FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

# A Macro Study of the Unequal Effects of Climate Change

Stephie Fried
Federal Reserve Bank of San Francisco
CEPR

May 2024

Working Paper 2024-18

https://doi.org/10.24148/wp2024-18

**Suggested citation:**

Fried, Stephie. 2024. "A Macro Study of the Unequal Effects of Climate Change." Federal Reserve Bank of San Francisco Working Paper 2024-18. https://doi.org/10.24148/wp2024-18

# A Macro Study of the Unequal Effects of Climate Change

Stephie Fried
Federal Reserve Bank of San Francisco and
CEPR

May 24, 2024

**Abstract**

This paper develops a macro heterogeneous-agent model to quantify the distributional impacts of higher temperatures in the US. Households adapt to temperature by using energy and equipment for heating and cooling. A key insight is that temperature acts as a transfer from nature, augmenting household income by the value of heating or cooling provided by nature. The welfare effects of climate change vary substantially with income, increasing welfare inequality in the colder parts of the US. This heterogeneity results from the effects of climate change on transfers from nature and on households' extensive-margin decisions to purchase heaters and air conditioners.

Email: stephie.fried@sf.frb.org. The views in this paper are those of the author and do not reflect the views of the Federal Reserve system or its staff. I thank Greg Casey, Betty Daniel, Kevin Novan, William Peterman, Brigitte Roth Tran, James Traina, and seminar and conference participants at George Mason University, The Federal Reserve Board of Governors, Duke, the University of Chicago, the University of Oregon, UC Davis, Cornell, UCSD, Federal Reserve Bank of Richmond, LSE environment week, Federal Reserve Bank of Minneapolis, Bank of France, University of Washington, LSE, Bank of England, Oxford, University of Glasgow, University of Edinburgh, and Middlebury College for their many helpful comments. For excellent research assistance, I thank Zoë Arnaut-Hull, River Bell and Ethan Goode. All potential errors are my own.

# 1. Introduction

Will climate change worsen US inequality? The effects of climate change could vary across households due to spatial heterogenity in exposure to climate damage and due to income heterogenity which affects households' ability to adapt. Recent macro climate-economy models either abstract from all heterogeneity (e.g., Acemoglu et al., 2012; Golosov et al., 2014; Barrage, 2020) or study spatial heterogeneity (e.g., Cruz and Rossi-Hansberg, 2024; Bilal and Rossi-Hansberg, 2023; Rudik et al., 2022). Instead, this paper studies income heterogeneity, and develops an Aiyagari (1994)-style heterogeneous agent model to analyze the distributional effects of climate change across income groups. I focus on one important type of climate damage; the direct effects of higher temperatures in the US. I find that the welfare consequences of higher temperatures caused by climate change vary considerably with income, increasing welfare inequality in the colder parts of the US.

I first develop a simple model to analytically explore the distributional consequences of changes in temperature. Households derive utility from consumption and housing services. I extend this standard framework by assuming that utility from housing services depends on the indoor temperature of the house. All else constant, households derive less utility from their house if it is too hot or too cold, than if it is a comfortable temperature. The indoor temperature depends on the outdoor temperature determined by nature and on any energy for heating or cooling the household purchases. The ability to purchase heating and cooling energy gives the household the option to adapt to the outdoor temperature by adjusting the indoor temperature of its house.[1]

A key insight from the simple model is that the outdoor temperature acts as a financial transfer from nature to households. For example, if the hottest possible outdoor temperature is 50°C and the actual outdoor temperature is 30°C, then nature gave all households 20°C of cooling for free, augment their income by the value of the cooling. Moving from moderate to extreme temperatures decreases these transfers from nature. Even though the decrease in transfers is the same for all households, it has higher welfare costs for lower income households because the value of the loss is larger in proportion to their income. In isolation, this intuition implies that climate change will be regressive in regions where it leads to more extreme temperatures, and thus fewer transfers from nature, and progressive in regions where it leads to less extreme temperatures.[2]

[1]Rode et al. (2021) and Jeon (2023) also consider the potential for households to adapt to temperature by using energy for heating and cooling.

[2]The notion that the outdoor temperature is a transfer from nature is closely related to the notion that environmental quality is a form of natural capital (Fenichel et al., 2018). The transfers from nature are similar to the income flows from the natural capital embodied in the temperature distribution. By shifting the temperature distribution, climate change affects the stock of natural capital and the associated income flows.

Expanding on the transfers-from-nature intuition, I develop a dynamic heterogeneous agent model to quantify the distributional effects of rising temperatures. An important feature of the quantitative model is that households need specialized temperature-control equipment as well as energy to produce heating and cooling. There are three types of temperature-control equipment; heaters which produce heat, air conditioners which produce cooling, and heat pumps which produce both heating and cooling. Each type of equipment has a fixed cost, implying that the household could optimally choose a corner solution with one or more types of temperature-control equipment equal to zero. The possibility of corner solutions allows the model to match the empirical fractions of households without heating or cooling.

The quantitative model includes several other features that are important for a numerical analysis. Low-income households receive energy assistance payments that broadly mimic federal and local programs for heating and cooling assistance in the US. Additionally, households experience an uncertain distribution of outdoor temperatures each period. I split the period into 100 sub-periods, where each sub-period corresponds to a different outdoor temperature. Finally, households vary in terms of their income, wealth, and location. The income and wealth distributions are endogenous; building from the benchmark models of income and wealth inequality (Bewley, 1977; Huggett, 1993; Aiyagari, 1994), households draw idiosyncratic labor-productivity shocks each period and endogenously adjust their consumption and savings decisions in response. The spatial distribution is exogenous; households live in one of five regions which differ based on the temperature distribution.

An important component of the model calibration is to determine the parameters in the production functions for heating and cooling. I assemble a unique data set on the price, capacity, and energy use of 1,125 heaters, air conditioners, and heat pumps. I use these data to estimate the heating and cooling production functions. I calibrate the remaining key parameters in the model to match average heating and cooling energy budget shares, the variation in energy budget shares and energy expenditures with income, and the fractions of households without heating or cooling. The calibrated model replicates the empirical patterns of rising energy expenditures and falling energy budget shares with income. Previous literature has added subsistence energy to the utility function to match these patterns (see e.g., Metcalf, 1999; Grainger and Kolstad, 2010; Fried et al., 2018). I show that directly modeling the effect of temperature on utility and allowing households to use energy to adapt to temperature enables the model to match these patterns without subsistence energy.

I use the model to study the distributional consequences of climate change, modeled as a rightward shift in the annual temperature distribution. I compare two stationary equilibria: (1) a contemporary equilibrium with the contemporary temperature distribution, and (2) a climate-change equilibrium with the projected temperature distribution in year 2100. To determine the 2100 temperature distribution, I use a climate projection designed to forecast carbon emissions in the absence of large-scale climate policy (RCP 8.5). I measure welfare using the consumption-housing equivalent variation, defined as the percent increase in consumption and housing the household would need in expectation in every period in the contemporary equilibrium so that they are indifferent between the contemporary and the climate-change equilibrium.

I find that the welfare impacts of higher temperatures from climate change vary substantially with income and region. In the colder regions, climate change creates welfare costs for the lowest income households, welfare benefits for the middle income households, and near-zero welfare effects for the highest income households. In contrast, in the hotter regions, climate change creates welfare benefits or relatively small welfare costs for the lowest income households, larger welfare costs for the middle income households, and again near-zero welfare effects for the highest income households. As an example of the magnitudes, in the cold region, the welfare impacts range from a welfare cost for households in the lowest income decile equal to 3 percent of consumption and housing expenditures to a welfare gain for households in the middle income deciles equal to 0.6 percent of consumption and housing expenditures.

The variation in the welfare impacts stems from two key mechanisms. First, drawing on the intuition from the simple model, climate change increases transfers from nature in the colder regions, creating welfare benefits, and decreases transfers from nature in the hotter regions, creating welfare costs. The changes in transfers from nature relative to income in all regions are larger for lower income households, implying that, all else constant, the magnitude of the welfare impact of climate change falls with income.

Second, moving beyond the intuition from the simple model, climate change affects households' choices of temperature-control equipment. All else constant, households would prefer to specialize in heater or air-conditioner equipment and avoid having to purchase large amounts of both heater and air-conditioner equipment or (more expensive) heat-pump equipment. Climate change affects households' ability to specialize. In the colder regions, the increase in hot days from climate change makes it harder for households to specialize in heater equipment, creating welfare costs. These costs are particularly acute for the lowest income households who move from only purchasing heater equipment in the contemporary equilibrium to purchasing both heater and air-conditioner equipment in the climate-change equilibrium. The welfare cost of this move away from complete specialization dominates the welfare benefit of higher transfers from nature, resulting in relatively high welfare costs of climate change for low-income households in the colder regions. In the hotter regions, the opposite effect leads to net welfare benefits of climate change among some low-income households.

I explore the potential for energy assistance policy to affect the distributional consequences of climate change. In the baseline simulation, I assume that the level of energy assistance does not change in response to climate change. A plausible alternative would be for policymakers to scale energy assistance in proportion to the changes in heating and cooling needs caused by climate change. This scaling partially offsets the welfare impacts stemming from the changes in transfers from nature among the low-income households that receive assistance. In particular, the scaling lowers energy assistance in the colder regions, where climate change increases transfers from nature, and raises energy assistance in the hotter regions, where climate change decreases transfers from nature. Consequently, among low-income households, scaling energy assistance in response to climate change leads to larger welfare costs of climate change in the colder regions and smaller welfare costs or larger welfare benefits in the hotter regions.

Motivated by local natural gas bans and the general push towards residential electrification, I study the effects of a mandate that all households use heat pumps for heating and cooling. Heat pumps have important implications for the distributional consequences of climate change because households with heat pumps are not affected by the impact of climate change on their ability to specialize in heater or air-conditioner equipment. Under the heat-pump mandate, all households in the colder regions still experience the welfare benefits from the higher transfers from nature, but low-income households no longer experience the welfare cost from purchasing their first air conditioner because the heat pump can also be used for cooling. Similarly, all households in the hotter regions still experience the welfare costs of lower transfers from nature but low-income households no longer experience the welfare benefits from not having to purchase a heater. In this case, climate change would create progressive welfare benefits across all income levels in the colder regions and regressive welfare costs across all income levels in the hot regions.

An important contribution of the paper is to study climate change in a setting with endogenous income and wealth heterogeneity. This heterogeneity matters not only for the distributional effects of climate change discussed above, but also for the aggregate welfare consequences. To highlight this point, I compare the aggregate welfare cost of climate change in the baseline simulation with the welfare cost from a representative agent version of the model. I find, for example, that climate change reduces aggregate welfare in the mild region by approximately 0.68 percent of consumption and housing expenditures in the baseline simulation compared to only 0.18 percent in the representative agent simulation. The differences in the other regions are similarly large in magnitude. These different predictions emphasize the importance of incorporating income heterogeneity for understanding aggregate outcomes, consistent with the findings from the macro literature on heterogeneous agents in other contexts (Krueger et al., 2016; Heathcote et al., 2009).

Finally, I consider an extension of the model in which low-income households' labor productivity partially depends on temperature. Many low-income workers work in outdoor sectors, like construction and agriculture, which could be particularly vulnerable to the effects of temperature on labor productivity. I apply a widely-used rule-of-thumb from the construction industry to determine how the changes in temperature from climate change affect labor productivity. The rule-of-thumb suggests that climate change will increase labor productivity by at most 2 percent in the colder regions and decrease labor by at most 2 percent in the hotter regions. Including these labor productivity effects for low-income households decreases the welfare cost of climate change for low-income households in the colder regions and increases the welfare cost for low-income households in hot regions, with the biggest effects in the cold and hot regions, where the changes in productivity are largest.

This paper contributes to the broader literature on the unequal consequences of climate change. Much of this literature focuses on region-level inequality and studies how the consequences of climate change vary across regions based on differences in regional income and climate damages (Nordhaus and Yang, 1996; Carleton et al., 2022; Nath, 2022; Krusell and Smith, 2022; Rudik et al., 2022; Cruz and Rossi-Hansberg, 2024; Bilal and Rossi-Hansberg, 2023). A more theoretical literature adds households with different fixed income types to integrated assessment climate-economy models to study the importance of within-region income inequality for optimal climate policy (Dennig et al., 2015; Kornek et al., 2021; Belfiori and Macera, 2022; Douenne et al., 2023).[3] Moving beyond fixed income types, Blanz (2023) extends a standard incomplete markets model of income heterogeneity, as in the present paper, to study the effects of food-price changes caused by climate change in developing countries. Finally, an empirical literature estimates how the impacts of natural disasters (Strobl, 2011; Deryugina et al., 2018; Roth Tran and Wilson, 2022) or temperature (Hsiang et al., 2017; Park et al., 2018; Behrer et al., 2021; Doremus et al., 2022) vary across income groups in the US.[4]

The paper proceeds as follows: Section 2 builds a simple model to develop the intuition that the outdoor temperature acts as a transfer from nature to households. Section 3 constructs a heterogeneous-agent model to quantify the distributional effects of climate change. Section 4 calibrates the model and Section 5 presents the quantitative results. Section 6 concludes.

[3] Malafry and Brinca (2022) study the importance of inequality for optimal climate policy in a stylized two-period model in which households draw idiosyncratic labor productivity shocks.

[4] An additional literature examines the distributional consequences of carbon prices instead of climate damage (see e.g., Parry, 2004; Fullerton and Heutal, 2007; Metcalf, 2007; Chiroleu-Assouline and Fodha, 2014; Parry and Williams, 2010; Williams et al., 2015; Fried et al., 2018, 2022). These papers generally find that how the government uses the carbon price revenue has a substantial impact on the distributional effects of the policy. In contrast to carbon taxes or permits, climate damage does not create a stream of revenue that the government can allocate to affect its equity implications.

## 2. Simple Model

I develop a simple model that incorporates the key relationships between temperature, utility, and adaptation to temperature through heating and cooling. Using the model, I show that outdoor temperature acts as a transfer from nature to households and derive the implications for the distributional consequences of a change in outdoor temperature. The model is intentionally stylized; it is static, deterministic, and abstracts from the equipment, such as heaters and air conditioners, necessary for heating and cooling. I relax these assumptions in the richer, quantitative model to follow.

**Environment.** Households derive utility from consumption, housing, and indoor and outdoor temperature, according to the function:

$$u(c_i, h_i, T_i) = \begin{cases} G(\zeta) c_i^{\alpha} [D(T_i) h_i]^{1-\alpha} & : 0 \leq \underline{\zeta} < T_i < \bar{\zeta} \\ -\Theta & : \text{otherwise} \end{cases}, \tag{1}$$

where $\alpha \in (0,1)$. Variables $c_i$ and $h_i$ denote household $i$'s choices of consumption and housing services, respectively. Variable $T_i$ denotes the household's choice of the indoor temperature of its house. If the indoor temperature is very cold, $T_i < \underline{\zeta}$, or very hot, $T_i > \bar{\zeta}$, then the household dies from exposure to temperature extremes and receives utility $-\Theta$.

Parameter $\zeta \in [0, 2\zeta^{\star}]$ is the outdoor temperature. Outdoor temperature is determined by nature and ranges from 0 to $2\zeta^{\star}$, where $\zeta^{\star}$ denotes the most comfortable outdoor temperature. Following Conte et al. (2022), function $G(\zeta)$ captures the amenity value of the outdoor temperature. While it is not necessary to specify a functional form, it is useful to think of $G(\zeta)$ as having a bliss point at the preferred outdoor temperature, $\zeta^{\star}$. In this case, nice days, i.e., days with temperatures near $\zeta^{\star}$, have higher amenity values, while cold or hot days, i.e., days with temperatures far from $\zeta^{\star}$, have lower amenity values.

Function $D(T_i) \in [0,1]$ is a damage function that reduces the utility the household receives from housing services based on the relationship between the indoor temperature of the house and the bliss point temperature, $\zeta^{\star}$, according to,

$$D(T_i) = \begin{cases} \frac{T_i}{\zeta^{\star}} & : 0 \leq T_i \leq \zeta^{\star} \\ \frac{2\zeta^{\star} - T_i}{\zeta^{\star}} & : \zeta^{\star} \leq T_i \leq 2\zeta^{\star}. \end{cases} \tag{2}$$

If the indoor temperature equals the bliss point, then the household receives full utility from its housing services; $D(T_i = \zeta^{\star}) = 1$. If instead, the house is cold, $T_i < \zeta^{\star}$, or hot, $T_i > \zeta^{\star}$, then the household receives less utility from the housing services; $D(T_i < \zeta^{\star}) < 1$ and $D(T_i > \zeta^{\star}) < 1$.

The indoor temperature, $T_i$, depends on the outdoor temperature, $\zeta$, and on any energy the household purchases for heating, $e_i^h$, or cooling, $e_i^c$: $T_i = \zeta + e_i^h - e_i^c$. Energy purchased for heating raises the indoor temperature relative to the outdoor temperature, while energy purchased for cooling reduces the indoor temperature relative to the outdoor temperature. The ability to use energy for heating and cooling allows the household to adapt to the outdoor temperature by changing the indoor temperature of its house.

The household chooses consumption, housing services, and heating and cooling energy to maximize utility (equation (2)) subject to the budget constraint, $y_i = c_i + p^h h_i + p^{eh} e_i^h + p^{ec} e_i^c$. Variables $p^h$, $p^{eh}$ and $p^{ec}$ denote the relative prices of housing, heating and cooling energy, respectively. Variable $y_i$ is the household's income. The consumption good is the numeraire.

