including geography, race, income, and building stock. Given this, addressing climate change requires a nuanced understanding of its economic impacts across different communities. However, comprehensive carbon footprint data do not exist at a national level. Therefore, a critical component of this challenge is accurately estimating household carbon footprints (HCFs) to understand the distributional effects of various climate policies.

This paper addresses these challenges by estimating household carbon footprints (HCFs) at a granular level, namely the census tract. We then use these data to model the impacts of several climate policies. Broadly speaking, this takes three steps. First, using data on sets of representative US households, we employ machine learning techniques to train prediction models that relate consumption of several products, such as electricity and natural gas, to variables that capture household demographics, location of the household, and weather, among others. The survey data employed for this step are drawn from the 2020 Residential Energy Consumption Survey (RECS), the 2017 National Household Transportation Survey (NHTS), and the 2022 Consumer Expenditure Survey (CEX), as well as our own survey on airline travel. (U.S. Energy Information Administration, 2020a; Federal Highway Administration, 2017; U.S. Bureau of Labor Statistics, 2022) Our models capture household energy consumption, transportation habits, and spending on goods and services to accurately show how consumption patterns vary across households. Our preferred method relies on the Least Absolute Shrinkage and Selection Operator (LASSO) algorithm to estimate emissions.

The second step is to use these prediction models to predict consumption at the census tract level by using values of the features chosen in the LASSO model from the census. This provides a prediction of the consumption of each product for an average household within a census tract. We discuss why, given our LASSO estimates, this yields an unbiased predictor of average consumption for a given tract. The third step is to measure the embedded emissions within the predicted consumption. We use emission factors from the Department of Energy's eGRID and the Environmental Protection Agency's Compilation of Air Pollutant Emissions Factors. (Argonne National Laboratory, 2019; U.S. Environmental Protection Agency, 2020)

Our analysis reveals substantial variation in carbon footprints driven by both income and geographic factors. Several stylized facts follow from the data. First, higher-income households typically have larger carbon footprints due to increased consumption of transportation and consumer goods. The HCF of the average household in the bottom decile of the income distribution is 15.25 tons per year, compared to 22 in the top decile. Using the EPA's proposed social cost of carbon value of $190 per metric ton, this translates to an over $1282.5 difference in climate damages.[1] Second, despite the positive correlation, carbon footprints grow less than one-for-one with income. While the average HCF between the first and tenth decile grows by about 44%, average incomes grow by about 415%.

Third, HCFs tend to be higher in the central part of the US, and this is due to both geographic differences in the carbon intensity of the local electricity grid and differences in consumption levels. Average HCFs in the four coastal Census divisions (New England, Mid-Atlantic, South Atlantic, and Pacific) are 17.32 tons per year, while the average everywhere else in the country is 20.2 tons. And, fourth, HCFs are lower in urban areas compared to suburban and rural areas. Each of these stylized facts has important implications for the distributional impacts of climate policies. On average, households in rural areas emit 21.14 tons of carbon dioxide per year, compared to 18.56 tons in urban areas. This disparity is largely due to rural households' reliance on private vehicles and fossil fuels for heating. Additionally, carbon footprints in the Midwest and central US are higher due to the carbon intensity of local electricity grids, whereas coastal regions benefit from cleaner energy sources.

We use our data on HCFs to evaluate the effects of nine different climate policy scenarios, focusing on carbon pricing schemes and performance standards such as the Corporate Average Fuel Economy (CAFE) standards and the Clean Energy Standard (CES). We model four carbon tax policies that vary the way revenues are recycled. Except for the scenario where we assume the revenues are not recycled, our results show that carbon pricing policies are progressive. For instance, under a $100-per-ton carbon tax with an evenly distributed dividend, low-income households receive a larger dividend than they pay in carbon taxes, leading to a net financial benefit. However, without geographic adjustments, this policy risks transferring wealth from the Midwest and Plains regions to coastal areas, where carbon footprints are generally lower. Our findings highlight the need for climate policies that account for both income and geographic disparities. A straightforward tax-and-dividend plan can be progressive, but geographic adjustments may be crucial to reduce transfers from rural and Midwest households. We show that policymakers can tailor approaches that balance the economic burden of climate action across different regions and income groups.

We also model a federal gasoline tax within the tax group. The tax payments under this program are similar in nature to the carbon tax payments discussed above. Therefore, as with carbon taxes, the use of the funds is crucial for understanding the regressivity of the taxes. Ignoring the revenues collected, gasoline taxes disproportionately impact rural households due to higher average vehicle usage and lower fuel economy, while urban households experience relatively lower burdens. We allocate the revenues based on the current allocation of the federal gasoline tax program, which uses the gasoline tax revenues for the interstate highway system. We use the state-level allocations and assume that households within the given state equally benefit from these state-level allocations. We find that north-central states disproportionately benefit from the program, as do Gulf Coast states.

[1] Source here.

We model three performance standards. These policies typically regulate carbon intensities or some close proxy, such as fuel economy, as opposed to carbon levels as in a carbon tax. Kwoka Jr (1983), Helfand (1991), and Holland et al. (2009) have shown that this translates into an implicit tax for any product that is worse than the standard and an implicit subsidy for products that are better than the standard. The degree of implicit taxation and subsidization is driven by how binding the standard is. We investigate three performance standards: fuel economy standards, clean energy standards, and carbon emission intensity standards applied to states' electricity sector. The performance standards have many of the same shortcomings as carbon pricing *without* dividends. The regulatory standards are regressive, imposing higher costs on low-income and rural households. For example, CAFE standards, which incentivize the production of fuel-efficient vehicles, disproportionately benefit high-income households that can afford newer, more efficient cars, while low-income households bear the brunt of higher vehicle costs. Similarly, a Clean Energy Standard increases electricity prices without providing revenue to offset these costs, further burdening households that rely on carbon-intensive grids.

We are not the first to explore HCFs and the distributional impacts of climate policy. (Jones and Kammen 2014; Jones and Kammen 2011) use regression methods to estimate household carbon footprints; they do not investigate how policies are likely to affect households. Using different methods, a separate literature has also explored the distributional impacts of climate policies, emphasizing the importance of designing policies that mitigate regressive outcomes. Metcalf et al. (2008) and Mathur and Morris (2014) have highlighted the potential regressive nature of carbon taxation. Research by Goulder et al. (2019) and others has shown that while emission reductions are relatively insensitive to the methods of revenue recycling, the welfare and distributional impacts vary significantly. Furthermore, geographic disparities play a crucial role in the effectiveness and fairness of climate policies, as demonstrated by studies from Davis and Knittel (2019) and Gillingham and Stock (2018).

The paper proceeds as follows: Section 2 outlines the methodology and data sources used in this study, including the machine learning models employed to estimate household carbon footprints (HCFs) and the census tract-level data that underpin our analysis. Section 3 presents the results, detailing the geographic and socioeconomic variation in HCFs and the distributional impacts of various climate policy scenarios. Section 4 discusses the policy implications of these findings, emphasizing the need for tailored solutions that address disparities across income groups, regions, and urbanization levels. Section 5 concludes by summarizing the key insights and offering recommendations for policymakers to ensure equitable and effective climate policies. Appendices provide additional details on data sources, modeling techniques, and supplementary results.

2 Methodology and data

2.1 Preliminaries

At a high level, our goal is to model the financial impact of a given climate policy on specific households. This requires three steps. First, we use machine learning techniques and data on a survey of representative households to predict a household's consumption of carbon-intensive goods based on household demographics, location of residence, and housing stock details. We model the consumption of nine goods: electricity, natural gas, propane, kerosene, gasoline, air travel, food, health care, and personal care. Collectively, these comprise virtually the entire footprint of households. Second, we combine the model trained on a survey of representative households and nationwide census tract data on the features used in the machine learning model to project the consumption of these goods across *all* census tracts. This yields an estimate of the average household-level consumption across the nearly 85,000 census tracts within the US.[2] Third, given consumption estimates of each of these goods, we use data on carbon intensities to calculate the household's total carbon footprint. Finally, with these "data" in hand, we can model particular climate policies using economic theory's predictions on what each policy will do to the prices of the goods (or carbon).

2.2 Representative survey data

We use three different household surveys, the 2020 Residential Energy Consumption Survey (RECS), the 2017 National Household Transportation Survey (NHTS), and the 2022 Consumer Expenditure Survey (CEX), for household-level data on energy consumption, vehicle miles traveled, and expenditure on goods and services, respectively. (U.S. Energy Information Administration (2020a), Federal Highway Administration (2017), U.S. Bureau of Labor Statistics (2022)) The 2020 Residential Energy Consumption Survey (RECS) collects household-level data on consumption of units of energy, including electricity, natural gas, propane, and kerosene, along with demographic and geographic characteristics. We add the unit price of energy to these data from the State Energy Data System (SEDS). (U.S. Energy Information Administration, 2020b) The 2017 National Household Transportation Survey (NHTS) records miles traveled by each household vehicle per trip. We aggregate the per-trip per-vehicle miles to obtain household vehicle miles.[3] The 2022 Consumer Expenditure Survey (CEX) provides data on household expenditure on food, health, and personal consumption.

For air travel, we surveyed 10,000 individuals, representative of the US population, about their air travel behavior.[4] We ask them to fill in the number of trips taken in the last year (2023) and include

[2] Census tracts average roughly 4,000 households. We discuss the implications of using averages at a census tract, as opposed to individual household data.

[3] We use the 2017 National Household Transportation Survey (NHTS) even though the 2022 version of the data is available as it was collected during the COVID-19 pandemic. Thus, it does not provide an accurate representation of household miles traveled in a year.

[4] We describe the survey in detail in Appendix G.

the IATA airport codes for origin and destination airports for these trips, along with their demographic characteristics. To calculate the distance between airport pairs, we used the latitude and longitudes of airports and the Haversine distance formula, also called the great-circle distance. This procedure gives us the total annual air miles traveled by individuals.

2.3 Machine learning models

For our main results, we rely on an adaptive lasso model, specifically a two-step lasso model, shown in equation 1. First, we split the household-level survey data into "training" and "testing" sets. On the "training" set, we run a lasso regression as the first step, with λ_{min} as the penalty choice. λ_{min} is the minimum possible value of λ. Running a lasso as the first step reduces error from overfitting as our model consists of highly correlated independent variables.[5]

Next, we run a K-fold ($K = 10$) cross-validation lasso regression with λ_{min} and λ_{1se}, respectively. λ_{1se} is the maximum possible value of λ within one standard error of λ_{min}. Cross-validation allows us to determine the optimal trade-off between including too few variables (with a large λ) and too many variables (with a small λ). The coefficients obtained from the first-step lasso regression make for the weights in this second-step lasso regression. We use the inverse of the absolute values of coefficients of the first-step lasso regression as the weights for the cross-validation second-step lasso regression. This lasso regression penalizes variables conforming to their importance to the model.

$$L(y, \lambda) = \arg\min_{\beta} \left\| y - \sum_{j=1}^{p} x_j \beta_j \right\|^2 + \lambda \sum_{j=1}^{p} w_j \left| \beta_j \right| \tag{1}$$

where, $w_j = \frac{1}{|\beta_j^{lasso_{fs}}|}$. $\beta_j^{lasso_{fs}}$ are the coefficients from first-step lasso regression, w_j are the weights added to cross-validation second-step lasso regression. The penalty term λ adjusts how restrictive the model will be, with larger values of λ leading to a more restrictive model selection.

We evaluate several predictions, varying the final prediction function once features are selected, the functional form of the model, and the possible set of features. We compare the results of two choices for the prediction functions—an OLS prediction function or a cross-validation lasso prediction function. The OLS prediction function takes the selected variables and uses the parameters from a final OLS regression for prediction. The lasso prediction function uses the (shrunk) lasso parameters directly. We consider two λ choices, λ_{min} or λ_{1se}. While the base set of possible features is dictated by those that are included in both the representative sample and the ACS data, we also vary the matrix of the explanatory variables by considering higher orders of the variables. In particular, we consider: (a) a base model with just the levels of the features, (b) a model that includes squares of each variable, (c) a model with first-order interactions, and (d) a model representing a second-order Taylor series expansion with squares

[5] We also tried other versions of adaptive lasso using ridge or OLS as the first step. Consistent with (Ballout et al., 2023), we find lasso outperforms the others.

and interactions. Finally, we consider two choices of functional forms, logs and levels. This yields 32 models to choose from.[6] The chosen model is the one with the minimum out-of-sample mean squared error (maximum out-of-sample R-squared).

We note that the Residential Energy Consumption Survey (RECS) is a mix of self-reported and provider-reported household energy consumption. Some households report zero energy use. For our purpose, we exclude the households with zero electricity usage. A fair proportion of households report zero natural gas, propane, or kerosene usage due to constraints on the natural gas distribution network. To overcome biased predictions, we divide the survey data into two parts, one with zero energy consumption and another with positive energy consumption for each of the three sources. We find corresponding census tracts in the ACS data using the variable reporting the proportion of households using a particular energy source (natural gas, propane, or kerosene) as the primary fuel for the household. This distinction allows us to use the characteristics of households with zero (and positive) energy consumption to predict energy use on the census tracts with zero (and positive) proportion of households using fuel as the primary energy source. For other models, we do not make such a distinction in reported zeros.

2.4 Tract-level data

Once we have our preferred prediction model, we use data at the census tract level to estimate consumption for the average household within each tract. We use the American Community Survey (ACS) to gather average household characteristics per census tract. (U.S. Census Bureau, 2022) We use the five-year estimates for 2022 as these provide detailed tables on an extensive set of variables.[7] A census tract is a subdivision in a county and consists of approximately 4,000 residents. We also supplement the ACS with weather information and building efficiency as these variables are in the survey data but not the census. We use the state-level cumulative number of cooling degree days (CDD) and heating degree days (HDD) from the National Oceanic and Atmospheric Administration (2022). Next, we add the International Energy Conservation Code (IECC) code to each census tract based on state and county. Finally, we also include data on average fuel economy, i.e., miles per gallon, for each zip code, matched with the census tracts.[8] Because the models are trained on individual household data, yet we evaluate the models using census tract average values for the features, we potentially suffer from Jensen's inequality bias. However, we note that in practice, very few non-linear functions of the features are selected. Furthermore, below, we compare our predictions to national data on actual outcomes.

