

The Distributional Consequences of Climate Change: Housing Wealth, Expectations, and Uncertainty

[\[Click here for latest version\]](#)

Jeffrey E. Sun*

March 24, 2025

Abstract

The median US household holds most of its wealth in one immobile, climate-exposed asset: the home. I study the implications of this fact for the distributional consequences of climate change, using a dynamic, stochastic, heterogeneous-agent model with 1713 locations. Forward-looking equilibrium real estate prices and rents endogenously respond to climate news shocks in spatially-segmented markets. Households participate on both sides of each market. To quantify climate impacts on economic fundamentals, I harmonize recent estimates of local productivity, amenities, energy costs, and disaster damage sensitivities. I find the model's global solution under aggregate climate uncertainty with a simple but general deep learning method introduced in a companion paper. In the calibrated model, a switch from widespread climate denial to widespread climate acceptance causes an effective transfer of housing wealth across regions of \$41bn immediately and \$507bn over the following century. Migration exacerbates this by amplifying housing price responses. The spatially-equalizing effects of migration only dominate for all households 50 years post-shock. Climate uncertainty causes ongoing regressive wealth transfers through higher equilibrium rents, borne mostly by poorer households.

*Princeton University. Email: jesun@princeton.edu.

I am deeply indebted to my advisors, Ezra Oberfield, Stephen Redding, Richard Rogerson, and Ishan Nath, for their thoughtfulness and generosity. This project was completely derailed (for the better) by conversations with Monika Mrázová, Dávid Nagy, Luigi Pascoli, Gianluca Violante, and Motohiro Yogo. I am grateful for the detailed comments of Leah Boustan, Levi Crews, José-Luis Cruz, Mayara Felix, Allan Hsiao, Ricardo Lagos, Lukas Mann, and John Sturm Becko. I have greatly benefited from conversations with Elena Aguilar, David Argente, Marnie Ginis, John Grigsby, Gene Grossman, Sebastián Guarda, Jacob Hartwig, Michael Jenuwine, Nobuhiro Kiyotaki, Benny Kleinman, Hugo Lhuillier, Mikhail Mamonov, Kristina Manysheva, Pablo Ottonello, Esteban Rossi-Hansberg, Anna Pestova, Thomas Sargent, Karthik Sastry, William Toms, Robert Wagner, Mark Watson, Sifan Xue, Yucheng Yang, and many seminar participants. I gratefully acknowledge The Simpson Center for the Study of Macroeconomics and the International Economics Section at Princeton for their generous funding.

1 Introduction

I study heterogeneity in the welfare impacts of climate change in a model in which households differ in their wealth and income, make forward-looking investment and migration decisions, are uncertain about the future of the climate, and hold an empirically realistic share of their wealth in housing. Housing prices and rents are endogenously priced in spatially-segmented markets in which households participate on both sides. A number of methodological contributions make this possible.

Climate change is one of the largest challenges currently facing humanity. A recent dynamic spatial climate economics literature studies the crucial question of who this will affect, taking into account that people and markets can adapt in various ways. In a seminal paper, Desmet et al. (2021) study climate-induced flooding and find that migration and local investment adaptation reduce overall productivity impacts by over 97%.¹ Bilal and Rossi-Hansberg (2023) make a substantial methodological contribution by incorporating forward-looking households and find that anticipated climate change induces significantly less migration.

Models in this literature typically incorporate household heterogeneity in location but abstract from other differences such as wealth, age, or income. This is likely due in part to its position at the intersection of multiple methodological challenges that each constrain broader literatures. In all heterogeneous-agent (HA) models, household state space constraints impose an often-onerous trade-off between different dimensions of household heterogeneity. Spatial models and transition dynamics each multiply this state space by the number of modeled units of space and time, respectively. With aggregate (e.g. climate) uncertainty, the aggregate state spaces of most HA models would be vast, generally precluding conventional solutions entirely.

Incorporating housing wealth compounds this complexity by introducing nontrivial housing asset and rental market equilibria within each location, complicating the household’s problem, and introducing household exposures to equilibrium housing prices. Housing wealth is empirically relevant nonetheless. In the U.S., housing wealth is approximately one third of the net value of all economic assets, or \$51.9tn² out of \$154.3tn.³ Over 80% of all housing units are owned by individual households, representing, for most households, a large and undiversified climate exposure, and for the median household, the majority of their wealth. Its monetary value, determined in equilibrium, primarily reflects expectations about the economic future of its location—a future that may be shaped by climate change.

I introduce a number of computational methods to solve my model, which features rich household heterogeneity, 1713 locations, transition dynamics, endogenous housing prices, and aggregate climate uncertainty. These methods do not impose common restrictions on the flexibility of household decisions or idiosyncratic

¹Relative to a version of their model that shuts down these adaptations.

²Zillow (2023)

³Excluding U.S. government wealth. Specifically, Q2 2023 household net worth (FRB), equal to \$154.3tn.

risk exposures,⁴ ex ante identicalness of households,⁵ or non-linearities in climate outcomes by scenario.⁶

An additional difficulty is that a spatial equilibrium analysis requires estimates of exogenous shocks which are spatially varying and specifically quantify impacts to economic fundamentals rather than picking up equilibrium adaptations. Incorporating forward-looking behavior and uncertainty requires also quantifying household expectations and potentially multiple sources of uncertainty. I quantify the climate impacts to economic fundamentals by identifying a set of studies from the literature that identify different climate impacts, such that, where applicable, each study controls in its methodology for the effects picked up by the others. The estimated impacts are plausibly exogenous to the economic model and can be harmonized and locally linearized well for use in the model.

In the calibrated model, housing wealth dramatically shapes where, when, and to whom climate impacts occur. A change in expectations from widespread climate denial to widespread climate acceptance induces a spatially-heterogeneous adjustment of housing prices equivalent to an effective transfer of housing value across space of \$41bn immediately and \$507bn over the following century. Within-location household heterogeneity—by wealth, income, housing wealth, and age—increases overall inequality in these impacts by 58%. In contrast to a version of the model without housing wealth, migration initially exacerbates spatial inequality in these impacts, as house prices in negatively-affected locations also respond to anticipated out-migration.

In the global solution, household heterogeneity, housing wealth, and uncertainty interact to exacerbate wealth inequality in expectation. Climate uncertainty introduces undiversifiable climate risk to housing assets, reducing demand for owning rental housing and thus reducing supply in the housing rental market. This increases equilibrium rents, effectively shifting the costs of climate risk from households who own housing to households who rent. Wealthier households with endogenously lower absolute risk aversion benefit from this climate uncertainty through increased risk premia on housing assets. For them, this more than counteracts the welfare costs of climate risk in housing prices.

Model

The model is essentially a macro-style incomplete-markets dynamic heterogeneous agent model with multiple locations, as opposed to a typical spatial model. It is a small open economy: interest rates and goods prices are exogenous and fixed over time. The household block is modeled after Kaplan et al. (2020), although the household’s housing choice is over continuous housing quality as in Favilukis et al. (2017). This is then embedded in a spatial equilibrium model with frictional migration and heterogeneous bilateral migration costs, similar to Borusyak et al. (2022) or Giannone et al. (2023), who similarly embed a version of Kaplan

⁴e.g., continuous-time models such as Greaney (2021) impose that household constraints never bind.

⁵e.g., Kleinman et al. (2023) impose that some households may own capital and others may migrate, but none may do both.

⁶e.g. Bilal and Rossi-Hansberg (2023) impose that climate impacts under different scenarios are linear in scenario severity.

et al. (2020) into a spatial equilibrium model. Heterogeneous households make forward-looking migration, consumption-savings, and housing rent and ownership decisions which collectively determine housing prices in each location by intersection with location-specific housing supply curves.⁷

I calibrate the model to the 1713 locations covering the contiguous U.S. reported in the Public Use Microdata Sample (PUMS) of the 1990 U.S. Census.⁸ I calibrate location-specific amenity values, housing supply cost shifters, and productivity shifters for each location, and exactly match empirical populations, housing prices, and mean wages in each location. I simultaneously calibrate the non-spatially-varying parameters of the model so that the stationary steady state matches key moments of the 1990-1992 U.S. economy computed from the PUMS microdata and the Survey of Consumer Finances (SCF).⁹

Model Solution

I cleanly separate the model solution into (1) the solution to the household’s problem, which does not use neural networks; and (2) the global solution strategy, which uses neural networks but requires no information about the details of the household’s problem. The global solution is intractable with conventional methods due to a vast state space.¹⁰

To solve the complex household’s problem, I temporally decompose each period into simple sub-period “stages” which can be implemented and optimized separately. I develop specialized algorithms to speed up costly root-finding, interpolation, and discrete-choice operations which bottleneck¹¹ the performance of discrete-time dynamic heterogeneous-agent (HA) models in general.¹² These reusable, high-performance building blocks are quite flexible and can allow complex frontier HA models to be rapidly developed in a modular fashion. The solution of the steady state and of transition dynamics without aggregate uncertainty do not use neural networks at all. The perfect-foresight transition path, with over 200 million idiosyncratic states per period, solves in a matter of hours on a single midrange CPU.¹³

The global solution strategy treats the household’s problem as a black-box operator taking an end-of-period value function and a beginning-of-period state distribution to a beginning-of-period value function and an end-of-period state distribution.¹⁴ The problem of finding the global solution essentially reduces to the problem of approximating these end-of-period value functions, for which I use a neural network. This is basically a form of supervised learning, which is a robust capability of neural networks.

⁷Housing supply is in the form of location-specific durable housing stocks which depreciate over time but are added to by construction along the spatially-heterogeneous housing supply curves estimated by Saiz (2010).

⁸The locations I use are Public Use Microdata Areas (PUMAs), which are the units reported in the PUMS. They generally nest counties, which number 3109 in the contiguous U.S. (including county equivalents).

⁹I choose this relatively early year to allow for the role of anticipatory adjustment to climate expectations in shaping 2020 spatial populations, housing stocks, and asset holdings.

¹⁰The model contains over 200m household states, so that the aggregate state space has over 200m dimensions.

¹¹See, e.g., Greaney (2021, p. 4) who uses continuous-time methods to overcome the “costly root-finding and interpolation operations that would be necessary in discrete time.”