**Analysis**. I consider a region of the parameter space in which the utility from death, $-\Theta$, is substantially less than the optimizing household's utility from life, implying that the household never chooses indoor temperatures outside of $(\underline{\zeta}, \bar{\zeta})$. I solve the model for a cold day, $\zeta < \zeta^\star$. The solution and implications for a hot day, $\zeta > \zeta^\star$, are symmetric and included in Appendix A. Focusing on an interior solution in which the household's optimal indoor temperature is less than the bliss point but greater than the outdoor temperature, $\zeta < T_i^\star < \zeta^\star$, the optimal levels of consumption, housing, and indoor temperature are

$$c_i^\star = (y_i + p^{eh}\zeta)\left(\frac{\alpha}{2-\alpha}\right), \quad h_i^\star = \left(\frac{y_i + p^{eh}\zeta}{p^h}\right)\left(\frac{1-\alpha}{2-\alpha}\right) \quad \text{and} \quad T_i^\star = \left(\frac{y_i + p^{eh}\zeta}{p^e}\right)\left(\frac{1-\alpha}{2-\alpha}\right). \tag{3}$$

The heating and cooling energy required to achieve indoor temperature $T_i^\star$ are $e_i^{h\star} = T_i^\star - \zeta$ and $e_i^{c\star} = 0$.

The optimized levels of consumption, housing, and indoor temperature in equation (3) reveal that the outdoor temperature, $\zeta$, acts as a transfer from nature to the household. Nature gives the household $\zeta$ degrees of heating for free, augmenting its income by $p^{eh}\zeta$. Importantly, all households, regardless of their income, receive the same transfer from nature. However, the transfer from nature has a larger percentage impact on low-income households optimal choices because it constitutes a larger share of their income.

To determine the distributional implications of a change in outdoor temperature, equation (4) derives the percent change in the household's optimized utility from a marginal increase in outdoor temperature,

$$\frac{\partial u(c_i^\star, h_i^\star, T_i^\star)/\partial \zeta}{u(c_i^\star, h_i^\star, T_i^\star)} = \frac{G'(\zeta)}{G(\zeta)} + \frac{p^{eh}\lambda^\star}{u(c_i^\star, h_i^\star, T_i^\star)}. \tag{4}$$

Variable $\lambda^\star$ denotes the Lagrange multiplier on the budget constraint. The welfare gains from a warmer day, measured as the percent change in utility, equal the relative increase in the amenity value of temperature plus the relative increase in utility from the higher income resulting from the larger transfer from nature. Substituting in the value for multiplier, $\lambda^\star$, from the first order conditions reveals that the welfare gains from a warmer day fall with income:

$$\frac{\partial u(c_i^\star, h_i^\star, T_i^\star)/\partial \zeta}{u(c_i^\star, h_i^\star, T_i^\star)} = \frac{G'(\zeta)}{G(\zeta)} + \frac{p^{eh}(2-\alpha)}{y_i + p^{eh}\zeta}. \tag{5}$$

The intuition for the progressivity of a warmer day parallels the intuition for the progressivity of lump-sum transfers from the public finance literature. Low-income households value the transfer from nature more than high-income households because it is a larger share of their budgets. Hence, an increase in the transfer from nature from a warmer day creates larger welfare gains for low-income households. More generally, in cold weather ($\zeta < \zeta^\star$), moving from cooler (more extreme) to warmer (less extreme) days is progressive, and in hot weather, ($\zeta > \zeta^\star$), moving from warmer (more extreme) to cooler (less extreme) days is progressive. All else constant, these results imply that the rightward shift in the temperature distribution caused by climate change will be regressive in regions in which it leads to more extreme temperature days and progressive in regions in which it leads to fewer extreme temperature days.

The heterogeneity in these welfare effects depends critically on households ability to adapt to temperature through heating and cooling. If households cannot adapt, then all households must choose the corner solution in which the indoor temperature equals the outdoor temperature, $T_i^\star = \zeta$. In this case, the model collapses to a standard two-good utility optimization problem. Outdoor temperature does not affect the household's consumption and housing choices, eliminating the transfers-from-nature channel. The welfare impact of a change in outdoor temperature depends only on the direct effect of temperature on utility, through both the amenity value and the damage function, and thus is the same for all households, regardless of income.[5]

The notion that extreme temperatures act like decreased transfers provides an interesting perspective on energy assistance policy. Much of this assistance takes the form of direct payments to low-income households, which, when viewed through the lens of the simple model, implicitly replace the lost transfers from nature on extreme temperature days. For example, the Low Income Home Energy Assistance Program (LIHEAP) provides payments that are largely based on households' heating or cooling expenses from the previous year, or on the average number of heating and cooling degree days. Both of these approaches are practical ways to mimic indexing payments to outdoor temperature, and thus compensate households for the reduced transfers from nature on extreme temperature days.

[5] See Appendix A for the model solution with infeasible adaptation and for the model solution when the household's optimal indoor temperature equals the bliss point, $T_i^\star = \zeta^\star$.

One implication of the simple model is that the amenity value has no effect on the distributional consequences of a change in temperature. Specifically, the derivative of equation (5) with respect to income does not depend on the amenity value. The implicit assumption is that the preferences underlying the amenity value are uncorrelated with income. I continue to make this assumption in the quantitative model and abstract from the amenity value. This abstraction affects level of the welfare impacts from a change in temperature. However, the analysis from the simple model suggests that it is not important for understanding the relative welfare effects of changes in temperature across income groups, which are paper's focus.

The analysis of the simple model focuses on a region of the parameter space in which the household's optimal indoor temperature is within $(\underline{\zeta}, \bar{\zeta})$, implying that the household does not die from exposure to extreme temperatures. Moreover, since the household's optimal indoor temperature is strictly less than $\bar{\zeta}$ and strictly greater than $\underline{\zeta}$, the possibility of death has no impact on the household's heating and cooling decisions. In the quantitative model, I continue to focus on this same region of the parameter space and I abstract from the utility cost of death from extreme temperature exposure.

This abstraction comes with the caveat that some people die each year from exposure to extreme temperatures in the US. For example, from 1999-2020, exposure to extreme temperatures caused approximately 0.04 percent of US deaths, equal to 1,145 people per year, on average.[6] Over half of these deaths are estimated to be among the un-housed population (Snow, 6/20/2022; NHCHC, 2022) which are outside the model. Some of the remaining deaths are likely due to behavioral factors unrelated to the temperature of the home, which are also outside of the model, such as preforming intense physical activities in extreme heat.[7] Even so, a small number of the temperature-related deaths are likely caused by the temperature of the home. It is important to acknowledge that the quantitative model omits this channel.

## 3. Quantitative model

Building on the transfers-from-nature intuition, I develop a dynamic, heterogeneous-agent model to quantify the distributional effects of climate change. Time is discrete and infinite. The economy is composed of $N$ regions which are differentiated by their temperature distributions. Each region contains a continuum of infinitely-lived heterogeneous households, and a continuum of perfectly competitive firms that produce the final good, energy, housing,

[6]Centers for Disease Control and Prevention, National Center for Health Statistics. National Vital Statistics System, Mortality 1999-2020 on CDC WONDER Online Database, released in 2021. Data are from the Multiple Cause of Death Files, 1999-2020, as complied from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program. Accessed at http://wonder.cdc.gov/ucd-icd10.html on Mar 3, 2023. ICD-10 codes X30 and X31 correspond to deaths from heat exposure and cold exposure, respectively.

[7]The Environmental Protection Agency (EPA) stresses the importance of early warning systems and educating the public on the dangers of extreme temperatures to reduce deaths from behavioral factors (EPA, 2021).

and temperature-control equipment. Households derive utility from consumption, housing services, and the temperature of their house. Households can affect the temperature of their house by producing heating or cooling from temperature-control equipment and energy. Households supply labor in a competitive regional labor market, and their labor productivity is subject to persistent, idiosyncratic shocks, as in Aiyagari (1994). Households can accumulate assets to self-insure against these shocks. A federal government taxes households and uses the revenue to provide energy assistance payments to low-income households. There is no migration between regions.

**Temperature distribution.** Households can experience $J$ outdoor temperatures during the period, ranging from $\zeta_1$ to $\zeta_J$. I divide the period into $J$ different sub-periods, one for each value of outdoor temperature, $\zeta_j \in [\zeta_1, \zeta_J]$. The temperature distribution is the set of weights $Q_{in} \equiv \{q_{ijn}\}_{j=1}^{J}$ which correspond to the fraction of the period household $i$ in region $n$ experiences outdoor temperature $\zeta_j$. To partially capture the uncertainty inherent in weather, I assume that the temperature distribution is a random variable. The household draws a temperature distribution each period from a distribution of temperature distributions, $\pi_n^Q(Q_{in})$. The realizations of the temperature distribution are i.i.d. across households within a region.

**Preferences.** Households derive utility from consumption and housing according to,

$$u_{in} = \sum_{j=1}^{J} q_{ijn} \left[ \frac{c_{ijn}{}^{1-\sigma}}{1-\sigma} + \psi \frac{(D(T_{ijn})h_{in})^{1-\sigma}}{1-\sigma} \right]. \tag{6}$$

The period utility of household $i$ in region $n$ is the sum of its utility in each sub-period, weighted by the fraction of the period the household spends in each sub-period, $q_{ijn}$. Following Gervais (2002), utility is separable in consumption, $c_{ijn}$ and housing services $h_{in}$. Subscript $j$ on variable $c$ denotes that consumption can vary across sub-periods. Consumption is a relatively flexible input and households can choose to consume different amounts based on the outdoor temperature. In contrast, housing services are less flexible and cannot vary by sub-period based on the outdoor temperature. Households choose housing at the start of the period and it is fixed for the entire period; there is no subscript $j$ on variable $h$ in equation (6).

Damage function $D(T_{ijn}) \in [0,1]$ depends on the difference between the household's indoor temperature in sub-period $j$, $T_{ijn}$, and the bliss point, $\zeta^\star$,

$$D(T_{ijn}) = \frac{1}{1+\chi(T_{ijn} - \zeta^\star)^2}. \tag{7}$$

As in the simple model, households receive full utility from their housing services in sub-periods in which the indoor temperature equals the bliss point, $D(\zeta^\star) = 1$. Households receive less than full utility from their housing services in sub-periods in which their house is hotter or colder than the bliss point, $D(T_{ijn} < \zeta^\star) < 1$ and $D(T_{ijn} > \zeta^\star) < 1$. Parameter $\chi$ controls the utility cost from being away from the bliss point. At the extreme, when $\chi = 0$, there is no decrease in utility from indoor temperatures that deviate from the bliss point. Similarly, as $\chi \to \infty$ households must choose indoor temperatures equal to the bliss point to receive any utility from their housing services.

The household's indoor temperature in sub-period $j$ depends on the outdoor temperature in that sub-period, $\zeta_j$, and on the heating and cooling the household produces. The household produces heating and cooling from energy and temperature-control equipment. I consider three types of temperature-control equipment: heater equipment, such as a furnace, that produces heat, air-conditioner equipment that produces cooling, and heat-pump equipment that produces both heating and cooling. In practice, the energy required to run the temperature-control equipment differs across the equipment types: air conditioners and heat pumps use electricity, while heaters use a variety of different fuels, such as electricity or natural gas, depending on the type of boiler or furnace. In the model, I assume that air conditioners and heat pumps require energy $e^e$ to operate, where superscript $e$ denotes electricity. I abstract from the heater fuel choice and assume that heaters require energy $e^m$ to operate, where superscript $m$ denotes a mix of fuel types. Households must pay a fixed cost for each equipment type, $\Omega^h$, $\Omega^c$, and $\Omega^p$, where superscripts $h$, $c$, and $p$ denote the fixed costs for heaters, air conditioners, and heat pumps, respectively. The fixed costs capture installation and maintenance costs and allow the model to match the empirical reality some households do not have temperature-control equipment.

The indoor temperature for household $i$ in sub-period $j$ in region $n$ equals:

$$T_{ijn} = \begin{cases} \zeta_j + \frac{1}{h^\gamma}\left[\underbrace{A^h(x^h_{in})^{\theta^h}(e^{mh}_{ijn})^{\eta^h}}_{\text{heating}} - \underbrace{A^c(x^c_{in})^{\theta^c}(e^{ec}_{ijn})^{\eta^c}}_{\text{cooling}}\right] & : \text{no heat pump} \\ \zeta_j + \frac{1}{h^\gamma}\left[\underbrace{A^p(x^p_{in})^{\theta^h}(e^{eh}_{ijn})^{\eta^p}}_{\text{heating}} - \underbrace{A^p(x^p_{in})^{\theta^p}(e^{ec}_{ijn})^{\eta^p}}_{\text{cooling}}\right] & : \text{heat pump.} \end{cases} \tag{8}$$

Variables $x^h$, $x^c$, and $x^p$ denote the household's levels of heater, air-conditioner, and heat-pump equipment. The additional superscripts $h$ and $c$ on the energy terms denote that the energy is used for heating or cooling, respectively. For example, $e^{ec}$ denotes energy type $e^e$ used for cooling. Parameters $A^h$, $A^c$, and $A^p$ denote the productivity of heating using a heater, of cooling using an air conditioner, and of heating and cooling using a heat pump, respectively.

The Cobb-Douglas functional forms in equation (8) have several properties that make them well-suited to describe the production of heating and cooling. First, both equipment and energy are essential inputs. For example, a household cannot produce heat without a furnace and energy to operate the furnace. Second, in the empirically relevant region of the parameter space ($\eta < 1$), the marginal product of energy decreases with energy use, all else constant. Hence the energy efficiency of heating or cooling falls as the difference between the indoor and outdoor temperature rises, consistent with the physical properties of heaters, air conditioners, and heat pumps.

As with housing services, the household's choices of temperature-control equipment are fixed for the entire period; there is no subscript $j$ on the temperature-control equipment variables, $x^h$, $x^c$ or $x^p$, in equation (8). Households cannot purchase a furnace each winter when it is cold and return it each summer when it is hot. This inflexibility implies that households optimally choose to idle their temperature-control equipment when the outdoor temperature makes the equipment unnecessary. While temperature-control equipment is fixed for the entire period, households can still adjust their energy use in each sub-period to achieve their desired indoor temperature; the energy variables $e^{mh}$, $e^{ec}$, and $e^{eh}$ in equation (8) all include a subscript $j$. For example, the household could use more energy for heating on a very cold day than on a moderately cold day.

The change in indoor temperature from given amounts of temperature-control equipment and energy is decreasing in the level of housing services, $h$. Larger levels of $h$ correspond to bigger houses which require more energy and equipment to heat and cool. However, this relationship is not necessarily one-for-one. First, doubling the square footage of a house typically less than doubles the energy necessary to heat and cool the house. This is because a major source of heat loss is due to air near the building's envelope and doubling the square footage of a house typically less than doubles the envelope. Second, higher levels of $h$ correspond not only to larger houses but also to higher quality houses, which are likely to be better insulated and more energy efficient.

**Labor endowment and productivity.** Each household is endowed with one unit of labor, which it supplies exogenously to firms in its region. The household earns labor income $wz_{in}$, where $w$ denotes the market wage and $z_{in}$ is the household's idiosyncratic labor productivity draw. The log of the household's idiosyncratic labor productivity is the sum of two components: $\log(z_{in}) = \nu_{in} + \xi_{in}$. Component $\nu_{in}$ is an idiosyncratic persistent productivity shock which follows a finite-state Markov chain with transition probabilities $\pi^{\nu}(\nu'_{in}|\nu_{in})$, and unique invariant distribution $\Pi^{\nu}(\nu_{in})$. I use 'prime' to denote next period's value of the variable. Component

$\xi \sim N(0, \sigma_\xi^2)$ is a household-specific fixed effect (i.e., ability) that is constant over time.

**Energy assistance.** The federal and local governments in the US provide aid to help low-income households with their heating and cooling energy bills. The federal government provides the majority of this aid through LIHEAP. While the exact formulas vary by state, most LIHEAP payments are direct transfers based on the household's expected heating and/or cooling expenses. I incorporate energy assistance into the model as direct transfers, $B_{in}$, to low-income households. The transfers vary by the household's region, based on the heating and cooling needs in the different climates. The government in the model finances energy assistance payments with a flat tax, $\tau$ on labor income.

**Recursive formulation of the household's problem.** The state variables are the household's assets, $a_{in}$ and the persistent value of its labor productivity shock, $\nu_{in}$. The model timing is as follows. The period starts and the household draws its labor productivity shock. The household chooses housing services and temperature-control equipment. After making these decisions, the household learns its temperature distribution for the current period. Finally, the household chooses heating and cooling energy and consumption in each sub-period, and its level of assets for the next period. I set up the timing in this way because, in practice, households make longer-run decisions for housing and temperature-control equipment based on their expectations of the weather. The assumption that households choose their housing and temperature-control equipment for the current period before the realization of the temperature distribution for that period partially captures this uncertainty.

