Riding the Waves:
Inequality and Adaptation to Extreme Temperatures in a Changing Climate
(Preliminary and incomplete)

Stephie Fried
Federal Reserve Bank of San Francisco

September 29, 2023

Abstract

Will climate change worsen U.S. inequality? Focusing on the direct effects of changes in temperature in the U.S., this paper develops an Aiyagari-style heterogeneous agent model to study the distributional impacts of climate change across income groups. Households can adapt to temperature by using capital and energy for heating and cooling. The model replicates empirical relationships between energy budget shares, energy expenditures, and income. A key insight from the model is that the outdoor temperature acts as a transfer from nature to households. Extreme temperatures correspond to reductions in transfers from nature and thus have higher welfare cost for lower income households. Consequently, climate change is generally regressive in hot regions of the U.S., where it leads to more extreme temperatures and progressive in cold regions, where it leads to fewer extreme temperatures. Households in the lower income deciles break this pattern because climate change affects whether these households purchase both heating and cooling capital or can specialize in a single type of energy capital.

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, Brigitte Roth Tran, Kevin Novan, William Peterman, 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, and LSE environment week for their many helpful comments. For excellent research assistance, I thank River Bell, Ethan Goode, and Zoe Arnaut-Hull. All potential errors are my own.

1. Introduction

Will climate change worsen U.S. 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 focus on spatial heterogeneity (e.g., Carleton et al., 2022; Nath, 2022; Cruz and Rossi-Hansberg, forthcoming). Instead, this paper focuses on income heterogeneity, and develops an Aiyagari (1994)-style model to analyze the distributional effects of climate change across income groups. I study one important type of climate damage: the rightward shift of the annual temperature distribution. Among the middle and high income deciles, I find that climate change creates progressive welfare benefits in the colder regions and regressive welfare costs in the hotter regions. The lowest income deciles break these patters; a portion of low income households in the colder regions experience substantial welfare costs while a portion in the hotter regions experience substantial welfare benefits.

I first develop a simple model to analytically explore the distributional consequences of changes in temperature. Following much of the literature, households derive utility from consumption and housing services. What I add to the standard framework is that the temperature of the house affects utility from housing services. The intuition is that, 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. Mathematically, the indoor temperature of the house affects utility from housing services with a bliss point at the "comfortable temperature." The indoor temperature depends on two factors: (1) the outdoor temperature determined by nature and (2) any heating or cooling energy 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.

A key insight from the analytic model is that the outdoor temperature acts as a transfer from nature to the household. 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. Extreme temperatures correspond to reduced transfers from nature. Continuing the same example, if the outdoor temperature is 40°C instead of 30°C, then nature only gave households 10°C of cooling for free. The decrease in transfers has higher welfare costs for lower income households because the value of the loss is larger in proportion to their income. The intuition for this result mirrors the intuition for the progressivity of lump-sum transfers from the public finance literature. Interpreting extreme temperatures as reductions in transfers implies that climate change will be regressive in regions where it leads more extreme temperatures and progressive in regions where it leads to fewer.

The insights from the simple model have implications for designing energy assistance policies for low income households. Viewing extreme temperatures as reduced transfers suggests that the most direct way to compensate low-income households for extreme temperatures is to provide direct payments that implicitly replace the lost transfers. Interestingly, this is exactly how a portion of energy assistance policy in the United States, provided by the Low-Income Home Energy Assistance Program (LIHEAP), is designed. While the exact formulas vary by state, LIHEAP payments typically depend on the household's heating or cooling expenses from the previous year, or on the average number of heating and cooling degree days. Both approaches are practical ways to mimic indexing payments to outdoor temperature, and thus directly compensate households for reduced transfers from nature on extreme temperature days.

To quantify the distributional effects of rising temperatures from climate change, I develop a dynamic, heterogeneous-agent model. The quantitative model incorporates the intuition from the simple model but adds several features that are important for a numerical analysis. First, households draw idiosyncratic labor-productivity shocks each period, which generate endogenous distributions of income and wealth. Second, larger houses require more energy to heat or cool. Third, to produce heating and cooling, the household much purchase heating capital (e.g., a furnace) and cooling capital (e.g., an air conditioner) in addition to heating and cooling energy. Fourth, I include federal energy assistance policy; low-income households receive energy-assistance payments that are broadly in line federal, state and local energy assistance policies in the US. Finally, households experience an uncertain distribution of outdoor temperatures each period. I split the period (a year) into 100 sub-periods, where each sub-period corresponds to a different outdoor temperature, ranging from -40°C to 59°C in one-degree increments.

To calibrate the model, I divide U.S. counties into five regions based on the historical annual average temperature in each county. There are three possible realizations of the temperature distribution within each region, cold, moderate, and hot. I calibrate the key parameters of the model to match average heating and cooling budget shares, and the variation in energy budget shares and energy expenditures with income, where energy is the sum of heating and cooling energy. 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 incorporating the impact of the temperature distribution on utility allows the model to match these patterns without subsistence energy.

I first use the model to study the distributional effects of weather in the current economy. In the model, weather corresponds to the particular realization of the temperature distribution (e.g., a hot year versus a cold year). I find that in the hotter regions, cold years make households better off and hot years make households worse off. The opposite patterns hold in the colder regions. In all regions, lower-income households' welfare is more sensitive to weather than higher income households' welfare. Relative to income, lower income households experience larger changes in transfers from nature from a hot or cold year compared to a moderate year, leading to larger welfare effects. Additionally, a portion of these households are liquidity constrained. Households adjust their savings to smooth the effects of the weather shock between current and future utility. This smoothing is not available to the lowest income, liquidity constrained households, magnifying the welfare consequences of weather shocks for this group.

I next use the model to study the distributional consequences of climate change, modeled as a rightward shift in the annual temperature distribution. The key difference between weather and climate change is that climate change alters households' expectations about the temperature distribution, thereby influencing their longer-run housing and heating and cooling capital decisions. To measure the rightward shift, I use county-level scientific projections of the annual temperature distribution in 2100, consistent with a no-policy emissions scenario, Representative Concentration Pathway (RCP) 8.5. I solve for a climate-change stationary equilibrium of the model with the 2100 temperature distributions in each region. I compare outcomes in the climate-change equilibrium with outcomes in a baseline equilibrium calculated under the historical temperature distributions.

Among the middle and high income deciles, climate change leads to welfare benefits in the colder regions, which fall with income, and welfare costs in the hotter regions, which also fall with income. Thus, across these income groups, climate change is progressive in the colder regions and regressive in the hotter regions. The welfare consequences for the lowest income deciles break these general patterns. Some of the lowest income households in the colder regions experience welfare costs from climate change, despite the welfare benefits experienced by the middle and high income households. Similarly, some of the lowest income households in the hottest regions experience welfare benefits from climate change.

This variation in the welfare impacts stems from the differential effects of climate change on transfers from nature and on households' ability to specialize in heating or cooling capital. In expectation, the rightward shifts in the temperature distributions in colder regions increase transfers from nature because they reduce the number of extreme temperature days, leading to welfare benefits. The opposite effect applies in the hotter regions. The changes in transfers from nature relative to income in all regions are largest for the lower income households, implying that, all else constant, the welfare consequences of climate change fall monotonically with income.

The effect of climate change on households' ability to specialize in heating or cooling capital breaks the monotonic relationship between the welfare impacts and income among the lowest income deciles. Since heating and cooling require separate types of capital, households could choose to purchase relatively more heating capital if they experience more cold days, or relatively more cooling capital if they experience more hot days. In the baseline, the temperature distributions in the colder regions are more favorable to specialization in heating capital because households only experience a small number of hot days. Climate change leads to more hot days, which reduces the specialization favorability in the colder regions. The opposite effect occurs in the hotter regions; climate change increases the favorability to specialization in cooling capital.

Across all income deciles, reductions in the specialization favorability in the colder regions decrease the welfare benefits of climate change, while increases in the specialization favorability in the hotter regions decrease the welfare costs of climate change. These effects are particularly acute for the lowest income households that move to, or away from, complete specialization. For example, a portion of the lowest income households in the colder region have a heater in the baseline, but do not have an air conditioner (e.g., they are completely specialized in heating). However, in the climate change equilibrium, these households choose to purchase both heater and an air conditioner. The welfare cost of this move away from complete specialization dominates the welfare benefit of higher transfers from nature, resulting in welfare costs of climate change for the lowest income households in the cold region. The opposite effect leads to the welfare benefits from climate change for the lower income households in the hot region.

Lastly I use the model to explore the distributional and welfare consequences of adaptation to temperature though heating and cooling, in both the current economy and in response to climate change. I solve for a no-climate-change and a climate-change equilibrium in which households cannot adapt to temperature. In these equilibria, households do not have access to heating and cooling capital or energy, implying that the indoor temperature always equals the outdoor temperature. As one might expect, the ability to adapt to temperature raises welfare in the current economy. More interestingly, the welfare benefits from this adaptation increase with income. While all households have the option to adapt, higher income households are able to purchase more heating and cooling capital and energy, resulting in indoor temperatures closer to the bliss point and higher welfare benefits from adaptation.

The distributional consequences of climate change depend critically on households' ability to adapt. Comparing the no-adaptation equilibria with and without climate change reveals that when households cannot adapt, climate change has identical welfare impacts across income deciles within a region. The rightward shift in the temperature distribution leads to the same change in indoor temperature for all households (within a region) and thus the same welfare costs. Hence, the differential welfare effects of climate change across income groups are entirely driven by households' differing abilities to adapt in response to the changes in temperature caused by climate change.

This paper contributes to the broader literature on the unequal consequences of climate change. Much of this literature is focused on country or region-level inequality and studies how the consequences of climate change differ around the world 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, forthcoming). A smaller, more theoretical, literature adds households with different 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).

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; Park and Stainier, 2021; Doremus et al., 2022) vary across income groups in the U.S. Consistent with the results in the present analysis, the studies on temperature tend to find that the changes in temperature caused by climate change have smaller consequences for higher income households, either because they are less exposed to the temperature changes or because adaptation is relatively less costly, or both.

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 revenue from the climate policy 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.

The paper proceeds as follows: Section 2 develops a simple model to analytically characterize the distributional effects of extreme temperatures. Section 3 builds a dynamic heterogeneous agent model to quantify the distributional effects of weather and climate change. Section 4 calibrates the model and section 5 presents the quantitative results.

2. Simple Model

I develop a simple model that incorporates the key relationships between temperature, utility, and adaptation through heating and cooling. The model is intentionally stylized to allow for an analytic characterization of the impact of extreme temperatures across the income distribution. The model is static, deterministic, and abstracts from the capital costs necessary for heating and cooling. I relax these assumptions in the richer, quantitative model to follow.

2.1. Environment

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

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

Variables c_i and h_i denote household i's choices of consumption and housing, 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 the indoor temperature is 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 based on the relationship between the indoor temperature of the house and the bliss point temperature, $\zeta^\star$, according to,

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

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, ζ, 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. Heating and cooling allow the household to adapt to the outdoor temperature by changing the indoor temperature of its house.

The household's optimization problem is to choose consumption, housing, heating, and cooling 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, respectively. Variable y_i is the household's income. The consumption good is the numeraire.

2.2. 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 the Appendix. Focusing first on an interior solution, in which the household's optimal indoor temperature is less than the bliss point but greater than the outdoor temperature, $\zeta < T_i^\star \leq \zeta^\star$, the optimal levels of consumption, housing, and indoor temperature are

$$c_i^\star = \frac{y_i + p^{eh}\zeta}{3}, \quad h_i^\star = \frac{y_i + p^{eh}\zeta}{3p^h}, \quad \text{and} \quad T_i^\star = \frac{y_i + p^{eh}\zeta}{3p^{eh}}. \tag{3}$$

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

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

The optimized levels of consumption, housing, and indoor temperature in equation (3) reveal that the outdoor temperature, ζ, acts as a transfer from nature to the household. Nature gives the household ζ 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, 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 warmer day creates larger welfare gains for lower income households, while a decrease in the transfer from nature from a colder day creates larger welfare losses for lower income households. This intuition for the distributional effects of changes in transfers from nature parallels the intuition for the progressivity of lump-sum transfers from the public finance literature.