¹²The largest gains are from sharing information across solutions for individual households, reducing duplication of work.

¹³I use GPU acceleration for the global solution.

¹⁴As an optimization, I also use neural networks to generate guesses for prices.

Existing deep learning techniques typically also use a neural network to approximate the household policy function. This generally imposes restrictions on the policy function¹⁵ as well as practical difficulties in training its approximator. With the optimized household problem solution, it becomes viable to re-solve for optimal policies every time the model is simulated. As an added benefit, the policy function is correct from the beginning.¹⁶ I outline these computational methods in Appendix A and introduce them more fully in a companion paper.

Quantification

Climate Damages to Economic Fundamentals—I combine and harmonize estimates from the literature to quantify the impact of global climate conditions on labor productivity, amenities, residential energy costs, and disaster damages. The overall strategy is to combine projections from climate models—of the impact of global climate conditions on local climate conditions—with estimates from environmental economics on the impacts of local climate conditions on local economic fundamentals. The primary challenge is avoiding double-counting: for each damage estimate I use, it is essential that the methodology used to generate it controls for the effects picked up by the other estimates. Another challenge is that estimates are reported in terms of different local climate variables. In the global temperature range of interest, these estimates are well-approximated¹⁷ by linear (or log-linear) functional forms taking (1) global temperatures to local annual heating- and cooling-degree-days¹⁸ and (2) local annual heating- and cooling-degree-days to economic fundamentals. This allows the modeled sensitivities of local economic fundamentals to global climate conditions to be flexibly location-specific while remaining functionally simple.¹⁹

Expectations and Uncertainty—To aid in calibrating overall climate uncertainty, I employ a richer auxiliary model of climate expectations which can interpretably match moments representing many sources of climate uncertainty. The simpler climate process in the main model is calibrated to match moments of the auxiliary model that summarize overall climate uncertainty.²⁰

This auxiliary model consists of a richer stochastic climate process and a representative Bayesian household who observes realized climate outcomes but has imperfect information about the parameters of the generating process. The auxiliary model matches empirical moments or estimates corresponding to the three major sources of climate uncertainty defined by the Intergovernmental Panel on Climate Change (IPCC): internal variability (natural fluctuations in climate conditions), scenario uncertainty (uncertainty in the future path of emissions), and model uncertainty over the contribution of emissions to temperatures.²¹ It

¹⁵For instance, that the continuation value be quasiconvex and differentiable in the policy variable.

¹⁶Up to interpolation error and as a function of continuation values, which may not be correct.

¹⁷For disaster damages, data limitations preclude this check.

¹⁸Heating- and cooling-degree-days are a measure of how cold and hot a location is, respectively, relative to 18.3°C (65°F). See Section 3.2 for a formal definition. Disaster damages are done differently.

¹⁹I do not, however, fully (log-)linearize the model as in Caliendo et al. (2019) and Kleinman et al. (2023).

²⁰Specifically, the quartiles of 2100 global temperatures and disaster damages given information in 2020.

²¹We might also expect considerable uncertainty over the economic model governing housing markets and equilibrium re-

also matches model uncertainty in the contribution of emissions to disaster damages, which I estimate by applying a form of “nested” Kalman filter to historical time series data on homeowners insurance premiums and disaster-related payouts.²² Intuitively, I identify the insurer’s level of model uncertainty using the responsiveness of premiums to realized payouts.

Related Literature

This paper connects to five literatures. First, an active, recent dynamic spatial environmental economics literature seeks to understand heterogeneity in the welfare impacts of climate change. A key theme is understanding how spatial climate damages are shaped by different climate impacts and adaptation mechanisms. A seminal paper without dynamics, Hsiang et al. (2017) consider spatial climate impacts on disasters, mortality, productivity, energy costs, and crime, but do not incorporate economic or migration adaptation. Desmet et al. (2021) study the impact of land loss due to coastal flooding and find that adaptation through the dynamic responses of investment and migration largely ameliorate losses to global GDP. Krusell and Smith (2022) consider productivity impacts due to temperature and adjustment via capital mobility but not migration, finding that capital mobility does not significantly change results. Bilal and Rossi-Hansberg (2023) introduce anticipation, allowing forward-looking migration in anticipation of future climate impacts on capital depreciation. They find that migration substantially reduces spatial inequality in climate damages. I introduce within-location household heterogeneity, aggregate uncertainty, and housing wealth, as well as an approach for incorporating external estimates from on climate impacts to economic fundamentals.

Second, a literature on flood policy studies the role of government in subsidizing flood insurance and investing in flooding mitigation. These include Balboni (2021), Hsiao (2023), Pang and Sun (2023), and Fried (2022). Whereas these typically focus on a smaller area with relatively few locations and reduce the number of dimensions of within-location household heterogeneity to one or two,²³ My model incorporates five dimensions of within-location household heterogeneity and 1713 locations. This level of spatial resolution covers the contiguous U.S. while still allowing a study of flooding impacts with fine spatial resolution.²⁴

Third, a new dynamic spatial literature seeks to understand how forward-looking agents make frictional decisions and/or respond to shocks in spatial equilibrium. These include Crews (2023), Komissarova (2022), and Kleinman et al. (2023). A subset of these papers specifically incorporate housing wealth, such as Giannone et al. (2023), Greaney (2021), and Fried (2022). I draw from this literature methodologically and introduce local real estate and rental markets, uncertainty in the form of a stochastic transition, and a

sponses. This is beyond the scope of this study.

²²Essentially, I assume that the insurer is applying a Kalman filter, then use a Kalman filter of my own whose unobserved state embeds the insurer’s learning process and whose observables include the insurance premiums resulting from that process.

²³Fried (2022) has three dimensions of household heterogeneity, including homeownership, with the trade-off of only having two locations.

²⁴e.g. Miami-Dade County alone has 11 locations.

solution method for finding the global solution under aggregate uncertainty.

Fourth, an emerging literature develops deep learning tools to solve dynamic models with aggregate uncertainty, which in turn can be connected to the literature spawned by Krusell and Smith (1998) on the solution of dynamic models with aggregate uncertainty through state space reduction. This literature is somewhat split between papers which solve continuous-time models, such as Fernandez-Villaverde et al. (2020) and papers which solve discrete-time models, such as Azinovic et al. (2022), Han et al. (2021), and Maliar et al. (2021). I draw from those using discrete time. My primary contribution is to develop algorithms which overcome bottlenecks in the solution of the household’s policy, allowing a reduction in the role of neural nets in the model solution. It also allows a separation of the solution to the intra-period household’s problem and the deep learning global solution method.

Fifth, I draw from a housing literature which provides the foundation for understanding and modeling the household’s housing and homeownership choice, either alone or in equilibrium. This literature includes Sinai and Souleles (2005), Kaplan et al. (2020), and Favilukis et al. (2017). Yao and Zhang (2005) specifically study the relationship between wealth, portfolio choice, and asset returns. I draw from this literature. Like Giannone et al. (2023), I extend a version of Kaplan et al. (2020) from one location to many, simplifying mortgages and general-equilibrium feedbacks.

Outline

The paper is structured as follows. In Section 2, I describe a dynamic spatial equilibrium model with forward-looking heterogeneous households and housing wealth. In Sections 3, 4, and 5, I quantify the model. In Section 3, I quantify the impact of global climate conditions on local economic fundamentals (labor productivity, amenities, residential, energy costs, and disaster risk) by combining and harmonizing estimates from the climate science and environmental economics literatures. In Section 4, I quantify expectations and uncertainty over the future course of climate change using CMIP5 projection ensembles and time series data on historical homeowners insurance premiums. In Section 5, I calibrate the economic model to match spatial and aggregate data on the contiguous U.S. from the 1990 PUMS and 1992 SCF. In Section 6, I explore how household heterogeneity, housing wealth, and uncertainty shape spatial and household heterogeneity in the welfare impacts of climate change in the calibrated model. Section 7 concludes.

2 A Dynamic Spatial Equilibrium Model

2.1 Overview

I describe a heterogeneous-agent dynamic spatial model with housing wealth traded in spatially-segmented housing markets. The 1713 locations in the model exist within a small open economy representing the

contiguous United States. Locations have systematically heterogeneous exposure to a stochastic global climate process.

Time is discrete. The economy is populated by overlapping generations of heterogeneous households who differ, within generations, in their bondholdings, income type, real estate holdings, location, and age, and make forward-looking location, saving, and real estate investment decisions subject to a number of spatial and financial frictions. Each period, each household selects a single location to live and work in, paying monetary and utility migration costs if they change locations. Bilateral utility migration costs vary by origin-destination location pair.

Households may save through a risk-free foreign-issued bond with fixed and exogenous interest rate or by purchasing residential real estate. When buying real estate, households may obtain a mortgage in the form of negative bondholdings subject to a collateral constraint and a higher interest rate than is paid on savings. When selling real estate, households pay an adjustment cost. Housing assets are priced in equilibrium by households and are thereby exposed to location-specific climate risk which is neither insurable nor diversifiable across locations. Households may only own housing in the location in which they live and, conversely, all housing in each location is owned by households who live there.

In equilibrium, due to this illiquidity, risk, and borrowing constraint, residential real estate offers a higher expected return than the safe bond.

A household may let out some of its owned housing to other households in the same location, generating a stochastic cash flow. Conversely, households may rent housing from other households in the same location instead of owning. Homeowners pay lower maintenance costs, with the maintenance differential representing any administrative, moral hazard, and utility costs of renting.

Housing is durable and the housing stock depreciates over time. New housing is constructed by a competitive construction sector which is static, building new houses instantaneously and immediately selling them at the prevailing market price until it is unprofitable to do so. In effect, the quantity of housing in each location is given by the cost curve of construction if prices are sufficiently high or by the depreciated stock of existing housing if prices are sufficiently low.

Housing rent and real estate prices are determined in equilibrium in each location and period in spatial and intertemporal equilibrium. By contrast, the labor and goods markets are stylized. Production functions are linear, though productivity differs by location. Households supply labor inelastically and are paid their marginal product in the form of the consumption good, a numeraire costlessly tradable with the rest of the world. There is no capital.

Climate change introduces aggregate uncertainty. Global climate conditions affect local productivity, amenities, energy costs, and disaster risk heterogeneously across locations. These climate damages affect forward-looking household location and real estate investment choices, which collectively determine equilib-

rium housing prices.