The dynamic programming problem for household $i$ in region $n$ equals:

$$V(a_{in}; \nu_{in}) = \max_{h_{in}, x_{in}^h, x_{in}^c, x_{in}^p} \sum_{Q_{in}} \pi^Q(Q_{in}) \left\{ \max_{\{e_{ijn}^{mh}, e_{ijn}^{ec}, e_{ijn}^{eh}, c_{ijn}\}_{j=1}^J} \left[ \sum_{j=1}^J q_{ijn} \left( \frac{c_{ijn}^{1-\sigma}}{1-\sigma} + \psi \frac{(D(T_{ijn}) h_{in})^{1-\sigma}}{1-\sigma} \right) \right. \right.$$
$$\left. \left. + \beta \sum_{\nu'_{in}} \pi^\nu(\nu_{in} | \nu'_{in}) \sum_{Q'_{in}} \pi^Q(Q'_{in}) V(a'_{in}; \nu'_{in}) \right] \right\}, \tag{9}$$

subject to the budget constraint:

$$(1-\tau) w z_{in} + (1+r) a_{in} + B_{in} = \sum_{j=1}^J q_{ijn} [c_{ijn} + p^{em} e_{ijn}^{mh} + p^{ee}(e_{ijn}^{ec} + e_{ijn}^{eh})]$$
$$+ p^s h_{in} + p^{xh} x_{in}^h + p^{xc} x_{in}^c + p^{xp} x_{in}^p + \Omega^h \mathbf{1}_{x^h > 0} + \Omega^c \mathbf{1}_{x^c > 0} + \Omega^p \mathbf{1}_{x^p > 0} + a'_{in},$$

and the non-negativity constraints:

$$a'_{in} \geq 0,\ h_{in} \geq 0,\ e^{mh}_{ijn} \geq 0,\ e^{ec}_{ijn} \geq 0,\ e^{eh}_{ijn} \geq 0,\ x^{c}_{in} \geq 0,\ x^{h}_{in} \geq 0,\ x^{p}_{in} \geq 0,\ c_{ijn} \geq 0.$$

The left-hand-side of the budget constraint is the household's cash-at-hand, equal to the sum of its after tax labor income, the gross value of its assets, and any energy assistance the household receives. Variable $r$ denotes the interest rate. The right-hand-side is the household's total expenses, equal to the sum of its expenditures on consumption and heating and cooling energy in each sub-period, housing services, temperature-control equipment and fixed costs, and any assets the household chooses to carry into the next period. Prices $p^{em}$ and $p^{ee}$, denote the relative prices of heating-specific energy and electricity, respectively. Prices $p^{xh}$, $p^{xc}$ and $p^{xp}$ denote the relative prices of heater, air-conditioner, and heat-pump equipment, respectively. Price $p^{s}$ denotes the relative price of housing services. The final good is the numeraire.

**Firms.** Each region contains unit masses of perfectly competitive firms that produce the final good, housing services, temperature-control equipment and energy. The final good, $y$, is produced from a Cobb-Douglas production function in capital, $k$, and labor, $l$, $y = A^{y}(k^{y})^{\alpha} l^{1-\alpha}$. Variable $A^{y}$ denotes total factor productivity (TFP) in the production of the final good. Final-good firm profits are given by, $\pi^{y} = A^{y}(k^{y})^{\alpha} l^{1-\alpha} - Rk^{y} - wl$. Variable $R \equiv r + \delta$ denotes the rental rate on capital, where $\delta$ is the depreciation rate.

Housing, energy, and temperature-control equipment are produced using linear technologies in capital:

$$h = A^{s}k^{s}, \quad e^{m} = A^{em}k^{em}, \quad e^{e} = A^{ee}k^{ee}, \quad x^{h} = A^{xh}k^{xh}, \quad x^{c} = A^{xc}k^{xc} \quad \text{and} \quad x^{p} = A^{xp}k^{xp}.$$

Variables $A^{s}$, $A^{em}$, $A^{ee}$, $A^{xh}$, $A^{xc}$, and $A^{xp}$ denote total factor productivity in the production of housing services, heating-specific energy, electricity, and heater, air-conditioner and heat-pump equipment, respectively. Profits are given by $\pi^{j} = p^{j}A^{j}k^{j} - Rk^{j}$, $j \in \{s, em, ee, xh, xc, xp\}$. I assume that a unit of capital, $k$, can be used to produce any of the goods and services in the economy, implying that the rental rate on capital does not vary across the different firms.

The solutions to the firms' profit-maximization problems yield the equilibrium wage, and relative prices

$$w = (1-\alpha)\left(\frac{\alpha}{R}\right)^{\frac{\alpha}{1-\alpha}}, \quad \text{and} \quad p^{j} = \frac{R}{A^{j}}, \quad j \in \{s, em, ee, xh, xc, xp\}. \tag{10}$$

I assume a small open economy with respect to capital, implying that the interest rate, $r$, is exogenous and equal to the world interest rate, $r^{\star}$.

## 4. Calibration

The time period is one year. I divide US counties into five regions based on their average annual temperature. I calibrate some of the parameters directly from the data or the existing literature. Given these parameters, I jointly calibrate the remaining parameters internally so that a set of moments in the model match their corresponding empirical targets. The calibration relies on four data sets: (1) the NIPA accounts from the BEA, (2) the 2015 Residential Energy and Consumption Survey (RECS), (3) PRISM climate data and (4) a unique product-level data set on heaters, air conditioners, and heat pumps collected from ecomfort.com. All values reported from the BEA data are historical averages over the 1997-2020 time period.[8]

**Temperature distribution.** I divide the counties in the continental US into five regions, cold, cool, mild, warm, and hot, based on the county's average annual temperature from 1950-2022. The regions have approximately equal population. County-level data on average daily temperature are from the PRISM climate group.[9] Appendix Figure C.1 shows a map of the US counties by region. The colder regions are in the northern and mountainous parts of the US.

Table 1: Average Annual Temperature Relative to a Moderate Year

| | Cold | Cool | Mild | Warm | Hot |
|---|---|---|---|---|---|
| Cold year | 0.91 | 0.93 | 0.95 | 0.97 | 0.96 |
| Moderate year | 1 | 1 | 1 | 1 | 1 |
| Hot year | 1.14 | 1.11 | 1.07 | 1.06 | 1.04 |

Note: In each region, the table reports the average annual temperature for each realization of the temperature distribution, relative to the value in a moderate year.

Temperature in each region can range from −40°C to 59°C in one-degree increments. Appendix Figure C.2 plots the average annual temperature distribution by region. To model uncertainty over the temperature distribution, I assume that there are three possible realizations of the temperature distribution, cold, moderate, and hot, with probabilities 0.1, 0.8, and 0.1, respectively. The temperature distribution for the hot realization equals the average of the temperature distribution in each region among the hottest 10 percent of years in the historical data. The moderate and cold realizations of the temperature distributions are similarly

[8]The EIA administers the RECS every five years. The most recent survey available is from 2020. I focus on the 2015 RECS instead of the 2020 RECS because of complications caused by COVID-19. Two key moments for the calibration are heating and cooling budget shares. COVID-19 began in March of 2020, after much of the heating season had passed, and so had relatively little effect on heating budget share. However, COVID-19 caused households to stay home for much of the cooling season, and thus pushed up cooling budget share. Average heating budget share is similar between the 2015 and 2020 RECS but cooling budget share is considerably higher in the 2020 RECS. To avoid the challenges introduced by these asymmetric behavioral changes in 2020, I use the data from the 2015 RECS for the analysis.

[9]See Appendix C.4 for additional details on the temperature data and the construction of the different regions.

constructed. Table 1 reports the average annual temperature for the cold and hot realizations relative to the moderate realization in each region. For example, in the cold region, the average annual temperature in a cold year is 9 percent below its value in a moderate year and the average annual temperature in a hot year is 14 percent above its value in a moderate year.

**Externally determined parameters.** I estimate the exponents in the heating and cooling production functions, $\theta^h$, $\eta^h$, $\theta^c$, $\eta^c$, $\theta^p$, and $\eta^p$, from product-level data on heaters, air conditioners, and heat pumps. The estimating equations are

$$\ln(y_i^j) = \kappa^j + \theta^j \ln(x_i^j) + \eta^j \ln(e_i^j) + \varepsilon_i^j \quad j \in \{h, c, p\}, \tag{11}$$

where subscript $i$ denotes the observation of a particular heater, air conditioner or heat pump. Variable $y$ is output of heating or cooling, variable $x$ is the quantity of the temperature-control equipment, and variable $e$ is the energy input.

Product-level data on heaters, air conditioners, and heat pumps are from ecomfort.com, a direct-to-consumer online store for heating and cooling equipment. The data provide information on the value and characteristics of the cross section of heaters, air conditioners and heat pumps for sale on the site during the fall of 2023. I focus exclusively on units that are for residential use.

I use the data to construct the variables in equation (11). I set the heating or cooling output, variable $y$, equal to the unit's capacity, i.e., the amount of heating or cooling the unit can produce in one hour. I assume that the quantity of temperature-control equipment, $x$, embodied in a heating or cooling unit is proportional to its price. Therefore, the level of $x$ in my data will be off by a constant factor across units, equal to the inverse of the unobserved price index for temperature-control equipment. This does not affect the estimates of $\theta^j$, but it does imply that the constant, $\kappa^j$, in equation (11) is not identified. Implicit in this approach is the assumption that the heater, air-conditioner, and heat-pump markets are competitive, so any price dispersion across units reflects differences in quality.[10]

Finally, I use data on efficiency and capacity to measure the energy input, $e$. The efficiency of the unit equals the heating or cooling output divided by the energy input. Thus, the energy input is the ratio of the unit's capacity to its efficiency. In both the model and the data, all air conditioners and heat pumps use the same energy input, electricity. However, the data contain heaters that use a variety of different energy inputs, such as natural gas, oil, and propane. The model abstracts from this choice over fuel types and assumes a single energy input for heaters, $e^m$. To align the data with the model, I follow Hassler et al. (2021), and multiply each heater's energy input by the relative price per BTU of that energy type. This adjustment

[10] This approach mimics the BEA's process for constructing data on the capital stock (Herman et al., 2003).

assumes that if a BTU of oil is twice as expensive as a BTU of natural gas, then it is twice as easy to obtain the energy from oil than from natural gas, and hence a BTU of oil is equivalent to two BTUs of natural gas. The data set has 137 air conditioners (including both window units and central air), 942 heaters (including both boilers and furnaces), and 46 heat pumps with unique observations of $y$, $x$ and $e$. Appendix C.3 reports the summary statistics and additional details on the data.

I estimate equation (11) using ordinary least squares. Table 2 reports the coefficient estimates with standard errors in parentheses. The sum of the coefficients for all three types of temperature control slightly exceeds unity, implying increasing returns to heating and cooling. The increasing returns could partially result from changes in technology with capacity. For example, window units and portable heaters have low capacity and low efficiency, while central heat and air systems (HVACs) have higher capacity and efficiency.

Table 2: Heating and Cooling Production Function Estimates

| | Air conditioners | Heaters | Heat pumps |
|---|---|---|---|
| Exponent on equipment: $\theta$ | 0.271*** | 0.350*** | 0.281*** |
| | (0.020) | (0.031) | (0.011) |
| Exponent on energy: $\eta$ | 0.853*** | 0.767*** | 0.858*** |
| | (0.029) | (0.024) | (0.011) |
| Constant: $\kappa$ | 0.326* | 2.814*** | 0.242** |
| | (0.195) | (0.161) | (0.093) |
| Observations | 137 | 942 | 46 |
| $R^2$ | 0.968 | 0.794 | 0.996 |

Note: The table reports coefficient estimates of equation (11), for air conditioners (column 1), heaters (column 2) and heat pumps (column 3) with standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

It is generally well known that estimating aggregate production functions, say between output, capital, and labor, can be challenging because unobserved changes in total factor productivity affect both the left- and right-hand side variables. However, in this case, there is less scope for unobserved variation in total factor productivity to create endogeneity concerns. Manufacturers measure the unit's capacity and efficiency under a standard set of temperature and humidity conditions, eliminating some of the main sources of variation in total factor productivity across heating and cooling outcomes. The high R-squares further indicate that the vast majority of the variation in heating and cooling output across units is explained by temperature-control equipment and energy.

I set capital's income share in the production of the final good, $\alpha$, equal to 0.26 (Kiyotaki et al., 2011; Nakajima, 2020). This value is lower than the typical value of capital share in a single-asset model (between 0.3 and 0.4) because final good production excludes the capital-intensive housing sector. I set the depreciation rate, $\delta$ equal to 0.066, the average depreciation rate in the national accounts data (NIPA Tables 1.1 and 1.3). I normalize total factor productivity in the production of the final good, $A^y$, to unity. I choose total factor productivity in the production of housing services, $A^s$, equal to 0.12, the average ratio of housing services to residential capital in the national accounts (NIPA Tables 2.4.5 and 1.1).

I choose productivity in the production of electricity, $A^{ee}$, equal to 0.14, the ratio of value-added in electricity production to electricity capital (NIPA Tables 2.1 and U.Value added by Industry). Heaters can use electricity, natural gas, oil or propane. To determine productivity in the production of heating energy, $A^{em}$, I compute the ratio of the average cost per British Thermal Unit (BTU) of electricity compared to the average cost per BTU of heating energy in RECS.[11] This ratio equals 1.93, implying that TFP in heating energy is 1.93 times higher than TFP in electricity: $A^{em} = 1.93 \times A^{ee} = 0.28$. Conditional on values for $A^{ee}$, $A^{em}$, and the heating and cooling productivites, $A^j$, $j \in \{c, h, p\}$, the productivities for the production of temperature-control equipment, $A^{xj}$, $j \in \{c, h, p\}$ amount to a choice of units (which are not relevant for the model). I normalize these three productivities to unity: $A^{xh} = A^{xc} = A^{xp} = 1$.

The bliss point temperature, $\zeta^\star$, equals 18°C (65°F). This is the average daily temperature the EIA uses to construct heating and cooling degree days, their primary measure of heating and cooling demand. Additionally, Albouy et al. (2016) estimate that households' optimal average daily temperature equals 18°C based on a hedonic analysis. The exogenous world interest rate, $r^\star$, equals 4 percent (McGrattan and Prescott, 2003). The coefficient of relative risk aversion, $\sigma$, equals 2. The persistent component of the labor productivity process, $\nu_{it}$, follows an AR(1) process of the form: $\nu_{i,t} = \rho \nu_{i,t-1} + \varepsilon_{it}$, with $\varepsilon_{i,t} \sim N(0, \sigma^2_\varepsilon)$. Parameter $\rho$ denotes the persistence and $\varepsilon_{it}$ is a white noise process with variance $\sigma^2_\varepsilon$. I take the values for $\rho = 0.97$, $\sigma^2_\varepsilon = 0.02$, and the variance of household-specific fixed effect, $\sigma^2_\xi = 0.66$, from Kaplan (2012). I use the Rouwenhorst method to approximate the AR(1) process with a five-state Markov chain. Appendix C.3 reports the Markov transition matrix, $\pi^\nu$, and the invariant distribution, $\Pi^\nu$. Appendix Table C.3 summarizes the values of the externally calibrated parameters.

**Internally determined parameters.** I jointly choose these parameters so that a set of moments in the model match their corresponding empirical targets. While all the parameters depend on all of the targets, some targets are more important for some parameters than others. I discuss each parameter and its primary target in turn.

[11] I calculate the average cost per BTU of electricity as the average cost of electricity used for cooling, among households with air conditioners, and electricity used for heating and cooling, among households with heat pumps. I calculate the average cost per BTU of heating energy the average ratio of total expenditures on heating energy (TOTALDOLSPH) to total BTUs of heating energy used (TOTALBTUSPH).

I choose TFP for heating produced from heater equipment, $A^h$, and TFP for cooling produced from air-conditioner equipment, $A^c$, to match average heating and cooling budget shares among households without heat pumps in RECS. Note that these averages include households that do not have any form of heating or cooling equipment. I choose TFP for heating and cooling from heat-pump equipment, $A^p$, to match the average combined heating and cooling budget share among households with heat pumps. All else constant, higher levels of heating and cooling TFP decrease the energy required to achieve a given indoor temperature, reducing heating and cooling budget shares.

Parameter $\gamma$ controls the effect of housing services on the amount of energy and temperature-control equipment necessary to change the indoor temperature. For example, if $\gamma = 1$, then, all else constant, the amount of required energy and equipment almost doubles when housing services double. Similarly, if $\gamma = 0$, then housing services have no effect on the level of energy and equipment required to change the indoor temperature. Since higher income households purchase more housing services, $\gamma$ is an important determinant of how energy expenditures vary with income. I choose $\gamma$ to match total heating and cooling energy expenditures in the fifth income quintile relative to the first quintile. This ratio equals 1.72, implying that the richest 20 percent of households spend 72 percent more on heating and cooling energy than the poorest 20 percent.

Damage coefficient, $\chi$, controls the variation in energy budget share with income. The relationship between energy budget share and income depends, in part, on how close households' optimal indoor temperatures are to the bliss point. Higher values of $\chi$ raise the utility cost of being away from the bliss point, making the bliss point temperature more of a necessity. As a result, households optimally choose indoor temperatures closer to the bliss point, which, in turn, causes heating and cooling budget shares to fall more steeply with income. I choose $\chi$ to match the ratio of energy budget share (the sum of heating and cooling budget shares) in the fifth income quintile relative to the first quintile. This ratio equals 0.12, implying that energy budget share in the fifth quintile is almost ten times smaller than in the first quintile.

The fixed costs for temperature-control equipment, $\Omega^h$, $\Omega^c$, and $\Omega^p$, determine the fractions of households that have any type of heating and any type of cooling, and the fraction of households that use a heat pump to produce their heating and cooling. In RECS, 95 percent of households have heating, 86 percent have cooling, and 9 percent produce both heating and cooling using heat pumps. I choose the fixed costs to match these targets.

I choose the level of energy assistance per recipient household in each region, $B_n$, such that it covers an average of 83 percent of recipient households' heating and cooling energy expenses. I calculate this target from data on LIHEAP payments for bill pay assistance, assistance programs by local utilities, and heating and cooling expenses among households that received assistance. Appendix C.2 provides more details on the energy assistance data. Consistent with the structure of LIHEAP (DHHS, 2020), the level of energy assistance in the model varies across regions based on the differences in heating and cooling expenses. I assume that energy assistance is received by all households who use some heating or cooling and have a combined capital and labor income below the bottom 5 percent of the income distribution, the average fraction of households that receive LIHEAP funds in the data.