To see the distributional implications of changes in outdoor temperature more formally, 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 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:

$$\frac{\partial u(c_i^\star, h_i^\star, T_i^\star)/\partial \zeta}{u(c_i^\star, h_i^\star, T_i^\star)} = \frac{G'(\zeta)}{G(\zeta)} + \left(\frac{3p^{eh}}{y + p^{eh}\zeta}\right)\left(\frac{1}{\ln(c_i^\star) + \ln(D(T_i^\star)h_i^\star)}\right). \tag{5}$$

Thus, in cold weather ($\zeta < \zeta^\star$), warmer (more moderate) days are progressive with larger welfare benefits accruing to lower income households. Likewise, colder (more extreme) days are regressive, with larger welfare costs accruing to lower income households. The parallel analysis for hot weather, ($\zeta > \zeta^\star$, see the Appendix) implies that cooler (more moderate) days are progressive while hotter (more extreme) days are regressive. 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 finding that extreme temperatures act like decreased transfers suggests that energy assistance for low income households should take the form of direct payments that implicitly replace the lost transfers from nature. Interestingly, this is exactly how the energy-bill assistance payments provided by the Low Income Home Energy Assistance Program (LIHEAP) are designed. While the exact formulas vary by state, LIHEAP payments often depend on the household's 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 directly compensate households for the reduced transfers from nature on extreme temperature days.

Importantly, for the functional form of the utility specified in equation (2), the effect of income on the welfare consequences of a change in temperature does not depend on the amenity value of temperature, $G(\zeta)$. More generally, the amenity value of temperature will not have a first order effect on distributional consequences of changes in temperature as long as the preferences underlying the amenity value are uncorrelated with income. Since the paper is focused on the distributional consequences of changes in temperature, and the simple model suggests that the amenity value is not critical to these distributional implications, I abstract from the amenity value of temperature in the quantitative model. The advantage of this abstraction is that it avoids parameterizing function $G(\zeta)$. The cost is that the welfare impacts of climate change in the quantitative model do not incorporate changes in the amenity value of temperature. While this abstraction matters for the interpretation of the level of the welfare impacts from a change in temperature, the analysis from the simple model suggests that it is not important for understanding the distributional implications.

The analysis thus far focused on an interior solution in which and $\zeta < T_i^\star < \zeta^\star$. I next consider what happens for each corner solution. 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

$$c_i^\star = \frac{y_i + p^{eh}\zeta - p^{eh}\zeta^\star}{2} \quad \text{and} \quad h_i^\star = \frac{y_i + p^{eh}\zeta - p^{eh}\zeta^\star}{2p^h}. \tag{6}$$

Again, we see that the outdoor temperature, ζ, 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.

If instead the household's optimal indoor temperature equals the outdoor temperature, $T_i^\star = \zeta$, then the household does not purchase any heating or cooling energy $e_i^h = e_i^c = 0$ and the model collapses to a standard two-good utility optimization problem.[1] The optimized levels of consumption and housing, are,

$$c_i^\star = \frac{y_i}{2} \quad \text{and} \quad h_i^\star = \frac{y_i}{2p^h}. \tag{7}$$

Optimal heating and cooling energy are both zero. In this case, the outdoor temperature has no effect on the household's decisions.

Lastly, I explore how energy budget share varies with income. Energy budget share for household i, EBS_i, for a cold day, $\zeta < \zeta^\star$, equals:

$$EBS_i = \begin{cases} 0 & : 0 < y_i \leq \frac{3p^h}{p^{eh}} \\ \frac{1}{3} - \frac{2}{3}\left(\frac{p^{eh}\zeta}{y_i}\right) & : \frac{3p^h}{p^{eh}} < y_i \leq \frac{3p^h\zeta^\star}{p^{eh}\zeta} \\ \frac{p^{eh}(\zeta^\star - \zeta)}{y_i} & : \frac{3p^h\zeta^\star}{p^{eh}\zeta} < y_i. \end{cases} \tag{8}$$

The lowest income households (first row of equation (8)) do not purchase any heating en-

---

[1] The analysis of this corner only applies for outdoor temperatures greater than the extreme lower bound necessary for survival, $\zeta > \underline{\zeta}$.

ergy. The energy budget share is constant and equal to zero for all households in this group. The middle income households (second row of equation (8)) purchase some heating energy, but choose indoor temperatures less than the bliss point. Energy budget share increases with income among these households. To understand the intuition, it is useful to decompose energy budget share into the energy cost necessary to attain $T_i^\star$ if the outdoor temperature is zero, minus the value of the transfer from nature, $p^{eh}\zeta$, all divided by income. As income rises, households choose indoor temperatures closer to the bliss point, reducing the relative importance of the transfer from nature in their energy expenditures. Consequently, energy expenditures increase faster than income, causing energy budget share to rise with income. In contrast, energy budget share falls with income among the highest income households (third row of equation (8)). These households all choose outdoor temperature equal to the bliss point. Since the energy required to attain $\zeta^\star$ does not change with income, energy expenses as a share of income fall as income rises.

In the quantitative model, the intuition for the effect on of income on energy budget share is similar, but the results are not as stark. Energy budget share falls with income in regions of the parameter space in which households choose indoor temperature close to (but not necessarily equal to) the bliss point and rises with income in regions of the parameter space in which households choose indoor temperature far from the bliss point. The calibration suggest that the empirically relevant region of the parameter space is one in which households choose indoor temperature relatively close to the bliss point and energy budget share falls monotonically with income.

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 $\underline{\zeta}$, death from exposure to temperature extremes 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 exposure to extreme temperatures.

This abstraction comes with the caveat that some people die each year from exposure to extreme temperatures in the U.S. For example, from 1999-2020, exposure to extreme temperatures caused approximately 0.04 percent of U.S. deaths, equal to 1,145 people per year, on average.[2] Over half of these deaths are estimated to be among the un-housed popula-

---

[2] 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.

tion (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.[3] Even so, a small number of the temperature-related deaths in the U.S. are likely caused by the temperature of the home. It is important to acknowledge that the quantitative model does not include this channel.

3. Quantitative model

I develop a dynamic, heterogeneous-agent model that incorporates the intuition from the simple model and adds several features that are important for a quantitative analysis of the distributional effects of the changes in temperature caused by climate change. First, households experience a range of outdoor temperatures each period. Incorporating the impact of the distribution of outdoor temperatures on utility is important because extreme temperatures are critical for understanding household's heating and cooling decisions and the resulting welfare impacts of climate change. Second, I develop a richer model of adaptation to temperature. Households produce heating and cooling from heating and cooling capital as well as from energy. Moreover, the energy and capital required to heat or cool a house by one degree depends on the size of the house, allowing the model to match the empirical regularity that energy expenses increase with income. Third, low-income households receive energy assistance payments that are indexed to expected heating and cooling degree days, broadly in line with the structure of LIHEAP payments in most states.

Finally, I incorporate two types of uncertainty. First, households face uncertainty over the temperature distribution. When households make their decisions for housing and heating and cooling capital, they must do so based on their expectations that it will be a hot year or a cold year, etc. Second, households face uncertainty over labor income. They draw idiosyncratic labor-productivity shocks each period, which generate endogenous distributions of income, earnings, and wealth.

3.1. Environment

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, and heating and cooling capital. Households derive utility from consumption, housing services, and the temperatures of their house. Households can affect the

---

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

temperature of their house by producing degrees of heating or cooling from heating or cooling capital 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). There is no migration between regions.

3.2. Temperature distribution

Households can experience J outdoor temperatures during the period, ranging from ζ_1 to ζ_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 ζ_j. Households know the range of possible outdoor temperatures (e.g., $[\zeta_1, \zeta_J]$), but, at the start of the period, they do not know the distribution, Q_{in}. Instead, after the household has made its longer term decisions (e.g., housing, heating and cooling capital), it draws a temperature distribution, Q_{in}, from a distribution of temperature distributions, $\pi_n^Q(Q_{in})$. For example, in the calibrated model, the households know that the outdoor temperature during a year can range from −40°C to 59°C. However, at the start of the year, when the household chooses what type of heater and air conditioner to purchase, it does not know whether it will be a hot year, a cold year, or a moderate year. The realizations of the temperature distribution are i.i.d. across households within a region.

3.3. Preferences

Households derive utility from consumption and housing according to the function,

$$u_{in} = \sum_{j=1}^{J} q_{ijn} \left[\frac{{c_{ijn}}^{1-\sigma}}{1-\sigma} + \psi \frac{(D(T_{ijn})h_{in})^{1-\sigma}}{1-\sigma} \right]. \tag{9}$$

The utility of household i in region n in a given period 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 c_{ijn} denotes the household's consumption in sub-period j, which could differ from the household's consumption in sub-period $j' \neq j$. Consumption is a relatively flexible input and households can choose to consume different amounts on hot and cold days. In contrast, housing services are much less flexible and cannot vary by sub-period based on the outdoor temperature. Instead, households choose housing at the start of the period and it is fixed for the entire period; there is no subscript j on h_{in} in equation (9).

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, $\zeta^\star$,

$$D(T_{ijn}) = \frac{1}{1+\chi(T_{ijn}-\zeta^\star)^2}. \tag{10}$$

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, $D(T_{ijn}<\zeta^\star)<1$ and $D(T_{ijn}>\zeta^\star)<1$, in sub-periods in which their house is hotter or colder than the bliss point. Parameter χ 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 temperature equal to the bliss point in order to receive any utility from their housing services.

The indoor temperature of household i in sub-period j in region n depends on the outdoor temperature in that sub-period, ζ_j, and on the heating and cooling the household produces,

$$T_{ijn} = \zeta_j + \frac{1}{h^\gamma}\left[\underbrace{A^h(x^h_{in})^{\theta^h}(e^h_{ijn})^{\eta^h}}_{\text{heating}} - \underbrace{A^c(x^c_{in})^{\theta^c}(e^c_{ijn})^{\eta^c}}_{\text{cooling}}\right]. \tag{11}$$

To produce heating, the household combines heating capital (e.g., a furnace), x^h_{in}, with heating energy, e^h_{ijn}. If the household does not have any heating capital, the marginal product of heating energy is zero. The household cannot produce heat without some type of heater. Similarly, to produce cooling, the household combines cooling capital (e.g., an air conditioner), x^c_{in}, with cooling energy, e^c_{ijn}. Parameters A^h and A^c determine overall productivity in the production of heating and cooling, respectively. Operating heating and cooling capital requires that the household purchase energy and that the household pay a fixed cost for each capital type, Ω^h and Ω^c. This fixed cost captures the costs associated with installing and maintaining heaters and air conditioners.

Importantly, the household's choice of heating and cooling capital is fixed for the entire period– there is no subscript j on x^h_{in} or x^c_{in} in equation (11). Households cannot purchase a furnace each winter when it is cold and return it each summer when it is warm. This inflexibility implies that households optimally choose to idle their heating capital and on hot days, so that they have it available for use on cold days, and likewise for cooling capital. While energy capital is fixed for the entire period, energy levels are not. Households can adjust their heating and cooling energy on hot and cold days to achieve their desired indoor temperature; there is a subscript j on e^h_{ijn} and e^c_{ijn} in equation (11). For example, on a very cold day, the household could use more heating energy then on a moderately cold day.

The marginal product of heating or cooling on indoor temperature is decreasing in the level of housing services, h. The intuition is that larger levels of h correspond to more expensive, and thus likely bigger, houses which require more energy 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 thus more energy efficient per square foot.