Households are ex ante identical with constant relative risk aversion. However, absolute risk aversion varies endogenously, decreasing in consumption. Climate uncertainty thus affects households heterogeneously by wealth and age. In equilibrium, climate uncertainty reduces demand for owning housing, reducing the supply of rental housing and increasing housing rents relative to prices.

The full set of frictions that households in the model are subject to are: migration costs, real estate transaction costs, borrowing limits, borrowing costs, uninsurable income risk, and location-specific climate-driven uninsurable risk in fundamentals ²⁵ and housing prices.

2.1.1 Timing

Time is discrete. At the beginning of each period, housing prices and rents are observed. Immediately afterward, the housing stock in each location depreciates and a construction sector builds new housing until the marginal cost of new construction in each location is equal to the housing price. Households then draw location preference shocks and choose locations, housing assets, consumption and, if they do not own a home, rental housing services. At the end of each period, households receive income, pay expenses, realize stochastic income shocks, and age. Then the global climate state evolves and a global climate news shock is drawn which affects the future trajectory of the global climate. The climate news shock does not immediately affect the physical state of the climate in the subsequent period, which is determined before the news shock is drawn.

That is, each period t :

- Housing prices $q_{\ell t}$ and rents $\rho_{\ell t}$ are observed in each location ℓ .
- Housing quantities $H_{\ell t}$ are determined in each location ℓ by depreciation and construction.
- Each household i inherits state $x_{it} = (b_{it}, z_{it}, h_{it-1}^{\text{live}}, h_{it-1}^{\text{let}}, \ell_{it-1}, a_{it})$.
- The household draws location preference shocks $\{\varepsilon_{\ell it}\}_{\ell}$.
- The household chooses location ℓ_{it} , real estate assets h_{it}^{live} and h_{it}^{let} , goods consumption g_{it} , and rental housing h_{it}^{rent} .
- The household receives income—earnings, interest, and rent—and pays expenses—goods consumption, interest, rent, and maintenance.
- The household realizes next-period human capital z_{it+1} .
- Global mean temperature SST_t and disaster intensity $\delta_t^{\text{disaster}}$ evolve.

²⁵Productivity, amenities, energy costs, and disaster risk.

- A global climate news shock ε_t^m is drawn.

2.2 Households

Households live for a_{max} periods, each representing one decade. There are six non-overlapping age groups representing adult households, with the youngest aged 21-30 and the oldest aged 71-80.

Individual households, indexed by $i \in I$, are atomistic, with beginning-of-period state

$$x_{it} = (b_{it}, z_{it}, h_{it-1}^{\text{live}}, h_{it-1}^{\text{let}}, \ell_{it-1}, a_{it}),$$

consisting of bondholdings b_{it} , human capital z_{it} , owner-occupied real estate h_{it-1}^{live} , owned rental real estate²⁶ h_{it-1}^{let} , location ℓ_{it-1} , and age a_{it} . Beginning-of-period owner-occupied real estate h_{it-1}^{live} , possibly zero, is housing which was owned and occupied by the household in the previous period. Beginning-of-period rental real estate h_{it-1}^{let} , possibly zero, is housing which was owned by the household and let out to other households in the same location in the previous period. Housing is perfectly divisible so that h_{it-1}^{live} and h_{it-1}^{let} will be chosen from a continuum as in Favilukis et al. (2017).

2.2.1 Household Choices

Location

Within each period, a household must live, work, and own real estate in the same location. At the beginning of each period, a household realizes preference shocks over moving to each location, chooses a location, then pays migration costs if they move across locations.

At the beginning of each period t , a household i in location ℓ_{it-1} draws an additive preference shock $\varepsilon_{\ell it}$ for each location ℓ ,²⁷

$$\varepsilon_{\ell it} \sim \text{Gumbel}(0, \psi^{-1}) \quad \text{i.i.d.}$$

The preference shock $\varepsilon_{\ell it}$ is added to the household's overall period- t utility, where ℓ_{it} is the chosen location.

The preference shock does not persist into future periods.

²⁶Which the household lets out to other households in the same location.

²⁷Technically, this means that the household's action depends on both its state and its draw of ε (as well as the aggregate state Γ_t):

$$\ell_{it}(x_{it}, \{\varepsilon_{\ell it}\}_{\ell}; \Gamma_t).$$

In effect, however, the properties of the Gumbel distribution imply that, conditional on x_{it} and Γ_t , the probability that a household chooses a location is proportional to the (exponentiated) continuation value resulting from choosing that location. Furthermore, as the location preference shocks are additive and i.i.d., households effectively receive the mean preference-shock-contingent continuation value with certainty before the preference shock is drawn. This is made precise in Appendix B.1.

Housing Investment, Housing Consumption, Goods Consumption, and Savings

The household then chooses owned housing h_{it}^{live} and h_{it}^{let} , location ℓ_{it} , goods consumption g_{it} , and rental housing h_{it}^{rent} :

$$\text{choice}_{it} = \{h_{it}^{\text{live}}, h_{it}^{\text{let}}, \ell_{it}, g_{it}, h_{it}^{\text{rent}}\}.$$

In contrast to a model of hand-to-mouth households, they need not spend all of their wealth each period. Any wealth not spent is saved in the form of a risk-free bond.

The household cannot occupy both a rented home and an owned one, so

$$h_{it}^{\text{live}} h_{it}^{\text{rent}} = 0.$$

2.2.2 Spatial and Financial Frictions

Households are subject to migration costs, an owned housing adjustment cost, borrowing costs, and a borrowing constraint.²⁸

Migration Costs

If a household's new location, ℓ_{it} , differs from its previous location, ℓ_{it-1} , they pay a fixed monetary²⁹ cost F^m and a utility cost proportional to the distance $d(\ell_{it-1}, \ell_{it})$ between the two locations:

$$F^u(\ell_{it-1}, \ell_{it}) = \tau d(\ell_{it-1}, \ell_{it}).$$

Housing Transaction Cost

If the household adjusts owned housing assets, either occupied, h_{it-1}^{live} , or let out, h_{it-1}^{let} , they pay a percentage of the sale price, $q_{\ell_{it-1}t} h_{it-1}^s$, as a realtor's fee:³⁰

$$\text{realtor's fee} = \phi q_{\ell_{it-1}t} h_{it-1}^s.$$

²⁸Households are also subject to uninsurable idiosyncratic income risk and aggregate climate risk to which they differ in their exposure. I defer the exposition of these risks to their respective sections.

²⁹Technically this is not a monetary economy and F^m is in terms of the numeraire good.

³⁰The relevant housing price is the *current* price in the *old* location, $q_{\ell_{it-1}t}$.

Borrowing Costs and Borrowing Constraint

After adjusting their location and housing assets (but before collecting income or paying other expenses), a household's new bondholdings are given by,

$$\tilde{b}_{it} = b_{it} - F^m \mathbb{1}_{it}^{\text{move}} + \sum_{s \in \{\text{live}, \text{let}\}} [(1 - \mathbb{1}^{\text{sell}, s} \phi) q_{\ell_{it-1}t} h_{it-1}^s - q_{\ell_{it}t} h_{it}^s],$$

$$\text{where } \mathbb{1}^{\text{move}} = \mathbb{1}[\ell_{it} \neq \ell_{it-1}]$$

Because households may only own housing in one location, they must sell all owned housing if they move across locations:

$$\mathbb{1}^{\text{sell}, s} = \mathbb{1}^{\text{move}} \vee \mathbb{1}[h_{it-1}^s \neq h_{it}^s] \quad \forall s \in \{\text{live}, \text{let}\}.$$

Bondholdings may be negative, representing borrowing. However, this borrowing must be collateralized by owned housing subject to a loan-to-value limit,

$$b_{it+1} \geq -\kappa q_{it\ell} (h^{\text{live}} + h^{\text{let}}),$$

and the interest rate which the household must pay when borrowing, r^m , is higher than the interest rate it receives when saving, r^f :

$$r^m > r^f.$$

2.2.3 Income

At the end of each period t , a household i living in location ℓ receives labor earnings equal to their marginal product under a linear production function with idiosyncratic component z_{it} and climate-impacted³¹ location-specific component $A_{\ell t}$. They receive interest income and rental income on post-adjustment liquid bondholdings, \tilde{b}_{it} , and rental real estate, h_{it}^{let} , respectively.³²

Given location-level productivity $A_{\ell t}$ and rent $\rho_{\ell t}$, the household receives income,

$$\text{Income}_{it} = \underbrace{A_{\ell t} z_{it}}_{\text{earnings}} + \underbrace{r^f \tilde{b}_{it} \mathbb{1}[\tilde{b}_{it} \geq 0]}_{\text{interest income}} + \underbrace{\rho_{\ell t} h_{it}^{\text{let}}}_{\text{rental income}}.$$

³¹Housing prices $q_{\ell t}$ and rents $\rho_{\ell t}$ are indirectly impacted by the climate state and expectations in equilibrium, but are not directly impacted by the climate state as productivity $A_{\ell t}$ is.

³²Note that this aspect of the model's timing is slightly unusual: households receive investment income within the same period that the investment is chosen. I choose this timing so that (1) renters, homeowners, and investors make housing decisions simultaneously (at the beginning of the period) and (2) rent paid by renters is received by landlords within (at the end of) the same period. To do otherwise imposes an unrealistic cost on landlords, who would otherwise have to wait ten years between the consumption foregone by purchasing a rental property and the consumption enjoyed out of rental income.

2.2.4 Expenditures

If the household rents, they pay rent at rate $\rho_{\ell t}$. If they own housing, they pay maintenance at rate $\chi_{\ell t}^{\text{live}}$ or $\chi_{\ell t}^{\text{let}}$ for owner-occupied or rental housing, respectively. Maintenance consists of a fixed cost depending only on whether the home is owner-occupied and two location-specific climate-impacted variables: energy costs $y_{\ell t}$ and disaster intensity $\delta_{\ell t}^{\text{disaster}}$:

$$\chi_{\ell t}^{\text{live}} = \chi^{\text{live}} + y_{\ell t} + \delta_{\ell t}^{\text{disaster}}$$

$$\chi_{\ell t}^{\text{let}} = \chi^{\text{let}} + y_{\ell t} + \delta_{\ell t}^{\text{disaster}}$$

where $\chi^{\text{live}} < \chi^{\text{let}}$.