To account for missing values in our data, we implement multiple imputations on the ACS data for the census tracts with zero population. The multiple imputation variables, such as tract population and number of households, are based on geography, i.e., state, county, and tract identifiers. We perform five

[6] Two prediction models times two λs times four sets of feature interactions times two functional forms.

[7] The `tidycensus` package in R is an easy-to-use tool to load the American Community Survey data from an API key. (Walker, 2019)

[8] To match zipcodes with census tracts, we use the `zipcode` package in R and USPS zipcode crosswalk files from the US Department of Housing and Urban Development. (Refer here)

iterations of multiple imputations and use the mean of all iterations to generate the data for our analysis.

We predict the average household consumption for each of the nine outlined models at the census tract level. Then, we multiply these estimates with the relevant carbon emissions factor to obtain the average household carbon footprint at the census tract level.

2.5 Carbon emissions factors

We use carbon emission factors to convert per-unit predictions to carbon footprint estimates from various sources. The Department of Energy's Emissions & Generation Resource Integrated Database (eGRID) provides the carbon intensity of the grid for each North American Electric Reliability Corporation (hereafter, NERC) region (see Figure 1).[9] (U.S. Environmental Protection Agency, 2020)[10] For other fuel types, we use the Greenhouse Gas Equivalencies Calculator provided by the US Environmental Protection Agency (EPA). U.S. Environmental Protection Agency (2023) Further, we obtain the carbon emission factors per dollar expenditure on goods and services from Ummel (2014), adjusting for inflation relative to 2021. For vehicle miles traveled, we use data on average miles per gallon at the census tract level from the IHS Markit report and the carbon emissions per gallon of gasoline from U.S. Environmental Protection Agency (2023). Finally, we use calculations from Center for Sustainable Systems, University of Michigan (2023) and International Civil Aviation Organization (2024) to calculate carbon emissions for domestic and international air miles, respectively. Appendix D elaborates on carbon emissions factor calculations. The following section expands on the model used to estimate household consumption and carbon footprints.

[9] The actual marginal emissions from electricity use for a given location at a given time depends on the level and distribution of demand at that time, as well as the flows and capacities of electricity across the transmission network. Therefore, it is overly simplistic to aggregate to any geographic area. However, the relevant carbon emission factor is not the *actual* carbon emission factor but the carbon emission factor that a given policy would use. For this reason, we feel using NERC subregions is accurate.

[10] Geo-spatial analysis allows us to match NERC subregions to census tracts

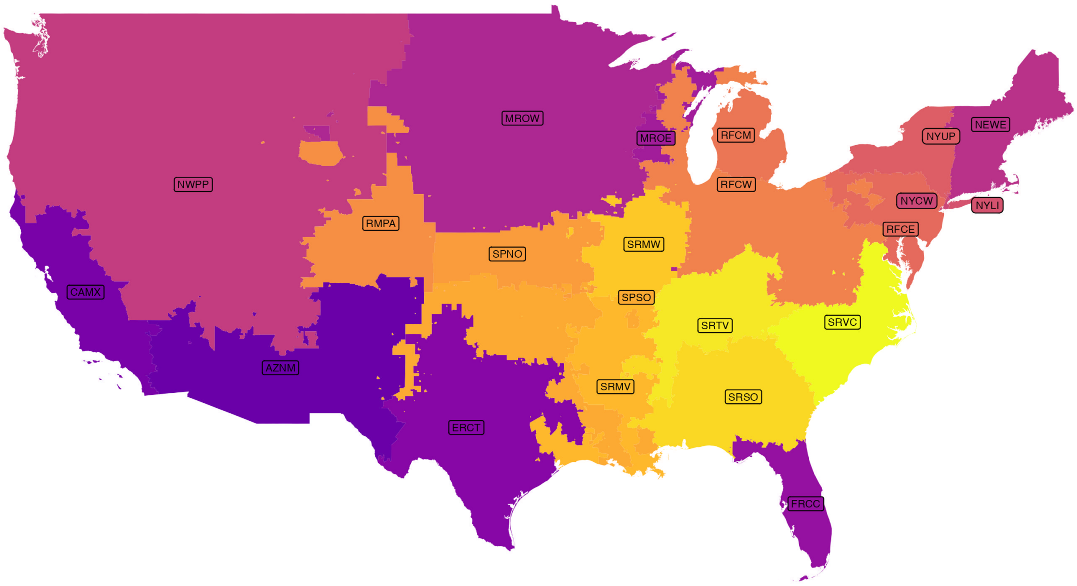

Figure 1: Subregions of the North American Electricity Reliability Corporation used to assign emissions factors to electricity consumption

2.6 Policy models

Given our data on household consumption and carbon footprints, the next step is to measure the effects of different climate policies on each household. As noted above, we consider carbon and gasoline taxes, as well as performance standards applied to vehicles and the grid. Importantly, for the carbon taxes, we further consider different ways to recycle the revenues. We model nine policies listed in Table 1. We describe them in more detail below.

Policy 1	Carbon tax with no revenue recycling
Policy 2	Carbon Price and Dividend (CPD)
Policy 3	CPD, adjusted for division and urbanity
Policy 4	Gasoline Tax
Policy 5	Corporate Average Fuel Economy (CAFE) standard
Policy 6	Electricity Clean Energy Standard (CES)
Policy 7	Electricity Carbon Emission Intensity Standard (CEIS)

Table 1: Policy models to estimate the distributed effects on US households.

To model the impact these policies have on households, we make two simplifying assumptions that, given that our focus is on the relative incidence of the policies across households, are unlikely to affect our conclusions. The first is that we focus only on the portion of a tax or performance standard that is passed through to consumers. For example, we do not attempt to measure the relative incidence across consumers and firms. Therefore, when we report the impact of, say, a $100/tonne carbon tax, the reader should interpret this as a carbon tax where consumers pay $100 per tonne. This may represent a $100 carbon tax that is fully passed through to consumers or a, for example, $200 per tonne carbon tax that is equally paid by firms and consumers. This does not detract from the focus of the paper since everything scales up and down with the value.[11] The second simplifying assumption is that we take existing consumption as given. Therefore, we do not allow consumers to reduce their consumption. In a standard demand curve analysis of consumer welfare, this implies we are effectively including the upper triangle defined by the change in consumption and the change in price. Once again, however, ignoring this triangle will only affect our results if the variation in demand elasticities is correlated with the group definitions discussed below.[12]

We analyze the impact of the policies across income deciles, urbanity, and geography. We discuss the average impacts as well as the distribution of impacts. We next turn to the specifics for each policy.

2.6.1 Policy specifics

A carbon tax with no revenue recycling is an edge-case scenario. In this case, the policymakers correct for the externality and commit the revenue to pay down the deficit or any other use of funds that do not benefit current households. It provides a baseline for other models. The carbon price and dividend plans impose a carbon price and pay back a fixed dividend, estimated using the following equation:

$$Dividend = \frac{\sum P_{CO_2} \times CO_2^{CT} \times HH^{CT}}{\sum HH^{CT}} - P_{CO_2} \times CO_2^{CT} \tag{2}$$

where P_{CO_2} is the price of CO_2, CO_2^{CT} is the predicted carbon emissions (in tonnes/household) for an average household in the census tract, and HH^{CT} is the total number of households in the census tract. For the CPD plans that adjust for household characteristics, such as division and urbanity, we average the carbon tax and dividends within these categories to estimate the net benefit to each household.

For the gasoline tax estimates, we use the current federal gasoline tax of \$0.184 per gallon and calculate:

$$Tax\ Paid_{GAS} = -0.184 \times \frac{1}{MPG_{CT}} \times VMT_{CT} \tag{3}$$

We also measure the household benefit from the gas tax. The Highway Trust Fund ensures that federal gas tax revenues are directed toward constructing and maintaining the interstate highway system and other federal and state road projects. We allocate the revenues from our tax simulations back to households using the per-capita state-level allocations from the 2023 Highway Trust Fund.[13] Each household is assumed to get an equal share within a state. This leads to the following measure of

[11] There is, admittedly, one caveat to this statement. If the relative incidence between firms and consumers varies across products, then variation in household expenditure weights will also drive the relative incidence across *consumers*. We expect this to be second order.

[12] In Appendix F, we do consider a positive income elasticity. This is potentially important since the change in disposable income from the tax-and-dividend programs is strongly correlated with income. However, our conclusions are not altered.

[13] These are available here: here.

incidence:

$$Incidence_{GAS} = Tax\ Paid_{GAS} + \frac{Share^{State} \times \sum Tax\ Paid_{GAS}{}^{CT} \times HH^{CT}}{HH^{State}} \tag{4}$$

where MPG_{CT} is the average miles per gallon for the census tract, VMT_{CT} is the annual vehicle miles traveled by an average household in a census tract, and \$0.183 is the federal gasoline tax.

For the Corporate Average Fuel Economy (CAFE) standard, we calculate the implicit tax (or subsidy) on household-owned vehicles (see, Kwoka Jr (1983) and Holland et al. (2009)). Vehicles with worse fuel economy than the standard get taxed, and vehicles with fuel economy better than the standard get subsidized following this equation:

$$Incidence_{CAFE} = \lambda \times \left(\frac{1}{MPG_{Tract}} - \frac{1}{MPG_{Standard}} \right) \times VMT \times g \tag{5}$$

where P_{CAFE} is the price of the vehicle, MC is the observed price of the vehicle, λ is the shadow value of the constraint at \$100/tonne, MPG_{Tract} is the miles per gallon for an average household in each census tract, $MPG_{Standard}$ is the US average miles per gallon, i.e., the assumed CAFE standard. We multiply this by VMT, which is the lifetime vehicle miles traveled, and g, which is the gallon to CO_2 metric tonnes conversion.[14]

We also model a "clean energy standard" (CES), which requires the regulated entity, typically utilities, to sell a percentage of energy from a "clean" source. The incentives of a CES are similar to those of a performance standard: clean electricity is subsidized, non-clean electricity is taxed. In practice, most CES and renewable portfolio standard (RPS) programs make these taxes and subsidies explicit; qualifying clean energy producers get credits that can be traded in the market, generating an implicit subsidy. Selling non-clean electricity is implicitly taxed since it requires the utility to procure additional clean electricity.

We must make an assumption as to how such a national program would be implemented in terms of the regulated entity. For simplicity, we assume that each state must meet a certain national target equal to the current share of carbon-free electricity. Energy providers with more carbon-intensive energy buy credits (which would cost \$100/tonne) from providers with less carbon-intense energy. Thus, areas with cleaner electricity will benefit through lower rates, while areas with dirtier electricity will pay higher rates. The household cost (or benefit) is the implicit tax (or subsidy) multiplied by the electricity consumed. The policy cost of the CES is calculated by:

$$Incidence_{CES} = \lambda \times CI^{US} \times \left(CFShare^{US} - CFShare^{State} \right) \times ELEC^{CT} \tag{6}$$

where $CFShare_{State}$ is the share of carbon-free energy for the state, $CFShare^{US} =$, is the US share of

[14]Average lifetime vehicle miles traveled is 160,000 miles over 12 years, or 13333.33 miles per year. $g = 0.0315$

carbon-free energy, λ is the shadow price at \$100/tonne, CI^{US} is 0.81 pounds/Kwh, and $ELEC^{CT}$ is the predicted value of household electricity consumption (in Kwh) per year at the census tract level. We treat as carbon-free electricity renewables, nuclear, and hydroelectric generation.

Finally, we model a carbon intensity standard applied to the electricity market. This differs slightly from the clean energy standard (CES). As noted, a CES creates an implicit subsidy on carbon-free generation and an implicit tax on any non-carbon-free generation. Importantly, however, the implicit tax on non-carbon-free electricity does not vary with the generation type; therefore, for example, coal is taxed by the same amount as natural gas despite having a carbon intensity of roughly twice that of a combined-cycle natural gas plant. The carbon intensity standard would, on the marginal, tax a state with a higher share of coal resources. The following equation represents the cost of this policy scenario:

$$Incidence_{CIS} = \lambda \left(CI^{State} - CI^{US} \right) \times ELEC^{CT} \tag{7}$$

where $\lambda = \$100/tonne$, CI^{State} is the carbon intensity of electricity (in tonne/Kwh) for the state, CI^{US} is the US average carbon intensity of electricity (in tonne/Kwh), and $ELEC^{CT}$ is the predicted value of household electricity consumption (in Kwh) per year at the census tract level. We use the CO_2 emissions (in metric ton) for 'Total Electric Power Industry' including all sources (includes coal, natural gas, petroleum, others).[15] We convert these CO_2 emissions to percentage share of state for the CI^{State} term. For CI^{US}, we use an average of state percentage shares.

3 Results

3.1 Household carbon footprint

Figure 2 shows the distribution of carbon footprints across the United States. The population-weighted average carbon footprint is 18.67 tons per household per year. Rural households generate more carbon emissions than urban households. While the average carbon footprint of urban areas is 18.56 tons per household per year, it is 21.14 tons per household per year for rural areas. We see a "donut" trend around major metropolitan areas, where the city center has low carbon emissions and the suburban areas outside the city have high emissions. We present zoomed-in versions of nine notable cities in Figure 3. An expanded view of New York City, for example, highlights this effect of urbanization: an average household on Long Island has a footprint nearly 1.8 times larger than that of an average household in Manhattan. We see similar trends for Boston (Massachusetts), Chicago (Illinois), Kansas City (Missouri), Jacksonville (Florida), Nashville (Tennessee), Houston (Texas), Phoenix (Arizona), and Los Angeles (California).