Table 3: Parameter Values: Internal Calibration

| Parameter | Value |
|---|---|
| *Utility* | |
| Weight on housing: $\psi$ | 0.20 |
| Discount factor: $\beta$ | 0.94 |
| Temperature damage coefficient: $\chi$ | 0.01 |
| *Heating and cooling* | |
| Housing exponent: $\gamma$ | 0.10 |
| Heater TFP: $A^h$ | 460 |
| Air conditioner TFP: $A^c$ | 555 |
| Heat pump TFP: $A^p$ | 600 |
| Heater fixed cost: $\Omega^h$ | 0.012 |
| Air conditioner fixed cost: $\Omega^c$ | 0.005 |
| Heat pump fixed cost: $\Omega^p$ | 0.018 |
| *Energy assistance* | |
| Cold region: $B_1$ | 6.9e-03 |
| Cool region: $B_2$ | 5.6e-03 |
| Mild region: $B_3$ | 4.3e-03 |
| Warm region: $B_4$ | 4.0e-03 |
| Hot region: $B_5$ | 4.5e-03 |

Note: The table reports the values of the parameters I choose internally so that a set of moments in the model match their corresponding empirical targets.

Lastly, the weight on housing in the utility function, $\psi$, is pinned down by the ratio of private residential assets to non-residential assets, 0.88 (NIPA Table 1.1). I determine the discount factor, $\beta$, to match the average ratio of US net wealth to output, 3.0, where net wealth equals the sum of fixed assets, consumer durables, and net foreign assets (NIPA Tables 1.1, 1.1.5, BEA International Data Table 1.1). Table 3 reports the values of the internally calibrated parameters.

Table 4 reports the values of the targeted moments in the model and in the data. Overall, the model fits the targeted moments quite closely. However, the model somewhat overstates the fraction of households with heat pumps. Until recently, heat pumps did not work well in the very cold. Depending on when households in the data purchased their temperature-control equipment, this drawback could have discouraged them from buying a heat pump.[12] The model abstracts from the inability of older heat pumps to operate in the very cold, causing it to overstate the fraction of households with heat pumps in the current economy. Importantly, this abstraction is less relevant for the model's predictions going forward because technological advances have made modern heat pumps effective in very cold weather. In the analysis, I show how variation in heat pump use in the current and future economy impacts the distributional effects of climate change.

Table 4: Model Fit

| Target | Model | Data |
|---|---|---|
| Wealth to output ratio | 3.0 | 3.0 |
| Housing to non-housing capital ratio | 0.88 | 0.88 |
| Cooling budget share: HHs w/o heat pumps | 0.0055 | 0.0055 |
| Heating budget share: HHs w/o heat pumps | 0.014 | 0.014 |
| Energy budget share: HHs w/ heat pumps | 0.014 | 0.014 |
| Fraction of HHs with heating | 0.97 | 0.95 |
| Fraction of HHs with cooling | 0.86 | 0.86 |
| Fraction of HHs with heat pumps | 0.18 | 0.09 |
| Energy budget share ratio: high to low income | 0.12 | 0.12 |
| Energy expenditure ratio: high to low income | 1.72 | 1.72 |
| $(\text{Energy assistance/energy expenditures})_n$ | 0.83 | 0.83 |

Note: The table reports the values of the targeted moments in the model and in the data.

It is useful to step back and compare heat pumps with heaters and air conditioners in the calibrated model. The advantage of a heat pump is that the household only needs to buy one type of equipment and that equipment is more efficient than both a heater, $A^p > A^h$, and an air conditioner, $A^p > A^c$. The disadvantage is that the heat pump is expensive. The fixed cost of a heat pump exceeds the combined fixed cost of an air conditioner and a heater $\Omega^p > \Omega^h + \Omega^c$. Moreover, the price of electricity to run the heat pump is greater than the price of energy to run the heater, $p^{ee} > p^{em}$. (Air conditioners also use electricity.)

The above comparisons imply that households will be more inclined to buy heat pumps if (1) they use both heating and cooling and (2) they use more cooling than heating, muting the drawback of the relatively high electricity price. In the calibrated model, heat pumps are primarily used by the richer households in the hot region. These households use both cooling and heating, but use substantially more cooling than heating. The data are largely consistent

[12] Consistent with this notion, 17 percent of households in RECS report using a heat pump as their main source of air conditioning, similar to the fraction as in the calibrated model, but only 9 percent report using a heat pump as their main source of both air conditioning and heating.

with this finding; approximately three quarters of households with heat pumps in the US live in the southern census region.

**External validation.** I compare the predictions of the model with the data for moments that I do not target. Much of the calibration focused on matching variation between income and households' heating- and cooling-related decisions. To validate the model, I instead focus on variation between climate and households' heating- and cooling-related decisions. One challenge with this approach is that the RECS 2015 does not report the household's county or state, making it difficult to match the households in the data to the regions in the model. However, RECS does report annual heating and cooling degree days, which can provide a mapping between the climate in the data and the model.

The number of heating degree days (HDDs) on a given day equals the difference between the mean temperature (the average of the high and the low) and 65$^o$F, if the mean temperature is below 65$^o$F and zero otherwise. Similarly, the number of cooling degree days (CDDs) equals the difference between the mean temperature and 65$^o$F, if the mean temperature is above 65$^o$F, and zero otherwise. Summing HDDs and CDDs over the year provides a weather-based measure of households' annual heating and cooling needs. Using these definitions, I calculate average heating and cooling degree days for each region in the model.

Table 5 reports the correlation (with standard errors in parentheses) between HDDs (panel A) or CDDs (panel B) and a number of important variables in both the model and the data. Households that experience more HDDs (i.e., those in locations with colder and/or longer winters) tend to have higher heating budget shares and are more likely to have heating equipment. Similarly, households that experience more CDDs (i.e., those in locations with warmer and/or longer summers) tend to have higher cooling budget shares are more likely have air-conditioning equipment. The magnitudes of these correlations are similar between the model and the data.

The third row of panel A reports the correlation between HDDs and the indoor temperature of the household's house during the winter months, among households that have heat. The third row of Panel B reports the correlation between CDDs and the indoor temperature of the household's house during the summer months, among households that have cooling.[13] In both the model and the data, higher levels of HDDs and CDDs are associated with households

[13]The winter indoor temperature in the data equals the average of the variables TEMPGONE, TEMPHOME, TEMPNITE, which are defined as the winter temperature when no one is home during the day, someone is home during the day, and at night, respectively. I calculate winter indoor temperature in the model as the household's average indoor temperature over the portion of the year with outdoor temperatures below $\zeta^{\star}$. The survey only asks households that have heating equipment about the indoor temperature during the winter. To ensure that the value in the model is comparable to the data, I only include households with heat in the model calculation. I use an analogues procedure to construct the values of summer indoor temperature in the model and the data. The relevant RECS variables are: TEMPGONEAC, TEMPHOMEAC, TEMPNITEAC.

choosing slightly warmer indoor temperatures. Again, the magnitudes of the correlations are similar between the model and the data.

Table 5: External Validation

| | Model | Data |
|---|---|---|
| A. *Heating degree days* | | |
| Heating budget share | 3.1e-06<br>( 5.8e-08) | 2.6e-06<br>( 2.0e-07) |
| Have heat | 2.0e-05<br>( 8.1e-07) | 3.6e-05<br>( 2.2e-06) |
| Winter indoor temperature ($^{o}$F) | 1.0e-05<br>( 5.0e-06) | 2.6e-05<br>( 4.6e-05) |
| B. *Cooling degree days* | | |
| Cooling budget share | 5.1e-06<br>( 6.8e-08) | 4.2e-06<br>( 2.1e-07) |
| Have air conditioning | 1.2e-04<br>( 3.4e-06) | 1.6e-04<br>( 7.5e-06) |
| Summer indoor temperature ($^{o}$F) | 1.1e-04<br>( 1.2e-05) | 3.8e-04<br>( 1.2e-04) |

Note: Panel A reports the correlations in the model (column 1) and the data (column 2) between annual heating degree days and households' heating budget share (row 1), whether the household has heating equipment (row 2) and the winter indoor temperature for households with heating equipment (row 3). Panel B reports the correlations in the model (column 1) and the data (column 2) between annual cooling degree days and households' cooling budget share (row 1), whether the household has cooling equipment (row 2) and the summer indoor temperature for households with cooling equipment (row 3). Standard errors are in parentheses.

Additionally, I evaluate the model's fit with regards to the income and wealth distributions. The model matches coarse measures of the income and wealth distributions relatively well. For example, the model predicts that 77 percent of income and 93 percent of wealth is held by the top 40 percent, compared to 75 percent and 97 percent respectively in the data. Appendix C.8 further disaggregates these statistics and reports the shares of income and wealth held by each quintile of the distribution in the model and the data. At this finer level of aggregation, the model continues to match the income distribution relatively well, but, as is common in this class of models, it struggles to match the high concentrations of wealth at the top of the wealth distribution (Quadrini and Ríos-Ruel, 1997). The model's quantitative results reveal the welfare consequences of climate change are relatively constant across the higher income (and wealth) quintiles, suggesting that missing the top of the wealth distribution is likely not critical for the paper's results.

## 5. Results

I use the model to run a series of counterfactual experiments. First, I analyze the distributional and welfare consequences of the rightward shift in the temperature distribution caused by climate change. Second, I show that these distributional impacts depend critically on households' ability to adapt to temperature by producing heating and cooling. Third, I study the potential for policy to interact with the distributional consequences of climate change. Fourth, I highlight the importance of income heterogeneity for understanding the aggregate welfare cost of climate change. Finally, I consider an extension of the model in which temperature affects the labor productivity of the lowest income workers.

**Distributional consequences of climate change.** Scientific models predict that climate change will cause a rightward shift in the temperature distribution. Rasmussen et al. (2016) report county-level projections of the annual distribution of average daily temperature through year 2100 under different representative concentration pathways (RCPs) for each climate model in the Coupled Model Inter-comparison Project (CMIP) archive. I use the RCP 8.5 projections, which were developed to approximate global emissions in the absence of large-scale climate policy. I take a probability-weighted average of the outcomes under RCP 8.5 for the different climate models, where the weights are designed to capture the relative probability that the climate model represents the true outcome (Rasmussen et al., 2016; Hsiang et al., 2017). The projections incorporate the effect of climate change on both the mean and the variance of temperature over a year.

I model uncertainty over the climate-change temperature distribution using the same approach as for the contemporary temperature distribution. I assume that there are three possible realizations of the temperature distribution, hot, cold, and moderate, with probabilities, 0.1, 0.1, and 0.8, respectively. The temperature distribution in the moderate year equals the temperature distribution from the scientific projections. To calculate the temperature distributions during a cold year, I assume the ratio of the mean and variance of the temperature distribution in the cold year relative to the moderate year is the same as in the historical data. I make the analogous assumption to determine the temperature distribution in the hot year. See Appendix C.5 for additional details on the climate projections.

Climate change shifts the temperature distributions in the cold, moderate and hot years in all regions to the right. As an example of these shifts, Figure 1 plots the contemporary (solid lines) and climate-change (dashed lines) temperature distributions in a moderate year in the cold (left panel) and hot (right panel) regions. (See appendix Figure C.3 for the analogous figure in the cool, mild, and warm regions.) Appendix Table C.2 reports the annual mean and variance of temperature in each region during a moderate year for the contemporary and

Figure 1: Effect of Climate Change on the Temperature Distribution

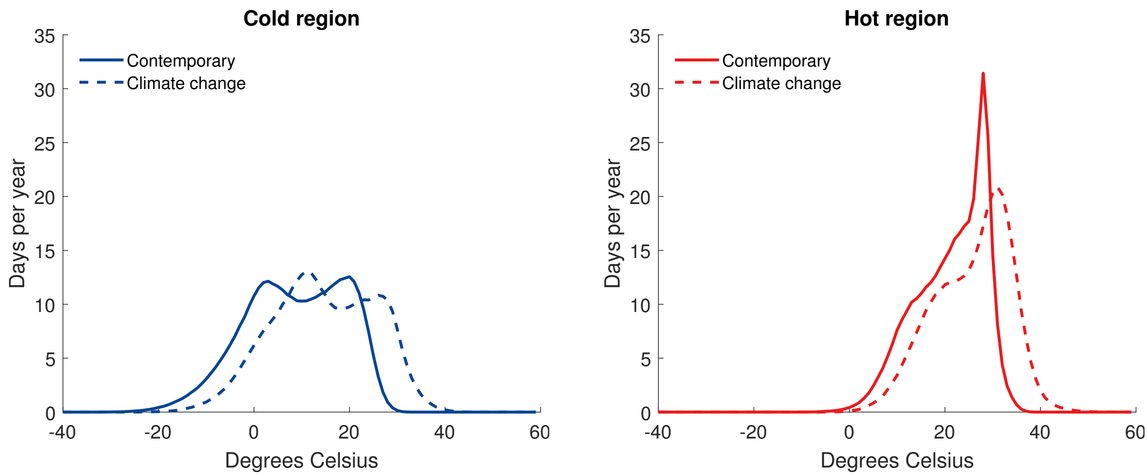


Note: The solid and dashed lines plot the contemporary and climate-change temperature distributions, respectively, in a moderate year in the cold (left panel) and hot (right panel) regions.

climate-change temperature distributions. Climate change increases the mean and variance.

To study the distributional impacts of climate change, I compare outcomes in a contemporary stationary equilibrium, in which I solve the model using the contemporary temperature distribution, to outcomes in a climate-change stationary equilibrium, in which I solve the model using the climate-change temperature distribution. I use the consumption-housing equivalent variation (CHEV) to measure the welfare consequences of climate change. The CHEV is the percent increase in consumption and housing a household would need in every period in the contemporary equilibrium, so that they are indifferent between the contemporary and the climate-change equilibrium. I measure the CHEV conditional the household's region and on their decile of the income distribution. Negative values indicate that climate change makes the household worse off.

I find that the welfare consequences of climate change vary considerably in both magnitude and sign across households. Figure 2 plots the CHEV by income decile in each region. In the colder regions, climate change leads to substantial welfare costs for low-income households, welfare benefits for middle-income households, and near zero welfare impacts for high-income households. In the hotter regions, climate change leads to welfare benefits or relatively small welfare costs for low-income households, larger welfare costs for middle-income households, and again near zero welfare impacts for high-income households.

This heterogeneity in the welfare impacts across income groups primarily results from two channels. First, returning to the intuition from the simple model, I look at the effects of climate change on transfers from nature. I show that the changes in transfers from nature caused by climate change are largely consistent with the welfare impacts among the middle and high

Figure 2: Welfare Consequences of Climate Change

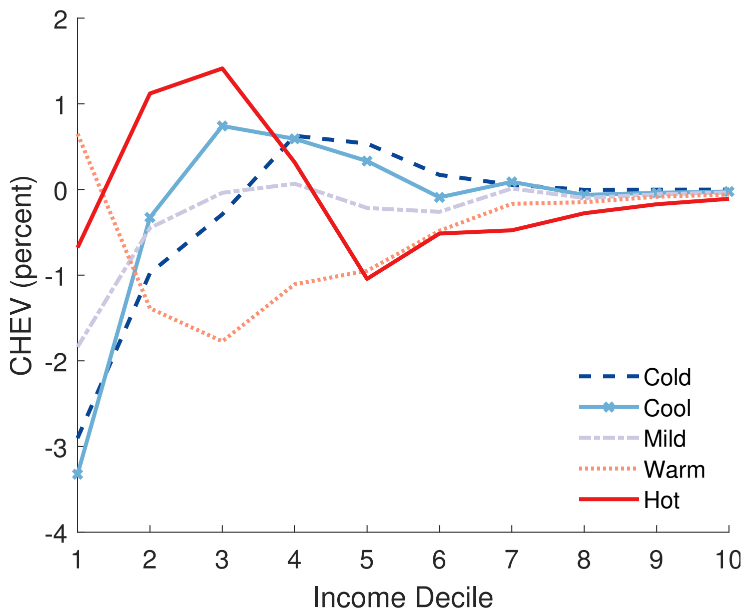


Note: Each line plots the welfare consequences of climate change conditional on the household's region and income decile, where decile one corresponds to the lowest income households. I measure welfare using the consumption-housing equivalent variation (CHEV) between the contemporary and climate-change equilibria. Negative values indicate that climate change makes the household worse off.

income households. Second, I consider how climate change affects households' choices of temperature-control equipment. I show that the effect of climate change on these equipment decisions is particularly important for households in the lower income deciles, and can lead to different welfare consequences for these groups.