Given location-level energy costs $y_{\ell t}$, disaster intensity $\delta_{\ell t}^{\text{disaster}}$, and rent $\rho_{\ell t}$, total expenditures are,

$$\text{Expenditures}_{it} = \underbrace{g_{it}}_{\text{goods}} + \underbrace{\rho_{\ell t} h_{it}^{\text{rent}}}_{\text{rent}} + \underbrace{\chi_{\ell t}^{\text{live}} h_{it}^{\text{live}} + \chi_{\ell t}^{\text{let}} h_{it}^{\text{let}}}_{\text{maintenance}} + \underbrace{r^m \tilde{b}_{it} \mathbb{1}[\tilde{b}_{it} < 0]}_{\text{interest expenses}}.$$

The maintenance cost $\chi^{\text{let}} > \chi^{\text{live}}$ is higher for rental properties and summarizes a number of reasons why owner-occupied housing might provide a lower user cost of housing services than rental housing for an unconstrained household, in equilibrium, before frictions are taken into account. These include the higher utility value of owner-occupied housing, the maintenance and administrative requirements of maintaining a rental property, and income tax on rental income.

Without this wedge, households would have no reason to own their home as homeownership would only impose borrowing constraints, illiquidity, and risk for no benefit. Without these frictions, however, households would have no reason to rent, as homeownership would provide unambiguously cheaper housing services. With these frictions, homeownership is more attractive in equilibrium for wealthier households who, needing to borrow less to buy housing, pay a lower effective user cost than less wealthy households for housing services from owned housing. Similarly, in equilibrium, rental housing investments provide higher expected returns than liquid assets at the cost of illiquidity and risk.

2.2.5 Idiosyncratic Shocks

At the end of the period, the household's income type evolves according to,

$$\log z_{it+1} = p \log z_{it} + \zeta_{a_{it+1}} + \nu_{it}$$

$$\nu_{it} \sim \mathcal{N}(0, \sigma_\nu) \quad \text{i.i.d.},$$

where $p \in [0, 1]$ is mean reversion and ζ_a is an age-specific shifter.³³

2.2.6 Utility

Similarly to Kaplan et al., 2020, a household i in location ℓ derives utility from goods g_{it} , housing services h_{it} , and location-level amenities $\alpha_{\ell t}$, with constant elasticity of substitution $1/\sigma$ between goods and housing services and constant elasticity of intertemporal substitution $1/\eta$. Location-specific and climate-impacted amenities $\alpha_{\ell t}$ enter as an effective multiplier on consumption spending. Formally,

$$u(g, h, \alpha) = \frac{\left(\alpha (g^{1-\sigma} + \gamma h^{1-\sigma})^{\frac{1}{1-\sigma}}\right)^{1-\eta} - 1}{1-\eta}.$$

Households i who have reached the maximum age $a_{it} = a_{\max}$ enjoy “warm-glow” bequest utility from the liquidation value of any assets which they own at the end of the period.³⁴

$$\begin{aligned} V_{a_{\max}+1,t}(b, z, h^{\text{live}}, h^{\text{let}}, \ell) &= Q \frac{\tilde{b}^{1-\eta}}{1-\eta} \\ \text{s.t. } \tilde{b} &= b + (1-\phi)(q_{\ell t} h^{\text{live}} + q_{\ell t} h^{\text{let}}). \end{aligned}$$

The full household’s problem is given in recursive form in Appendix B.

2.2.7 Household Entry

Each period, a new cohort, representing age 20-29, enters. Each household i in the new cohort draws individual state

$$x_{it} = (b_{it}, z_{it}, 0, 0, \ell_{it-1}, 1),$$

where $h_{it-1}^{\text{live}} = h_{it-1}^{\text{let}} = 0$ and $a_{it} = 1$. Bondholdings b_{it} and human capital z_{it} are drawn from an initial distribution

$$(b_{it}, z_{it} \mid a_{it} = 1) \sim \Lambda^{\text{init}}.$$

Initial location ℓ_{it-1} is equal to the population distribution of households in the fourth age group, representing age 50:

$$\Pr_t(\ell_{it} \mid a_{it} = 1) \sim \int_{x \in \mathbf{X}} \mathbb{1}[\ell(x; \Gamma_t) = \ell \wedge a(x; \Gamma_t) = 4] d\lambda_t(x),$$

³³This is isomorphic to the idiosyncratic human capital process of Kaplan et al. (2020), except that I pull out the aggregate productivity shifter $A_{\ell t}$ into the income equation and make it location-specific.

³⁴This model does not consider the general equilibrium implications of these bequests (or indeed of any goods consumption). In particular, there is no explicit intergenerational wealth transmission and a reduction in end-of-life saving would not feed back into lower initial wealth for new households.

where $\lambda_t : \mathcal{P}(\mathbf{X}) \rightarrow \mathbb{R}$ is the measure over individual states $x \in \mathbf{X}$ in period t , and Γ_t is the aggregate state in period t .

2.3 Housing

At each time t , the stock of housing $H_{\ell t} \in \mathbb{R}^+$ in each location ℓ is a state variable. Housing assets are traded in local residential real estate markets and rented in local housing rental markets. Markets are spatially segmented and all housing must be owned by households in the same location. A household may only own or rent housing in the location in which they live.

Housing depreciates over time and new housing is constructed by a static, competitive construction sector which produces using a location-specific housing production function.

2.3.1 Residential Real Estate Market

Let $\lambda_t : \mathcal{P}(\mathbf{X}) \rightarrow \mathbb{R}$ be the measure over individual states $x \in \mathbf{X}$, and Γ_t the aggregate state, in period t .

Demand In any location ℓ and time t , the stock of housing $H_{\ell t}$ is entirely owned by households who live in that location. Denote demand for residential real estate by³⁵

$$H_{\ell t}^D(q) \equiv \int_{x \in \mathbf{X}} (h^{\text{live}}(x; q, \Gamma_t) + h^{\text{let}}(x; q, \Gamma_t)) \mathbb{1}[\ell(x; q, \Gamma_t) = \ell] d\lambda_t(x).$$

Supply Before real estate is traded, a competitive construction sector constructs a quantity $(H_{\ell t} - \tilde{H}_{\ell t})$ of new housing, where $\tilde{H}_{\ell t} \equiv (1 - \delta^h)H_{\ell t-1}$ is the quantity of housing carried over from the previous period, and then sells it onto the market. The construction sector observes prices then builds instantaneously and immediately sells at those prices, making its decision static. The marginal cost of housing construction is a location-specific convex function of the (new) housing stock $H_{\ell t}$, as in Greaney (2021) and Saiz (2010).

$$c_\ell(H) = \Pi_\ell H^{\beta_\ell}.$$

I assume that the construction sector is competitive, so that if $H_{\ell t} > \tilde{H}_{\ell t}$, then $q_{\ell t} = c_\ell(H_{\ell t})$. Housing supply is thus given by

$$H_{\ell t}^S(q) = \begin{cases} c_\ell^{-1}(q) & q \geq c_\ell(\tilde{H}_{\ell t}) \\ \tilde{H}_{\ell t} & q < c_\ell(\tilde{H}_{\ell t}) \end{cases} \quad (1)$$

where $H_{\ell t}^D(q)$ denotes housing demand given housing price q .

³⁵I omit the dependence of choice variables on housing rent $\rho_{\ell t}$ in this subsection to focus on the dependence on housing price $q_{\ell t}$. In a sense, $\rho_{\ell t}$ is a function of Γ_t anyway.

The equilibrium condition for each residential real estate market is

$$H_{\ell t}^S(q) = H_{\ell t}^D(q). \quad (2)$$

2.3.2 Housing Rental Market

In each location ℓ , demand for rental housing is given by³⁶

$$H_{\ell t}^{\text{rent}}(\rho) = \int_{x \in \mathbf{X}} (h^{\text{rent}}(x; \rho, \Gamma_t)) \mathbb{1}[\ell(x; \rho, \Gamma_t) = \ell] d\lambda_t(x)$$

and demand for rental real estate is given by

$$H_{\ell t}^{\text{let}}(\rho) = \int_{x \in \mathbf{X}} (h^{\text{let}}(x; \rho, \Gamma_t)) \mathbb{1}[\ell(x; \rho, \Gamma_t) = \ell] d\lambda_t(x).$$

Rental housing markets are also segmented by location. The equilibrium condition for each rental housing market is

$$H_{\ell t}^{\text{rent}}(\rho) = H_{\ell t}^{\text{let}}(\rho). \quad (3)$$

2.4 Climate

Locations are heterogeneously exposed to global climate conditions, affecting local productivity $A_{\ell t}$, amenities $\alpha_{\ell t}$, residential energy costs $y_{\ell t}$, and disaster damages $\delta_{\ell t}$. Global climate conditions evolve according to a stylized process whose purpose is to match key moments of climate uncertainty and to allow for a news shock affecting the future global climate trajectory without affecting the current physical climate state.³⁷

2.4.1 Productivity, Amenities, and Energy Costs

Productivity $A_{\ell t}$, amenities $\alpha_{\ell t}$, and residential energy costs $y_{\ell t}$ are fully flexible by location both in their level and in their response to global mean temperature SST_t . I linearize (or log-linearize) this relationship:

$$\log A_{\ell t} = \log \bar{A}_{\ell} + A_{\ell}^g \text{SST}_t$$

$$\log \alpha_{\ell t} = \log \bar{\alpha}_{\ell} + \alpha_{\ell}^g \text{SST}_t$$

$$y_{\ell t} = \bar{y}_{\ell} + y_{\ell}^g \text{SST}_t$$

I argue in Section 3 that this linearization is reasonable around the $\sim 4^{\circ}\text{C}$ range of global temperatures

³⁶I omit the dependence of choice variables on housing price $q_{\ell t}$ in this subsection to focus on the dependence on housing rent $\rho_{\ell t}$. In a sense, $q_{\ell t}$ is a function of Γ_t anyway.