Some rural and suburban areas have lower than average emissions, such as, through the Carolinas,

[15]Source: EIA (2022)

southern Mississippi Valley, and parts of the Pacific Coast. This is driven by different factors. For one, more households in the former two regions have a lower income and, thus, a lower carbon footprint. Also, the latter has households with higher incomes but operates on a grid with lower emissions intensity and in a climate that does not necessitate the same level of energy required for cooling and heating.

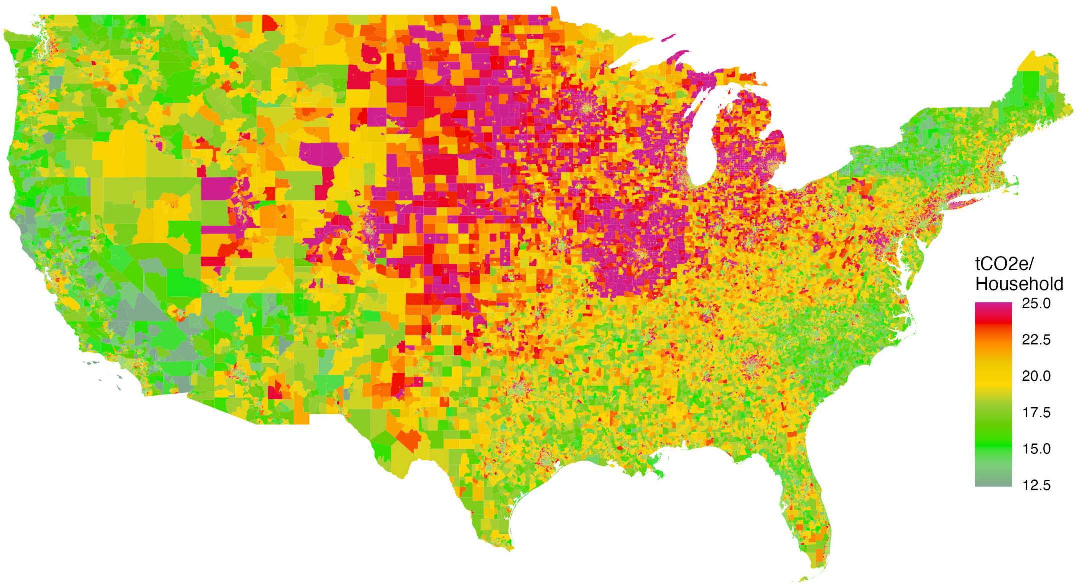

Figure 2: Total Household Carbon Footprints for the Continental United States

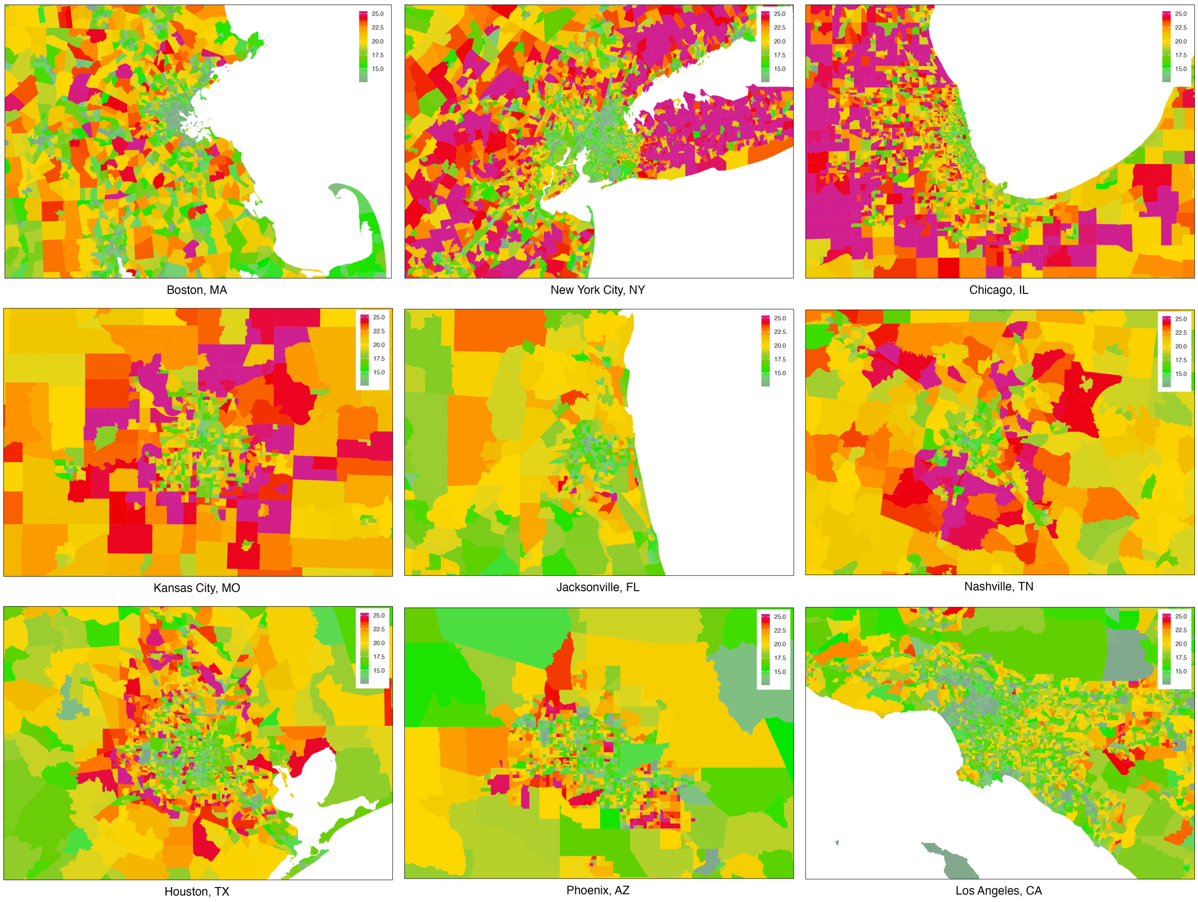

Figure 3: Total Household Carbon Footprints for the Most Populous City in Each Census Division

Table 2 highlights the percentage contribution of each consumption category in the household's total average carbon footprint. Vehicle miles driven and electricity usage are the top two contributors to household carbon footprint. Figure 5 shows spatial differences in carbon footprints of different consumption categories.

Figure 4 displays general patterns in the components of the HCFs across Census regions, income quintiles, urbanization, and NERC regions. We then dive deeper into each of these below. Consistent with the map above, we find much higher HCFs in the Midwest compared to the other regions of the country. Much of this difference is coming from energy consumption. Fuel oil- and methane-related carbon emissions are concentrated in regions that rely on those fuels for home heating. One such notable region is the Northeast, where fuel oil is heavily relied upon. Southern homes have higher HFCs than the northeastern and western homes. HCFs are increasing in incomes driven by energy use, goods and services, and air travel.[16] As noted above as well, we see a substantial increase in HCFs among rural homes and a smaller increase for suburban homes. Transportation emissions are greater in the suburban regions due to longer commutes and households owning multiple cars. The footprint associated with electricity consumption is heavily influenced by the emissions intensity of the associated North American

[16] These results are consistent with Ummel (2014) and Jihoon Min et al. (2010)).

Electric Reliability Corporation (NERC) subregions. The transportation footprint is concentrated in the Midwest, where fuel economy for private vehicles tends to be lower.

Consumption Category	Household's Average Carbon Footprint ($tCO2e$/Household)	Percentage of Household's Total Average Carbon Footprint (%)
Electricity	5.78	30.97
Natural Gas	1.04	5.55
Propane	0.05	0.29
Kerosene	0.22	1.15
Energy	7.09	37.96
Food	1.56	8.33
Health	0.22	1.20
Personal	0.53	2.84
Goods and Services	2.31	12.37
Vehicle Miles	8.13	43.52
Domestic Airmiles	0.64	3.44
International Airmiles	0.51	2.71
Total Airmiles	1.15	6.15

Table 2: Average household carbon footprint by consumption categories

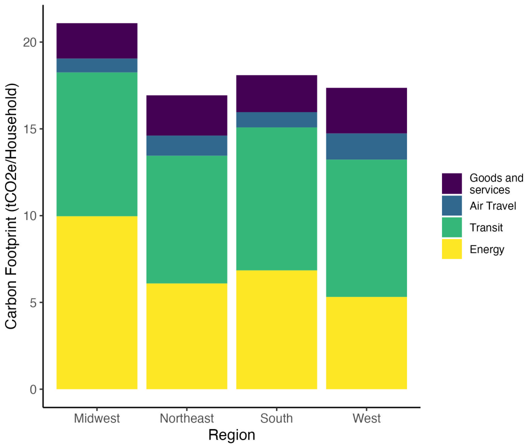

(a)

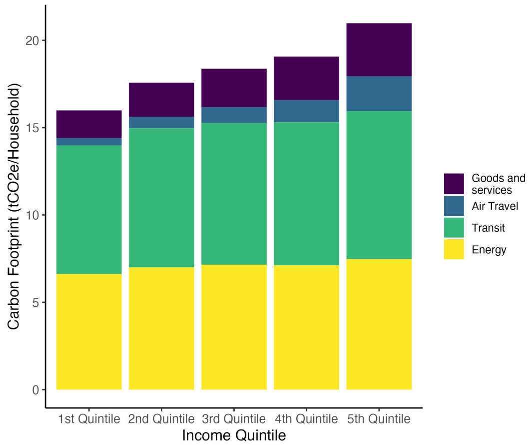

(b)

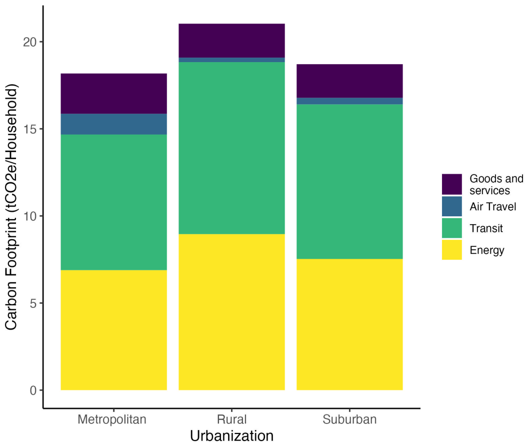

(c)

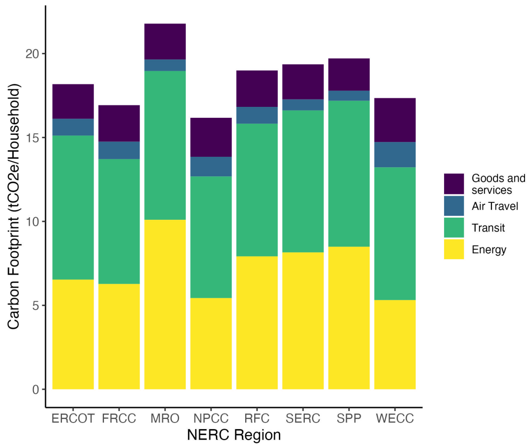

(d)

Figure 4: Total Household Carbon Footprints according to Census Region, income quintile, NERC region, and urbanization, by footprint contribution