I measure the transfers from nature for household $i$ in region $n$, $\Upsilon_{in}$, as the expected annual energy expenditures required to attain the outdoor temperature in each sub-period,

$$\Upsilon_{in} = \begin{cases} E\left[p^{em}\sum_{j=1}^{59} q_{ij}\left(\frac{(\zeta_j - -40)h_{in}^{\gamma}}{A^h(x_{in}^h)^{\theta^h}}\right)^{\frac{1}{\eta^h}} + p^{ee}\sum_{j=60}^{100} q_{ij}\left(\frac{(59-\zeta_j)h_{in}^{\gamma}}{A^c(x_{in}^c)^{\theta^c}}\right)^{\frac{1}{\eta^h}}\right] & : \text{no heat pump} \\ E\left[p^{ee}\sum_{j=1}^{59} q_{ij}\left(\frac{(\zeta_j - -40)h_{in}^{\gamma}}{A^h(x_{in}^p)^{\theta^p}}\right)^{\frac{1}{\eta^p}} + p^{ee}\sum_{j=60}^{100} q_{ij}\left(\frac{(59-\zeta_j)h_{in}^{\gamma}}{A^c(x_{in}^p)^{\theta^p}}\right)^{\frac{1}{\eta^p}}\right] & : \text{heat pump.} \end{cases} \tag{12}$$

I assume a minimum possible outdoor temperature of −40°C and maximum possible outdoor temperature of 59°C, the same range as used in the scientific projections. The first summation in equation (12) is the energy expenditures required to heat the household's house on cold days to the outdoor temperature of $\zeta_j$ from the minimum temperature of −40°C. The second summation is the energy expenditures required to cool the household's house on hot days to

Figure 3: Effect of Climate Change on Transfers From Nature Relative to Income

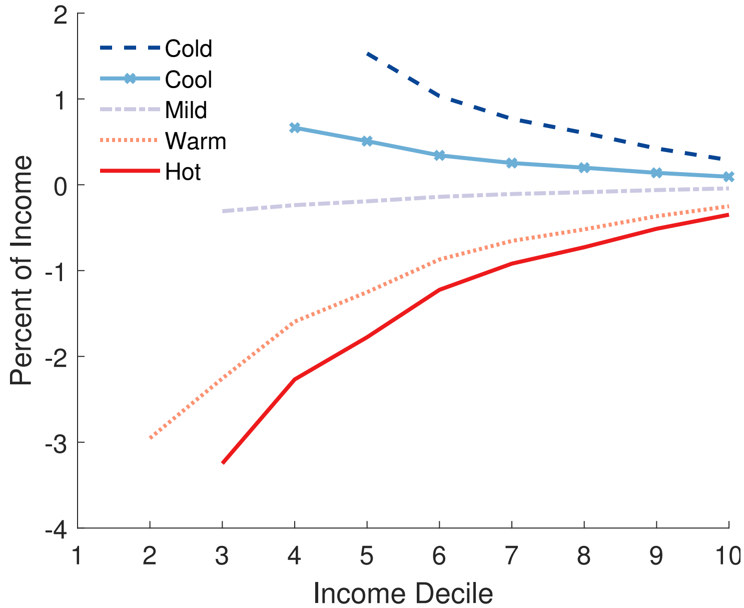


Note: Each line plots the change in transfers from nature relative to income caused by climate change, conditional on the household's region and income decile. Decile one corresponds to the lowest income households. Transfers from nature are defined in equation (12).

the outdoor temperature of $\zeta_j$ from the maximum outdoor temperature of 59°C. Transfers from nature are not defined for households that do not purchase any heating or cooling equipment.

Figure 3 plots the change in transfers from nature caused by climate change as a percent of the household's income.[14] For example, climate change reduces transfers from nature by slightly more than 3 percent of income among households in the third decile in the hot region. Overall, the figure reveals that climate change increases transfers from nature in the colder regions where it leads to more moderate temperatures, and decreases transfers from nature in the hotter regions, where it leads to more extreme temperatures. The magnitude of the change in transfers from nature relative to income falls as income rises. These patterns are consistent with the progressive welfare benefits among the middle- and high-income deciles in the colder regions in Figure 2 and the regressive welfare costs among the middle- and high-income deciles in the hotter regions.

Turning to the second channel, I examine the interaction between the temperature distribution and households' choices of temperature-control equipment. To understand the intuition, consider two temperature distributions: one with only cold days and one with only hot days. If there are only cold days, then households choose to fully specialize in heater equipment

[14] To isolate the direct effect of climate change on transfers from nature, I hold housing and temperature-control equipment constant at their values in the contemporary equilibrium and calculate $\Upsilon_{in}$ using the sub-period weights (the $q's$) from the contemporary temperature distribution and from climate change temperature distributions.

and not purchase any air-conditioner or heat-pump equipment. Similarly, if there are only hot days, then households choose to fully specialize in air-conditioner equipment. All else equal, households benefit from being able to specialize in either heater or air-conditioner equipment because it reduces the total cost of temperature control. Climate change affects how favorable the temperature distribution is to this specialization.

To quantify these effects, I define a specialization favorability index, $S_n \in [0,1]$ in each region $n$,

$$S_n \equiv \mathbb{E}\left( \frac{\left| \sum_{\zeta_j < \zeta^\star} q_{jn}(\zeta^\star - \zeta_j) - \sum_{\zeta_j > \zeta^\star} q_{jn}(\zeta_j - \zeta^\star) \right|}{\sum_{\zeta_j < \zeta^\star} q_{jn}(\zeta^\star - \zeta_j) + \sum_{\zeta_j > \zeta^\star} q_{jn}(\zeta_j - \zeta^\star)} \right). \tag{13}$$

The expectation is taken with respect to the different realizations of the temperature distribution. The numerator of $S_n$ is the difference between the degrees of heating and the degrees of cooling necessary to attain the bliss point in every sub-period. The denominator is the total degrees of heating and cooling necessary to attain the bliss point every sub-period of the year. If $S_n = 0$, then degrees of heating exactly equal degrees of cooling, implying that the temperature distribution is not at all favorable to specialization. Higher values of $S_n$ indicate larger differences between degrees of heating and cooling, implying that the temperature distribution is more favorable to specialization. If $S_n = 1$, then degrees of either heating or cooling are zero. In this case, the household can fully specialize in heating or cooling equipment and still attain the bliss point temperature in every sub-period, if they choose to do so.

Figure 4 plots the specialization favorability index in the contemporary (blue) and the climate-change (orange) equilibrium for each region. In the contemporary equilibrium, the temperature distribution in the colder regions is relatively favorable to specialization. Households in these regions experience only a small number of hot days, allowing households to specialize more in heater equipment. Climate change leads to more hot days, reducing specialization favorability. The opposite effect occurs in hotter regions. Specialization favorability is low in the contemporary equilibrium but increases with climate change.

All else constant, the decrease in specialization favorability in the colder regions reduces the welfare benefits from climate change and the increase in specialization favorability in the hotter regions reduces the welfare costs of climate change. While changes in specialization favorability affect the welfare of all income groups, they have the most pronounced effects among the low-income households because it causes a portion of these households to move to, or away from, complete specialization. Intuitively, this is because these households move between an equilibrium in which they only pay for heater or air-conditioner equipment (complete specialization) and an equilibrium in which they pay for both heater and air-conditioner

Figure 4: Specialization Favorability Index

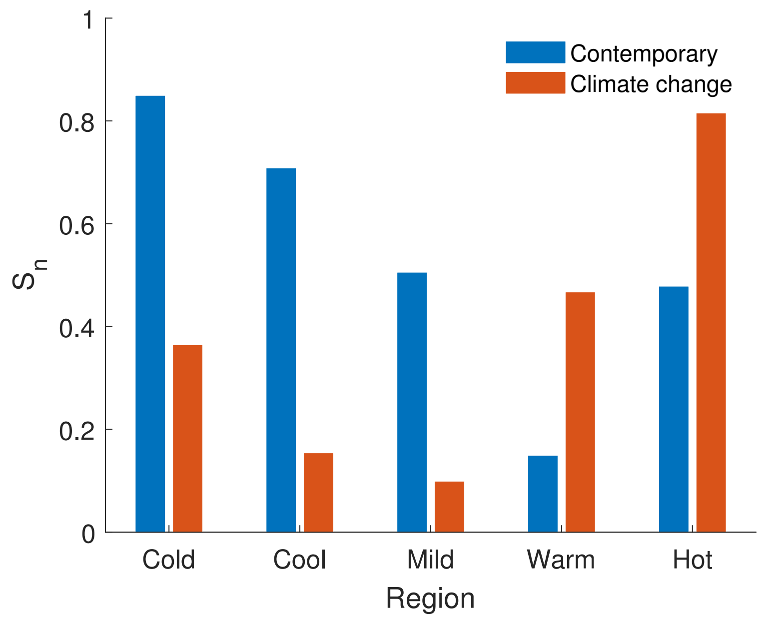


Note: The figure plots the specialization favorability index, $S_n$, defined in equation (13), in each region, for the contemporary (blue) and climate-change (orange) temperature distributions. Higher values of $S_n$ indicate that the temperature distribution is more favorable to specialization.

equipment or heat-pump equipment (incomplete specialization). Mathematically, this is because, all else constant, the welfare impact of moving between a corner solution, where the first order condition for heater or air-conditioner equipment does not hold, and an interior solution, where the first order conditions do hold, tends to be larger than the welfare impact from moving between two interior solutions where the first order conditions always hold.

To understand which households move between complete and incomplete specialization in response to climate change, the left panel of Figure 5 plots the fraction of households in the cold region without air conditioning, meaning they do not have air-conditioner or heat-pump equipment, in the contemporary (blue) and climate-change (orange) equilibria. All households in the cold region have heat in both equilibria. Comparing the blue and orange lines reveals that climate change causes the majority of households in the bottom three deciles of the income distribution to move from a corner solution in which they do not have air conditioning to an interior solution in which they have air conditioning. Among these households, the welfare costs from lower specialization favorability dominate the welfare benefits from higher transfers from nature, leading to the negative welfare effects of climate change in Figure 2.

We see the opposite patterns in the hot region. Here, all households have air conditioning in both equilibria but not all households have heat. The right panel if Figure 5 plots the fraction of households in the hot region without heat, meaning they do not have heater or heat-pump equipment, in the contemporary (blue) and climate-change (orange) equilibria. Most house-

Figure 5: Changes in Specialization

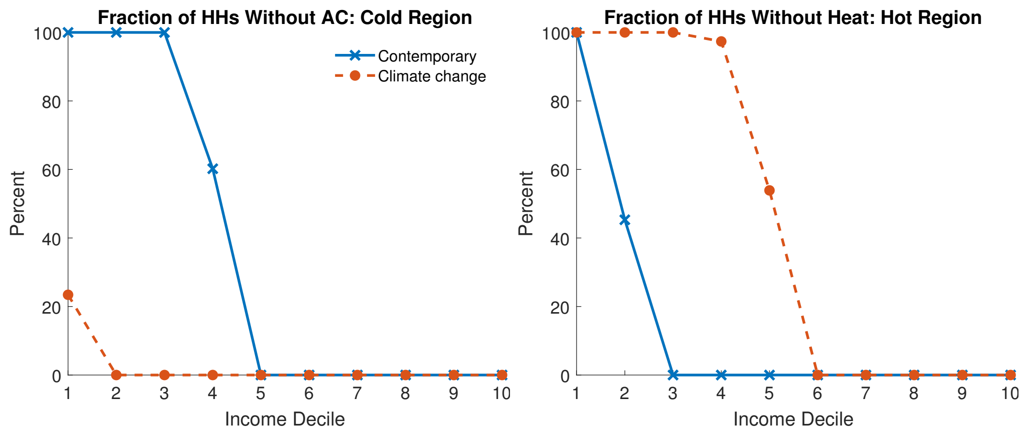


Note: The figure plots the fraction of households without air conditioning in the cold region (left panel) and without heat in the hot region (right panel) in the contemporary (blue) and climate-change (orange) equilibria.

holds in the second through fourth deciles move from an interior solution in the contemporary equilibrium in which they had heat, to a corner solution in the climate-change equilibrium in which they do not have heat. Among these households, the welfare benefits from the move to complete specialization dominate the welfare cost from lower transfers from nature, leading to the positive welfare effects of climate change Figure 2. Appendix D reports the specialization results in the other three regions.

Overall, the heterogeneous effects of climate change increase welfare inequality in the cold, cool, and mild regions and decrease welfare inequality in the warm and hot regions. Table 6 reports the effect of climate change on the Gini coefficient on lifetime welfare. Increases in the Gini coefficient correspond to increases in inequality. Climate change raises inequality in the colder regions despite the progressive increase in transfers from nature because it causes many low-income households in to purchase an air conditioner for the first time. Similarly, climate change lowers inequality in the hotter regions despite the regressive decrease in transfers from nature because it causes many low-income households to not purchase a heater for a first time.

Table 6: Impact of Climate Change on Within Region Welfare Inequality
(Percent change in the Gini Coefficient)

| Cold | Cool | Mild | Warm | Hot |
|---|---|---|---|---|
| 1.00 | 1.01 | 0.51 | -0.18 | -0.06 |

Note: The table reports the effect of climate change on the Gini coefficient. Increases in the Gini correspond to increases in inequality.

In sum, the distributional impacts of climate change in Figure 2 result from the different effects of changes in transfers from nature and specialization favorability across income groups.

All else constant, the negative relationship between income and changes in transfers from nature relative to income in Figure 3 implies that the welfare benefits of climate change fall with income in the colder regions and the welfare costs of climate change fall with income in the hotter regions. However, the comparatively large welfare impact from changes in specialization favorability for low-income households that move between complete and incomplete specialization breaks the monotonicity in this relationship. Overall, climate change increases welfare inequality in the colder regions and decreases welfare inequality in the hotter regions.

**Adaptation to temperature.** The welfare impacts of changes transfers from nature and specialization favorability are closely linked to households' ability to adapt to temperature though heating and cooling. If households cannot adapt to temperature, then, as discussed in the simple model (Section 2), transfers from nature have no effect on household decisions. Moreover, the specialization favorability of the temperature distribution is irrelevant if the household does not have the option to purchase temperature-control equipment. Without these two channels, the welfare consequences of climate change are the same for all households within a region.

To see this result quantitatively, I solve for a contemporary and a climate-change equilibrium without adaptation to temperature. Since households cannot produce heating and cooling, the indoor temperature always equals the outdoor temperature in both equilibria. I measure the welfare effects of climate change without temperature adaptation as the CHEV between the contemporary equilibrium without temperature adaptation and the climate-change equilibrium without temperature adaptation. Appendix Figure D.2 plots the distributional effects of climate change for this case. In contrast to the baseline results in Figure 2, the lines in Figure D.2 are flat, implying that climate change would have identical welfare effects across income groups if households could not adapt to temperature.

**Role of policy.** I study the effects of two policies on the welfare impacts of climate change: (1) energy assistance for low-income households and (2) a mandate that all households use heat pumps for heating and cooling. Beginning with energy assistance, in the baseline simulation, I assume that the level of assistance does not change with climate change. A plausible alternative to this assumption is that policymakers adjust the level of energy assistance in proportion to the changes in heating cooling needs caused by the change in climate. To explore the implications of this potential policy response, I solve for a new climate-change equilibrium in which I scale the level of energy assistance in each region such that it also averages 83 percent of recipient households' energy expenditures in the climate-change equilibrium, the same as in the contemporary equilibrium. This adjustment reduces energy assistance in the cold and cool regions by 24 and 20 percent, respectively, and increases energy assistance in the mild, warm, and hot regions by 9, 36, and 79 percent respectively. Total government spending on energy

Figure 6: Impact of Energy Assistance on the Welfare Impact of Climate Change: Households in the First Income Decile

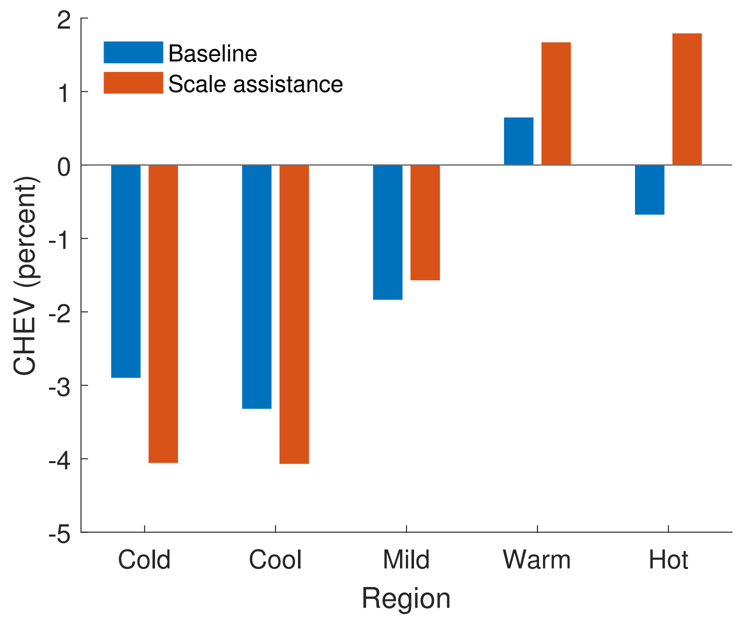


Note: The bars plot the welfare consequences of climate change for households in the first income decile in the baseline simulation (blue) and when energy assistance scales with climate change (orange).

assistance increases by 20 percent, though the total size of the program is small in comparison to GDP, equal to less than 0.02 percent of output in both simulations.

Figure 6 plots the welfare impacts of climate change for households in the first income decile when policymakers scale energy assistance in response to climate change (orange bars) compared to the welfare impacts in the baseline simulation (blue bars). Overall, scaling energy assistance partially offsets the welfare impacts stemming from changes in transfers from nature. Energy assistance falls in the colder regions (where transfers from nature increase), raising the welfare cost of climate change. In contrast, energy assistance rises in the hotter regions (where transfers from nature decrease), reducing the welfare cost or magnifying the welfare benefits from climate change.

Turning to the heat-pump mandate, policymakers have designed a range of incentives to encourage households to switch from traditional air conditioners and heaters to heat pumps. Some of the most extreme versions are bans on natural gas or natural gas appliances. For example, New York and Washington DC have banned the use of most natural gas in newly constructed buildings while California has banned the sale of new natural gas furnaces after 2030. In many places, heat pumps are the most viable alternative to natural gas heating, making a ban on natural gas isomorphic to a mandate that households use heat pumps for heating and cooling.

To explore how a heat-pump mandate would affect the distributional consequences of climate change, I solve for a contemporary and a climate-change equilibrium in which households' only option for temperature-control equipment is a heat pump. I measure the welfare consequences of climate change under a heat-pump mandate as the CHEV between these two equilibria. Figure 7 plots this CHEV by income decile in each region.