³⁷This stylized process will be calibrated to match moments generated by a richer, quantified model of climate change and Bayesian climate expectations in Section 4.

of interest, for both the sensitivity of local temperatures to global temperatures and the sensitivity of these economic fundamentals to local temperatures.³⁸

In each location ℓ and period t , local disaster intensity (damages per unit of housing) $\delta_{\ell t}^{\text{disaster}}$ is a (temporally-constant, spatially-varying) multiple of the aggregate disaster intensity $\delta_t^{\text{disaster}}$ for that period:

$$\delta_{\ell t}^{\text{disaster}} = \delta_{\ell}^g \delta_t^{\text{disaster}}.$$

2.4.2 Global Climate State

The global climate state evolves according to the following stylized process. This is an approximant to the richer climate process which I take to the data, described in Section 4.

Each period t , the global climate state is given by

$$\mathbf{D}_t = \left\{ w_t, m_t, \underbrace{\text{SST}_t}_{\text{Global Temperature}}, \underbrace{\delta_t^{\text{disaster}}}_{\text{Global Disaster Intensity}} \right\},$$

where w_t and m_t are state variables loosely corresponding to global energy generation and emissions intensity, respectively, though in their calibration they will be effectively adjusted to account for scenario uncertainty and model uncertainty.

An annual index of global anthropogenic climate forcing rates (loosely representing annual emissions) e_t is constructed as the product of a factor w_t that grows at a constant rate and a stochastic factor m_t that shrinks in expectation faster than m_t grows, so that e_t starts near zero, grows once w_t becomes sufficiently large, then almost surely eventually converges to zero as m_t extinguishes w_t . I am interested in the region of the process, away from the stationary distribution, where e_t is still positive. Formally,³⁹

$$e_t = w_t m_t \tag{4}$$

$$\log w_{t+1} = \log w_t + g^w$$

$$\log m_{t+1} = \log m_t - \beta_m w_t + \varepsilon_t^m$$

$$\varepsilon_t^m \sim \mathcal{N}(0, \sigma_m^2).$$

Global temperature SST_t and disaster intensity δ_t (expressed as the deviation from pre-industrial levels⁴⁰)

³⁸When local temperatures are measured in heating- and cooling-degree-days.

³⁹Note that this particular process does not admit simulation from the infinitely distant past because m_t is nonstationary. This is easily solved by substituting $\log m_{t+1} = \log m_t - (\rho_m + \varepsilon_t^m)w_t$, but this is unnecessary for the purposes of this study.

⁴⁰Disaster damages at the pre-industrial baseline level can be considered a form of maintenance, whose spatial variation can be approximately soaked up by $\Pi_{\ell t}$.

follow an autoregressive process impacted by forcings e_t :

$$\begin{aligned} \text{SST}_{t+1} &= \rho_c \text{SST}_t + \beta_{\text{SST}} e_t \\ \log \delta_{t+1} &= \rho_\delta \log \delta_t + \beta_\delta e_t. \end{aligned}$$

The evolution of m_t is stochastic and summarizes all uncertainty about the future trajectory of the global climate. Note that ε_t^m represents a persistent shock to the growth rate of temperatures and disaster damages and impacts the future trajectory of the global climate state but not its current physical state.

2.5 Recursive Stochastic Equilibrium

2.5.1 State Space

I now drop time t subscripts and define the evolution of the aggregate state in recursive form.

The beginning-of-period aggregate state Γ is,

$$\Gamma = (\lambda, \mathbf{H}, \mathbf{D}),$$

where $\lambda : \mathcal{P}(\mathbf{X}) \rightarrow \mathbb{R}$ is the measure over individual states $x \in \mathbf{X}$, $\mathbf{H} = \{\tilde{H}_\ell\}_\ell$ is the set of beginning-of-period location-level housing stocks, and \mathbf{D} is the aggregate climate state.

2.5.2 Equilibrium

A recursive stochastic equilibrium is

- A set of household policy and value functions solving Equation (11),
- The aggregate transition function

$$\Omega(\Gamma' | \Gamma),$$

over

$$\Gamma = \{\lambda, \mathbf{H}, \mathbf{D}\}$$

composed of the transition functions over λ , \mathbf{H} , and \mathbf{D} (formally defined by Equation (12) in Appendix B, and Equations (1) and (4), respectively), and

- Prices $q_\ell(\Gamma)$, $\rho_\ell(\Gamma)$ satisfying the market clearing conditions in Equations (2) and (3), respectively.

2.6 Shock Concept

I consider an economy subject to two types of shock: (1) a shock to an initial stationary steady state corresponding to the unexpected onset of a stochastic climate transition and (2) subsequent per-period news shocks from an anticipated distribution about the future trajectory of the global climate.

First, I consider an economy in an initial stationary steady state which unexpectedly begins to experience climate change. Second, as soon as the stochastic climate process begins, households correctly anticipate that it will continue according to the stochastic climate process described in Section 2.4.2.

Note that climate change is highly persistent but not permanent. Climate forcings e_t almost surely eventually converge to zero, as do, consequently, all temperature and disaster intensity deviations-from-baseline,

$$\begin{aligned}\lim_{t \rightarrow \infty} e_t &= 0 \\ \lim_{t \rightarrow \infty} \text{SST}_t &= 0 \\ \lim_{t \rightarrow \infty} \delta_t &= 0,\end{aligned}$$

returning the economy to the stationary steady state.

I use this stationary steady state as the initial aggregate state which experiences an unexpected climate transition. Specifically, I consider an initial stationary steady state in which $m_t = m = 0$. (The value of $w_t \in \mathbb{R}$ affects nothing.) After the shock, the values (m, w) are then unexpectedly set to (m_0, w_0) , which I calibrate in Section 4.

3 Climate Damages

I spend the next three sections quantifying the model. In this section, Section 3, I combine and harmonize estimates from the climate science and environmental economics literatures to quantify the impact of global climate conditions on local economic fundamentals: productivity, amenities, energy costs, and disaster risks.

In Section 4, I calibrate the global climate process. I do this by constructing and quantifying a richer model of stochastic global climate change and Bayesian climate expectations, then calibrating the simpler climate process in the economic model to match key moments of the richer, quantified model.

In Section 5, I calibrate the remaining parameters of my economic model using data from the Survey of Consumer Finances and the Public Use Microdata Sample of the U.S. Census.

3.1 Global Climate Conditions \rightarrow Local Economic Fundamentals

To calibrate the impact of global climate conditions on local economic fundamentals, I combine estimates of the impact of global climate conditions on local climate conditions with estimates of the impact of local climate conditions on local economic fundamentals:

Global Climate Conditions \rightarrow Local Climate Conditions \rightarrow Local Economic Fundamentals.

As my spatial unit, I use the Public Use Microdata Areas (PUMAs), which are the smallest spatial unit available in publicly available U.S. Census microdata.⁴¹

Available in the literature are estimates of local climate conditions under specific climate scenarios and estimates of the effects of local temperature metrics on local economic fundamentals. Furthermore, estimates of different local economic impacts taken from different studies may overlap in terms of the effects they pick up.

To generalize estimates of local temperature responses to global climate conditions, I linearize between estimates for 2000 and 2080 under RCP4.5.

To combine and generalize estimates of the response of local economic fundamentals to local temperature metrics, I compile a set of studies whose methodologies, where applicable, each control for the effects measured by the other studies. Where available, I use nonparametric or overparameterized estimates and show that the effects measured are approximately linear in the temperature metric I choose. I therefore argue that I am able to linearize all the way through, obtaining reasonably linearized estimates of the responses of local economic fundamentals to global climate conditions.⁴²

Explicitly, I allow local labor productivity $A_{\ell t}$, amenities $\alpha_{\ell t}$, residential energy costs $y_{\ell t}$, and disaster damages to housing $\delta_{\ell t}$, to depend on global climate conditions. I allow these variables to be fully-flexibly location-specific in their levels and sensitivity to global climate conditions:⁴³

$$\log A_{\ell t} = \log \bar{A}_{\ell} + A_{\ell}^g \text{SST}_t$$

$$\log \alpha_{\ell t} = \log \bar{\alpha}_{\ell} + \alpha_{\ell}^g \text{SST}_t$$

$$y_{\ell t} = \bar{y}_{\ell} + y_{\ell}^g \text{SST}_t$$

$$\delta_{\ell t} = \delta_{\ell}^g \bar{\delta}_t.$$

I define $A_{\ell t}$ as a multiplier on the effective quantity of labor supplied by each household in a location. I define $\alpha_{\ell t}$ as a multiplier on the effective quantity of consumption enjoyed by each household in a location.

⁴¹In particular, I use the 1990 PUMA boundaries used in the 1990 U.S. Census Public Use Microdata Sample.

⁴²For impacts on disaster damages, data limitations preclude this check.

⁴³The intercept for $\delta_{\ell t}$ is zero because $\delta_{\ell t}$ and $\bar{\delta}_t$ are defined as the deviation from the same pre-industrial baseline.

I define $y_{\ell t}$ as the overall cost of energy required to heat and cool a typical owner-occupied two-bedroom in location ℓ over the course of period t . I define $\delta_{\ell t}$ as the excess disaster damages due to climate change suffered by a typical owner-occupied two-bedroom housing unit in location ℓ in period t .