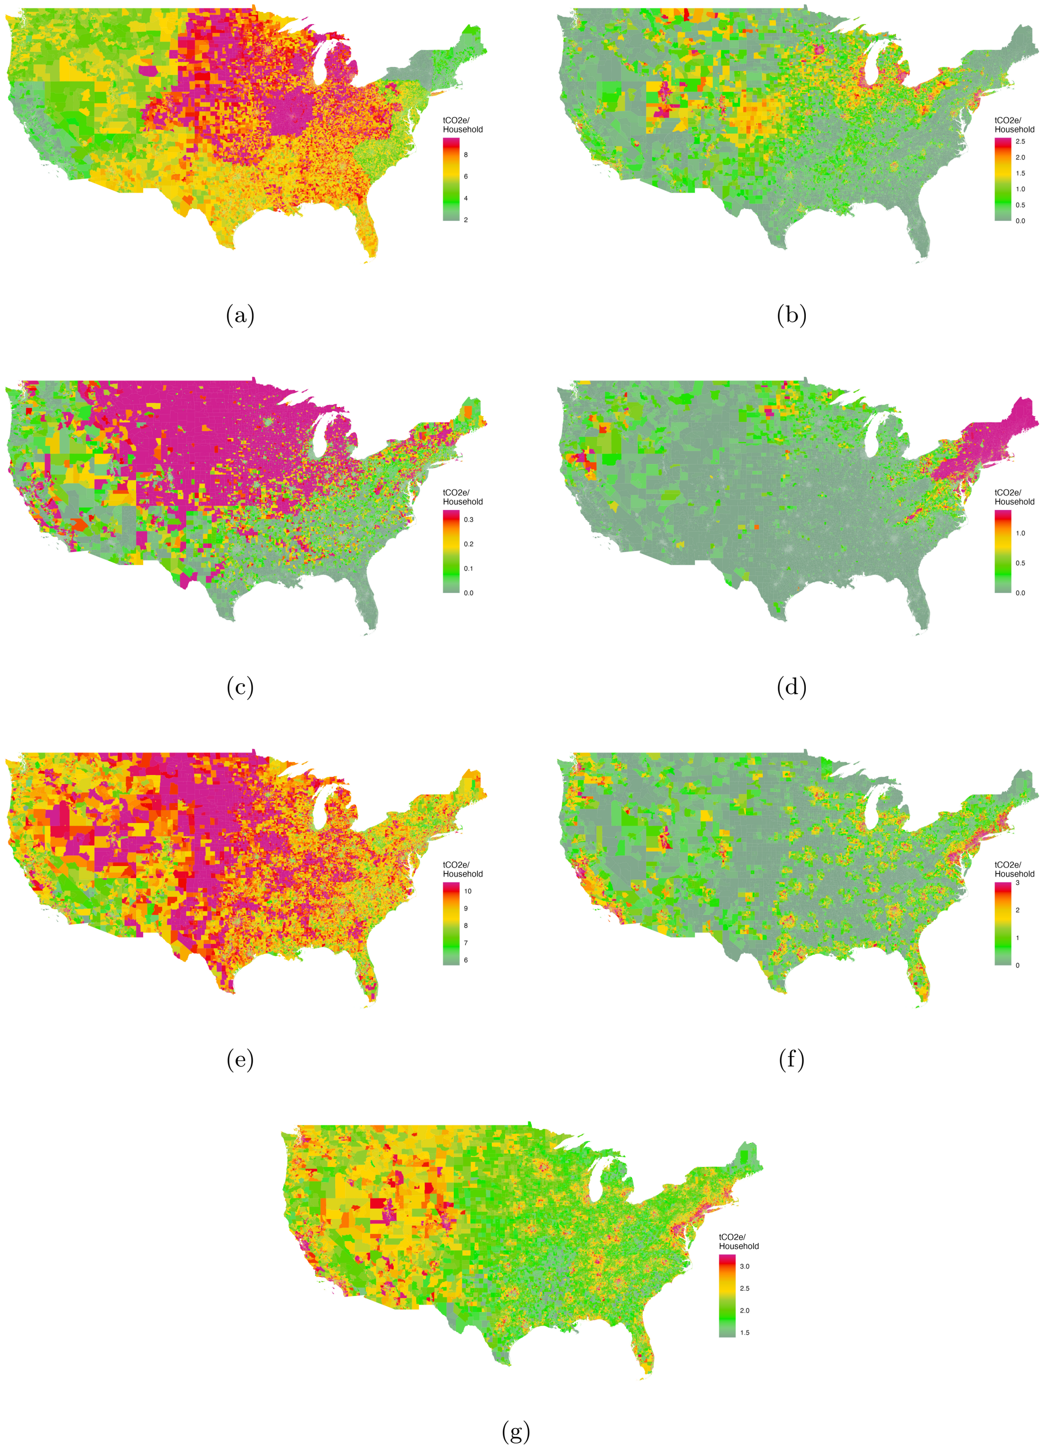

Figure 5: Carbon Footprint by Category: (a) Electricity, (b) Natural Gas, (c) Propane, (d) Kerosene, (e) Vehicle Miles, (f) Air Miles, (g) Good and Services.

Figure 6 provides further insights into how HCFs vary with income and urbanity, plotting the distribution of household carbon footprints across income and urbanity. The dashed line represents the average household carbon footprint across all households. The distribution shifts rightward with an increase in income. Further, for a given income group, increased urbanization shifts the distribution to the left. Figure 4 reiterates that rural households have a higher carbon footprint, primarily arising from higher energy use and greater vehicle miles driven. Metropolitan areas have a higher proportion of carbon footprint coming from air travel compared to suburban and rural areas. In a similar fashion, the top two income quintiles generate a comparatively larger proportion of household carbon footprint from air travel. Household carbon footprints also differ by geography and NERC region.

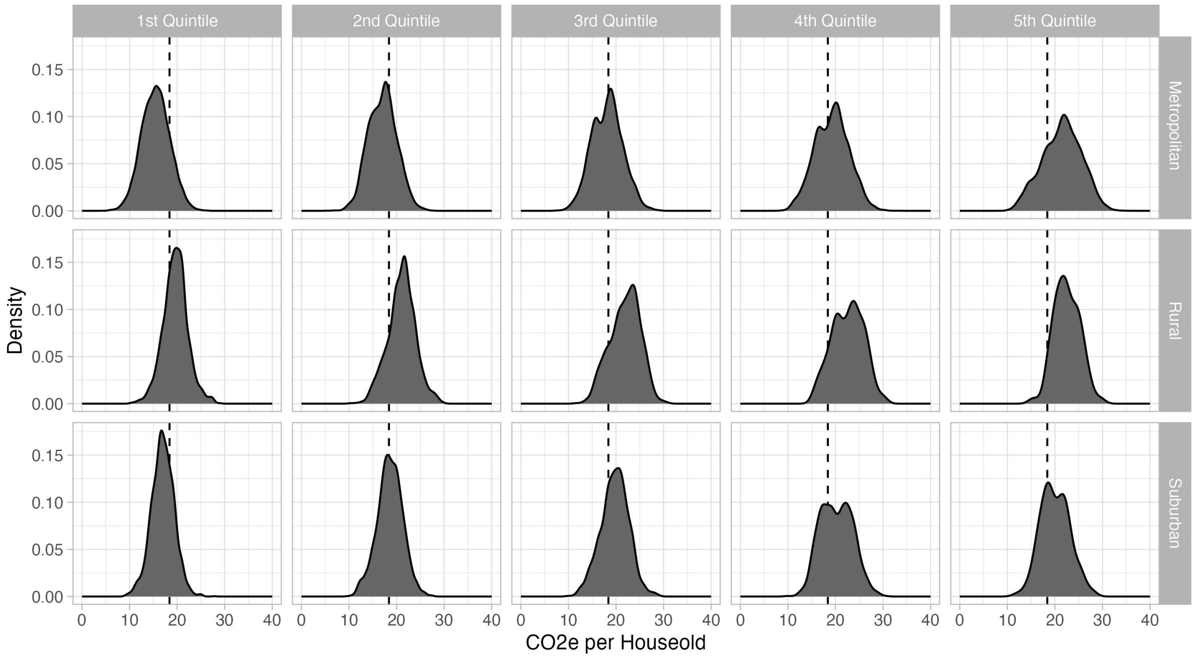

Figure 6: Total Household Carbon Footprints (in tons) across income quintiles and urbanity, compared to US average (represented by the dashed line)

Figure 7 shows the variation in household carbon footprint across party affiliations.[17] While the Republican households have slightly higher footprints than average and the Democratic households have slightly lower footprints than average, there is a large overlap between the two distributions.

[17]We match data on Congressional Districts and legislators of the 117th Congress, including their party affiliations, to the American Community Survey census tracts to do so.

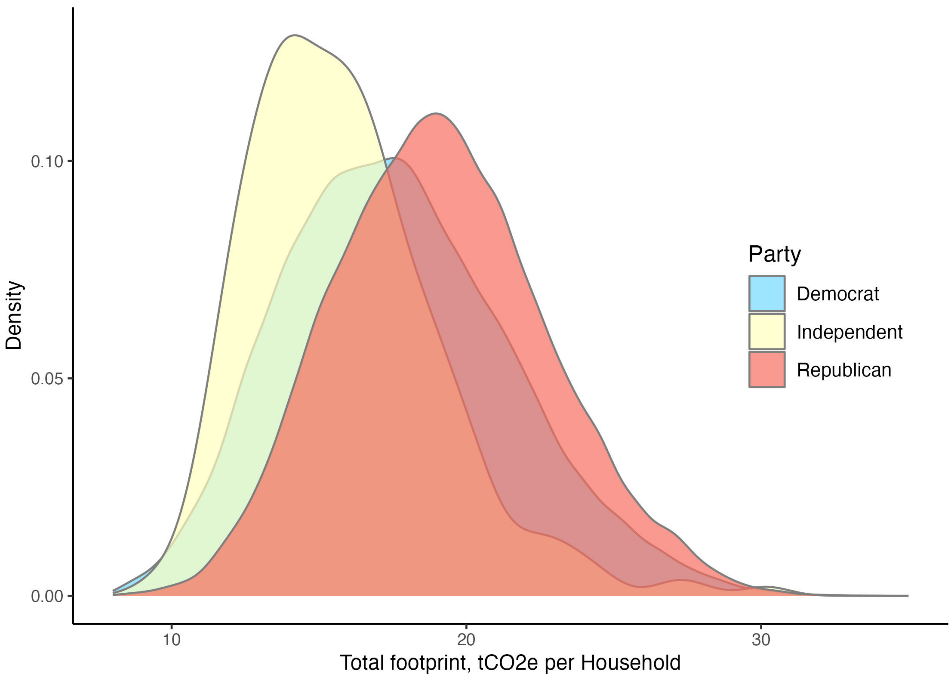

Figure 7: Distribution of Household Carbon Footprints across political parties (according to party affiliation of the House Members in the 118th Congress)

To further understand how carbon footprints correlate with household characteristics and geography, Table 3.1 reports the results of a regression of the log of carbon footprints on demographics, weather, and housing characteristics. We report results with and without state-fixed effects. While these estimates do not reflect causal relationships, looking at the conditional correlations can be useful for our purpose.

We find that cooling degree days (CDD) and heating degree days (HDD) are not associated with higher footprints when accounting for within-state variation. Household income is positively correlated with carbon footprints, but this association is substantially lower than one-for-one. Consistent with the discussion above about urbanity, population density is negatively correlated with carbon footprints. Not surprisingly, the carbon intensity of the grid is positively correlated with carbon footprints.

Turning to household demographics and socio-economic status, older households have lower footprints. Native Hawaiian and Pacific Islander, mixed race, and Hispanic households have higher carbon footprints, conditional on the other variables listed, compared to Caucasian households. No other race or ethnicity shows a statistically significant difference. Carbon footprints fall with education (note we are also conditioning on income), with the omitted group being less than a high school education. We also include several variables capturing the built environment. Interestingly, newer homes, despite having stricter building codes, are not associated with lower carbon footprints. The omitted category is the proportion of homes built prior to 1940. This may reflect the fact that newer homes tend to be larger, counterbalancing the higher efficiency levels. Finally, fossil-fuel-based heating sources are correlated with higher carbon footprints.

	$ln(CO_2)$ (1)	$ln(CO_2)$, with State FEs (2)
Log of HDD	0.005	-0.003
	(0.005)	(0.006)
Log of CDD	-0.013***	-0.007
	(0.004)	(0.008)
Log of Income	0.211***	0.239***
	(0.016)	(0.017)
Log of Household Age	-0.537***	-0.440***
	(0.048)	(0.036)
Log of Population Density	-0.016***	-0.013***
	(0.001)	(0.002)
Log of Electricity Emissions Rate (lb/mWh)	0.322***	0.204***
	(0.033)	(0.035)
Proportion Race African American	0.046**	0.060***
	(0.018)	(0.012)
Proportion Race Asian	-0.070	0.024
	(0.065)	(0.047)
Proportion Race Native Hawaiian & Pacific Islander	0.458***	0.310***
	(0.170)	(0.074)
Proportion Race Mixed	-0.035	-0.047
	(0.057)	(0.042)
Proportion Hispanic	0.060*	0.111***
	(0.034)	(0.038)
Proportion Education High School	-0.208***	-0.175***
	(0.020)	(0.013)
Proportion Education Some College	-0.208***	-0.201***
	(0.031)	(0.015)
Proportion Education Bachelors or Higher	-0.354***	-0.348***
	(0.038)	(0.032)
Proportion Owner Occupied	0.442***	0.385***
	(0.026)	(0.010)
Proportion Home Built 2010 to 2019	-0.025	-0.061***
	(0.016)	(0.012)
Proportion Home Built 2000 to 2009	0.076***	0.046***
	(0.019)	(0.013)
Proportion Home Built 1980 to 1999	0.054***	0.024**
	(0.019)	(0.012)
Proportion Home Built 1960 to 1979	0.059***	0.029**
	(0.018)	(0.014)
Proportion Home Built 1940 to 1995	0.064***	0.044**
	(0.022)	(0.021)
Proportion Heat with Propane	0.057**	0.066**
	(0.027)	(0.026)
Proportion Heat with Electricity	0.009	-0.018
	(0.034)	(0.026)
Proportion Heat with Kerosene	0.122***	0.341***
	(0.039)	(0.028)
Proportion Heat with Natural Gas	0.116***	0.123***
	(0.029)	(0.021)
Observations	83,507	83,507
R^2	0.755	0.787

Table 3: OLS regressions of log of the average household carbon footprints at the census tract level on demographics, housing characteristics, and weather conditions. The second column includes state-fixed effects. Standard errors are clustered at the state level.

3.1.1 Model selection, fit, and confidence intervals

Here, we review model selection and performance, estimating the models described in Section 2.3, generating 16 models for each consumption variable. Appendix A Table A1 provides an overview of the "best" fit models for all consumption categories. It reports the choices of the dependent variable, independent variable matrix, λ, the predict function, and out-of-sample R-squared) for these "best" fit models.

For each of the nine consumption categories, we report the out-of-sample test and train R-squared values.[18] For example, Appendix A Table A2 reports the out-of-sample R-squared and out-of-sample adjusted R-squared for household electricity consumption model. In the case of the same out-of-sample test R-squared and adjusted R-squared values, we choose the more conservative model, i.e, the base matrix is chosen over the squares matrix of independent variables.[19] Appendix A Table A3 provides the summary statistics for the "test" dataset, which is a split sample from RECS 2020 on which the model is predicted, predicted data from the lasso predict function, and predicted data from the OLS predict function of the "best" fit model described above.

Similarly, Appendix Tables A4 to A23 report the out-of-sample test R-squared values and the summary statistics for the actual and predicted data for other household consumption models. We split the RECS 2020 into "test" and "train" datasets for household energy consumption models (including electricity, natural gas, propane, and kerosene consumption). (Appendix Tables A3, A5, A7, A9) For expenditures on food, health, and personal care, we use CEX 2022 data to train and test our models. (Appendix Tables A11, A13, A15) We use NHTS 2017 for household vehicle miles traveled, as shown in Appendix Table A17. Finally, we use our survey data to train and test the household air miles as shown in Appendix Tables A19, A21, A23.