3.4. 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 ν_{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 $\xi \sim N(0, \sigma^2_{\xi})$ is a household-specific fixed effect (i.e., ability) that is constant over time.

3.5. Energy assistance

The federal government provides aid to help low-income households adapt to extreme temperatures through heating and cooling. The majority of this aid is provided through the Low Income Home Energy Assistance Program (LIHEAP) and is used to assist low-income households with their energy bills. While the exact formulas vary by state, most LIHEAP payments are direct transfers based on the household's expected heating and/or cooling expenses.[4] In addition to LIHEAP, low-income households can also receive energy assistance through the Weatherization Assistance Program (WAP) funded by the Department of Energy. WAP assistance provides low-income households with energy capital-related capital, such as a more efficient furnace. At the state and local level, approximately one quarter of utilities provide energy assistance to low income households which supplements the federally provided assistance. 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

---

[4] The LIHEAP benefits in each state are available for download from: https://liheapch.acf.hhs.gov/delivery/benefits.htm.

climates.[5]

3.6. Recursive formulation of the households' problem

I describe the optimization problem for a household i in region n. The state variables are the household's assets, a_{in} and the persistent value of its labor productivity shock, ν_{in}. Following the realization of the labor productivity shock, the household chooses housing services, and heating and cooling capital, to maximize expected lifetime welfare, where the expectation is taken with respect to the realization of the temperature distribution in the current period, and with respect to all future realizations of the temperature distribution and labor productivity shocks. After the temperature distribution realizes, the household chooses consumption and heating and cooling energy in each sub-period, and savings for the next period to maximize expected lifetime welfare, where now the expectation is taken with respect to future realizations of the temperature distribution and labor productivity shocks.

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

$$V(a_{in};\nu_{in}) = \max_{h_{in},x^h_{in},x^c_{in}} \sum_{Q_{in}} \pi^Q(Q_{in}) \left\{ max_{\{e^h_{ijn},e^c_{ijn},c_{ijn}\}^J_{j=1}} \left[\sum_{j=1}^{J} q_{ijn} \left(\frac{c_{ijn}^{1-\sigma}}{1-\sigma} + \psi \frac{(D(T_{ijn})h_{in})^{1-\sigma}}{1-\sigma} \right) \right.\right. \tag{12}$$

$$\left.\left. + \beta \sum_{\nu'_{in}} \pi^{\nu}(\nu_{in}|\nu'_{in}) \sum_{Q'_{in}} \pi^Q(Q'_{in}) V(a'_{in};\nu'_{in}) \right] \right\} \tag{13}$$

subject to the budget constraint:

$$(1-\tau)wz_{in} + (1+r)a_{in} + B_{in} =$$
$$\sum_{j=1}^{J} q_{ijn}(c_{ijn} + p^{eh}e^h_{ijn} + p^{ec}e^c_{ijn}) + p^h h_{in} + p^{xh}x^h_{in} + p^{xc}x^c_{in} + \Omega^h \mathbf{1}_{x^h>0} + \Omega^c \mathbf{1}_{x^c>0} + a'_{in},$$

and the non-negativity constraints:

$$a'_{in} \geq 0,\ h_{in} \geq 0,\ e^c_{ijn} \geq 0,\ e^h_{ijn} \geq 0,\ x^c_{in} \geq 0,\ x^h_{in} \geq 0,\ c_{ijn} \geq 0.$$

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

---

[5] I do not model transfers of energy capital separately from the direct transfers. In the calibration, the WAP transfers are small and always less than the household's optimal level of energy capital, making them equivalent to direct transfers in the model.

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, heating, and cooling in each sub-period, housing services, heating and cooling capital and fixed costs, and any assets the household chooses to carry into the next period. Variables p^{eh}, p^{ec}, p^{xh}, p^{xc}, and p^h denote the relative prices of heating and cooling energy, heating and cooling capital, and housing services, respectively. The final good is the numeraire.

3.7. Firms

Each region contains unit masses of perfectly competitive firms that produce the final good, housing services, heating and cooling energy, and heating and cooling capital. The final good, y, is produced from a Cobb-Douglas production function in capital and labor, $y = A^y (k^y)^\alpha l^{1-\alpha}$. In equilibrium, parameter α equals capital's share of production. Variable A^y denotes total factor productivity (TFP) in the production of the final good. Housing, energy, and energy capital are all produced using a linear production technologies in capital:

$$h = A^h k^h, \quad e^h = A^{eh} k^{eh}, \quad e^c = A^{ec} k^{ec}, \quad x^h = A^{xh} k^{xh}, \quad \text{and} \quad x^c = A^{ec} k^{ec}.$$

Variables A^h, A^{eh}, A^{ec}, A^{xh} and A^{xc} denote total factor productivity in the production of housing, heating and cooling energy, and heating and cooling capital, respectively.

The solution to the firms' profit-maximization problems yields the equilibrium wage and relative prices:

$$\begin{aligned} & w = (1-\alpha)\left(\frac{\alpha}{R}\right)^{\frac{\alpha}{1-\alpha}}, \\ & p^h = R/A^h, \quad p^{eh} = R/A^{eh}, \quad p^{ec} = R/A^{ec}, \quad p^{xh} = R/A^{xh}, \quad p^{xc} = R/A^{xc}. \end{aligned} \tag{14}$$

Variable $R \equiv r + \delta$ denotes the rental rate on capital, where δ is the depreciation rate. 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, $r^\star$.

3.8. Stationary equilibrium

I define a stationary recursive competitive equilibrium. Throughout the definition of the equilibrium, I suppress the individual household subscripts. Let $g = (a, \nu, n)$ denote the vector of household states and characteristics at the start of the period. Let μ 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, τ, and an international interest rate, $r^\star$, a *stationary recursive competitive equilibrium* consists of time-invariant value and policy functions for households $\{V(g), h(g), x^h(g), x^c(g)\}$, $\{e_j^h(g,Q), e_j^c(g,Q), c_j(g,Q)\}_{j=0}^J$, production plans for firms, $\{k_n^y, l_n, k_n^h, k_n^{eh}, k_n^{ec}, k_n^{xh}, k_n^{xc}\}_{n=1}^N$, prices, $\{w, R, p^h, p^{eh}, p^{ec}, p^{xh}, p^{xc}\}$ and stationary distribution, μ, such that:

1. Given prices and policies, the household value function, V, solves the optimization problem in equation (13) and $h(g), x^h(g), x^c(g)$, $\{e_j^h(g,Q), e_j^c(g,Q), c_j(g,Q)\}_{j=0}^J$, are the associated policy functions.

2. Prices satisfy equation (14).

3. The markets for housing services and heating and cooling capital clear in each region:

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

4. The markets for heating and cooling energy clear in each region in each period:

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

5. The federal budget constraint clears:

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

6. The labor market clears in each region:

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

4. Calibration

The time period is one year. I divide U.S. 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. I take all historical averages over the 1997-2020 time period.

4.1. Temperature distribution

I divide the U.S. economy into five regions, cold, cool, mild, warm, and hot, based on the county's average annual temperature from 1950-2019. County-level data on average daily

Figure 1: Map of U.S. Counties by Region

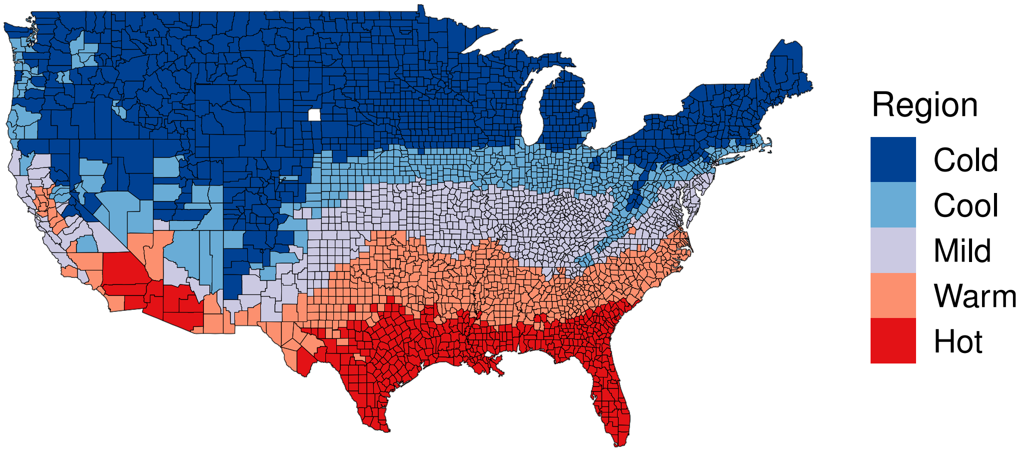

temperature are from PRISM climate group, and compiled by Wolfram Schlenker.[6] Figure 1 shows a map of the U.S. counties by region. As one would expect, the colder counties are located in the more northern and more mountainous parts of the U.S.

Figure 2 plots the average annual temperature distribution in each region. Temperature ranges from −40°C to 59°C in one degree increments. 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. To determine the temperature distribution in the hot realization, I calculate the average of the temperature distribution in each region among the hottest 10 percent of years in the historical data. The procedure to determine the temperature distributions in the moderate and cold years is analogous. 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.

---

[6] The data are available for download from Wolfram Schlenker's website: http://www.columbia.edu/ ws2162/links.html

Figure 2: Temperature Distribution in an Average Year

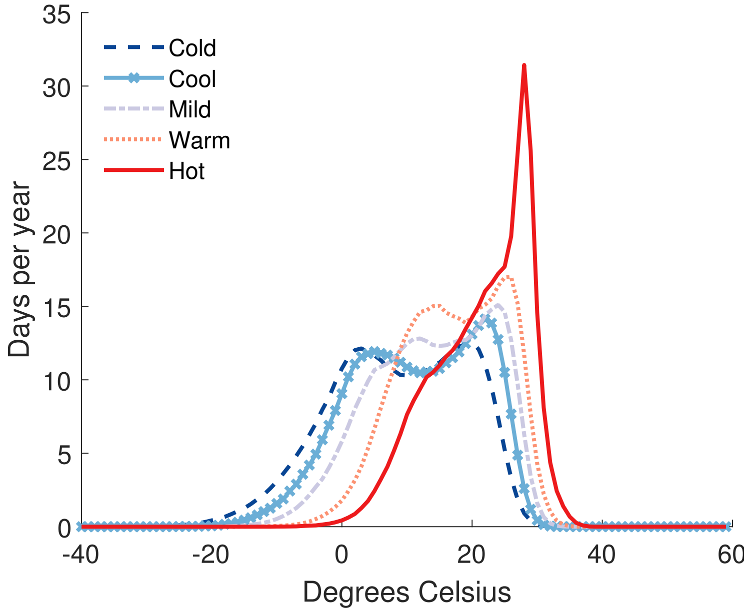

Table 1: Average Annual Temperature Relative to a Moderate Year

	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.

4.2. Externally determined parameters

4.2.1. Production functions for heating and cooling

I estimate heating and cooling production parameters θ^h, η^h, θ^c and η^c from product-level data on heaters and air conditioners. The estimating equations are

$$\ln(y_i^h) = \kappa^h + \theta^h \ln(x_i^h) + \eta^h \ln(e_i^h) + \varepsilon_i^h \quad \text{and} \quad \ln(y_i^c) = \kappa^c + \theta^c \ln(x_i^c) + \eta^c \ln(e_i^c) + \varepsilon_i^c,$$

where subscript i denotes the observation of a particular heater or air conditioner. Variable y is the heating capacity (measured in BTUs) of the heater or the cooling capacity (measured in tonnage) of air conditioner. Heating and cooling capacity refer to the amount of heating or cooling that could be produced from running the unit at its maximum level for one hour. Variable x denotes the price of the heater or air conditioner, and hence the value of the heating or

Table 2: Heating and Cooling Production Function Estimates

	Heating	Air conditioning
Exponent on capital: θ	0.09*** (0.004)	0.27*** (0.020)
Exponent on energy: η	0.91*** (0.003)	0.85*** (0.029)
Constant: κ	0.13*** (0.024)	0.33* (0.20)
N	1358	137
R^2	0.99	0.97

air conditioning capital. Variable e the energy required to operate the heater or air conditioner at its maximum level for one hour.