In Section 3.4, I report the general magnitudes and economic and spatial heterogeneity of these estimated climate damages, by naively dollarizing and summing them for each household in the 2020 American Community Survey (ACS). By far the most severely affected individual locations are those exposed to severe flooding, primarily along the East Coast and in Central Appalachia. However, a far larger area is exposed to temperature changes, such that overall damages due to temperature are likely much larger than those due to flooding.⁴⁴

3.2 Global Climate Conditions \rightarrow Local Climate Conditions

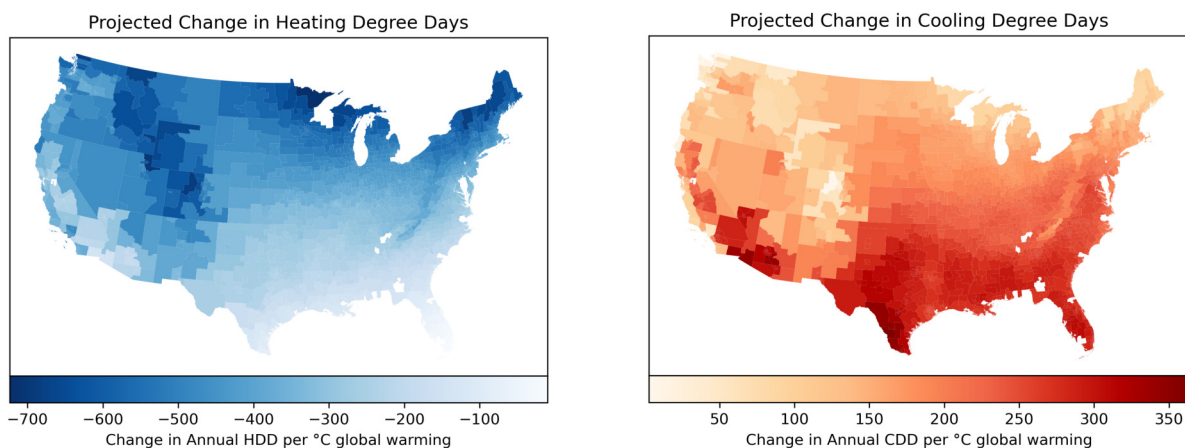
I use climate projections from the Coupled Model Intercomparison Project 5 (CMIP5) to estimate the impacts of global climate conditions on local climate conditions. The relevant measures of local climate conditions that I use are heating degree days (HDD) and cooling degree days (CDD).

The number of annual heating (cooling) degree-days in a location is the sum over the days in a year of the number of degrees by which daily temperatures are below (above) 18.3°C (65°F). For location i , year t , days d ,

$$\begin{aligned} \text{HDD}_{it} &= \sum_{d \in D_t} \max(18.3 - T_{idt}, 0) \\ \text{CDD}_{it} &= \sum_{d \in D_t} \max(T_{idt} - 18.3, 0). \end{aligned}$$

For each location, I linearize between CMIP5 ensemble model median projections of 2000 HDDs and 2080 HDDs under RCP8.5 to obtain the projected change in HDD and CDD per °C global warming. I do the same for CDD. Figure 1 displays the results.

⁴⁴The comparison of overall magnitudes at the end of this section is only meant to be descriptive and does not take into account the behavioral and price responses whose quantification is the purpose of the full model.



(a) Projected change in heating degree days per $^{\circ}\text{C}$ global warming.

(b) Projected change in cooling degree days, per $^{\circ}\text{C}$ global warming.

Figure 1

Projected change in heating and cooling degree days (HDD and CDD) per $^{\circ}\text{C}$ global warming. I use CMIP5 ensemble median estimates of projected HDD and CDD under RCP8.5 for 2000 and 2080 (Gassert et al., 2021). I average over each location for each scenario. For each location, I divide the projected 2020-2080 increase in HDD (CDD) by the corresponding CMIP5 ensemble median projected change in global sea surface temperatures (SST) under RCP8.5 (Hsiang et al., 2017). Locations are Public Use Microdata Areas, 1990 boundaries.

3.3 Local Climate Conditions \rightarrow Local Economic Fundamentals

I compile and harmonize estimates from the environmental economics literature on the impact of local climate conditions on local economic fundamentals: productivity, amenities, energy costs, and disaster risks. For each impact, I use estimates from a study which, where applicable, controls in its methodology for effects picked up by each other study. This is necessary to prevent double-counting.

3.3.1 Labor Productivity

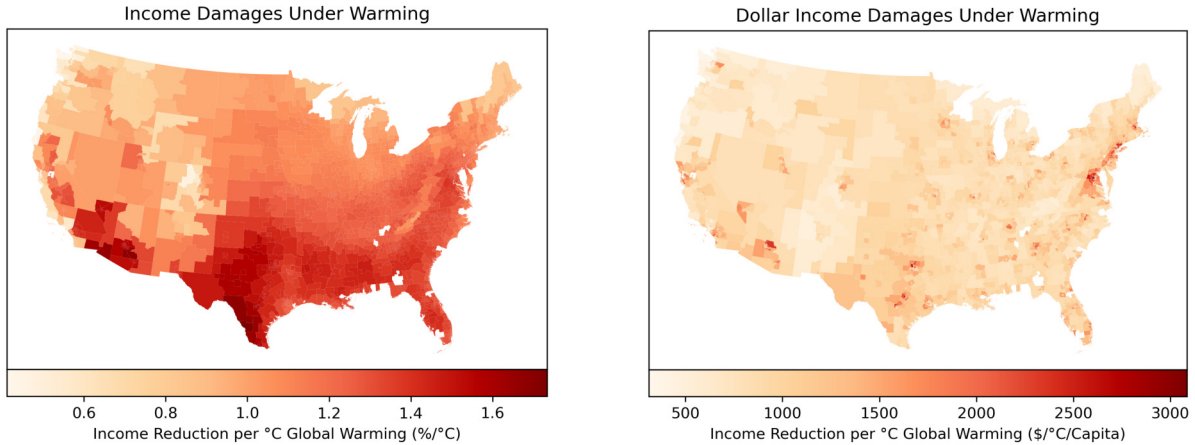
I take estimates of the effect of temperature on income from Deryugina and Hsiang (2017). They use spatial panel data on weather variation and income for U.S. counties in the period 1969-2014 to estimate the effect of temperature on income.⁴⁵

⁴⁵See Appendix Section C.1 for details.

Parameter	Productivity Effect per DD (% of Annual Income)	Productivity Effect per 365 DD (% of Annual Income)
HDD	6.57×10^{-6}	0.240
CDD	4.50×10^{-5}	1.66

Table 1: Productivity Damages per Heating and Cooling-Degree Day

Productivity damages per single ($^{\circ}\text{C}$) degree-day, and per 365 degree-days, relative to the 18.3°C (65°F) baseline, as a percentage of annual income. The second column, representing the effect of 365 additional CDDs (HDDs), is the implied marginal damage (benefit) of a one-degree increase in temperature for each day of the year, in a location where every day is already above (below) 18.3°C .



(a) Percentage reduction in productivity per capita, per $^{\circ}\text{C}$ global warming. Estimates are from Deryugina and Hsiang (2017). I linearize and multiply their estimates by the estimated change in HDD and CDD per $^{\circ}\text{C}$ global warming described in Figure 1. Further details are in Appendix Section C.1.

(b) Naively dollarized reduction in productivity per capita, per $^{\circ}\text{C}$ warming. I merely multiply the percentage reduction in productivity reported in Figure 2a by observed annual wages for each household in the 2020 ACS and mean across households by location.

Figure 2: Productivity Damages per $^{\circ}\text{C}$ Global Warming

3.3.2 Heat and Cold Disamenities

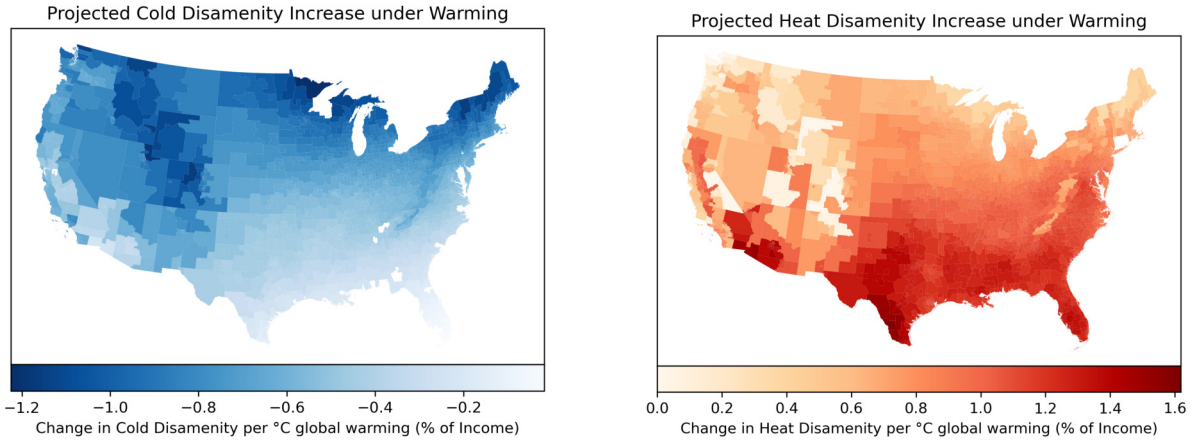
I take estimates of heat and cold disamenities from Albouy et al. (2016). The authors estimate the largest percentage of one’s income one would be willing to forego in order to avoid one marginal day at a certain temperature. I introduce this damage into my model as an effective multiplier on consumption spending in utility. As described in Appendix C.3, their estimates are close to log-linear in the *difference* between the daily temperature and 18.5°C (65°F), suggesting that they can be well-approximated by a linear function of HDD and CDD. I report these estimated coefficients on HDD and CDD in Table 2.

Figure 3 shows the implied change in heat and cold disamenities per $^{\circ}\text{C}$ global warming. Figure 4 shows the implied net change in location-level amenities per $^{\circ}\text{C}$ global warming.

Parameter	WTP per DD (% of Income)	WTP per 365 DD (% of Income)
HDD	3.05×10^{-3}	1.11
CDD	4.40×10^{-3}	1.61

Table 2

Amenity Damages per Heating- and Cooling-Degree Day. The first column represents the willingness-to-pay to avoid a single ($^{\circ}\text{C}$) degree-day, relative to the 65°F baseline, as a percentage of annual income. The second column, representing the willingness to pay to avoid 365 CDD (HDD), is the marginal damage (benefit) of a one-degree increase in temperature for each day of the year, in a location where every day is already above (below) 18.3°C .



(a) Projected change in cold disamenities per $^{\circ}\text{C}$ global warming.

(b) Projected change in heat disamenities per $^{\circ}\text{C}$ global warming.

Figure 3

Heat and Cold Disamenities per $^{\circ}\text{C}$ Global Warming. Panel (3a) reports the estimated maximum percentage of income an average household in a location would be willing to forego in order to avoid the change in local amenities resulting from 1°C global warming. Estimates are obtained by, first, linearizing the estimates of Albouy et al., 2016 of the amenity impacts of one additional annual cooling-degree-day (the results of this linearization are reported in Table 2), then multiplying by the estimates of the change in local cooling-degree-days induced by 1°C global warming reported in Figure 1b. Panel (3b) reports the same, but for heating-degree-days.