The fit of the models varies significantly across the consumption category. Our models do best in predicting electricity consumption, followed by natural gas and propane consumption. The models do less well with the other consumption categories. Fortunately, these products represent smaller shares of a typical household's carbon footprint. Most importantly, however, despite the relatively poor predictive power of some of the models, the 95% confidence intervals in our predictions tend to be below 5% for the average household. We report the 95% confidence intervals for our predictions for all consumption categories and model choices in Appendix B. Appendix Table B24 reports the prediction confidence intervals (ΔCI_P) for the household electricity consumption model. We also calculate the prediction confidence intervals as a percentage of the prediction means in Columns (4) and (8) of Appendix Table B24 for predictions from lasso and OLS predict functions. Similarly, Appendix Tables B25 to B34 report prediction confidence intervals for other household consumption models.

Using our "best" fit models, we outline the coefficients of the variables selected when predicted on ACS 2022 in Appendix C.[20] Appendix Table C35 reports the coefficients of the variables selected

[18]For the models predicting household air miles, we report model evaluations for total air miles and domestic and international air miles separately. As shown in Appendix A Table A1, our machine learning algorithm chooses the same model for all three sub-categories of household air miles. We use predictions from household domestic air miles and household international air miles for all our analysis throughout the paper.

[19]In other words, the first incidence of the maximum out-of-sample test R-squared in Appendix Table A2 for the household electricity consumption model is chosen to be the "best" fit model.

[20]Note that Urban and Log of income are forced variables in the first step lasso. We believe these two variables are important predictors of household consumption, and, thus, should be included in our predictive models. We

for models predicting household energy consumption. Appendix Tables C36, C37, and C38 report the coefficients of variables selected for models predicting household food, health, and personal consumption, vehicle miles traveled, and air miles traveled, respectively.

Finally, we also report our out-of-sample mean, median, standard deviation, and min/max consumption levels for each product as well as the mean level of consumption in the test data. In general, our mean predictions are close to those of the test data; however, are not generally able to capture the extreme fat tails, especially the right tail, in our predicted consumption levels.

3.2 Policy impacts

Next, we analyze the impacts of different policies on households. We begin by discussing the results for each policy and then summarize across policies.

3.2.1 Simple tax and dividend plan

We start with a carbon tax and dividend plan where there is a single-sized check sent to households (Policy Scenario 2 in Table 1). Plans similar to this have received a lot of support from organizations like the Climate Leadership Counsel and prominent policymakers such as James Baker and, the late, George Shultz. In 2019, over 3,500 economists signed a letter advocating such a plan.[21]

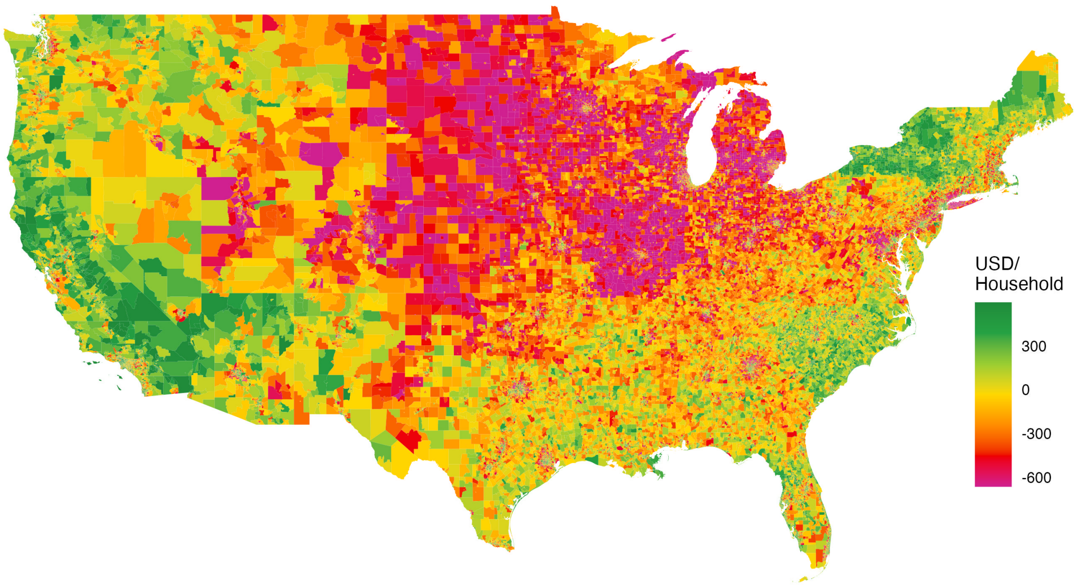

Figure 8: Net Impact of $100 Carbon Price and Dividend

Figure 8 shows the spatial distribution of household impacts for the simple CPD policy. There are three major stylized facts of such a policy that immediately follow from the discussion of HCFs above.

force them irrespective of the first step lasso including or excluding these variables.

[21]This is available here.

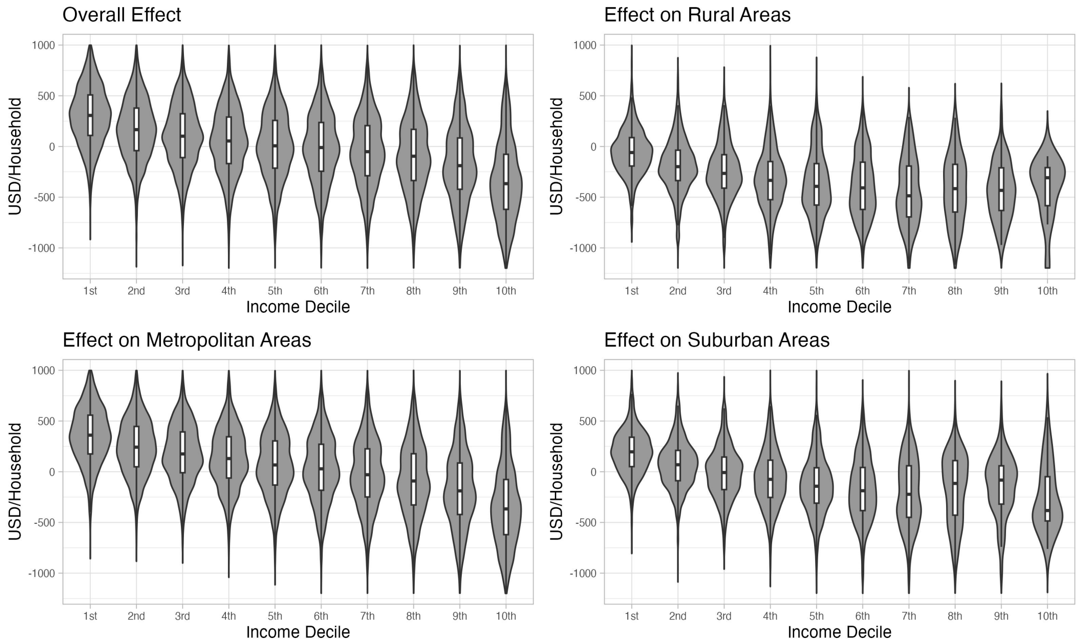

Figure 9: Net Impact of $100 Carbon Price and Dividend, by income decile and urbanity

Two of these facts are likely to be politically challenging, while one is likely to be politically convenient. The politically convenient stylized fact is represented in the upper left panel of Figure 9, which is the distribution of winners and losers across income deciles. A simple carbon tax and dividend plan is progressive, with the average net tax paid increasing as households get wealthier. The bottom 50% of households, on average, receive a dividend check larger than their tax liability. The average (median) net payment from the policy across the bottom five deciles is: $311.91 ($309.57), $171.64 ($166.95), $107.55 ($101.66), $57.79 ($54.80), $10.87 ($7.04), respectively. This is represented in Figure 10.[22] When we view this as a share of income, the results are even more substantial. These net payments represent 0.96% (0.92%), 0.37% (0.36%), 0.19% (0.18%), 0.09% (0.08%), and 0.01% (0.01%) of average decile income, respectively. We do note that a significant share of the below-median income earners pay more in carbon taxes than they receive from the dividend. This share of households with negative net transfers from the policy is 15%, 29%, 36%, 43%, and 49% for deciles one through five, respectively.

The first politically challenging result of this simple tax-and-dividend plan is represented in the other three panels of Figure 9 and flows from the fact that HFCs are lowest in cities. The policy is most progressive in cities, and low-income households are much more likely to gain from such a policy than rural and suburban households. As an example, 90% of bottom-decile households in cities receive a dividend check larger than their tax liability, but only 40% and 82% of rural and suburban bottom-decile households, respectively. Furthermore, for rural and suburban households, the policy is not progressive for the upper three income deciles. This leads to large amounts of money flowing from rural areas to the cities. Figure 11 shows the average rural household transfers roughly $260 to urban households, while

[22]Of course, the specific decile-to-decile transfers are arbitrary.

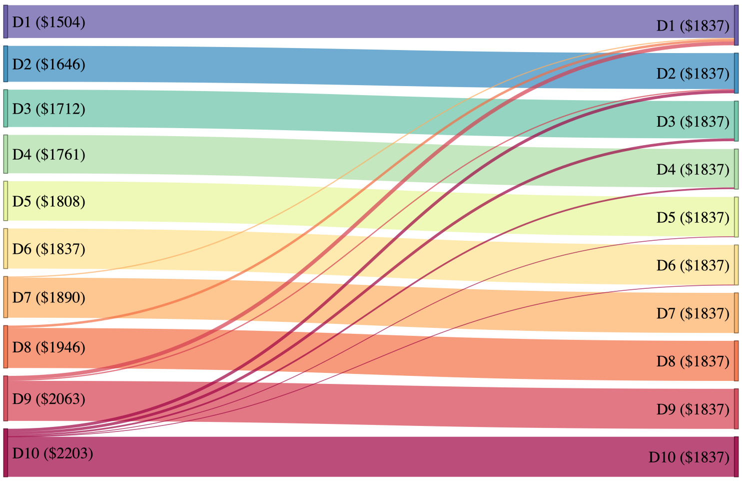

Figure 10: Average transfers between households by income deciles for a Carbon Price and Dividend

the average suburban household transfers $39.

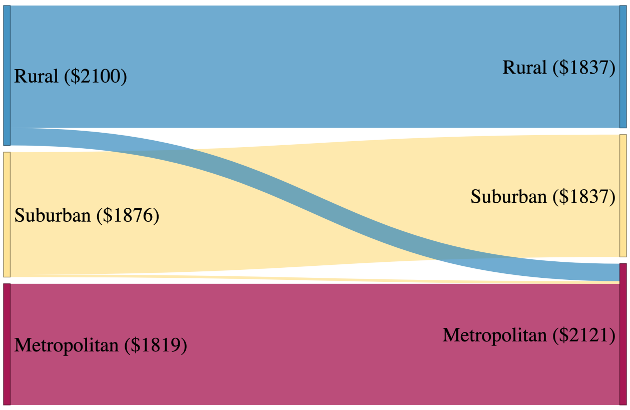

Figure 11: Transfers between metropolitan, suburban, and rural households for a Carbon Price and Dividend

The second politically challenging result of a simple tax-and-dividend plan comes from HCFs being higher in middle America compared to coastal areas. This is represented in Figure 12. We find that the Midwest transfers nearly $7B to the other regions of the US, with $3 billion going to the Northeast, $3 billion to the West, and about $1 billion to the South.

3.2.2 Geographic-dependent dividends

An alternative to the simple carbon tax-and-dividend plan analyzed above is to vary the size of the dividend check to ensure budget neutrality across the nine Census divisions and urbanity. In this scenario, the tax revenue collected would stay within Census divisions and urban/suburban/rural areas within those regions. We calculate the household gains and losses across such a policy. The map by Census

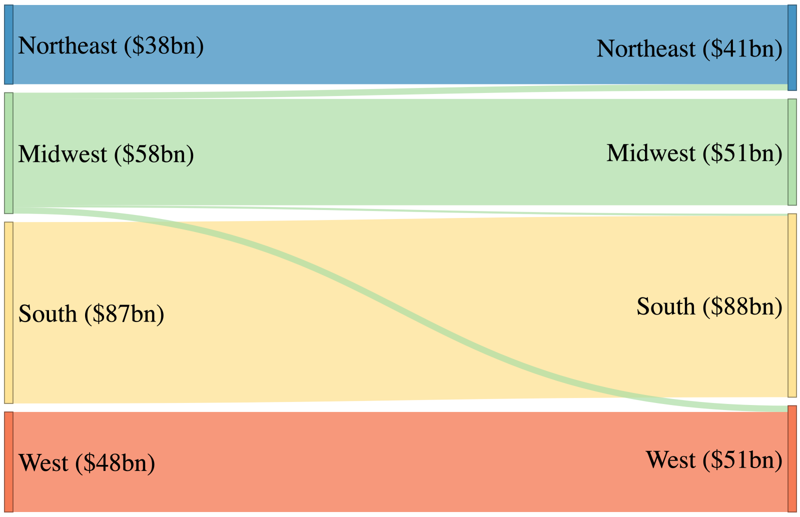

Figure 12: Transfers between households in each Census region for a Carbon Price and Dividend (the left side represents relative tax paid through a carbon price for the average household in each region, the right side represents the dividend received; the flows between regions reflect the transfers between those regions)

tract is in Figure 13. A carbon tax-and-dividend plan structured in this way directly "fixes" the issues illustrated in Figures 12 and 11, in the sense that the policy would guarantee money does not flow from one region of the country to another and from rural areas to cities. This results in a map that looks much more uniform.

The differences in the size of the dividend checks are substantial. Table 4 reports the 27 different dividend check levels. The smallest check size is $1536 (suburban Pacific), while the largest is about $2300 in Mountain North Central areas.[23] One could argue that having 27 different check levels would be too complicated or generate resistance. We do the same exercise where we use the four Census regions instead. In this, the range of the check sizes is from $1653 (suburban Northeast) to $2324 (rural Midwest).