Data on the price, capacity, fuel type, and efficiency of heating and cooling units are from ecomfort.com, a direct-to-consumer online store for heating and cooling equipment. The capacity of a heater, measured in BTUs, is the maximum amount of heat the unit can produce per hour. Similarly, the capacity of an air conditioner, measured in tonnage, is the maximum amount of heat the unit can remove per hour. The efficiency of the heater, measured by the Annual Fuel Utilization Efficiency (AFUE), and the efficiency of the air conditioner, measured by the Seasonal Energy Efficiency Ratio (SEER), are the heating or cooling output divided by the energy input. Variable e, the energy required to operate the unit at maximum capacity for one hour, equals the heater or air conditioner capacity divided by its efficiency. The data are based on product performance tests for each heater and air conditioner, implying the usual endogenity concerns with production function estimation do not apply. The resulting data set has 137 air conditioners (including both window units and central air) and 1,358 heaters (including both boilers and furnaces).

Columns (1) and (2) of Table 2 report the coefficient estimates with standard errors in parentheses. The R-square is near unity on both regressions, suggesting that the production functions for heating and cooling in the model provide a reasonably good fit to the data. The heating coefficients sum to one, implying constant returns to scale in heating. The sum of the air conditioning coefficients is slightly greater than one, implying increasing returns in cooling. These increasing returns implicitly capture the switch in technology as tonnage increases from smaller, less efficient window units to larger, more efficient central air systems.

4.2.2. Other parameters

I set capital's income share in the production of the final good, α, equal to 0.26 (Kiyotaki et al., 2011; Nakajima, forthcoming). 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 $\delta = 0.066$, 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, heating capital, A^{xh}, and cooling capital, A^{xc}, to unity. I choose total factor productivity in the production of housing, A^h, equal to 0.12, the ratio of housing services to residential capital in the national accounts (NIPA Tables 2.3.5 and 1.1). Similarly, I choose productivity in the production of cooling energy, A^{ec}, 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). Virtually all energy used for cooling is electricity. In contrast, energy used for heating can be electricity, natural gas, oil, or another source, such as wood. To determine productivity in the production of heating energy, I compute the ratio of the cost per British Thermal Unit (BTU) of heating energy compared to cooling energy in the 2015 Residential Energy Consumption Survey (RECS). This ratio equals 0.37, implying that TFP in heating energy is $1/0.37 = 2.67$ times higher than TFP in cooling energy, $A^{eh} = 2.67 \times A^{ec} = 0.38$.

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 from 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, σ, equals 2. The persistent component of the labor productivity process, ν_{it}, follows an AR(1) process of the form: $\nu_{i,t} = \rho \nu_{i,t-1} + \varepsilon_{it}$, with $\varepsilon_{i,t} \sim N(0, \sigma_\varepsilon^2)$. Parameter ρ denotes the persistence and ε_{it} is a white noise process with variance σ_ε^2. I take the values for ρ, σ_ε^2, and the variance of household-specific fixed effect, σ_ξ^2, from Kaplan (2012). I use the Rouwenhorst method to approximate the AR(1) process with a five-state Markov chain. Table 3 reports the values of the externally calibrated parameters.

Table 3: Parameter Values: External Calibration

Parameter	Value
Firms	
Depreciation rate: δ	0.07
Capital share in final-good production: α	0.26
TFP in housing services: A^h	0.12
TFP in cooling-energy production: A^{ec}	0.14
TFP in heating-energy production: A^{eh}	0.38
Households	
International interest rate: r	0.04
Coefficient of relative risk aversion: σ	2
Bliss point temperature: $\zeta^\star$	18
Heating capital exponent: θ^h	0.09
Cooling capital exponent: θ^c	0.27
Heating energy exponent: η^h	0.91
Cooling energy exponent: η^c	0.85
Labor productivity	
Persistent shock persistence: ρ	0.97
Persistent shock innovation variance: σ^2_ε	0.02
Fixed-effect variance: σ^2_ξ	0.66

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

4.3. Internally calibrated parameters

Table 4 reports the values of the internally calibrated 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.

Table 4: Parameter Values: Internal Calibration

Parameter	Value
Utility	
Weight on housing: ψ	0.20
Discount factor: β	0.94
Temperature damage coefficient: χ	0.01
Heating and cooling	
Housing exponent: γ	0.12
TFP in the production of heating: A^h	215
TFP in the production of cooling: A^c	685
Energy assistance	
Average assistance per recipient: $\bar{B}$	0.01

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.

I choose TFP in the production of heating, A^h, and cooling, A^c, to match average heating and cooling budget shares from the 2015 RECS. All else constant, higher levels of heating and cooling TFP reduce the flow energy required to achieve a given indoor temperature. Consequently, households choose indoor temperatures closer to the bliss point and have lower heating and cooling budget shares.

I choose the exponent on housing services in equation (11), γ, to match total heating and cooling energy expenditures in the fifth income quintile relative to the first quintile. In the 2015 RECS, this ratio equals 1.7, implying that the richest 20 percent of households spend 70 percent more on heating and cooling energy than the poorest 20 percent. Parameter γ controls how the marginal product of heating and cooling energy changes with the level of housing services. For example, if $\gamma = 1$, then, all else constant, the amount of energy required to attain a given indoor temperature doubles when housing services double. Similarly, if $\gamma = 0$, then housing has no effect on the level of energy required to attain a given indoor temperature. Since higher income households purchase more housing services, γ is an important determinant of how energy expenditures vary with income.

Damage coefficient, χ, controls the variation in energy budget share with income. As discussed in the simple model, the relationship between energy budget share and income depends, in part, on how close the households' optimal indoor temperatures are to the bliss point. Parameter χ governs the utility cost of choosing indoor temperatures not equal to the bliss pint. As χ increases, the bliss temperature becomes closer to a "necessity." Households optimally choose indoor temperatures closer to the bliss point, causing heating and cooling budget shares to fall more steeply with income. I choose χ to match the ratio of energy budget share (the sum of heating and cooling budget shares) in the fifth income quintile relative to its value in the first income quintile. In the 2015 RECS, this ratio equals 0.12, implying that energy budget share in the fifth quintile is almost ten times smaller than in the first quintile.

I choose the average value of energy assistance received by households who receive assistance, $\bar{B}$, to target the total funds for energy assistance relative to GDP, equal to 0.00027. The total funds include the sum of LIHEAP, WAP, and state and local energy assistance. I assume that energy assistance is received by all households with combined capital and labor income below the bottom 5.5 percent of the income distribution, the maximum of the fractions of households that received LIHEAP or WAP funds in the data. Empirically, energy assistance varies across states based on differences in expected heating and cooling expenditures (DHHS, 2020). To capture this fact, I assume energy assistance varies across regions in proportion to the sum of the heating and cooling degree days in each region, scaled by the average heating or cooling cost per degree day. Appendix 1.1 provides more details on the energy assistance policies and data.

Lastly, the weight on housing in the utility function, ψ, 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, β, to match the average ratio of U.S. net wealth to output, 3.1, 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 5 reports the values of the targeted moments in the model and in the data. Overall, the model fits the targeted moments quite closely.

Table 5: Model Fit

Target	Model	Data
Wealth to output ratio	3.2	3.1
Housing to non-housing capital ratio	0.89	0.88
Cooling budget share	0.0058	0.0056
Heating budget share	0.014	0.014
Energy-budget-share ratio: high to low income	0.12	0.12
Energy expenditure ratio: high to low income	1.72	1.72
Energy assistance to output	0.00027	0.00027

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

4.4. External validation

I compare the predictions of the model with the data for moments that are less aggregated than the calibration targets. The top panels of Figure 3 plot heating and cooling budget shares by income quintile in the model (solid blue line) and in the data (dashed orange line). In the

Figure 3: External Validation

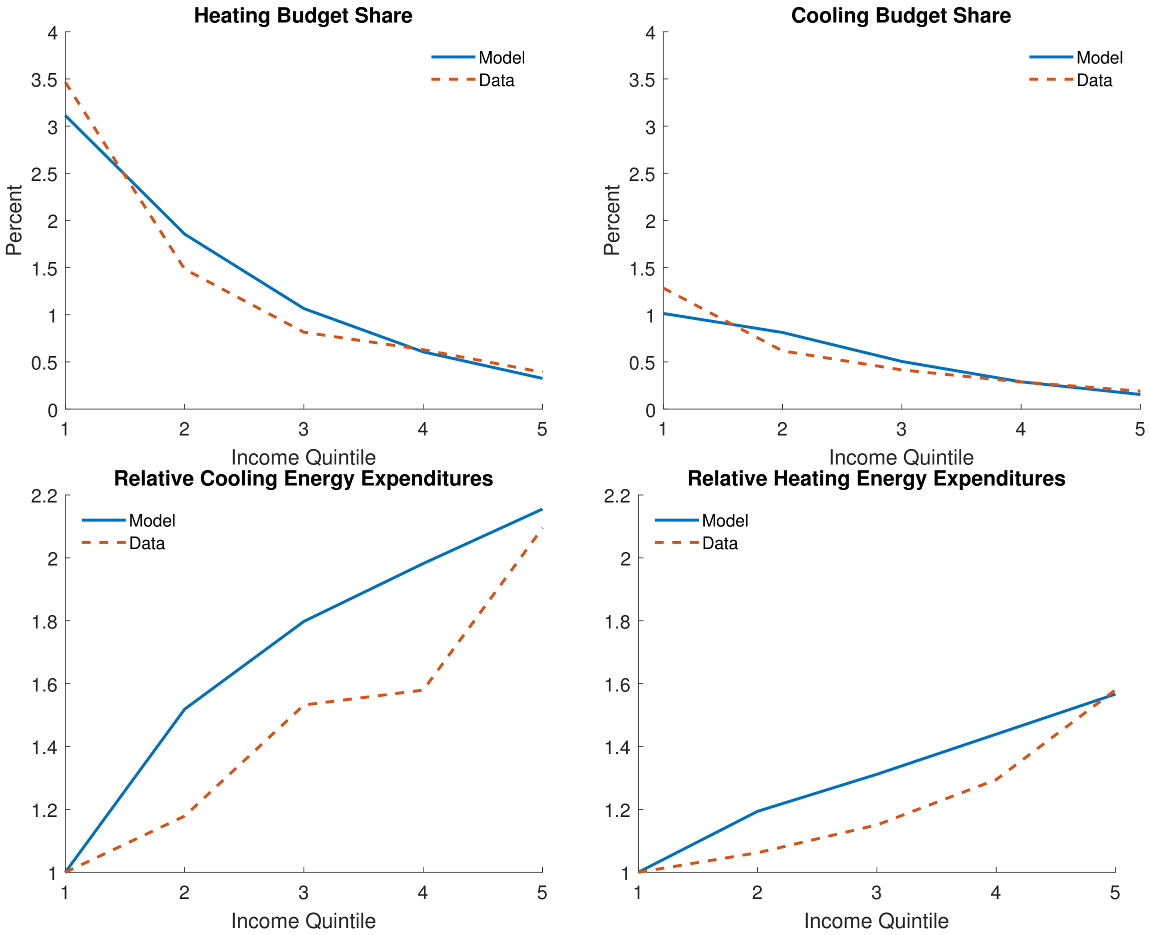

calibration, I target the average of heating and cooling budget shares across income groups, and the ratio of energy share (sum of heating and cooling budget shares) in the fifth quintile relative to the first quintile. Overall, the model does reasonably well matching the less-aggregated changes in heating and cooling budget shares with income plotted in Figure 3.