Figure 7: The Welfare Impacts of Climate Change Under a Heat-Pump Mandate

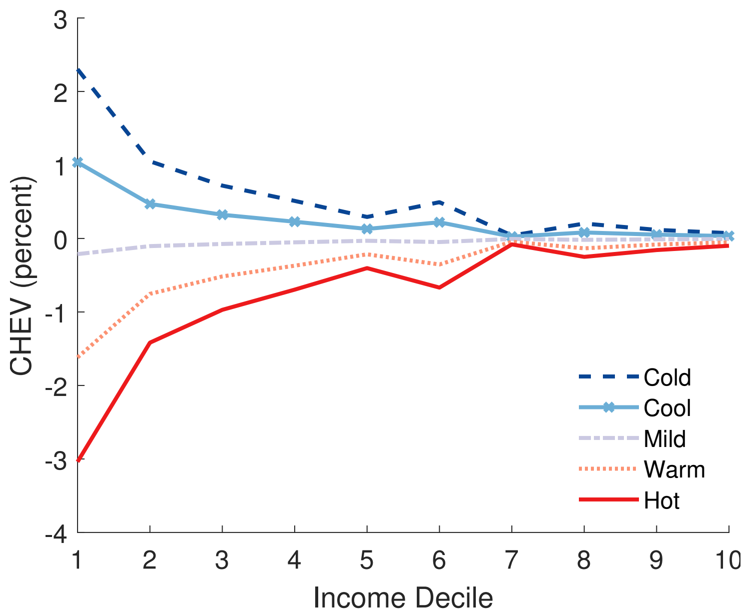


Note: Each line plots the welfare consequences of climate change under a heat-pump mandate conditional on the household's region and income decile, where decile one corresponds to the lowest income households. I measure welfare using the consumption-housing equivalent variation (CHEV) between the contemporary equilibrium under a heat-pump mandate and climate-change equilibrium under a heat-pump mandate. Negative values indicate that climate change makes the household worse off.

The welfare effects of climate change under a heat-pump mandate are entirely determined by the changes in transfers from nature. Changes in households' ability to specialize in heater or air-conditioner equipment no longer matter because all households have one piece of equipment, the heat pump, that can be used for both heating and cooling. Consequently, climate change leads to progressive welfare benefits across all income deciles in the colder regions, where it increases transfers from nature, and to regressive welfare costs across all income deciles in the hotter regions, where it decreases transfers from nature.

**Importance of income heterogeneity for aggregate outcomes.** Aggregate outcomes, such as the social cost of carbon and the effects of climate change on capital, labor, and output, are key inputs into the design of climate policy. To determine these aggregates, economists and policymakers typically simulate a climate-economy model with a representative agent and no within-region income heterogeneity, such as DICE (Barrage and Nordhaus, 2023) or Golosov et al. (2014). A key insight from the macroeconomic literature on heterogeneous agents is that the underlying income heterogeneity matters for aggregate outcomes, such as, in the macroeconomic context, the welfare cost of business cycles, inflation, and asset pricing (Krueger et al., 2016; Heathcote et al., 2009). Here, I demonstrate that income heterogeneity matters for understanding the aggregate welfare cost of the climate change within each region of the US.

Table 7: Welfare Impact of Climate Change (CHEV, percent)

| | Cold | Cool | Mild | Warm | Hot |
|---|---|---|---|---|---|
| Baseline simulation: aggregate | -0.99 | -0.87 | -0.68 | -0.61 | 0.13 |
| Baseline simulation: average HH | 0.05 | 0.09 | 0.01 | -0.17 | -0.48 |
| Representative agent simulation | -0.01 | -0.11 | -0.18 | -0.26 | -0.43 |

Note: The table reports the welfare cost of climate change for each region in the baseline model (row 1), for households in the income decile with average income (row 2) and from a representative agent version of the model (row 3).

The first row of Table 7 reports the aggregate (average) welfare cost of climate change in each region, measured as the CHEV conditional on the household's region. To understand the importance of income heterogeneity for the aggregate welfare cost, I consider two different ways to abstract from income heterogeneity in rows two and three of Table 7. Row two reports the welfare impact for the average household in each region, measured as the CHEV conditional on region for households in the seventh income decile, which contains average income. Row three reports the welfare impact from a representative agent version of the model in which all households receive the average asset and labor income from the contemporary equilibrium in the baseline simulation in every period.

Comparing row one with either row two or three suggests that abstracting from income heterogeneity reduces the welfare cost of climate change in all regions except for the hot region. For example, in the cold region, the aggregate welfare cost is -0.99 (row 1) even though the household with the average income actually experiences small welfare benefits from climate change (row 2). The aggregate welfare cost incorporates the large welfare costs to households at the bottom of the distribution, which are not offset by large welfare gains for households at the top. Focusing on a representative household (row two or three) misses this heterogeneity thus biases the welfare cost of climate change down. Similarly, in the hot region, the aggregate welfare benefit (row 1) incorporates the large benefits to the households at the bottom of the distribution that forgo heating, which are not offset by large welfare costs for households at the top. Again, focusing on a representative household (row two or three) misses this heterogeneity, and, in this case, biases the welfare costs of climate change up.

**Climate change and labor productivity.** I assume that temperature does not affect labor productivity in the main analysis. While this assumption is plausible for workers who work indoors where firms have access to heating and cooling, it is perhaps less plausible for workers who work outdoors in sectors like construction and agriculture. According to the Occupational Requirements Survey, approximately eight percent of US workers were constantly or frequently exposed to the outdoors in 2023, and thus could experience changes in labor productivity as a result of climate change. Many of these outdoor workers are likely to be low income, suggesting that the effect of temperature on labor productivity could have important implications for the distributional impacts of climate change. To explore this possibility, I design an experiment in which climate change reduces labor productivity for low-income households.

Importantly, both hot and cold temperatures affect labor productivity. To quantify these effects, I use a rule-of-thumb estimate for the affect of temperature on labor productivity from Richardson Products, a construction cost database. The Richardson database has been widely used for over 50 years by US construction companies to estimate project costs. The database includes estimates of the labor hours required to complete a project and an adjustment factor for these estimates based on temperature. The adjustment factor increases hours by 1 percent for every degree outside of 40°F - 85°F. Using this adjustment factor, I calculate the effect of climate change on labor productivity in each region. Climate change increases labor productivity by 2 and 0.9 percent in the cold and cool regions respectively, and decreases labor productivity by 0.1, 1.1, and 2.2 percent in the mild, warm, and hot regions, respectively.

I solve the model for a new climate-change equilibrium in which some workers are affected by these labor productivity changes. I consider two cases: (1) the change in labor productivity affects all workers with the lowest labor income shock, $z = z^1$, which corresponds to the bottom three percent of workers and (2) the change in labor productivity affects all workers with the the lowest two values of the labor productivity shock, $z \in \{z^1, z^2\}$, which corresponds to the bottom 16 percent of workers. The two cases bound the fraction of workers who are exposed to the outdoors reported in the Occupational Requirements Survey (eight percent).

Figure 8 reports the welfare consequences of climate change for households in the first income quintile, in the baseline simulation (blue), and when climate change affects labor productivity for households with the lowest value (orange) and lowest two values (yellow) of the labor income shock. I focus on households in the first income quintile, since, by assumption, climate change directly affects labor productivity for many of these households. The higher labor productivity from climate change reduces the welfare cost of climate change for households in the first income quintile in the cold and cool regions. Similarly, the lower labor productivity from climate change increases the welfare cost of climate change for households in the first quintile in the mild and warm regions and eliminates the welfare benefits for households in the first quintile in hot region. The magnitude of the effects are largest in the cold and hot regions because these regions experience the largest changes in labor productivity.

Figure 8: Effect of Changes in Labor Productivity on the Welfare Impacts of Climate Change: Households in the First Income Quintile

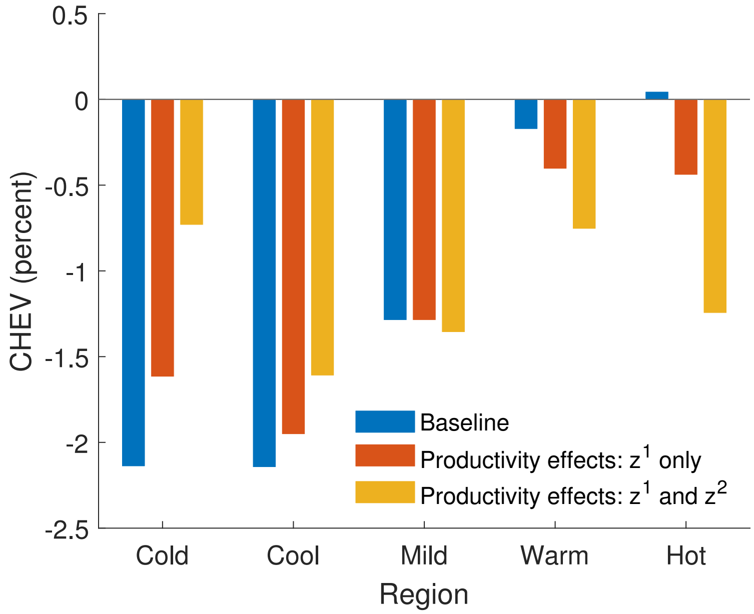


Note: The bars plot the welfare impacts of climate change for households in the first income quintile in the baseline simulation (blue), when climate change reduces lowest value (orange) and the lowest two values (yellow) of the labor productivity shock.

## 6. Conclusion

The welfare consequences of rising temperatures from climate change vary considerably across the income distribution. I show that one can view the outdoor temperature as a transfer from nature to households. These transfers are received by all households, but they have a larger impact on welfare for low-income households because they are a larger share of their income. Climate change increases transfers from nature in the colder regions because it leads to a more moderate temperature distribution, and it decreases transfers from nature in the hotter regions because it leads to a more extreme temperature distribution. In isolation, the changes in transfers from nature imply that climate change creates progressive welfare benefits in the colder regions and regressive welfare costs in the hotter regions.

In addition to transfers from nature, changes in specialization favorability also have important implications for the distributional effects of climate change. All else equal, households would prefer temperature distributions that allow them to fully specialize in heater or air-conditioner equipment and still attain indoor temperatures near the bliss point in every sub-period. Climate change makes it harder for households to specialize in heater equipment in the colder regions, causing some low-income households to purchase their first air conditioner. The welfare cost of this move away from complete specialization dominates the welfare benefits from higher transfers from nature. Similarly, climate change makes it easier for households to specialize in air-conditioner equipment in the hotter regions, causing some low-income households to forgo heaters. The welfare benefit from this move to complete specialization dominates the welfare cost of lower transfers from nature.

In sum, the analysis highlights that the welfare consequences of climate change vary substantially across income groups. The paper focuses explicitly on the direct effects of higher temperatures in the US. The quantitative importance of income heterogeneity in this context suggests that exploring the distributional impacts of higher temperatures in other countries or the distributional impacts of other types of climate damage would be promising avenues for future work. An interesting extension of the paper would be to combine this work focused income heterogeneity with the literature focused on spatial heterogeneity (see the citations in the introduction). Such an extension would allow one to consider how income and wealth interact with migration decisions in response to climate change.

## References


Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012): "The Environment and Directed Technical Change," *American Economic Review*, 102, 131–166.

Administration for Children and Families (2020): "Report to Congress Low Income Home Energy Assistance Program (LIHEAP). An Assessment of the Program's Formula and Allocations of Funding Among States: Appendices," *Department of Health and Human Services*.

Aiyagari, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109, 659–684.

Albouy, D., W. Graf, R. Kellogg, and H. Wolff (2016): "Climate Amenities, Climate Change, and American Quality of Life," *Journal of the Association of Environmental and Resource Economists*, 3, 205–246.

Barrage, L. (2020): "Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy," *The Review of Economic Studies*, 1–39.

Barrage, L. and W. D. Nordhaus (2023): "Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model," *NBER Working paper w31112*.

Behrer, A., R. Park, G. Wagner, C. Golja, and D. Keith (2021): "Heat Has Larger Impacts on Labor in Poorer Areas," *Environmental Research Communications*, 3, 095001.

Belfiori, M. E. and M. Macera (2022): "Climate Policy and Inequality," *Working paper*.

Bewley, T. (1977): "The Permanent Income Hypothesis: A Theoretical Formulation," *Journal of Economic Theory*, 16, 252–292.

Bilal, A. and E. Rossi-Hansberg (2023): "Anticipating Climate Change Across the United States," *NBER working paper 31323*.

Blanz, A. (2023): "Climate-related Agricultural Productivity Losses through a Poverty Lens," *Working paper*.

Carleton, T., A. Jina, M. Delgado, M. Greenstone, T. Houser, S. Hsiang, A. Hultgren, R. E. Kopp, K. E. McCusker, I. Nath, J. Rising, A. Rode, H. K. Seo, J. Simcock, A. Viaene, J. Yuan, and A. Zhang (2022): "Valuing the Global Mortality Consequences of Climate Change Accounting for Adaptation Costs and Benefits," *The Quarterly Journal of Economics*, 137, 2037–2105.

Chiroleu-Assouline, M. and M. Fodha (2014): "From Regressive Pollution Taxes to Progressive Environmental Tax Reforms," *European Economic Review*, 69, 126–142.

Conte, B., K. Desmet, and E. Rossi-Hansberg (2022): "On the Geographic Implications of Carbon Taxes," *NBER Working Paper 30678*.

Cruz, J.-L. and E. Rossi-Hansberg (2024): "The Economic Geography of Global Warming," *Review of Economic Studies*, 91, 899–939.

Dennig, F., M. B. Budolfson, M. Fleurbaey, A. Siebert, and R. H. Socolow (2015): "Inequality, Climate Impacts on the Future Poor, and Carbon Prices," *Proceedings of the National Academy of Sciences*, 112, 15827–15832.

Deryugina, T., L. Kawano, and S. Levitt (2018): "The Economic Impact of Hurricane Katrina on its Victims: Evidence From Individual Tax Returns," *American Economic Journal: Applied Economics*, 10, 202–233.

DHHS (2020): *Report to Congress, Low Income Home Energy Assistance Program (LIHEAP); An Assessment of the Program's formula and Allocations of Funding Among States*.

Doremus, J. M., I. Jacqz, and S. Johnston (2022): "Sweating the Energy Bill: Extreme Weather, Poor Households, and the Energy Spending Gap," *Journal of Environmental Economics and Management*, 112, 102609.

Douenne, T., A. J. Hummel, and M. Pedroni (2023): "Optimal Fiscal Policy in a Climate-Economy Model with Heterogeneous Households," *SSRN working paper 4018468*.

EPA (2021): "Climate Change Indicators: Heat-Related Deaths," *Environmental Protection Agency*.

Fenichel, E. P., J. K. Abbott, and S. Do Yun (2018): "The Nature of Natural Capital and Ecosystem Income," *Handbook of Environmental Economics*, 4, 85–142.

Fried, S., K. Novan, and W. B. Perterman (2018): "The Distributional Effects of a Carbon tax on Current and Future Generations," *Review of Economic Dynamics*, 30, 30–46.

——— (2022): "Understanding the Inequality and Welfare Impacts of Climate Policy," *Working paper*.

Fullerton, D. and G. Heutal (2007): "The General Equilibrium Incidence of Environmental Taxes," *Journal of Public Economics*, 91, 571–591.

Gervais, M. (2002): "Housing Taxation and Capital Accumulation," *Journal of Monetary Economics*, 49, 1461–1489.

Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2014): "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 82, 41–88.

Grainger, C. A. and C. D. Kolstad (2010): "Who Pays a Price on Carbon?" *Environmental and Resource Economics*, 46, 359–376.

Hassler, J., P. Krusell, and C. Olovsson (2021): "Directed Technical Change as a Response to Natural Resource Scarcity," *Journal of Political Economy*, 129, 3039–3072.

Heathcote, J., K. Storesletten, and G. L. Violante (2009): "Quantitative Macroeconomics With Heterogeneous Households," *Annu. Rev. Econ.*, 1, 319–354.

Herman, S., A. Katz, L. Loebach, and S. McCulla (2003): "Fixed Assets and Consumer Durable Goods in the United States, 1925-97," *Bureau of Economic Analysis*.

HSIANG, S., R. KOPP, A. JINA, J. RISING, M. DELGADO, S. MOHAN, D. RASMUSSEN, R. MUIR-WOOD, P. WILSON, M. OPPENHEIMER, ET AL. (2017): "Estimating Economic Damage from Climate Change in the United States," *Science*, 356, 1362–1369.

HUGGETT, M. (1993): "The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies," *Journal of Economic Dynamics and Control*, 17, 953–969.

JEON, W. (2023): "Pricing Externalities in the Presence of Adaptation," *Working paper*.

KAPLAN, G. (2012): "Inequality and the Life Cycle," *Quantitative Economics*, 3, 471–525.

KIYOTAKI, N., A. MICHAELIDES, AND K. NIKOLOV (2011): "Winners and Losers in Housing Markets," *Journal of Money, Credit and Banking*, 43, 255–296.

KORNEK, U., D. KLENERT, O. EDENHOFER, AND M. FLEURBAEY (2021): "The Social Cost of Carbon and Inequality: When Local Redistribution Shapes Global Carbon Prices," *Journal of Environmental Economics and Management*, 107, 102450.

KRUEGER, D., K. MITMAN, AND F. PERRI (2016): "Macroeconomics and Household Heterogeneity," *Handbook of Macroeconomics*, 2, 843–921.

KRUSELL, P. AND A. A. SMITH (2022): "Climate Change Around the World," *NBER working paper 30338*.

KUHN, M. AND J.-V. RIOS-RULL (2016): "2013 Update on the U.S. Earnings, Income, and Wealth Distributional Facts: A View From Macroeconomics," *Federal Reserve Bank of Minneapolis Quarterly Review*, 37, 1–73.