Urbanity and geographic-dependent dividends also make the policy more progressive. Figure 14 plots the violins graphs by income decile for all tracts and by urbanity. While the simple tax and dividend policy was progressive when we focused on the income deciles of all census tracts, within rural households, the simple plan was regressive across the top four income deciles and non-monotonic in suburban areas for the top four income deciles. Targeting dividend checks based on urbanity and geography implies the policy is progressive across all urbanity categories and "more" progressive across all census tracts; that is, the difference between the dividend check for the average first and tenth decile tracts is larger.

[23]We note that suburban areas on the coasts do follow the national pattern of having higher HCFs than cities.

	Metropolitan	Rural	Suburban
New England	1623.96	1791.55	1659.65
Middle Atlantic	1711.88	1854.21	1649.67
South Atlantic	1759.79	1915.57	1736.16
East North Central	2080.75	2324.94	2136.99
West North Central	2105.60	2323.88	2115.14
East South Central	1914.24	2057.41	1897.45
West South Central	1830.83	1995.75	1821.19
Mountain	1924.71	2120.67	1853.14
Pacific	1621.87	1666.34	1536.32

Table 4: Dividend Check by US Census Division and Urbanity

	Metropolitan	Rural	Suburban
Midwest	2087.76	2324.30	2129.90
Northeast	1689.25	1813.37	1652.71
South	1802.42	1995.07	1803.13
West	1718.89	1976.29	1709.68

Table 5: Dividend Check by US Census Region and Urbanity

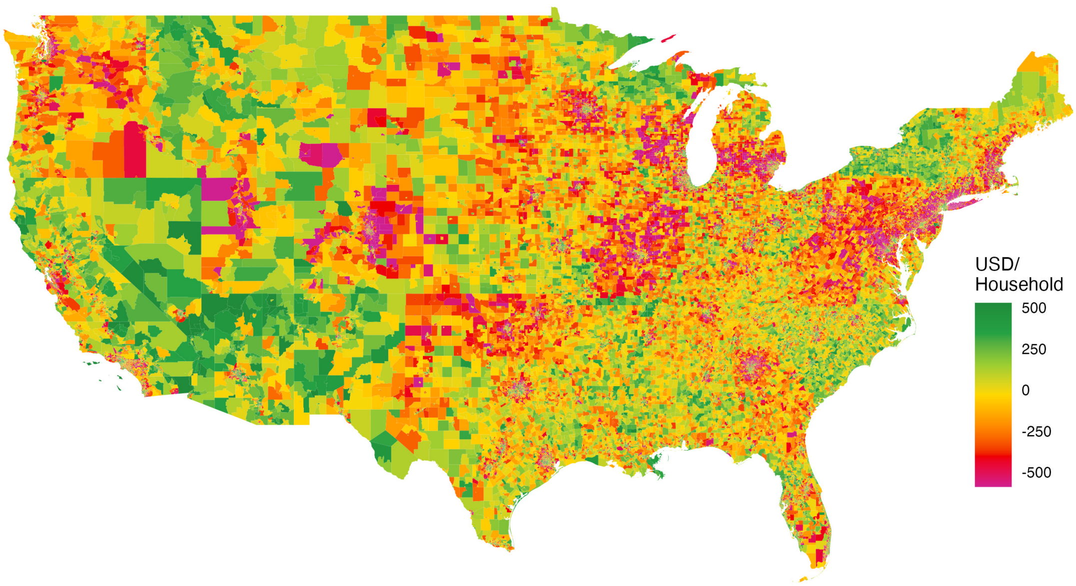

Figure 13: Net Impact of $100 Carbon Price, Dividend adjusted for Urbanity and Geography

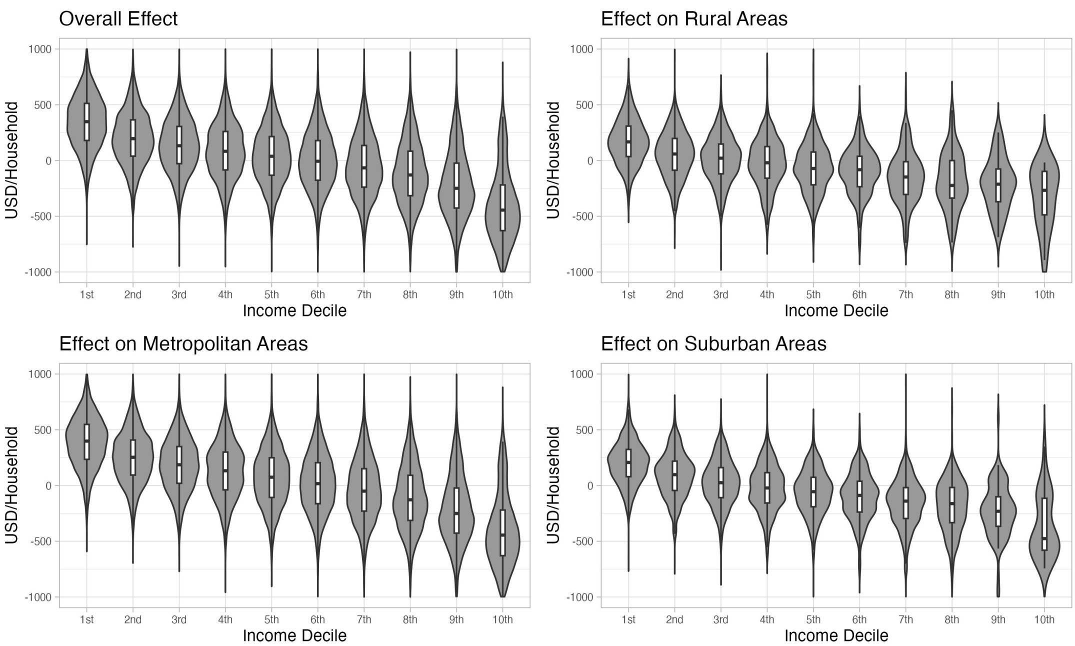

Figure 14: Net Impact of $100 Carbon Price, Dividend adjusted for Urbanity and Geography, by income decile and urbanity

3.2.3 Transportation-specific policies

We also analyze the impact of two transportation-specific policies: a gasoline tax and a fuel economy standard. While not strictly climate policies, gasoline taxes play a critical role in financing road infrastructure in the US, and because they tax driving, they have climate benefits. Furthermore, several policy discussions around carbon taxes have openly discussed exempting gasoline because of the critical role it plays in household finances.

The Gasoline Tax (Policy Scenario 4 in Table 1) imposes a gasoline tax at the current Federal gas tax level. As noted above, we recycle the revenue back to the states on a per-capita basis based on the current state-level allocations of the Highway Trust Fund; therefore, our tax is, by definition, revenue neutral. This leads to very visible state-level winners and losers.

Figure 16 shows the income distribution of household impacts. Given current allocation methods, the Federal gasoline tax is progressive. Households in the bottom three deciles receive net benefits while households in income deciles D4 to D10 pay taxes. (See Table 6). The Gasoline Tax benefits households in rural areas (See Table 7). Figure 15 shows the spatial distribution of household net impacts for the Gasoline Tax. Evident from the map is that the states Montana, North and South Dakota, and Wyoming receive a larger share of the Highway Trust Fund than they contribute, as do West Virginia and Vermont.

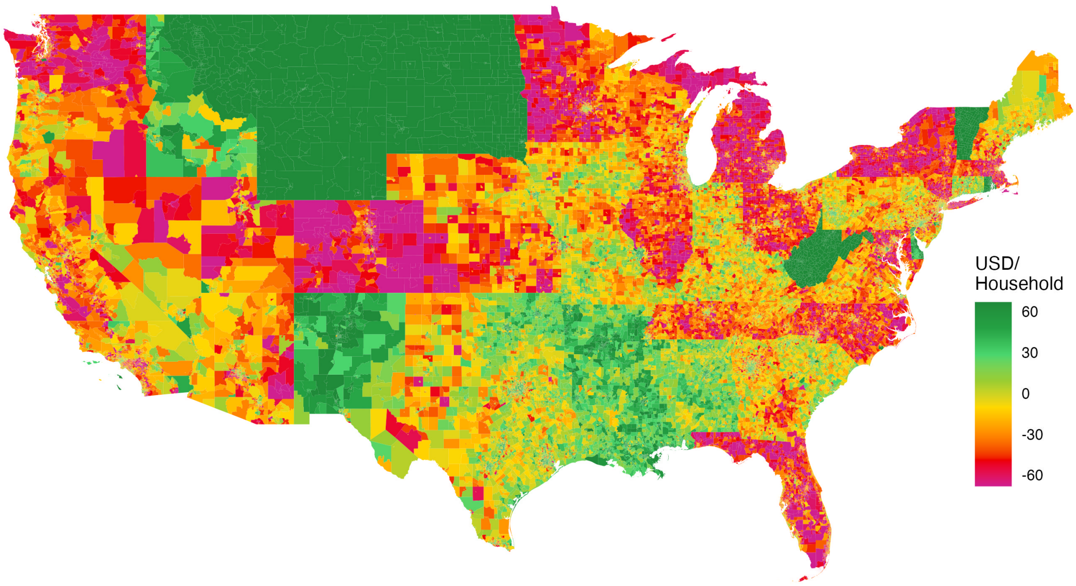

Figure 15: Net Impact of Gas Tax

While gasoline taxes operate on the price of fuel, fuel economy standards affect the price of new vehicles. The Corporate Average Fuel Economy (CAFE) standard imposes an implicit tax (or subsidy) on household-owned vehicles as shown by Equation 5 in Section 2.6. Using data on the purchases of new vehicles, Figure 17 displays the geographic distribution for the CAFE standard. We find the CAFE

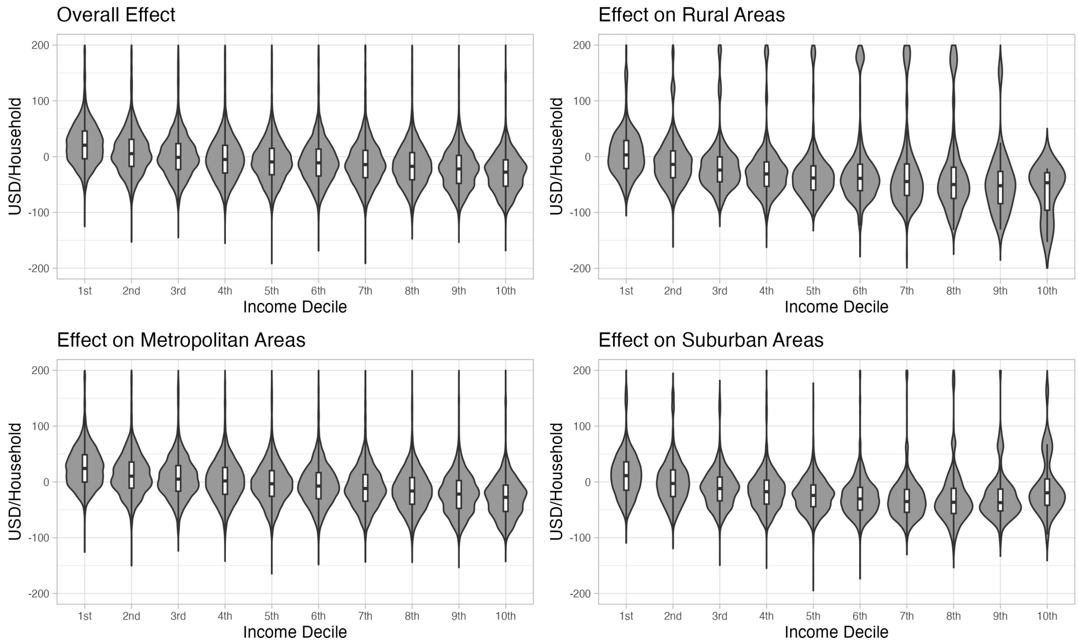

Figure 16: Net Impact of Gas Tax, by income decile and urbanity

standard to be progressive, as shown in Table 6. The bottom income decile earns a net benefit of $47 per household per year while the top income decile loses $104 per household per year. Moreover, the CAFE standard benefits rural and suburban households with an average net impact of $180 and $144 per year, respectively. Figure 18 shows the distribution of net effects across income groups and urbanity for the CAFE standard.