The bottom panels of Figure 3 plot expenditures on heating and cooling energy relative to their values in the first quintile in the model (solid blue line) and in the data (dashed orange line). In the calibration, I target the ratio of total energy expenditures (sum of heating and cooling energy expenditures) in the fifth quintile relative to the first quintile. The model captures the positive and monotonic relationship between both heating and cooling energy expenditures and income, though with somewhat different convexity.

5. Results

I first use the calibrated model to quantify the distributional effects of weather (e.g., the realization of a hot, cold, or moderate year) in the current US economy. Next, I run a series of counterfactual experiments. In the first experiment, I analyze the distributional and welfare cost of the rightward shift in the temperature distribution caused by climate change. In the second experiment, I analyze the welfare and distributional impacts of adaptation to temperature through heating and cooling in both the current economy, and in response to climate change.

5.1. Distributional consequences of weather

I study the welfare consequences of hot and cold years relative to moderate year. I measure the welfare effects using the *current-period* consumption-housing equivalent variation (CHEV). The current-period CHEV is the percent increase in consumption and housing a household would need in the current period so that they are indifferent between living in the moderate year and the hot or cold year. Positive values imply that the hot or cold year makes the household better off relative to the moderate year.

Figure 4 plots the welfare effects across income deciles and regions of a cold year (left panel) and a hot year (right panel) relative to a moderate year. In the hotter regions, cold years make households better off and hot years make households worse off. The opposite patterns hold in the colder regions. In all regions, the welfare of lower-income households is more sensitive to weather than higher income households, as evidenced by the fact that the magnitude of the welfare impacts in Figure 4 fall with income. Comparing across regions and deciles, the lowest income deciles in hot region experience considerably larger changes in welfare from hot and cold years than any other sub-group.

Figure 4: Welfare Effects of Extreme Years Relative to a Moderate Year (current-period CHEV)

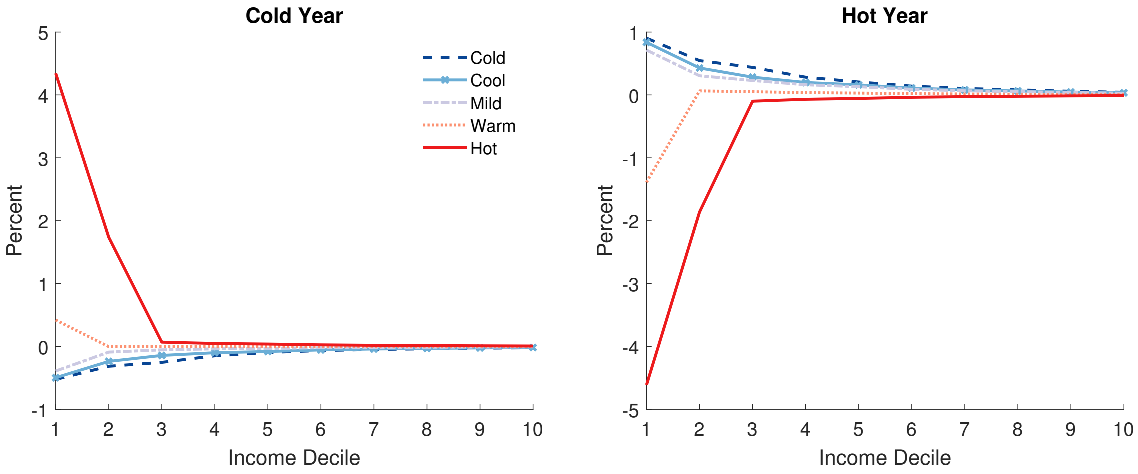

To understand these patterns, I return to the intuition from the simple model and quantify the effect of cold and hot years on transfers from nature. I calculate the transfer from nature for each household as the energy cost to produce the heating or cooling provided by nature, given the household's chosen levels of housing and heating and cooling capital. I measure the heating or cooling provided by nature relative to the minimum and maximum temperatures used in the scientific models, -40^{o}C and 59^{o}C respectively.

Figure 5: Change in Transfers From Nature as a Percent of Income From Extreme Years

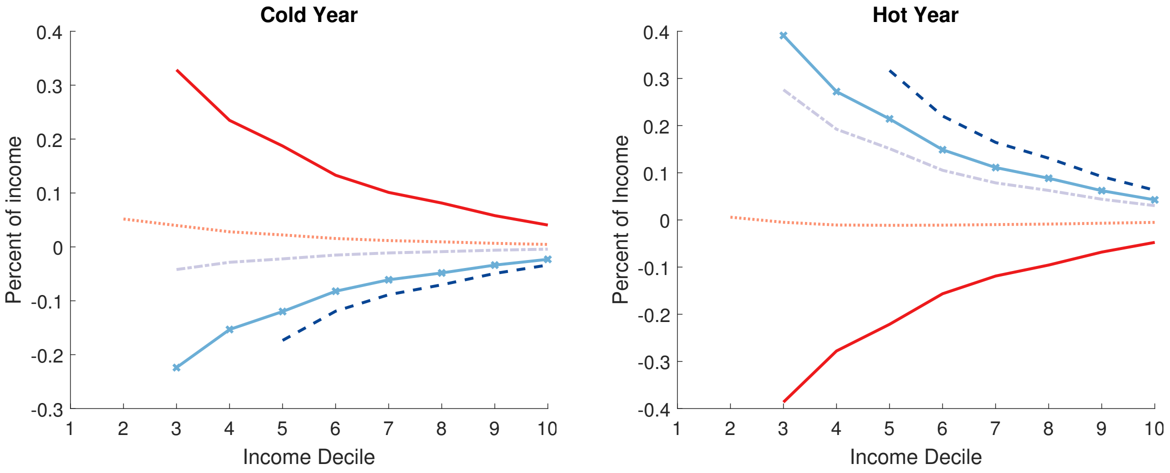

Figure 5 plots the change in transfers from nature relative to income for the cold (left panel) and hot (right panel) years relative to a moderate year. For example, the increase in transfers from nature in a cold year relative to a moderate year for the average household in the fourth decile in the hot region equals 0.85 percent of income. The transfers from nature are infinite (or not well-defined) among the lower income households who choose not to purchase any heating or cooling capital.

Overall, the effects of hot and cold years on transfers from nature in Figure 5 mirror the welfare effects in Figure 4. In the hotter regions, cold years increase transfers from nature, leading to welfare benefits. Hot years decrease transfers from nature, leading to welfare costs. Relative to income, the changes in the transfers from nature are larger for lower-income households, leading to the progressive welfare effects of cold years in the hotter regions in Figure 4 and the regressive effects of hot years. The opposite patterns hold in the colder regions.

Households have less options to adapt to weather (e.g., the realization of a hot, cold, or moderate year) than to climate change. Since households learn the realization of weather after they have made their heating and cooling capital decisions for the period, they can only adapt by adjusting their heating and cooling energy, not their heating and cooling capital. The particular weather realization has no impact of the households' expectations of future weather realizations, and thus no (direct) effect on their future heating and cooling capital decisions. In contrast, climate change does affect households' expectations about future weather realizations (e.g., the temperature distributions in hot, cold, and moderate years all shift to the right) and hence impacts both their heating and cooling capital and their heating and cooling energy decisions.

5.2. 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 take a probability weighted average of the different climate models where the weights equal the relative probability that the climate model represents the true outcome (Rasmussen et al., 2016; Hsiang et al., 2017). I use the RCP 8.5 projections, which are designed to approximate global emissions in the absence of large-scale climate policy. The projections incorporate the effect of climate change on both the mean and the variance of temperature over the course of a year.

I model uncertainty over the temperature distribution using the same approach as in the baseline. 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 mod-

Figure 6: Effect of Climate Change on the Temperature Distribution

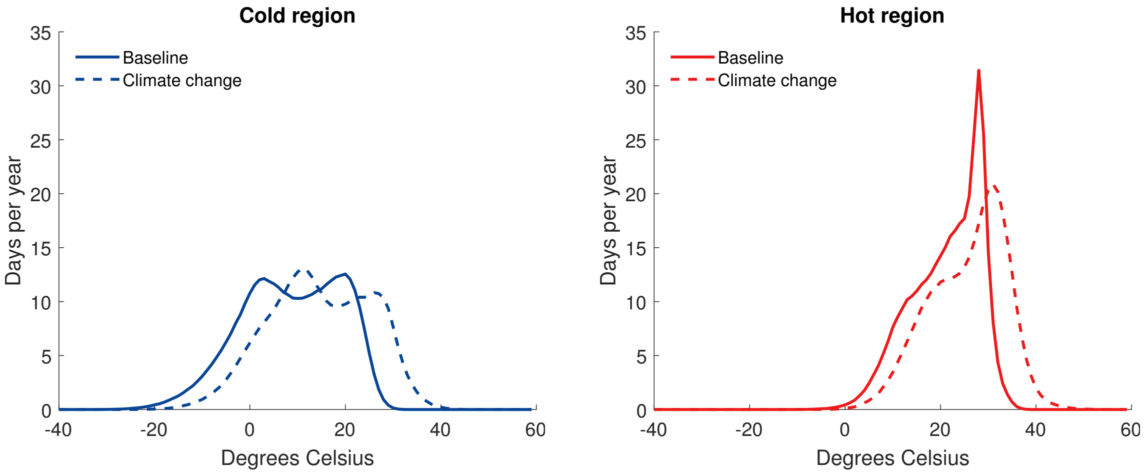

erate year is the same as in the baseline. I make the analogous assumption to determine the temperature distribution in the hot year.

Figure 6 plots the temperature distributions in a moderate year in the baseline (solid lines) and with climate change (dashed lines) in the cold (left panel) and hot (right panel) regions. Climate change shifts all realizations of the temperature distribution in all regions to the right. Appendix Table 8 reports the mean and the variance of temperature over the course of a year in each region in the baseline and climate change equilibrium. Both the mean and the variance of temperature during a year increase in response to climate change.

To study the distributional impacts of climate change, I solve for a new stationary equilibrium with the climate-change temperature distribution. I compare this climate-change equilibrium to the baseline equilibrium with the historical 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 baseline equilibrium, so that they are indifferent between the baseline and the climate-change equilibrium. Note that this welfare measure differs from the current-period CHEV used to analyze the distributional effects of weather in Section 5.1. To quantify the distributional impacts of climate change, I measure the CHEV conditional on the household being in each decile of the income distribution. Negative values indicate that climate change makes the household worse off.