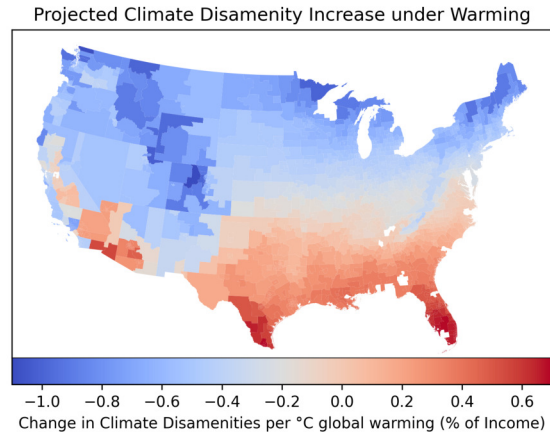
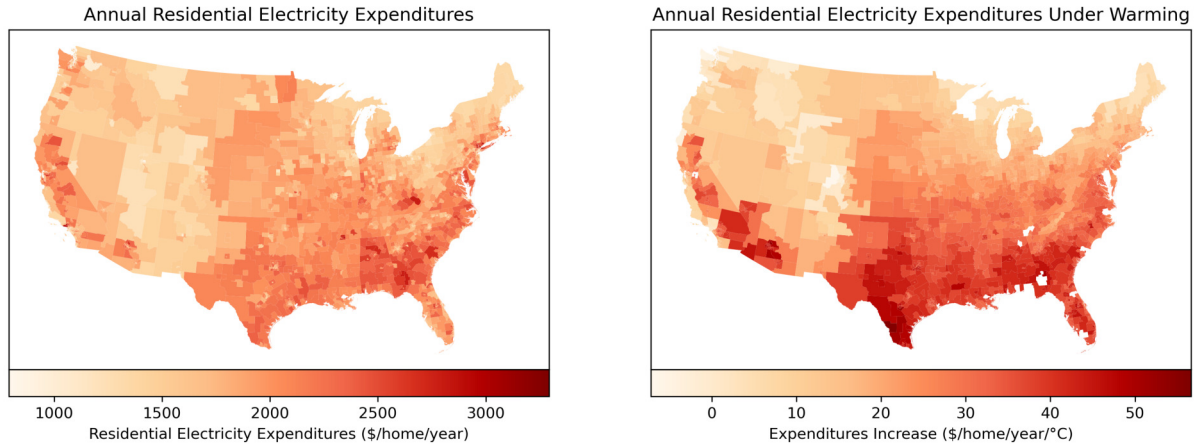


Figure 4: Net Climate Amenity Damages per °C Global Warming

Net projected change in cold and heat disamenities, as an equivalent percentage of annual income, per °C global warming. Estimates are obtained by adding the estimated local cold disamenity and heat disamenity impacts of 1°C global warming, reported in Figure 3.

3.3.3 Energy Costs

I take projections of the effect of warming on energy costs from Rode et al. (2021). Figure 5a describes the current levels of residential electricity consumption per capita in the ACS. Figure 5b describes the projected change in energy expenditure per capita per °C global warming.



(a) Annual residential electricity expenditure per capita, 2021.

(b) Projected change in residential electricity expenditure per capita, per °C global warming.

Figure 5

Residential energy expenditure damages per °C global warming. Estimates are obtained by dividing projected percentage electricity consumption growth between 2020 and 2100 under RCP8.5 for each county from Rode et al. (2021) by the ensemble median projected increase in mean global temperatures in that period and scenario from Rasmussen and Kopp (2017). I then apply a population-based crosswalk to PUMAs and multiplying by residential electricity expenditure per capita from the 2020 ACS.

3.3.4 Disaster Damages

I use projections of climate impacts on property damages due to flooding from Bates et al. (2021), as reported in Wing et al. (2022). The authors use a high-resolution inundation model to compute comprehensive coastal, fluvial, and pluvial flood risk. Theirs is the first large-scale inundation model to incorporate all three forms of flooding.

Flooding represents the large majority of projected climate damages through natural disasters. Figure 6 shows the relative property damages from flooding between 1975 and 2020 by FEMA hazard type from the Spatial Hazard Events and Losses Database for the United States (SHELDUS). Damages which involve both storms and flooding are split evenly between the relevant flooding-related category and storm-related category. This may thus represent a significant underestimate of total damages from flooding, as flooding damages often occur in conjunction with storms. This leaves wind-related damages and wildfires as quantitatively large, plausibly climate-impacted hazards. However, recent climate literature projects only small (on average) and highly uncertain impacts of climate change on wind-related damages. Wildfires remain unaccounted for, but I nevertheless conclude that the comprehensive flooding damage projections of Bates et al. (2021) likely capture a substantial share of overall excess disaster damages due to climate change.⁴⁶

Figure 7a displays their model-predicted Average Annual Losses (AAL), as a percentage of residential

⁴⁶See Appendix Section C.2 for details.

structure value, under 2020 climate conditions. Figure 7b shows projected flooding damage growth, as a percentage of current AAL, per °C global warming.


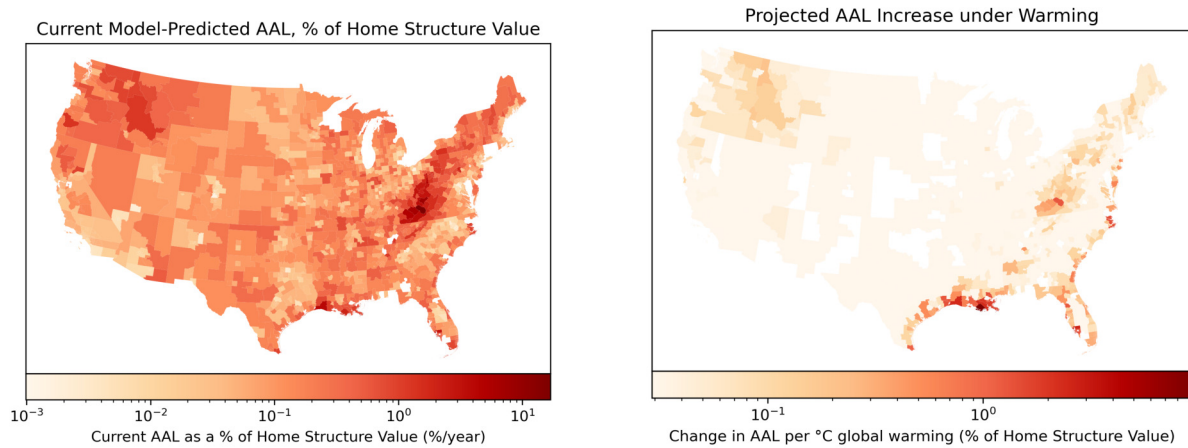


Figure 6: Total disaster property damage by FEMA hazard type, 1975-2020. From SHELDUS.



(a) Model-predicted 2020 Average Annual Losses (AAL), as a % of housing structure value, from the model of Bates et al. (2021) as reported by Wing et al. (2022).

(b) Projected flooding damage growth as a percentage of housing structure value in the model of Bates et al. (2021) as reported by Wing et al. (2022). I divide their projections, which are changes from 2020-2050 under RCP4.5, by the ensemble median of global temperature change in that period and scenario from Rasmussen and Kopp (2017). Units are percentage *points*, representing the expected percentage of total home structure value destroyed by flooding in each location following 1°C global warming. Log scale.

Figure 7

3.4 Damage Magnitudes and Variation

In order to roughly gauge the relative magnitudes and spatial distribution of these damages, I naively dollarize estimates of annual damages from each source, per 1°C global warming, for every household in the 2020 ACS. Specifically, I compute damages incurred by 1°C global warming, in the absence of behavioral or equilibrium responses, as a percentage of income. Damages are dollarized as follows. Productivity and amenity damages are both multiplied by current income. Residential energy expenditure damages are multiplied by current residential electricity expenditures. Disaster damages are multiplied by current home value, if the home is owned, or set to zero if the home is rented.

Figure 8 shows net damages from all sources combined. Figure 8a shows damages on a log scale to highlight the severity of damages in highly flooding-exposed counties. Figure 8b is on a linear scale to illustrate more widespread but less localized damages, mostly temperature-related.

Figure 9 compares damages as a percentage of income by income and by latitude. Damages are relatively flat by income, partly by construction: productivity and amenity damages are constructed as a percentage of income. Nevertheless, flooding and energy cost damages are proportionally larger for lower-income households, due to larger housing price exposure and energy expenditure for these households as a proportion of income.

Spatial inequality in damages is quite dramatic. Figure 9b shows damages by latitude. Households at lower latitudes experience significantly greater heat disamenities and productivity damages and smaller decreases in cold disamenities. Flooding damages are relatively small as a proportion of total damages, but are highly spatially concentrated.

Figure 18 displays mean damage estimates—in absolute dollar terms and as a percentage of income—per °C global warming in the absence of behavioral or equilibrium responses, along other spatial margins: distance to coast, latitude, and 2020 mean annual temperature. Damages are decreasing in distance from the coast. This is primarily driven by climate amenities and income impacts, which operate through temperature rather than flood risk, though flood risk is a contributor. Damages are strongly increasing in current mean annual temperature, due largely to climate impacts on amenities and income, though flood risk is a contributor in this case too. Overall damages are positive for most households, with households in the bottom decile of current mean annual temperature still experiencing positive damages (net harms) on average and households in the top decile of current mean annual temperature experiencing mean damages of almost 3% of income per °C global warming.