MALAFRY, L. AND P. BRINCA (2022): "Climate Policy in an Unequal World: Assessing the Cost of Risk on Vulnerable Households," *Ecological Economics*, 194, 107309.

MCGRATTAN, E. R. AND E. C. PRESCOTT (2003): "Average Debt and Equity Returns: Puzzling?" *American Economic Review*, 93, 392–397.

METCALF, G. (2007): "A Proposal for a US Carbon Tax Swap: An Equitable Tax Reform to Address Global Climate Change," *Hamilton Project, Brookings Institute*.

METCALF, G. E. (1999): "A Distributional Analysis of Green Tax Reforms," *National Tax Journal*, 655–681.

NAKAJIMA, M. (2020): "Capital Income Taxation with Housing," *Journal of Economic Dynamics and Control*, 115, 103883.

NATH, I. (2022): "Climate Change, the Food Problem, and the Challenge of Adaptation Through Sectoral Reallocation," *Working paper*.

NHCHC (2022): "Cold-Related Injuries," *National Health Care for the Homeless Council*.

NORDHAUS, W. D. AND Z. YANG (1996): "A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies," *The American Economic Review*, 741–765.

PARK, J., M. BANGALORE, S. HALLEGATTE, AND E. SANDHOEFNER (2018): "Households and Heat Stress: Estimating the Distributional Consequences of Climate Change," *Environment and Development Economics*, 23, 349–368.

PARRY, I. (2004): "Are Emissions Permits Regressive?" *Journal of Environmental Economics and Management*, 47(2), 364–387.

PARRY, I. AND R. WILLIAMS (2010): "What Are the Costs of Meeting Distributional Objectives for Climate Policy?" *B.E. Journal of Economic Analysis & Policy*, 10(2).

QUADRINI, V. AND J.-V. RÍOS-RUEL (1997): "Understanding the US Distribution of Wealth," *Federal Reserve Bank of Minneapolis Quarterly Review*, 21, 22–36.

RASMUSSEN, D., M. MEINSHAUSEN, AND R. E. KOPP (2016): "Probability-Weighted Ensembles of US County-Level Climate Projections for Climate Risk Analysis," *Journal of Applied Meteorology and Climatology*, 55, 2301–2322.

RODE, A., T. CARLETON, M. DELGADO, M. GREENSTONE, T. HOUSER, S. HSIANG, A. HULTGREN, A. JINA, R. E. KOPP, K. E. MCCUSKER, ET AL. (2021): "Estimating a Social Cost of Carbon for Global Energy Consumption," *Nature*, 598, 308–314.

ROTH TRAN, B. AND D. J. WILSON (2022): "The Local Economic Impact of Natural Disasters," *Federal Reserve Bank of San Francisco Working Paper 2020-34*.

RUDIK, I., G. LYN, T. WEILIANG, AND A. ORTIZ-BOBEA (2022): "The Economic Effects of Climate Change in Dynamic Spatial Equilibrium," *Working paper*.

SNOW, A. (6/20/2022): "Sweltering Streets: Hundreds of Homeless Die in Extreme Heat Each Year," *LA Times, Associated Press*.

STROBL, E. (2011): "The Economic Growth Impact of Hurricanes: Evidence From US Coastal Counties," *Review of Economics and Statistics*, 93, 575–589.

WILLIAMS, R., H. GORDON, D. BURTRAW, AND J. CARBONE (2015): "The Initial Incidence of a Carbon Tax Across Income Groups," *National Tax Journal*, 68(1), 195–214.

# Appendix

## A. Simple Model

I first describe the model solution when adaptation in infeasible. I assume the outdoor temperature is greater than the extreme lower bound necessary for survival, $\zeta > \underline{\zeta}$. In this case, the indoor temperature equals the outdoor temperature $T_i = \zeta^\star$. The optimized levels of consumption and housing are, $c_i^\star = \alpha y_i$ and $h_i^\star = (1-\alpha)y_i/(2p^h)$. The percent change in utility from a marginal increase in outdoor temperature equals:

$$\frac{\partial u(c_i^\star, h_i^\star, T_i^\star)/\partial \zeta}{u(c_i^\star, h_i^\star, T_i^\star)} = \frac{G'(\zeta)}{G(\zeta)} + \frac{1-\alpha}{\zeta^\star}$$

Thus, the welfare gains from an increase in outdoor temperature do not depend on household income.

I next describe the model solution for the opposite corner, when households choose indoor temperature equal to the bliss point, $T_i^\star = \zeta^\star$. If the household's optimal indoor temperature equals the bliss point, $T_i^\star = \zeta^\star$, then the optimal levels of consumption and housing on a cold day, $\zeta < \zeta^\star$, are $c_i^\star = \alpha(y_i + p^{eh}\zeta - p^{eh}\zeta^\star)$ and $h_i^\star = (1-\alpha)(y_i + p^{eh}\zeta - p^{eh}\zeta^\star)$. Again, the outdoor temperature, $\zeta$, acts as a transfer from nature, augmenting the household's income by $p^{eh}\zeta$. The same analysis as for the interior solution reveals that moving towards more extreme temperatures has higher welfare costs for lower income households.

The analysis of the simple model in the main text focused on a cold day, $0 < \zeta < \zeta^\star$. I provide the parallel analysis for a hot day, $\zeta^\star < \zeta < 2\zeta^\star$. Focusing on an interior solution, the optimal levels of consumption, housing, and indoor temperature are

$$\begin{aligned} c_i^\star &= (y_i + p^{ec}(2\zeta^\star - \zeta))\left(\frac{\alpha}{2-\alpha}\right), \quad h_i^\star = \left(\frac{y_i + p^{ec}(2\zeta^\star - \zeta)}{p^h}\right)\left(\frac{1-\alpha}{2-\alpha}\right) \\ \text{and } \; 2\zeta^\star - T_i^\star &= \left(\frac{y_i + p^{ec}(2\zeta^\star - \zeta)}{p^e}\right)\left(\frac{1-\alpha}{2-\alpha}\right). \end{aligned} \tag{A.1}$$

The heating and cooling energy required to achieve indoor temperature $T_i^\star$ are:

$$e_i^{h^\star} = 0 \quad \text{and} \quad e_i^{c\star} = \zeta - T_i^\star.$$

As was the case for cold days, the optimized levels of consumption, housing, and indoor temperature in equation (A.1) reveal that the outdoor temperature, $\zeta$, acts as a transfer from nature to the household. Nature gives the household $2\zeta^\star - \zeta$ degrees of cooling for free, augmenting its income by $p^{ec}(2\zeta^\star - \zeta)$. Importantly, all households, regardless of their income, receive the same transfer from nature. However, lower-income households value the transfer more than higher-income households because it is a larger share of their budgets. Thus, an increase in the transfer from nature from a colder day creates larger welfare gains for lower income households, while a decrease in the transfer from nature from a hotter day creates larger welfare losses for lower income households.

## B. Quantitative Model

I define a stationary recursive competitive equilibrium. Throughout the definition of the equilibrium, I suppress the individual household subscripts. Let $g=(a,v,n)$ denote the vector of household states and characteristics at the start of the period. Let $\mu$ be the invariant, cross-sectional distribution over the household states, characteristics, and realizations of the temperature distribution.

Given an energy assistance policy, $B(g)$, a labor-income tax, $\tau$, and an international interest rate, $r^\star$, a *stationary recursive competitive equilibrium* consists of time-invariant value and policy functions for households $\{V(g),h(g),x^h(g),x^c(g),x^p(g)\}$, $\{e_j^{mh}(g,Q),e_j^{ec}(g,Q),e_j^{eh}(g,Q),c_j(g,Q)\}_{j=0}^J$, production plans for firms, $\{k_n^y,l_n,k_n^h,k_n^{em},k_n^{ee},k_n^{xh},k_n^{xc},k_n^{xp}\}_{n=1}^N$, prices, $\{w,R,p^h,p^{em},p^{ee},p^{xh},p^{xc},p^{xp}\}$ and stationary distribution, $\mu$, such that:

1. Given prices and policies, the household value function, $V$, solves the optimization problem in equation (9) and $h(g),x^h(g),x^c(g),x^p(g)$, $\{e_j^{mh}(g,Q),e_j^{ec}(g,Q),e_j^{eh}(g,Q),c_j(g,Q)\}_{j=0}^J$, are the associated policy functions.

2. Prices satisfy equation (10).

3. The markets for housing services and temperature-control equipment clear in each region:

$$\int h(g)d\mu_{g|n}=A^h k_n^h,\quad \int x^h(g)d\mu_{g|n}=A^{xh}k_n^{xh},\quad \int x^c(g)d\mu_{g|n}=A^{xc}k_n^{xc},$$
$$\int x^p(g)d\mu_{g|n}=A^{xp}k_n^{xp}.$$

4. The energy markets clear in each region in each period:

$$\sum_{j=1}^{J}\int q_j e_j^m(g,Q)d\mu_{g|n}=A^{em}k_n^{em},\quad \text{and}\quad \sum_{j=1}^{J}\int q_j(e_j^{ec}(g,Q)+e_j^{eh}(g,Q))d\mu_{g|n}=A^{ee}k_n^{ee}.$$

5. The federal budget constraint clears:

$$\int \tau wz d\mu = \int B(g) d\mu.$$

6. The labor market clears in each region:

$$\int z d\mu_{g|n} = l_n.$$

Note that the small-open-economy assumption implies that the capital market does not need to clear domestically. Households' asset holdings do not necessarily equal the sum of firms' capital demands because firms can buy capital from abroad and household assets can be used for capital abroad.

## C. Data and Calibration

### C.1. RECS

RECS 2015 reports information on annual household income divided into eight bins, ranging from less than $20,000 to over $140,000. I use Census income Table HINC-06 to assign the average value of income to each bin. I code a household has having a heat pump if they report that both their heating and air-conditioning equipment are a heat pump (EQUIPM =4, and CENACHP = 1).

### C.2. Energy assistance

I focus explicitly on federal and local assistance for heating and cooling energy. This assistance is primarily provided through LIHEAP and through separate state and local utility programs. Data on LIHEAP are from the Administration for Children and Families Report to congress (Administration for Children and Families, 2020).[15] The data provide information on the number of households that receive assistance with their energy bills and the average level of bill-pay assistance from 2008 - 2019. Data on additional utility-specific programs can be downloaded from the LIHEAP Clearing House.[16] Since many of the utilities use the same eligibility criteria as for LIHEAP, I assume that the same fraction of households receive the supplemental utility insurance as receive LIHEAP (5 percent).

[15]The report can be downloaded from: www.acf.hhs.gov/sites/default/files/documents/ocs/rpt_liheap_congressional_request_for_formula_analysis_appendices.pdf. Information on how LIHEAP benefits are allocated in each state can be downloaded from: https://liheapch.acf.hhs.gov/delivery/benefits.htm.

[16]The data can be downloaded from https://liheapch.acf.hhs.gov/snapshots.htm.

Approximately one quarter of households are served by utilities that offer utility-specific assistance.[17] Utility-specific assistance typically takes the form of either a dollar reduction in the monthly base charge (base discount) or a percentage reduction in the marginal cost of energy (percentage discount). To construct the calibration target, I need to determine the average dollar values of the percentage and base discounts across all utilities, including those that do not offer assistance. I do this calculation by census division, the smallest geographic unit in the RECS. I take the following steps:

1. Calculate the weighted average base discount across all utilities, including those that do not offer any assistance (and thus have a base discount of zero). The weights are proportional to the number of households served by each utility.

2. Calculate the weighted average percentage discount across all utilities, including those that do not offer any assistance (and thus have a percentage discount of zero). The weights are proportional to the number of households served by each utility.

3. Calculate the average heating and cooling energy bill among households that receive assistance in the 2015 RECS.

4. Letting $x_j$ be the percentage discount in census division $j$, I calculate the annual dollar value of the percentage discount for the average (eligible) household in division $j$ according to:

$$\text{Dollar value of percentage discount in division j} = x_j \left( \frac{(\text{annual energy bill})_j}{1 - x_j} \right).$$

5. The average utility-specific assistance per recipient household equals the sum of the average base discount and the average dollar value of the percentage discount.

The calibration target is the fraction of the average energy bill covered by LIHEAP and utility-specific assistance. To determine the target, I calculate total assistance as the sum of LIHEAP bill-pay assistance and the average utility-specific assistance. I divide total assistance in each census division by the average energy bill, adjusted to reflect the fact that the reported energy bill incorporates the percentage and base discounts. Averaging across divisions implies that energy assistance covers approximately 83 percent of heating and cooling energy expenses among recipient households.

[17] Data on the number of households served by each utility is from the utility's website. When the utility does not report the number of households served on their website, I obtain the information from findenergy.com. Data on the total number of households by state is from the American Community Survey.

### C.3. Production functions for heating and cooling

I collect data on the price, capacity, and efficiency of all heaters, air conditioners, and heat pumps for sale on ecomfort.com during the fall of 2023. I focus the analysis exclusively on units that are for residential use. I drop all heaters with capacity greater than 200,000 BTUs as well as heaters coded as pump house, wash down, industrial, explosion resistant, freeze protection, or heavy duty. The air-conditioner and heat pump units on the site could all be used for residential use; I do not drop any of these units. The efficiency measure for heaters is the Annual Fuel Utilization Ratio (AFUE) and for air conditioners and heat pumps is the Seasonal Energy Efficiency Ratio (SEER). Table C.1 reports the summary statistics.

Table C.1: Summary Statistics

| Variable | Min | Max | Mean | SD | Obs |
|---|---|---|---|---|---|
| *Air conditioners* | | | | | |
| Price | 209 | 6572 | 2077.96 | 1377.46 | 137 |
| Efficiency (SEER, BTU/Wh) | 6.51 | 20 | 13.65 | 2.57 | 137 |
| Capacity (tonnage) | 0.42 | 5 | 2.61 | 1.39 | 137 |
| *Heaters* | | | | | |
| Price | 125.4 | 7069.5 | 2693.24 | 1340.75 | 942 |
| Efficiency (AFUE, percent) | 58 | 100 | 87.22 | 9.04 | 942 |
| Capacity (BTU, thousands) | 0.72 | 200 | 87.61 | 51.39 | 942 |
| *Heat pumps* | | | | | |
| Price | 1716 | 7522 | 3029.72 | 1336.29 | 46 |
| Efficiency (SEER, BTU/Wh) | 14 | 19.2 | 15.09 | 1.56 | 46 |
| Capacity (tonnage) | 1.5 | 5 | 3.08 | 1.04 | 46 |

Note: The table reports the min, max, mean, and standard deviation for the price, efficiency, and capacity of air conditioners, heaters, and heat pumps. The units of the Seasonal Energy Efficiency Ratio (SEER), the efficiency measure for air conditioners and heat pumps, are BTUs of heat removed per watt hour of energy consumed. The Annual Fuel Utilization Efficiency (AFUE), the efficiency measure for heaters, is a percent equal to the ratio of BTUs of heat produced to BTUs of energy consumed. These are the standard efficiency measures used by the industry.

### C.4. Historical temperature

Daily data on county temperature from 1950-1980 are from Wolfram Schlenker's fine-scaled weather data set[18], which is derived from the PRISM climate data set. Daily data on county temperature from 1981-2022 are from PRISM climate data set, accessed via Google Earth Engine. The resulting temperature distributions are similar if I instead use Wolfram Schlenker's data set for 1950-2019 (its available range). I choose to use the PRISM climate data directly for the more recent years because Schlenker's data set ends in 2019 and it is easy to access the PRISM climate data post 1980 through Google Earth Engine.

[18] The data are available for download at: http://www.columbia.edu/~ws2162/links.html.

Figure C.1: Map of US Counties by Region

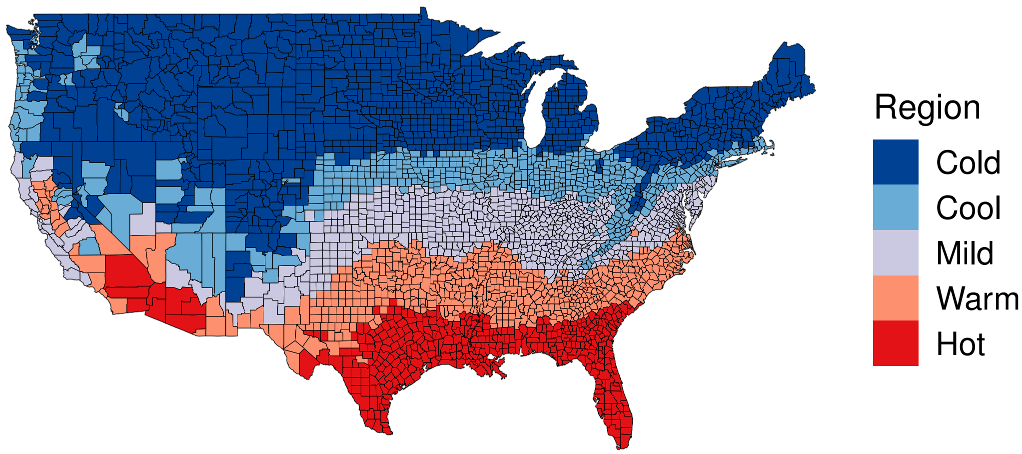


Note: The map shades US counties according to their temperature region.

I divide US counties into five regions based on the average of the historical annual average temperature in each county. I determine the cutoff temperatures for each region so that the 2020 population is approximately equal across the regions. Figure C.1 shows a map of US counties by region. To create a temperature distribution in each region and year, I bin the daily data into 100 one-degree bins, ranging from −40°C to 59°C, the same range as in the climate projections. The cold realization of the temperature distribution in each region equals the average temperature distribution over the years with annual average temperature in the bottom ten percent. Similarly, the moderate and hot realizations of the temperature distribution in each region equal the average temperature distribution over the years with annual average temperature in the middle 80 percent and in the top ten percent, respectively. Figure C.2 plots the average annual temperature distribution by region.