Figure 7 plots the CHEV by income decile in each region. For most income deciles, climate change leads to welfare costs in the hotter regions (the CHEV is negative) and welfare benefits in the colder regions (the CHEV is positive). To understand these regional differences, it is useful to examine how climate change affects the transfers from nature in each region. Fig-

Figure 7: Welfare Consequences of Climate Change

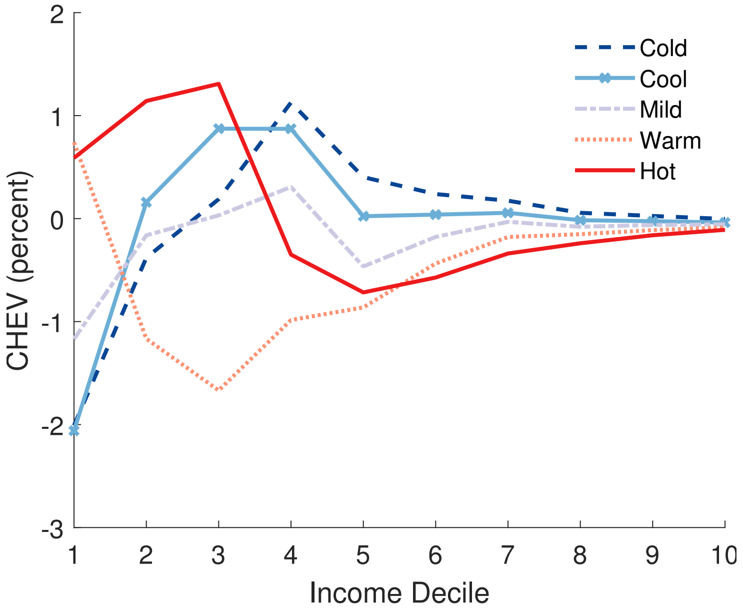

ure 8 plots the change in transfers from nature caused by climate change as a percent of the household's income. In hot regions, climate change reduces transfers from nature because it leads to more extreme temperature days. Similarly, in cold regions, climate change increases transfers from nature because it leads to fewer extreme temperature days. The lower transfers in hot regions contribute to the overall welfare costs from climate change in hot regions in Figure 7, while the higher transfers in cold regions contribute to the overall welfare benefits from climate change in the cold regions.

Among the medium- and high-income groups, the effect of climate change on transfers from nature relative to income falls as income rises. All else constant, the smaller impact of climate change on transfers from nature for higher income groups implies that the welfare costs of climate change fall with income in the hot regions and the welfare benefits from climate change fall with income in the cold regions. Referring back to Figure 7, these patterns generally hold. For all but the lowest income deciles, the lines in the colder regions are positive and downward sloping, implying that the welfare benefits fall as income rises. Similarly, the lines in the hotter regions are negative and upward sloping, implying that the welfare costs fall as income rises.

However, changes transfers from nature are not sufficient to explain the relationships between welfare and income for households in the lowest income deciles. To understand these impacts, we also need to consider how climate change affects the household's choice to specialize in heating or cooling capital. Since heating and cooling require separate types of capital (e.g. a gas furnace versus a central air unit), households choose to idle cooling capital on cold

Figure 8: Effect of Climate Change on Transfers From Nature Relative to Income

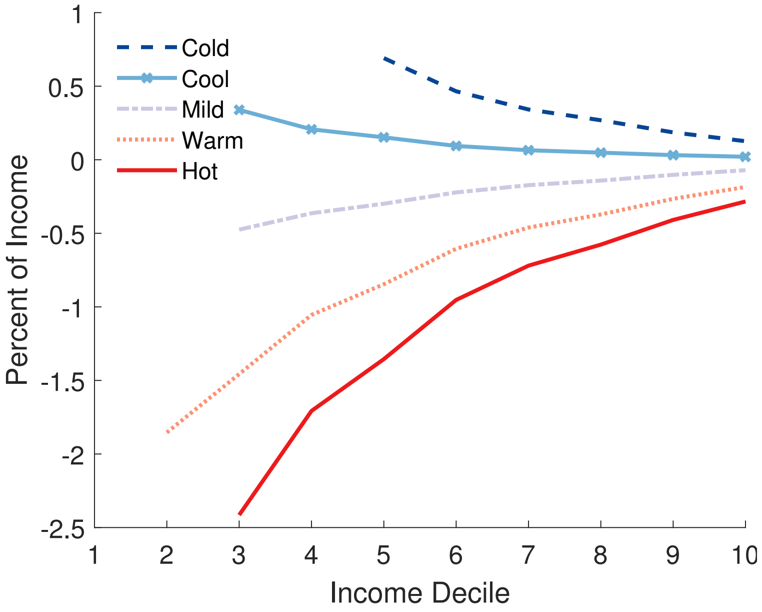

days and heating capital on hot days. If the temperature distribution was such that there were only hot or cold days, then households could fully specialize in heating or cooling capital, and avoid idling any capital.

To quantify the effect of climate change on how favorable the temperature distribution is for specialization, I calculate a specialization favorability index, $S_n \in [0,1]$ in each region n,

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

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 capital and still attain the bliss point temperature in every sub-period, if they choose to do so.

Figure 9 plots the specialization favorability index in the baseline (blue bars) and in the

Figure 9: Specialization Favorability Index

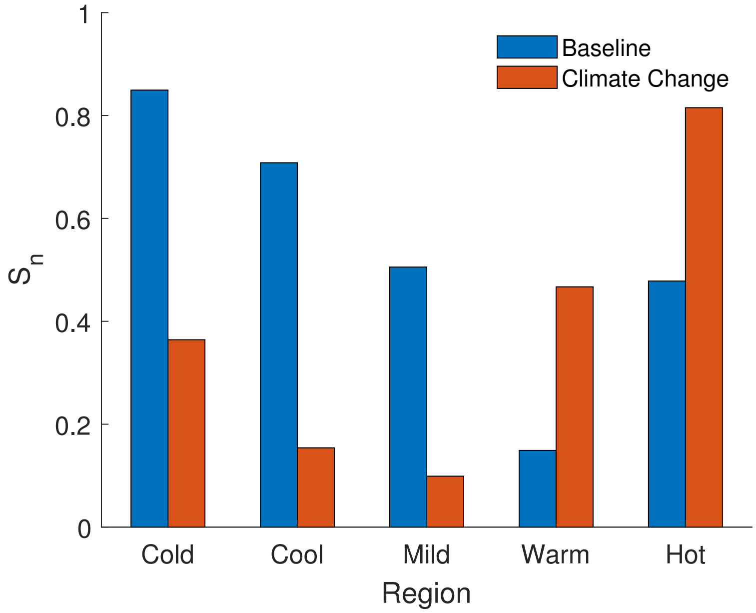

climate change equilibrium (orange bars) for each region. In the baseline, the temperature distribution in the cold and cool regions is relatively favorable to specialization; the specialization favorability index equals 0.85 and 0.71, respectively. The cold and cool regions experience only a small number of hot days in the baseline, allowing households to specialize more in heating capital. The rightward shift in the temperature distribution caused by climate change leads to more hot days in these regions, reducing the favorability of the temperature distribution to specialization.

The opposite effect occurs in the warm and hot regions. In the baseline, the temperature distribution is not particularly favorable to specialization. While households experience more hot days than cold days, they still experience a substantial number of cold days. The rightward shift in the temperature distribution caused by climate change reduces the number of cold days, increasing the specialization favorability of the temperature distribution. For example, in the hot region, climate change increases the specialization favorability index from 0.48 to 0.82.

All else constant, the decrease in specialization favorability in the colder regions reduces the welfare benefits from climate change. Similarly, the increase in specialization favorability in the hotter regions reduces the welfare costs of climate change. While changes in the specialization favorability affect all income groups, they are most pronounced for the lowest income deciles 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 heating or cooling capital and an equilibrium in which they pay for both heating and cooling capital. Mathematically, this is because, all else constant, the welfare consequences of moving between a corner solution, where the first order condition for heating or cooling capital does not hold, and an interior solution, where the first order conditions do hold, are larger than the welfare consequences from moving between two interior solutions where the first order conditions always hold.

Figure 10 plots the fraction of households in the hot and cold regions without heat (left panels) and without AC (right panels) in the baseline (blue) and climate change (orange) equilibria. Appendix B reports the results in the other regions. Focusing first on the cold region, all households have heat in both equilibria, but not all households have air conditioning. Comparing the blue and orange lines, 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 the lowest two deciles, the welfare costs of the decrease in the favorability of the temperature distribution to specialization dominate the welfare benefits from higher transfers from nature, leading to the overall negative welfare effect of climate change in Figure 7.

We see the opposite patterns in the hot region. Here, all households have air conditioning in both equilibria but not all households have heat. Again comparing the blue and orange lines, climate change causes many households in the second through fifth decile to move from an interior solution in which they have heat to a corner solution in which they do not have heat. The welfare benefits from the increase in specialization partially offset, or in some cases dominate, the welfare costs of the reduced transfers from nature, leading to the relatively smaller welfare costs of climate change, in the fourth and fifth deciles and the welfare benefits in the second and third deciles plotted in Figure 7.

Households in the first decile in the hot region do not have heat in either equilibrium. Since climate change does not cause the these households to move from an interior to a corner solution, the welfare benefits from the increase specialization favorability are smaller than for the slightly richer households who do switch from having heat in the baseline to not having heat in the climate change equilibrium.

In sum, the distributional impacts of climate change in Figure 7 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 8 implies that the welfare benefits of climate change fall with income in the cold regions and the welfare costs of climate change fall with income in the hot regions. However, the comparatively large welfare impact from changes in specialization favorability for low-income households that move between complete and incomplete specialization

Figure 10: Changes in Heat and AC Specialization

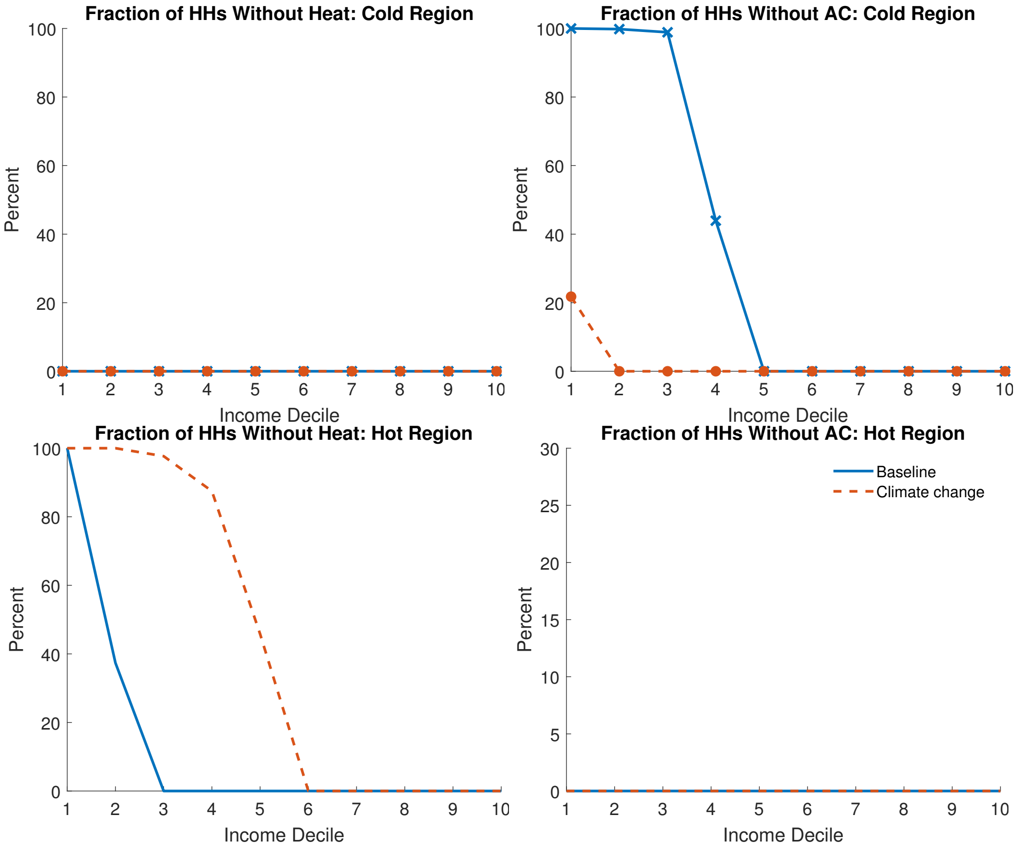

Figure 11: Welfare Effects of Adaptation in the Current Economy

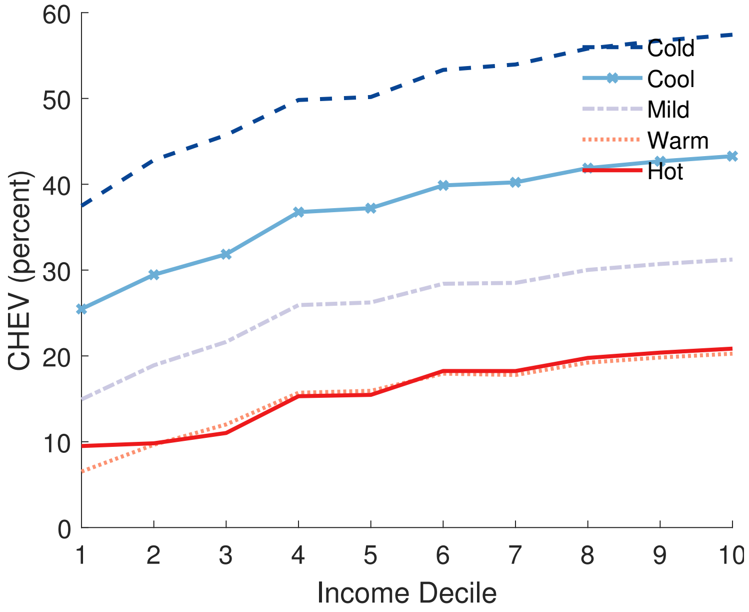

breaks the monotonicity in this relationship.