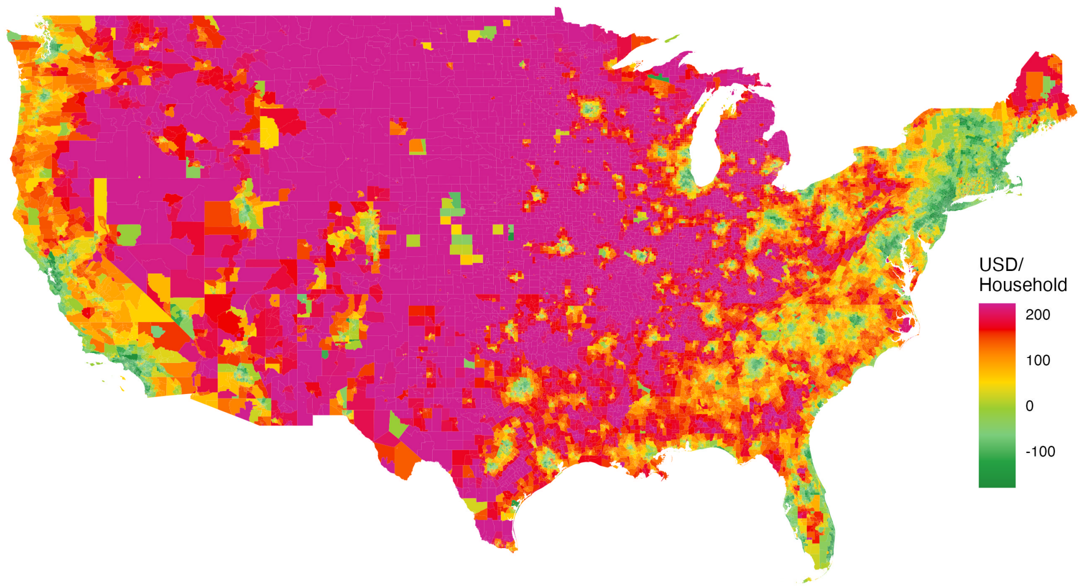

Figure 17: Net Impact of CAFE Standard

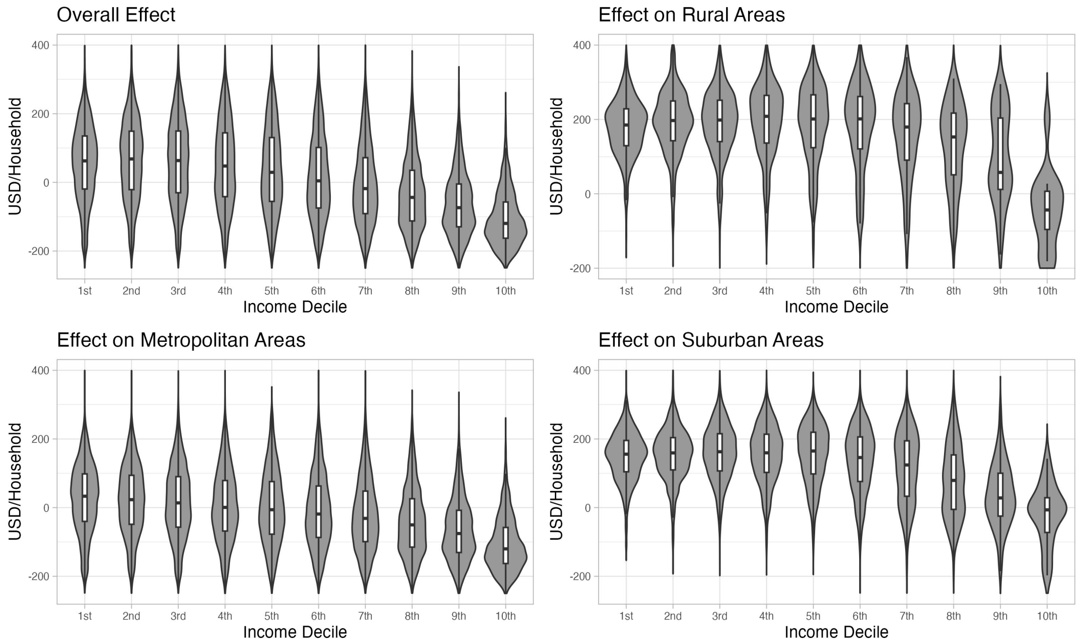

Figure 18: Net Impact of CAFE Standard, by income decile and urbanity

3.2.4 Electricity-specific policies

As with the transportation sector, individual climate policies often target the electricity sector separately. Common among these are either Renewable Portfolio Standards (RPSs) that require a certain share of electricity bought and sold to be from renewables or Clean Energy Standards (CESs) that require a certain share of electricity to be from carbon-free generation. CESs add nuclear generation and certain carbon-capture technologies to the regulated share. These are popular policies at the state level, with over 30 states having either a RPS or CES. A criticism of these policies is that they do not differentiate among the carbon-producing generation. That is, two states with the same level of renewable generation will be viewed equally even if one of the states uses only coal generation for the non-renewable share. An alternative policy would be to set a limit on the carbon intensity generation, similar to the fuel economy standard that sets a floor on average fuel economy.

We analyze the impact of a national clean energy standard and a national carbon emission intensity standard (CEIS). To do so requires us to define the electricity generation mix of an individual household. This is a regulatory choice. In practice, electrons cross state borders based on physical constraints in the network, so a household's actual consumption mix is not easily defined. With that said, one recent proposal for a national policy (the "Clean Energy Performance Program" under the Biden Administration) similar to a clean energy standard defined the generation mix at the state level. Therefore, we use the state generation mix. We use the national average share of clean energy and carbon intensity, respectively, as the baseline. Therefore, households that live in states with low (high) clean energy generation or a high carbon intensity will be implicitly taxed (subsidized), and the higher their electricity consumption, the greater the tax (subsidy) will be. This, necessarily, generates a high degree of correlation of the impact within a state.

Figure 19 shows the spatial distribution across US census tracts. States like Maine, Vermont, and New Hampshire in the New England region and California, Washington, and Oregon constituting the West Coast reap benefits from this policy. Figure 20 shows the distribution of net effects across income groups and urbanity. However, for this policy, the main driver of variation in the incidence is across states rather than income. The households in the bottom four deciles gain \$22 (D1), \$14 (D2), \$6 (D3), and \$5 (D4) on average per year, respectively, while households in the top six deciles lose \$2 to \$36 on average per year. Households in metropolitan areas lose \$8 on average per year, while households in suburban areas gain \$11 on average per year.

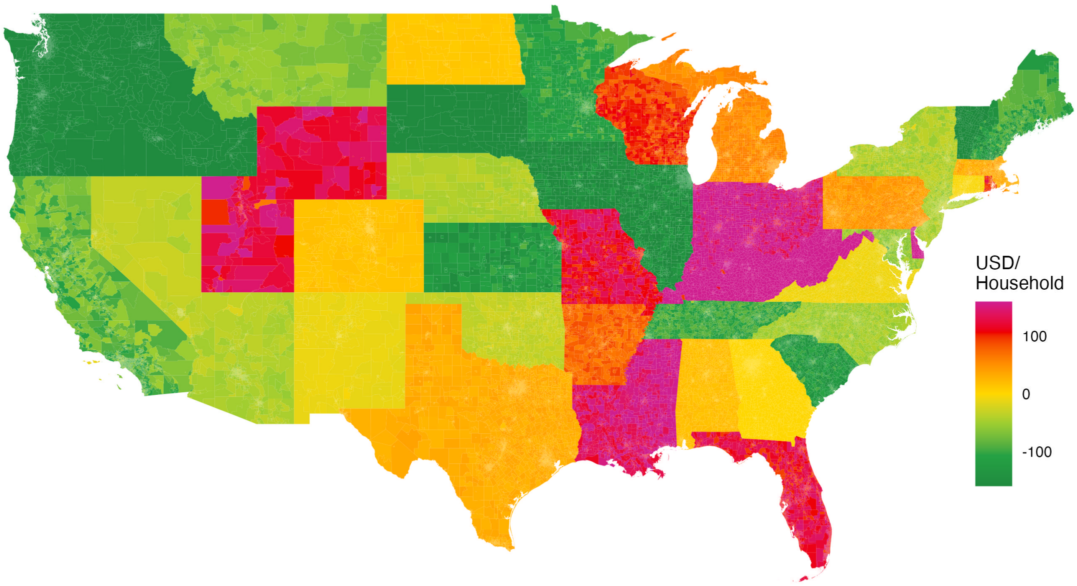

Figure 19: Net Impact of Clean Energy Standard

Finally, we calculate a Carbon Emission Intensity Standard as described in Section 2.6 Equation 7. This standard penalizes the states with a higher carbon share of electricity. The geographic distribution is shown in Figure 21. States like Wyoming and West Virginia are penalized for higher carbon content in their grids. Figure 22 displays the distribution across income deciles and urbanity. Households in the bottom decile lose \$23 per year on average from this policy, while the households in the top decile lose \$71 per year on average. About 70% of households in the bottom decile are penalized under this policy. Rural households, on average, gain \$160 per year from this policy. 40% of rural households incur penalties.

Interestingly, the incidence of the CES and CEIS is not very correlated. Figure 23 plots a scatterplot of the average household incidence of the CES for each state compared to the incidence under the CEIS. They are essentially uncorrelated, pointing to the fact that states with large renewable and nuclear generation levels also tend to have large amounts of coal generation. This also underscores the potential

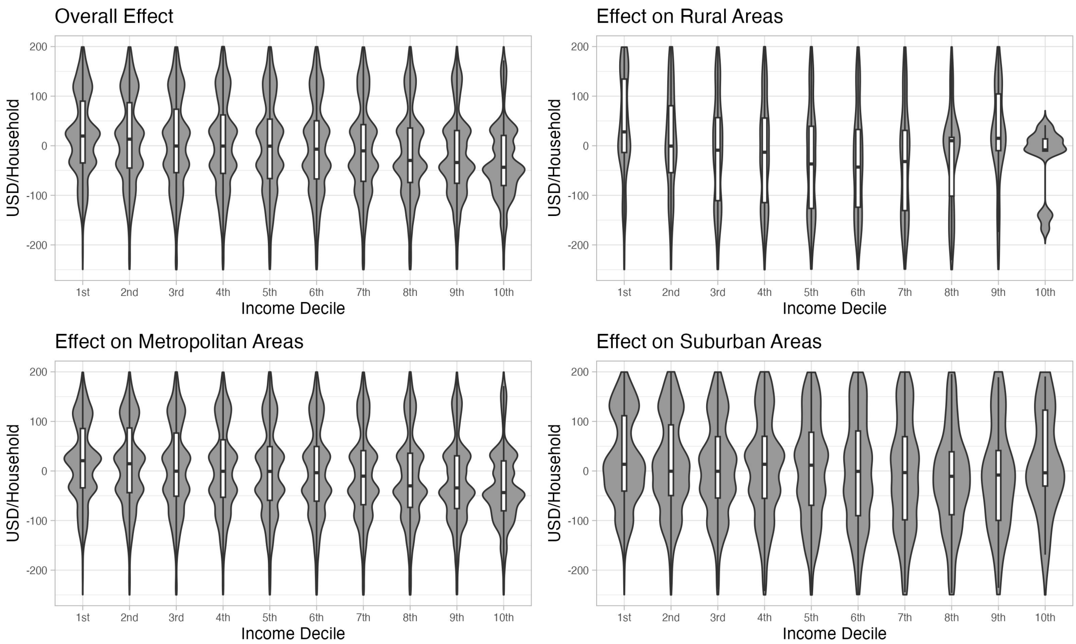

Figure 20: Net Impact of Clean Energy Standard, by income decile and urbanity

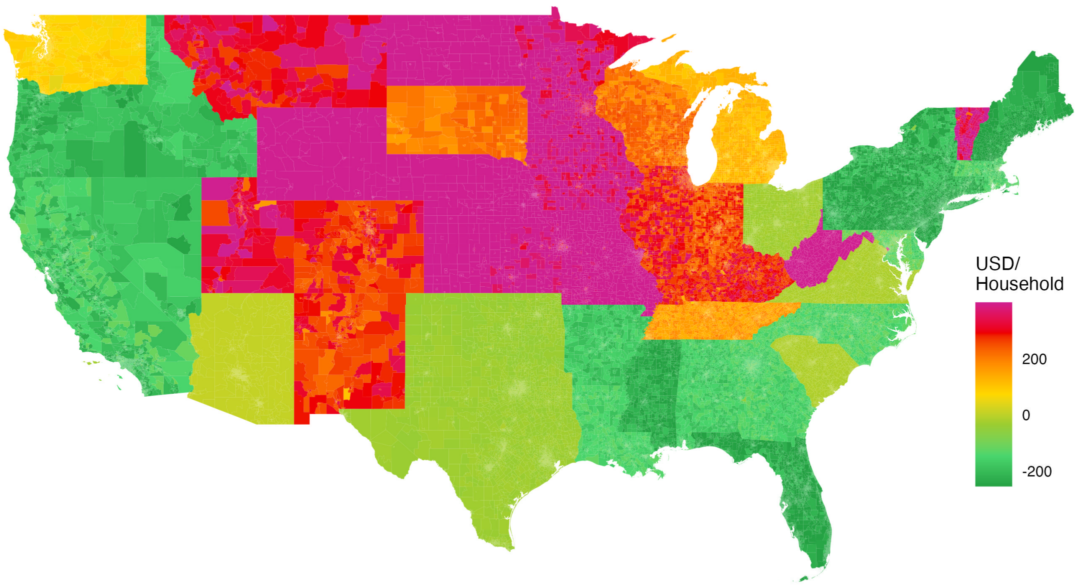

Figure 21: Net Impact of Carbon Emission Intensity Standard

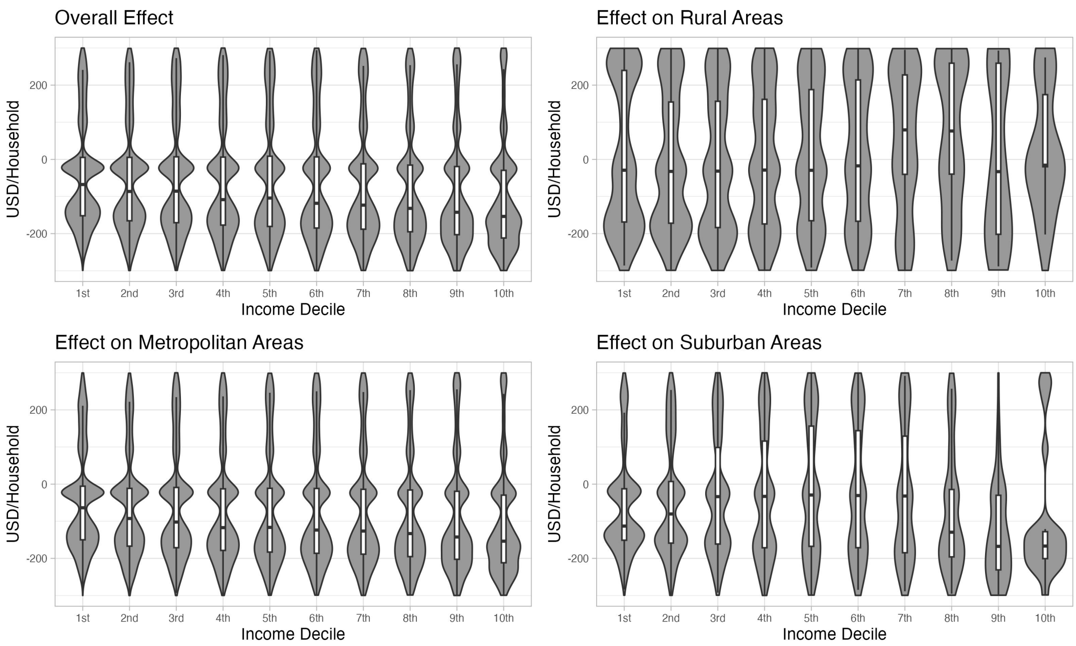

Figure 22: Net Impact of Carbon Emission Intensity Standard, by income decile and urbanity

inefficiency of clean energy standards since, for climate change, carbon emissions are what matters.