Figure C.2: Average Annual Temperature Distribution

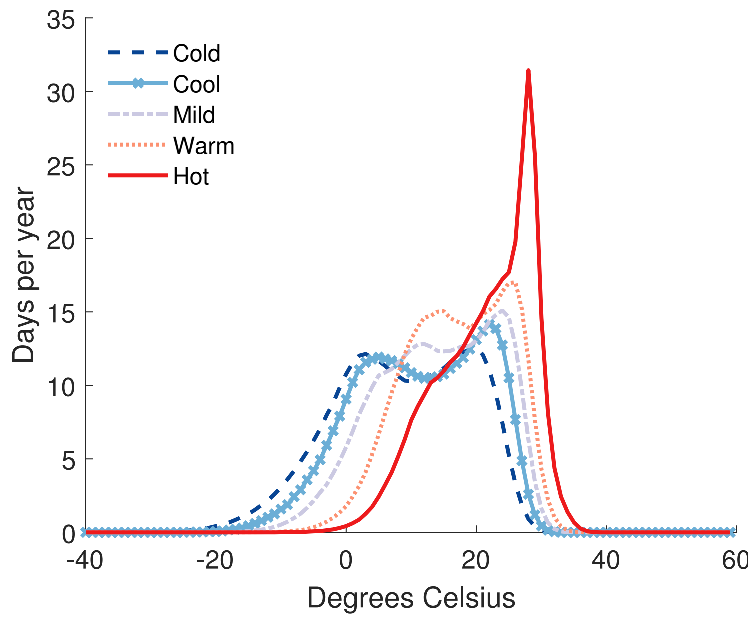


Note: The figure plots the average annual temperature distribution in each region.

## C.5. Climate change projections

As described in the main text, I take a probability-weighted average of the county-level projections of the annual distribution of average daily temperature in year 2100 under RCP 8.5 for each climate model in the CMIP archive (Rasmussen et al., 2016). The weights are designed to capture the relative probability that the climate model represents the true outcome (Rasmussen et al., 2016; Hsiang et al., 2017). I assume that there are three possible realizations of the temperature distribution, hot, cold, and moderate, with probabilities, 0.1, 0.1, and 0.8, respectively. I assume that the temperature distribution in the moderate year equals the temperature distribution from the scientific projections. To determine the cold realization of the temperature distribution, I scale the mean and the variance of the temperature distribution in the moderate year by the ratios of mean and variance of temperature in the moderate realization relative to the cold realization in the historical data. I use a similar process to determine the hot realization of the temperature distribution.

Table C.2 reports the mean and the variance of temperature for a moderate year in each region in the contemporary and climate-change equilibria. Both the mean and the variance of temperature during a year increase in response to climate change. Figure C.3 plots the contemporary (solid lines) and climate-change (dashed lines) temperature distributions in a moderate year in the cool (left panel), mild (right panel) and warm (bottom panel) regions.

Figure C.3: Effect of Climate Change on the Temperature Distribution

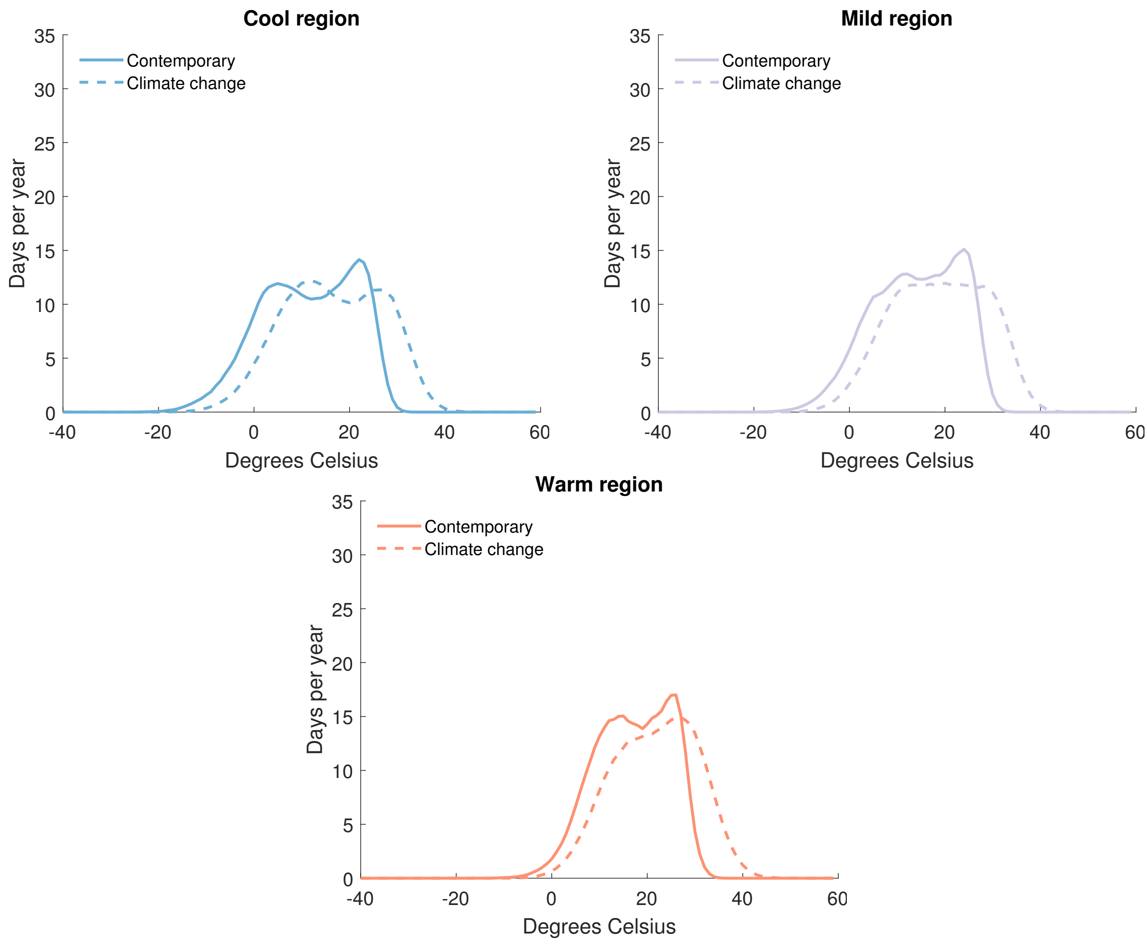


Note: The solid and dashed lines plot the contemporary and climate-change temperature distributions, respectively, in a moderate year in the cool (left panel), mild (right panel), and warm (bottom panel) regions.

Table C.2: Mean and Variance of Annual Temperature in a Moderate Year

| | Cold | Cool | Mild | Warm | Hot |
|---|---|---|---|---|---|
| *Mean* | | | | | |
| Contemporary | 8.50 | 11.20 | 13.90 | 16.98 | 21.30 |
| Climate change | 14.53 | 16.56 | 18.78 | 21.63 | 25.71 |
| *Variance* | | | | | |
| Contemporary | 109.01 | 97.54 | 82.93 | 60.86 | 52.75 |
| Climate change | 115.38 | 107.80 | 97.08 | 76.33 | 63.53 |

Note: The table reports the mean (panel 1) and variance (panel 2) of the contemporary and climate change temperature distributions in a moderate year in each region.

### C.6. Labor productivity process

I use the Rouwenhorst method to calculate the Markov transition matrix for the persistent shock:

$$\pi^{\nu} = \begin{pmatrix} 0.9413 & 0.0573 & 0.0013 & 0.0000 & 0.0000 \\ 0.0143 & 0.9420 & 0.0430 & 0.0007 & 0.0000 \\ 0.0002 & 0.0287 & 0.9422 & 0.0287 & 0.0002 \\ 0.0000 & 0.0007 & 0.0430 & 0.9420 & 0.0143 \\ 0.0000 & 0.0000 & 0.0013 & 0.0573 & 0.9413 \end{pmatrix}.$$

The corresponding invariant distribution equals: $\Pi^{\nu} = (0.0625, 0.2500, 0.3750, 0.2500, 0.0625)$.

### C.7. Parameter values

Table C.3 reports the values of the externally calibrated parameters.

Table C.3: Parameter Values: External Calibration

| Parameter | Value |
|---|---|
| *Firms* | |
| Depreciation rate: $\delta$ | 0.07 |
| Capital share in final-good production: $\alpha$ | 0.26 |
| TFP in housing production: $A^s$ | 0.12 |
| TFP in electricity production: $A^{ee}$ | 0.14 |
| TFP in heating-energy production: $A^{em}$ | 0.28 |
| *Households* | |
| International interest rate: $r$ | 0.04 |
| Coefficient of relative risk aversion: $\sigma$ | 2 |
| Bliss point temperature: $\zeta^{\star}$ | 18 |
| Heater-equipment exponent: $\theta^h$ | 0.35 |
| Air-conditioner-equipment exponent: $\theta^c$ | 0.27 |
| Heat-pump-equipment exponent: $\theta^p$ | 0.28 |
| Heater-energy exponent: $\eta^h$ | 0.77 |
| Air-conditioner-energy exponent: $\eta^c$ | 0.85 |
| Heat-pump-energy exponent: $\eta^p$ | 0.86 |
| *Labor productivity* | |
| Persistent shock persistence: $\rho$ | 0.97 |
| Persistent shock innovation variance: $\sigma_{\varepsilon}^2$ | 0.02 |
| Fixed-effect variance: $\sigma_{\xi}^2$ | 0.66 |

Note: The table reports the values of the parameters I take directly from the data and the existing literature.

### C.8. External validation

Data on the share of aggregate income received by each quintile of the income distribution are from Table H-2 of the US Census Bureau Current Population Survey (CPS). A household's income in the model equals its labor income plus its asset income. Data on the share of aggregate wealth held by each quintile of the wealth distribution are from the Survey of Consumer Finances, as reported in Kuhn and Rios-Rull (2016). Tables C.4 and C.5 report the model and empirical values of the shares of income and wealth for each quintile of the income and wealth distributions, respectively. The model matches the overall income and wealth distributions reasonably well, but, as discussed in the main text, it does not capture the high concentrations of wealth at the top of the distribution.

Table C.4: Share of Income Recieved by Each Qunitile of the Income Distribution

| | Quintile | | | | |
|---|---|---|---|---|---|
| | First | Second | Third | Fourth | Fifth |
| Model | 3.5 | 6.8 | 12.9 | 24.5 | 52.5 |
| Data | 3.1 | 8.3 | 14.1 | 22.6 | 52.0 |

Note: This table reports the share of income received by each quintile of the income distribution in the model (row 1) and in the CPS (row 2).

Table C.5: Share of Wealth Held by Each Qunitile of the Wealth Distribution

| | Quintile | | | | |
|---|---|---|---|---|---|
| | First | Second | Third | Fourth | Fifth |
| Model | 0.0 | 1.2 | 5.6 | 17.5 | 75.8 |
| Data | -0.7 | 0.6 | 3.2 | 9.8 | 87.0 |

Note: This table reports the share wealth held by each quintile of the wealth distribution in the model (row 1) and in the SCF (row 2).

## D. Additional Results

Figure D.1 plots the fraction of households that have only heat or have only air conditioning the contemporary and climate-change equilibria. Figure D.2 plots the distributional effects of climate change when households cannot use heating and cooling to adapt to temperature.

Figure D.1: Changes in Heat and AC Specialization

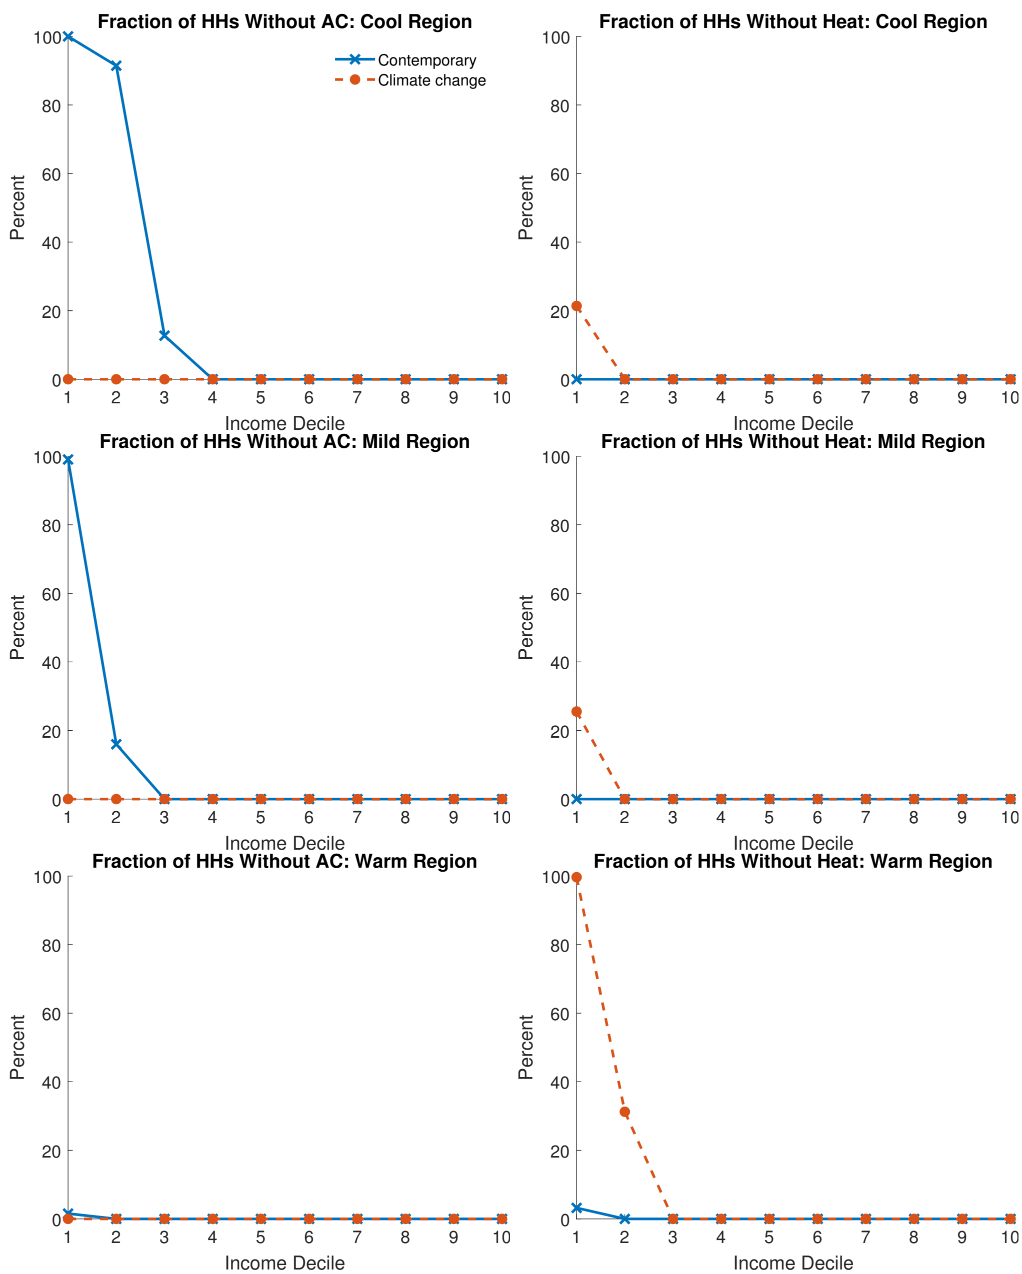


Note: The figure plots the fraction of households without air conditioning (left panels) and the fraction of households without heat (right panels) in the cool (row 1), mild (row 2), and warm (row 3) regions in for the contemporary (blue line) and climate-change (orange line) equilibria.

Figure D.2: Welfare Impacts of Climate Change When Households Cannot Adapt to Temperature (CHEV)

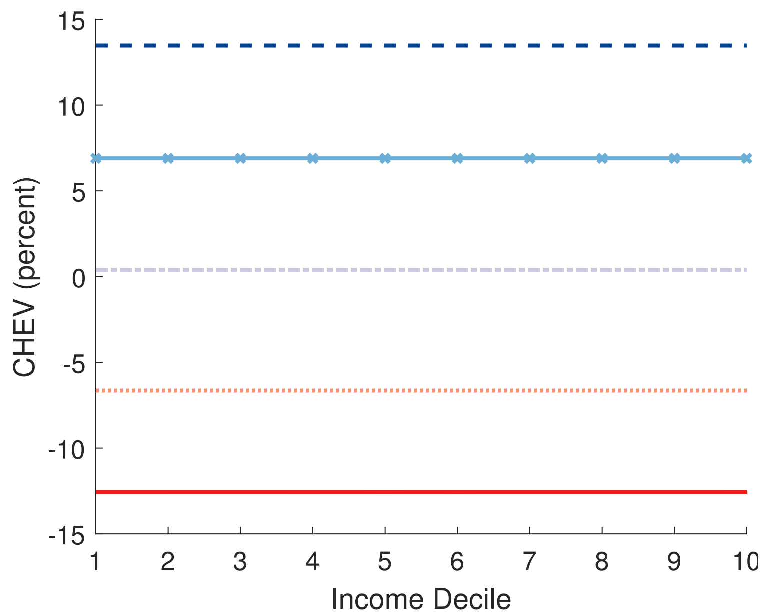


Note: Each line plots the welfare consequences of climate change without adaptation conditional on the household's region and income decile, where decile one corresponds to the lowest income households. I measure welfare using the consumption-housing equivalent variation (CHEV) between the contemporary and climate-change equilibria. Negative values indicate that climate change makes the household worse off.