5.3. Adaptation to temperature and climate change

Households adapt to outdoor temperature in the model primarily by using energy and capital for heating and cooling. I explore the effects of adaptation along three dimensions. First, I study the welfare and distributional consequences of adaptation to temperature in the current economy, without any additional climate change. Second, I analyze how climate change affects households' temperature adaptation decisions. Finally, I examine how this adaptation affects the welfare and distributional consequences of climate change.

To understand the role of adaptation in the current economy, I solve for a no-climate change equilibrium in which households cannot adapt to temperature. This equilibrium has the same temperature distribution as the baseline equilibrium. Since households cannot adapt to temperature, the indoor temperature equals the outdoor temperature in every sub-period. Figure 11 plots the CHEV between the baseline equilibrium without adaptation and the baseline equilibrium with adaptation. Positive values imply that adaptation makes households better off.

Adaptation improves the welfare of all households, with the largest gains in the colder regions. The utility cost of the damage from the baseline temperature distribution is highest in the coldest region, and falls monotonically across the cold, cool, mild, and warm regions, causing the welfare benefits from adaptation to increase monotonically across these regions. The hottest region breaks this pattern; the utility cost of the damage caused by the temperature distribution is almost as high as in the cold region, yet the welfare benefits from adaptation are substantially lower. Compared to the cold region the hot region has a more extreme temperature distribution, measured by the maximum number of degrees away from the bliss point in any sub-period. Additionally, the diminishing returns to cooling energy are larger than for heating energy ($\eta^c < \eta^h$). Combined, these two factors imply that adaptation is relatively costly in the hot region, leading households to choose indoor temperatures farther from the bliss point. The relatively high adaptation costs offset the relatively high utility damages, muting the overall welfare benefits from adaptation in the hot region.

Within each region, the welfare benefits from adaptation increase with income. To understand the intuition, consider the extreme case of a very low-income household who does not purchase a heater or an air conditioner. The welfare gains from adaptation for this household are zero. As income rises, households optimally choose to purchase more heating and air conditioning, leading to indoor temperatures closer to the bliss point and larger welfare benefits.

Climate change shifts the temperature distributions in all regions to the right, causing households to change their heating and cooling adaptation decisions. Figure 12 plots the percent change in households' average heating (left panel) and cooling (right panel) expenditures between the baseline and climate change equilibria. Across all regions, households reduce their heating expenditures and increasing their cooling expenditures in response to climate change. The magnitude of the percentage change in heating expenditures is largest in the hot region. This region has the lowest heating expenditures in the baseline and climate change causes many low and middle income households in this region to forgo all heating. Similar, but opposite reasoning implies that the magnitude of the change in cooling expenditures is largest in the cold region. Within any region, the percentage changes in heating and cooling expenditures are larger for lower income households, consistent with the notion that climate change leads to larger changes in the transfers from nature for these groups.

Households have the technology to fully adapt to climate change by attaining the same indoor temperature as they choose in the baseline. However, this full adaptation is not necessarily optimal because climate change affects the expected marginal benefits of heating and cooling capital and energy. To understand how much households choose to adapt, Figure 13 plots the change in indoor temperature (in degrees) when the outdoor temperature is 13°C, 5 degrees below the bliss point (left panel) and when the outdoor temperature is 23°C, 5 degrees above the bliss point (right panel).

In response to climate change, households choose colder indoor temperatures on cold days. Climate change reduces the marginal benefit of heating capital, causing households to purchase less heating capital. In any given sub-period, the lower level of heating capital reduces the marginal product of heating energy, leading households to choose indoor temperatures farther

Figure 12: Effect of Climate Change on Expenditures for Heating and Cooling Capital and Energy

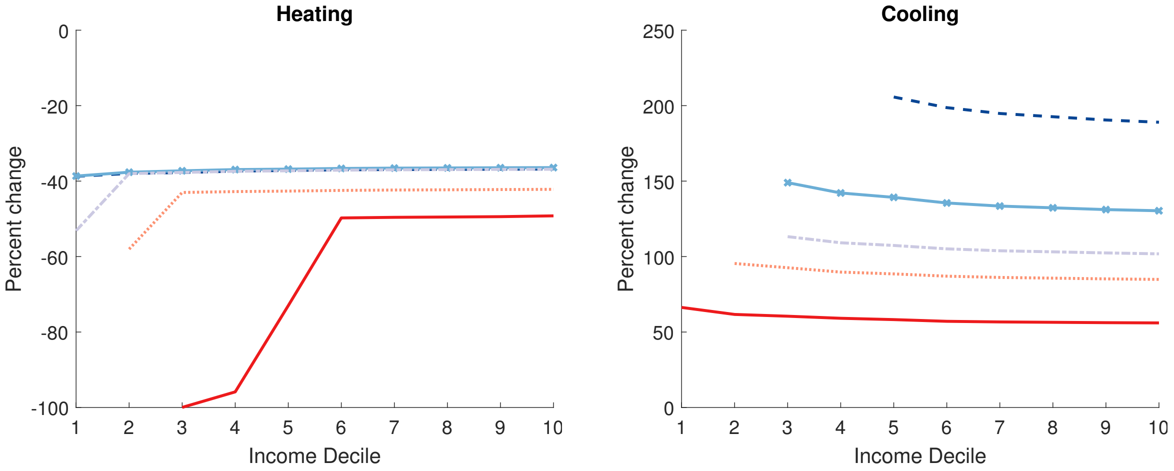

Figure 13: Effect of Climate Change on Indoor Temperature

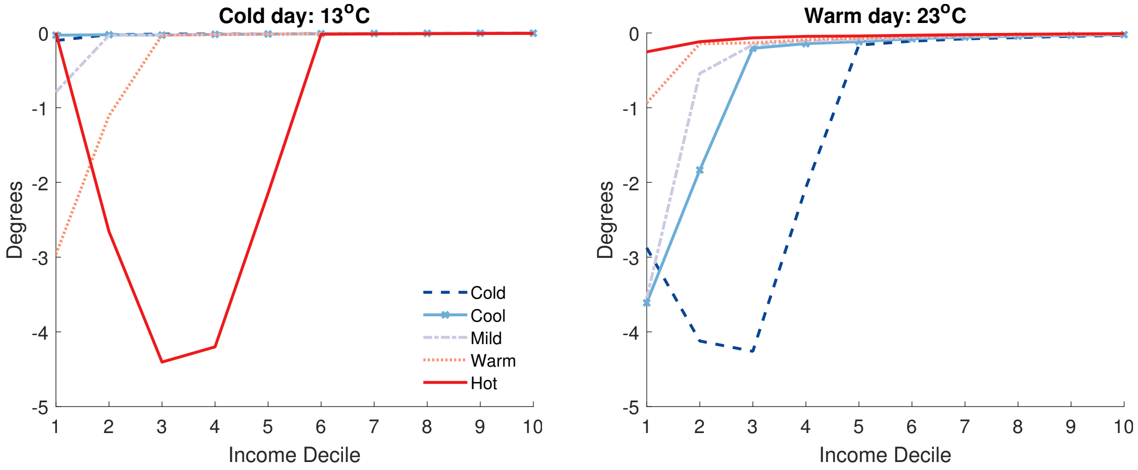

Figure 14: Welfare Impacts of Climate Change When Households Cannot Adapt to Temperature (CHEV)

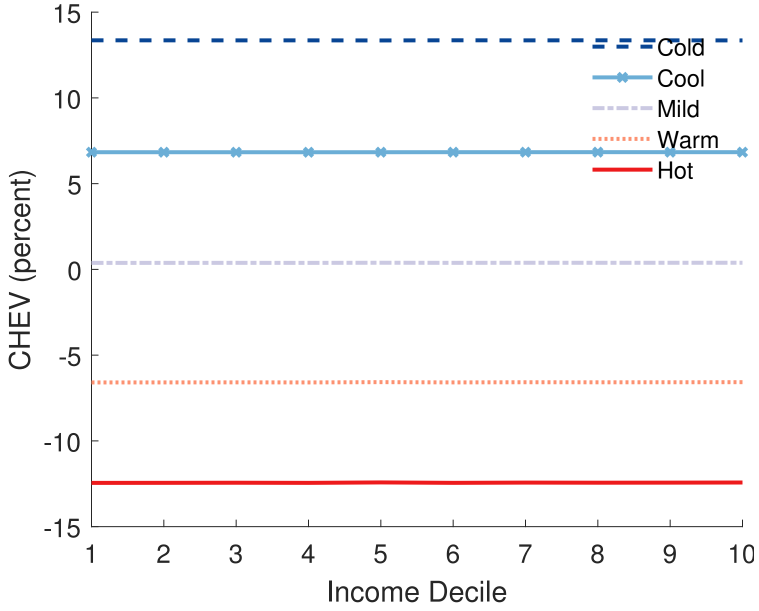

from the bliss point. This effect is particularly large among households in the hotter regions that switch from using heating in the baseline to not using heating in the climate change equilibrium. Parallel but opposite reasoning implies that climate change moves households closer to the bliss point on warm days.

Looking across income groups, the effect of climate change on households' optimal indoor temperature generally falls with income. Climate change leads to larger relative changes in transfers from nature for lower income households, generating larger responses across all of the choice variables. Compounding on this, climate change has the strongest effect on indoor temperature among the households that move into or out of complete specialization, and these households tend to be lower income.

Lastly, I consider how households' ability to adapt to temperature affects the distributional and welfare consequences of climate change. I solve for a no-adaptation climate-change equilibrium which has the same temperature distribution as the climate change equilibrium from Section 5.2. Since households cannot adapt, the outdoor temperature equals the indoor temperature. I measure the welfare effects of climate change when there is no adaptation as the CHEV between the baseline equilibrium without adaptation and the climate change equilibrium without adaptation.

Figure 14 plots the welfare effects of climate change without adaptation. Comparing Figures 7 and 14 reveals that adaptation mutes the magnitude of the welfare effects of climate change across all income groups and regions. Furthermore, adaptation substantially increases the variance in the welfare effects across income groups; in contrast to Figure 7 the lines in Figure 14 are virtually flat. All households within a region experience the same shift in the temperature distribution. If households cannot adapt, then they all experience the same welfare cost. The ability to adapt varies with income, and thus leads to the variance in the welfare effects of climate change across income groups in Figure 7.

5.4. Importance of 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 U.S.

To understand the importance of income heterogeneity, I simulate climate change in a counterfactual economy in which I halve the variance of the persistent labor productivity shock. Mechanically, the lowest income households in low variance economy are richer than their baseline counterparts and the highest income households are poorer. Average labor income in the low variance economy is the same as in the baseline. Table 6 reports the aggregate welfare cost of climate change within each region in the baseline and low-variance simulations.