Carbon Policy	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10
Carbon Tax	-1525.21	-1665.48	-1729.56	-1779.32	-1826.24	-1858.61	-1902.07	-1949.54	-2042.63	-2204.20
CPD	311.91	171.64	107.55	57.79	10.87	-21.50	-64.95	-112.43	-205.51	-367.09
CPD, adjusted for urbanity and division	336.95	196.56	131.80	79.56	32.26	-13.34	-70.37	-138.01	-246.63	-438.43
Gasoline Tax	21.90	10.63	4.49	-0.11	-3.63	-4.93	-9.61	-15.04	-21.91	-29.83
CAFE standard	46.91	56.06	51.87	42.84	30.80	12.32	-6.65	-33.05	-60.54	-104.41
Clean Energy Standard	21.13	13.73	5.72	4.36	-2.13	-4.27	-9.18	-16.25	-21.43	-35.61
Carbon Emission Intensity Standard	-22.59	-13.96	-7.02	-8.85	-1.38	-11.52	-19.46	-32.45	-40.73	-71.46

Table 6: Net Impact of Policy Scenarios by Income Decile (USD/Household Annually)

Carbon Policy	Metropolitan	Rural	Suburban
Carbon Tax	-1852.35	-2114.27	-1889.63
CPD	-15.24	-277.16	-52.51
CPD, adjusted for urbanity and division	-38.14	-15.07	-15.03
Gasoline Tax	-6.61	4.87	-9.20
CAFE standard	-26.22	179.49	143.70
Clean Energy Standard	-8.16	-2.05	11.27
Carbon Emission Intensity Standard	-39.34	160.55	28.16

Table 7: Net Impact of Policy Scenarios by Urbanity (USD/Household Annually)

4 General comparisons across all policies

The analysis of policy impacts above reveals substantial variation in how different climate policies affect households based on income, geography, and urbanity. Broadly, carbon pricing policies that generate revenue offer greater flexibility in mitigating inequities, while performance-based regulations tend to

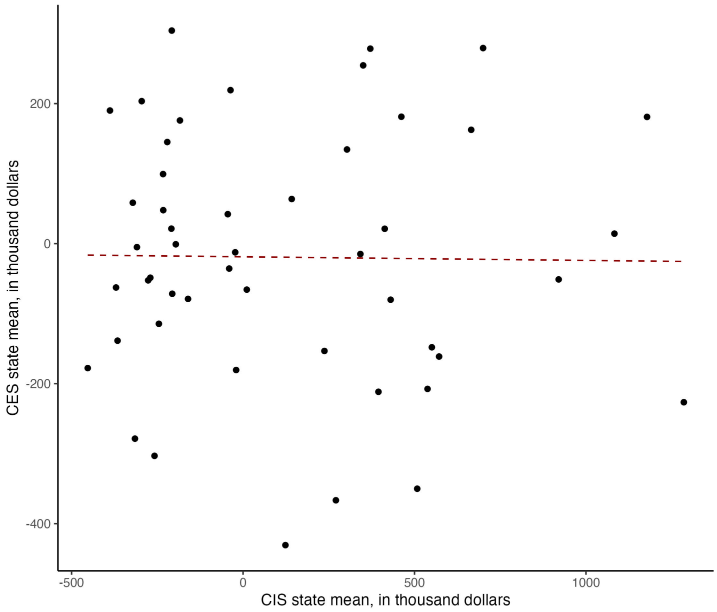

Figure 23: Scatter plot between state-level household mean impacts from the Clean Energy Standard and Carbon Emissions Intensity Standard

impose higher burdens on low-income and rural households without compensatory mechanisms.

A simple carbon tax and dividend approach, where all households receive an equal per-capita rebate, emerges as a generally progressive policy at the national level. Households in the lowest income deciles receive more in dividends than they pay in carbon taxes, while higher-income households bear a net cost. However, this broad progressivity masks significant regional disparities. Rural and suburban households, which tend to have larger carbon footprints due to higher energy consumption and greater reliance on personal vehicles, are disproportionately impacted. This results in net financial transfers from the Midwest and Plains states to urban areas on the coasts. Adjusting the dividend structure to account for regional and urban-rural differences helps to alleviate these imbalances. When dividends are recalibrated at the census division level and adjusted by urbanity, the redistributive effects of the policy become more neutral across regions, reducing opposition from rural and Midwestern households while maintaining progressivity within income groups.

Among transportation policies, a federal gasoline tax disproportionately affects rural households, which tend to drive longer distances and own less fuel-efficient vehicles. However, when revenues from the tax are allocated using the existing distribution of the Highway Trust Fund, the net effect varies by state. Some states, such as those in the Northern Plains and Appalachia, receive more in highway funding than their residents pay in gasoline taxes, leading to net benefits for households in those regions. The Corporate Average Fuel Economy (CAFE) standard, which effectively subsidizes fuel-efficient vehicles while penalizing less efficient ones, proves to be somewhat progressive. Lower-income households, which are more likely to purchase used vehicles, benefit from reduced fuel costs without facing higher upfront vehicle prices, while wealthier households bear greater costs due to their greater tendency to buy new vehicles. Notably, the CAFE standard also benefits rural and suburban households more than urban households, as they tend to own and use vehicles for longer periods, maximizing the benefits of improved fuel economy.

Policies targeting the electricity sector introduce additional disparities. A Clean Energy Standard (CES), which mandates a minimum share of electricity generation from renewable or nuclear sources, leads to geographic winners and losers. States with already high shares of low-carbon electricity, such as those in the Northeast and on the West Coast, benefit from lower compliance costs, while states reliant on coal face higher electricity prices. Similarly, a Carbon Emission Intensity Standard (CEIS), which penalizes states based on the carbon intensity of their electricity mix, disproportionately impacts states with heavy coal dependence, such as Wyoming and West Virginia. Interestingly, the incidence of the CES and CEIS is not highly correlated, as some states with high renewable generation also maintain significant coal capacity, leading to unexpected redistributive effects.

Taken together, these findings highlight key trade-offs in climate policy design. Carbon pricing mechanisms offer greater flexibility by generating revenue that can be used to mitigate regressive impacts, but their effectiveness depends on how the revenue is redistributed. Policies such as a simple tax-and-dividend scheme can be broadly progressive, but without geographic adjustments, they risk shifting economic burdens from urban to rural areas and from the Midwest to the coasts. Performance standards, including fuel economy mandates and clean energy requirements, tend to be more rigid and can impose substantial costs on lower-income and rural households without offering direct financial compensation. As a result, policymakers must carefully consider both the income-based and geographic distributional effects of climate policies to ensure they do not exacerbate existing economic disparities while working toward emissions reductions.

5 Conclusions

This study develops machine learning models that allow for the estimation of household carbon footprints and uses economic theory to measure the impact of several climate policies on households. It highlights the significant variation in household carbon footprints and the impacts of carbon policies across different regions of the United States. These findings underscore the complex interplay between regional characteristics—such as energy sources, transportation patterns, and income levels—and the effectiveness of various carbon policies. The variation we observe has critical implications for policymakers. First, it underscores the need for detailed regional assessments when formulating national carbon policies. Uniform policy approaches may inadvertently exacerbate disparities, placing disproportionate burdens on certain regions or demographics. For example, states heavily reliant on coal for electricity or regions with limited public transportation infrastructure may face significantly higher costs under clean energy standards or fuel economy standards compared to regions already investing in renewable energy and efficient transit systems. This highlights the importance of tailoring policy interventions to align with local economic and infrastructural realities.

Our analysis also underscores a key advantage of carbon pricing, above its economic efficiency. Unlike regulatory measures such as fuel economy standards or clean energy mandates, carbon taxes generate revenue that policymakers can redistribute to address equity concerns and promote social welfare. This revenue offers an opportunity to cushion low-income households from potential adverse effects, incentivize clean technology adoption, or invest in community-specific mitigation strategies. The ultimate net impact of a carbon tax on a given household depends heavily on the choices policymakers make about how to allocate its revenues, enabling a level of customization that regulatory approaches lack.

By addressing distributional concerns, policymakers can potentially build broader public support for climate policies, ensuring their effectiveness over time. Additionally, the ability to target specific regional challenges makes carbon taxes an invaluable tool for addressing the uneven economic and environmental landscape across the country. In contrast, fuel economy standards and clean energy mandates, while popular policy choices, lack the capacity to generate direct financial benefits for households and require substantial upfront investments, which may be more difficult to implement equitably.

References

Argonne National Laboratory (2019). GREET Model.

Ballout, N., Etievant, L., and Viallon, V. (2023). On the use of cross-validation for the calibration of the adaptive lasso. *Biometrical Journal*, 65(5):e2200047.

Center for Sustainable Systems, University of Michigan (2023). Carbon footprint factsheet.

Davis, L. W. and Knittel, C. R. (2019). Are Fuel Economy Standards Regressive? *Journal of the Association of Environmental & Resource Economists*, 6:S1.

Federal Highway Administration (2017). National Household Travel Survey.

Gillingham, K. and Stock, J. H. (2018). The Cost of Reducing Greenhouse Gas Emissions. *Journal of Economic Perspectives*, 32(4):53–72.

Goulder, L. H., Hafstead, M. A. C., Kim, G., and 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.

Helfand, G. E. (1991). Standards versus standards: the effects of different pollution restrictions. *The American Economic Review*, 81(3):622–634.

Holland, S. P., Hughes, J. E., and Hughes, C. R. (2009). Greenhouse Gas Reductions under Low Carbon Fuel Standards? *American Economic Journal: Economic Policy*, 1(1):106–146.

International Civil Aviation Organization (2024). Carbon emissions calculator.

Jihoon Min, Hausfather, Z., and Qi Feng Lin (2010). A High-Resolution Statistical Model of Residential Energy End Use Characteristics for the United States. *Journal of Industrial Ecology*, 14(5):791–807.

Jones, C. and Kammen, D. M. (2014). Spatial Distribution of U.S. Household Carbon Footprints Reveals Suburbanization Undermines Greenhouse Gas Benefits of Urban Population Density. *Environmental Science & Technology*, 48(2):895–902.

Jones, C. M. and Kammen, D. M. (2011). Quantifying Carbon Footprint Reduction Opportunities for U.S. Households and Communities. *Environmental Science & Technology*, 45(9):4088–4095.

Kwoka Jr, J. E. (1983). The limits of market-oriented regulatory techniques: The case of automotive fuel economy. *The Quarterly Journal of Economics*, 98(4):695–704.

Mathur, A. and Morris, A. C. (2014). Distributional effects of a carbon tax in broader U.S. fiscal reform. *Energy Policy*, 66:326–334.

Metcalf, G. E., Paltsev, S., Reilly, J. M., Jacoby, H. D., and Holak, J. F. (2008). Analysis of U.S. Greenhouse Gas Tax Proposals. SSRN Scholarly Paper ID 1131633, Social Science Research Network, Rochester, NY.

National Oceanic and Atmospheric Administration (2022). National Centers for Environmental Prediction.

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

U.S. Bureau of Labor Statistics (2022). Consumer Expenditure Surveys (CE) Public Use Microdata Data Files. Library Catalog: www.bls.gov.

U.S. Census Bureau (2022). American Community Survey 5-Year Data (2018-2022). Library Catalog: www.census.gov Section: Government.

U.S. Energy Information Administration (2020a). Residential Energy Consumption Survey (RECS) - Data - U.S. Energy Information Administration (EIA).

U.S. Energy Information Administration (2020b). United States - SEDS - U.S. Energy Information Administration (EIA).

U.S. Environmental Protection Agency (2020). Emissions & Generation Resource Integrated Database (eGRID). Library Catalog: www.epa.gov.

U.S. Environmental Protection Agency (2023). Greenhouse gases equivalencies calculator - calculations and references.

Walker, K. (2019). tidycensus: Load US Census Boundary and Attribute Data as 'tidyverse' and 'sf'-Ready Data Frames. R package version 0.9.2.

MIT CEEPR

MIT Center for Energy and Environmental Policy Research

Since 1977, the Center for Energy and Environmental Policy Research (CEEPR) has been a focal point for research on energy and environmental policy at MIT. CEEPR promotes rigorous, objective research for improved decision making in government and the private sector, and secures the relevance of its work through close cooperation with industry partners from around the globe. Drawing on the unparalleled resources available at MIT, affiliated faculty and research staff as well as international research associates contribute to the empirical study of a wide range of policy issues related to energy supply, energy demand, and the environment.

An important dissemination channel for these research efforts is the MIT CEEPR Working Paper series. CEEPR releases Working Papers written by researchers from MIT and other academic institutions in order to enable timely consideration and reaction to energy and environmental policy research, but does not conduct a selection process or peer review prior to posting. CEEPR's posting of a Working Paper, therefore, does not constitute an endorsement of the accuracy or merit of the Working Paper. If you have questions about a particular Working Paper, please contact the authors or their home institutions.

MIT Center for Energy and Environmental Policy Research
77 Massachusetts Avenue, E19-411
Cambridge, MA 02139
USA

Website: ceepr.mit.edu

MIT CEEPR Working Paper Series is published by the MIT Center for Energy and Environmental Policy Research from submissions by affiliated researchers.

For inquiries and/or for permission to reproduce material in this working paper, please contact:

Email ceepr@mit.edu
Phone (617) 253-3551
Fax (617) 253-9845