Table 6: Welfare Impact of Climate Change (CHEV, percent)

	Cold	Cool	Mild	Warm	Hot
Baseline simulation	-0.65	-0.54	-0.50	-0.53	0.52
Low variance simulation	-0.51	-0.37	-0.49	-0.45	0.35

In all regions, the degree of income heterogeneity matters for the aggregate welfare cost of climate change within the region. The effect of income heterogeneity on the welfare cost of climate change operates through two key channels. First, the range of income across households is larger in the baseline simulation. The heterogeneous welfare effects of climate change across income groups implies that these range differences could impact the aggregate welfare cost. For example, in the baseline simulation, the lowest income households in the coldest regions experienced substantial welfare costs. These households are not included in the low-variance simulation, implying that, all else constant, climate change will have smaller welfare costs in the low-variance simulation. Similarly, the richest households in the baseline simulation are also not included in the low-variance simulation. However, the variation in the welfare cost of climate change among the richer households in Figure 7 is relatively small, implying that differences in the income range matter less at the top than at the bottom.

The second channel is that, conditional on income, more households choose corner solutions without heat or air conditioning in the low-variance simulation than in the baseline simulation. Consequently, the welfare impacts of changes in specialization favorability are larger in the low-variance simulation. Again, using the cold region as an example, without climate change, a larger number of households do not have heat in the low variance simulation. Climate change moves all of these households out of the corner solution, raising the welfare cost of climate change relative to the baseline simulation.

More households choose corner solutions conditional on income in the low variance simulations because they have less savings; aggregate savings in the no-climate change equilibrium in the baseline simulation equals 3.61 compared to 0.97 in the no-variance simulation. Households' primary motive for saving in the model is to self-insure against labor income risk. The lower labor income risk in the low-variance simulation causes them to save less. Consequently, low realizations of the labor-income shock reduce the benefit of a heater or air conditioner more than in the baseline simulation because households' are less able to use the savings to purchase energy to operate the heater or air conditioner. As a result, more low-income households choose the corner solutions in the low-variance simulation.

Appendix

A. Data

1.1. Energy assistance

To determine the average value of energy assistance received by households who receive assistance, I calculate total funds for energy assistance relative to GDP in 2015, equal to 0.00027. Energy assistance, B_{in}, varies with the household's income and region. I use data on LIHEAP assistance (DHHS, 2020), WAP assistance (NASCSP, 2019) and state and local energy assistance (cite) to calculate total funds for energy assistance relative to GDP, equal to 0.00027. In praci

Data on LIHEAP expenditures for energy assistance by state are from the Administration for Children and Families Report to congress (Administration for Children and Families, 2020).[7] The data provide information on the number of households that receive assistance with their energy bills and the amount of funds allocated to this assistance from 2008 - 2019. Data on WAP funding from are from the WAP annual funding surveys.[8] The surveys provide information on total WAP funds from 2009-2019. The WAP fact sheet (US Department of Energy, 2021) reports that 35,000 households receive WAP assistance each year. To convert transfers of capital stock, such as a more efficient furnace, to the flow transfers in the model, I assume that the investment's average lifespan is 16 years (Fowlie et al., 2018).

I take the following steps to determine the total amount of additional funds for energy assistance (i.e., separate from LIHEAP and WAP) provided by individual utilities. First, I identify the utilities that provide additional discounts and the size of the discounts from the LIHEAP Clearing House.[9] Second, I determine the number of households the utility serves from the utility's official website.[10] Approximately one quarter of US households are served by utilities that offer additional energy assistance beyond LIHEAP and WAP. The utility-specific discounts take the form of either a dollar reduction in the monthly base charge (base discount) or a percentage reduction in the marginal cost of electricity (percentage discount). To calculate total expenditures on utility-specific assistance, I need to convert the percentage discount into a dollar value. To do this conversion, I take the following steps:

1. Calculate the average percentage discount across all utilities by census division, including

---

[7] The report can be downloaded from: `www.acf.hhs.gov/sites/default/files/documents/ocs/rpt_liheap_congressional_request_for_formula_analysis_appendices.pdf`.

[8] The surveys can be downloaded from: `https://nascsp.org/wap/weatherization-publications/wap-annual-funding-surveys/`.

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

[10] When the utility does not report the number of households served on their website, I obtain the information from `findenergy.com`.

those that do not offer percentage discounts (and thus have a percentage discount of zero).

2. Calculate the average electricity bill among households that receive assistance by census division in the 2015 RECS.

3. 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:

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

Similarly, I calculate the average base discount across all utilities, including those that do not offer base discounts. Total utility-specific expenditures equal the average dollar value of the percentage discount plus the average base discount multiplied by the number of eligible households. 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 (5.5 percent). The sum of LIHEAP, WAP and utility-specific assistance equals 0.027 percent of GDP in 2015. Approximately 60 percent of the total assistance is from LIHEAP, 20 percent is from WAP, and the remaining 20 percent is from the utility-specific assistance.

I determine the variation in energy assistance across regions and with climate change based on changes in heating and cooling degree days, scaled by the average cost of heating per heating degree day, 0.15 dollars per HDD, and the average cost of cooling per cooling degree day, 0.13 dollars per CDD, from the 2015 RECS. Table 7 reports the cost-weighted degree days relative the cold region in the baseline equilibrium. The cold region in the baseline has the largest value of the cost-weighted degree days hence receives the highest energy assistance payments. Climate change reduces the cost-weighted degree days, and hence energy assistance in the colder regions and increases the cost-weighted days in the warmer regions.

Table 7: Cost Weighted Degree Days

	Cold	Cool	Mild	Warm	Hot
Baseline	1	0.85	0.71	0.56	0.56
Climate change	0.81	0.75	0.70	0.65	0.75

1.2. Climate change temperature

Table 8: Mean and Variance of Annual Temperature in a Moderate Year

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

B. Additional Results

Figure 15: Changes in Heat and AC Specialization

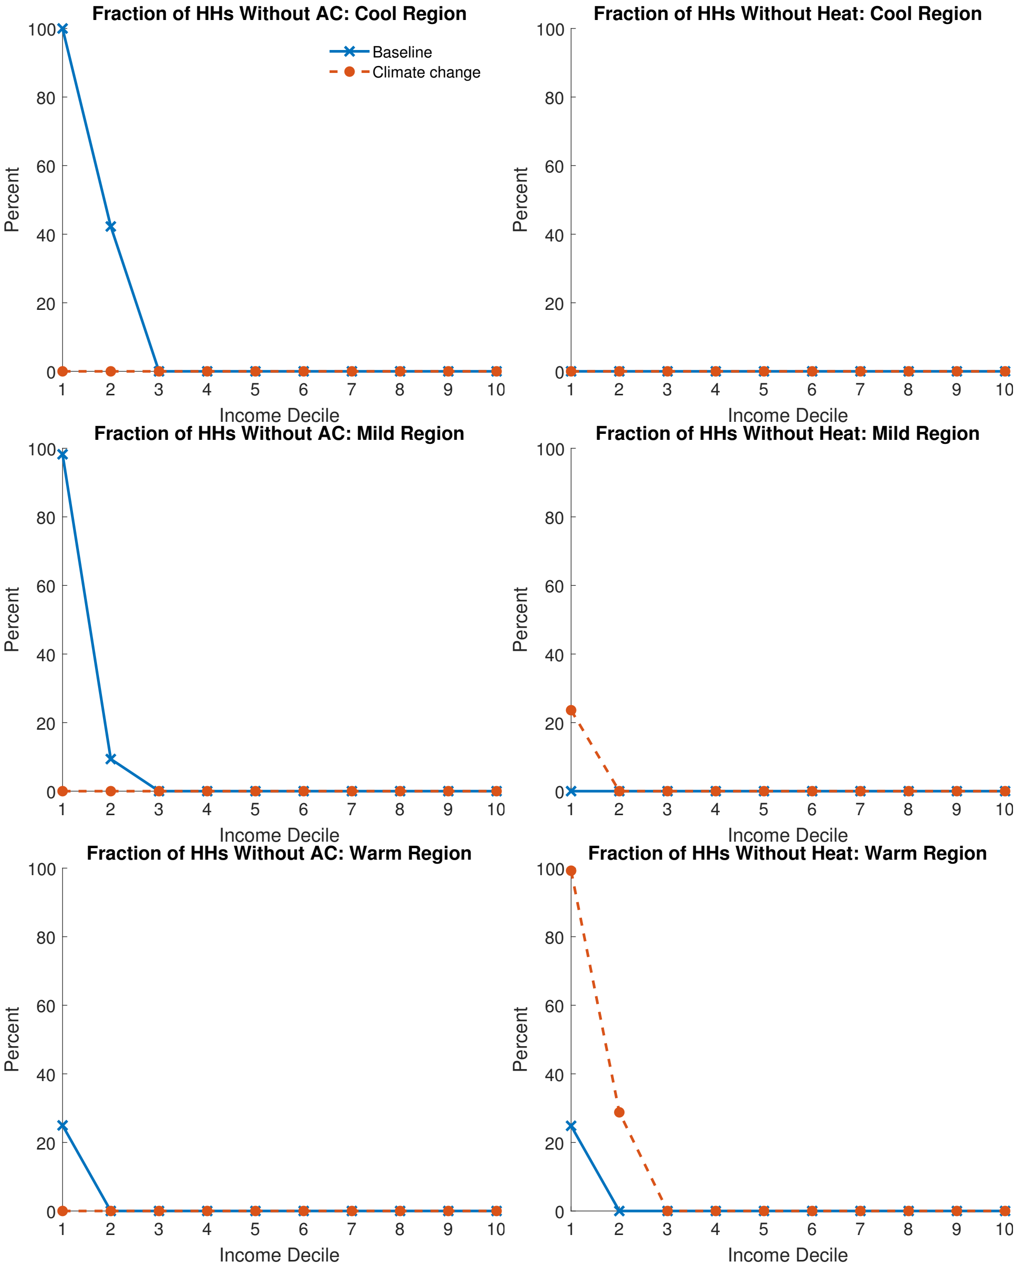

References

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 hte Program’s Formula and Allocations of Funding Among States: Appendices,” .

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*.

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,” .

CRUZ, J.-L. AND E. ROSSI-HANSBERG (forthcoming): “The Economic Geography of Global Warming,” *Review of Economic Studies*.

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.

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

FOWLIE, M., M. GREENSTONE, AND C. WOLFRAM (2018): “Do energy efficiency investments deliver? Evidence from the weatherization assistance program,” *The Quarterly Journal of Economics*, 133, 1597–1644.

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.

Heathcote, J., K. Storesletten, and G. L. Violante (2009): "Quantitative macroeconomics with heterogeneous households," *Annu. Rev. Econ.*, 1, 319–354.

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.

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*.

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. (forthcoming): "Capital Income Taxation with Housing," *Journal of Economic Dynamics and Control*.

NASCSP (2019): *Weatherization Assistance Program Funding Report*.

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.

PARK, J. AND P. STAINIER (2021): "Inequality in workplace climate exposure," *Working paper*.

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).

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.

ROTH TRAN, B. AND D. J. WILSON (2022): "The Local Economic Impact of Natural Disasters," *Federal Reserve Bank of San Francisco Workin gPaper 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," *Associated Press*.

STROBL, E. (2011): "The economic growth impact of hurricanes: Evidence from US coastal counties," *Review of Economics and Statistics*, 93, 575–589.

US DEPARTMENT OF ENERGY (2021): *Factsheet: Weatherization Assistance Program*, Office of Energy Efficiency and Renewable Energy, US Department of Energy, Washington DC.

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.