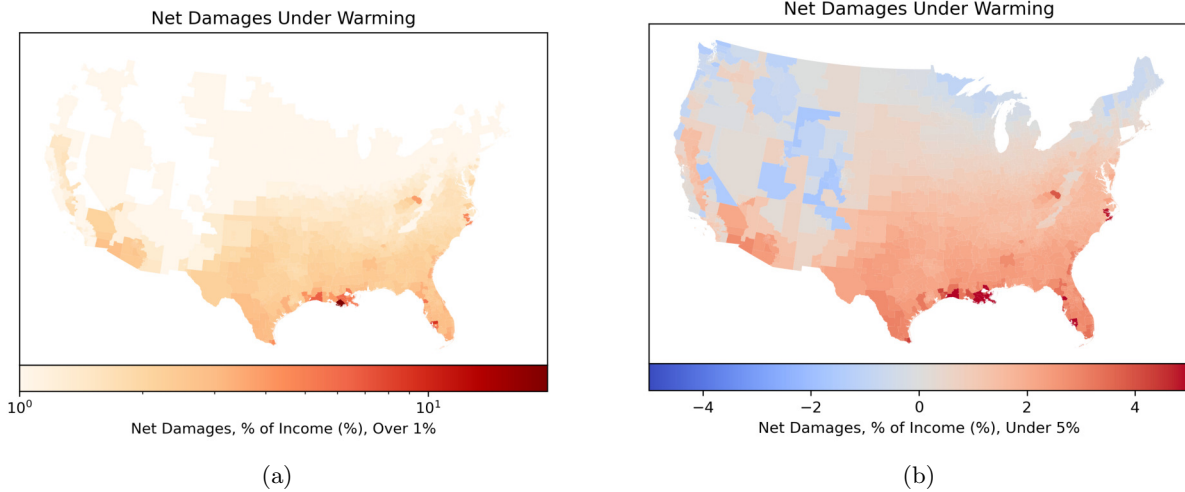


Figure 8

Mean damages per $^{\circ}\text{C}$ global warming, as an equivalent percentage of income, in the absence of behavioral or equilibrium responses. Estimates are obtained by summing over damages to labor income, location amenities, residential energy costs, and residential disaster damages, as computed in this section, for each household in the 2020 American Community Survey. Damages are then meaned within each PUMA. Panel (8a) is on a log scale and bottom coded at 1%. Panel (8b) is on a linear scale and top coded at 5%.

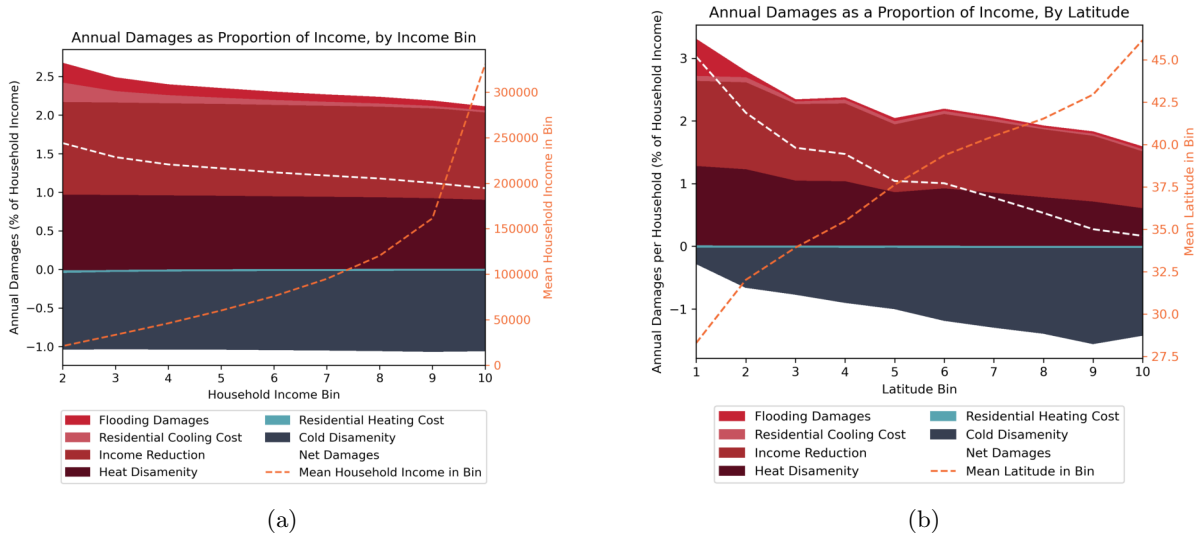


Figure 9

Mean damages per $^{\circ}\text{C}$ global warming, as an equivalent percentage of income, in the absence of behavioral or equilibrium responses. Estimates are obtained by summing over damages to labor income, location amenities, residential energy costs, and residential disaster damages, as computed in this section, for each household in the 2020 American Community Survey. Panel (9a) reports the mean for each income bin. Panel (9b) reports the mean for each latitude bin.

4 Calibrating Climate Expectations

In this section I calibrate the climate process by mapping data on climate expectations and uncertainty to an enriched auxiliary model of global climate and climate expectations. The auxiliary model explicitly incorporates multiple climate variables and sources of climate uncertainty and can exactly match moments from the data and existing climate projections on the expected evolution of the global climate, as well as each source of uncertainty. I then calibrate the simplified climate process in the main model to match key moments in the auxiliary model representing overall uncertainty: in particular, I match the model-predicted quartiles, given 2020 information, of 2100 global temperatures and disaster damages.


The auxiliary model consists of a richer stochastic global climate process and a representative Bayesian household who is uncertain about the sensitivities of global temperatures and disaster damages to emissions and learns about these over time by observing realized global temperatures and disaster damages.

The auxiliary model explicitly incorporates: internal variability, in that annual realized global temperatures and disaster damages have a random component; scenario uncertainty, in that the path of global emissions is stochastic; and climate model uncertainty, in that the household learns over time about the unobserved sensitivities of global temperatures and disaster damages to emissions.

4.1 Stochastic Global Climate Process

I describe a process in which the evolution of global climate conditions is influenced by greenhouse gas (GHG) emissions. The purpose of this model is not to be a realistic representation of the climate system but to parsimoniously encode, at any point in time, a plausible distribution of future climate paths, such that (1) the model can match key moments from data and climate projections on the evolution of, and uncertainty in the evolution of, the global climate; (2) news shocks affect the distribution of future paths in a straightforward and quantitatively plausible way, impacting the anticipated trajectory of the global climate without immediately impacting the physical climate state.

The enriched stochastic global climate process consists of a nonlinear process on the following variables, where arrows indicate dependence within each period:



Conceptually, global power generation grows at a constant rate. The annual level of GHG emissions is the product of global power generation and the emissions intensity of power generation. The emissions

intensity decreases over time, eventually overtaking power generation, so that emissions eventually vanish.⁴⁷ Emissions contribute to a stock of atmospheric GHGs which is gradually absorbed out of the atmosphere by natural processes. Annual mean temperatures and disaster damages are drawn each year from a distribution whose mean depends on the stock of atmospheric GHG.

Formally, time is discrete. Global power generation w_t grows deterministically at constant rate g^w :

$$\log w_{t+1} = \log w_t + g^w.$$

The log emissions intensity of power generation, m_t , decreases in proportion to current power generation:

$$\begin{aligned} \log m_{t+1} &= \log m_t - \beta_m w_t + \varepsilon_t^m \\ \varepsilon_t^m &\sim \mathcal{N}(0, \sigma_m). \end{aligned}$$

Annual emissions are simply the product of power generation and emissions intensity:

$$e_t = w_t m_t. \quad (5)$$

The stock of atmospheric greenhouse gases, GHG_t follows an autoregressive process. GHGs are added to by emissions and subtracted from by natural absorption at rate $(1 - \rho_{\text{GHG}})$:

$$\text{GHG}_{t+1} = c_{\text{GHG}} + \rho_{\text{GHG}} \text{GHG}_t + e_t. \quad (6)$$

Each year, realized temperatures and logged disaster damages are independently drawn from Normal distributions with means $\overline{\text{SST}}_t$ and $\bar{\delta}_t$ respectively,

$$\begin{aligned} \text{SST}_t &\sim \mathcal{N}(\overline{\text{SST}}_t, \sigma^d) \\ \log \delta_t &\sim \mathcal{N}(\log \bar{\delta}_t, \sigma^d). \end{aligned}$$

The mean of the temperature distribution $\overline{\text{SST}}_t$ is linear in the stock of atmospheric GHG:

$$\overline{\text{SST}}_t = \tilde{c}_{\text{SST}} + \tilde{\beta}_{\text{SST}} \text{GHG}_t.$$

⁴⁷This process remains agnostic on the extent to which these emissions intensity reductions are due to policy responses or “natural” technological progress.

Substituting in Equations 5 and 6 and simplifying,

$$\overline{\text{SST}}_{t+1} = c_{\text{SST}} + \rho_c \overline{\text{SST}}_t + \beta_{\text{SST}} w_t m_t. \quad (7)$$

Without loss of generality, I define $\overline{\text{SST}}_t$ as the deviation from the stationary steady state baseline, so that,

$$\overline{\text{SST}}_{t+1} = \rho_c \overline{\text{SST}}_t + \beta_{\text{SST}} w_t m_t.$$

Similarly, the mean of the annual disaster damage distribution $\bar{\delta}_t$ is log-linear in the stock of atmospheric GHG,

$$\log \bar{\delta}_t = \tilde{c}_\delta + \tilde{\beta}_\delta \text{GHG}_t.$$

Simplifying similarly to Equation (7),

$$\log \bar{\delta}_{t+1} = \rho_\delta \log \bar{\delta}_t + \beta_\delta w_t m_t.$$

Altogether, the global climate process of the auxiliary model is as follows:

$\log w_{t+1} = \log w_t + g^w$	Global Energy Generation
$\log m_{t+1} = \log m_t - \beta_m w_t + \varepsilon_t^m$	Emissions Intensity
$\overline{\text{SST}}_{t+1} = \rho_{\text{SST}} \overline{\text{SST}}_t + \beta_{\text{SST}} w_t m_t$	Global Temperature
$\log \bar{\delta}_{t+1} = \rho_\delta \log \bar{\delta}_t + \beta_\delta w_t m_t$	Global Disaster Risk
$\varepsilon_t^m \sim \mathcal{N}(0, \sigma_m^2)$	Scenario Uncertainty
$\text{SST}_t \sim \mathcal{N}(\overline{\text{SST}}_t, \sigma_{\text{SST}}^2)$	Internal Temperature Variability
$\log \delta_t \sim \mathcal{N}(\log \bar{\delta}_t, \sigma_\delta^2)$	Internal Disaster Variability

Two sources of uncertainty are already present: internal variability, in the variability of annual temperature and disaster realizations given global climate conditions; and scenario uncertainty, in that emissions intensity evolves stochastically. The following section adds climate model uncertainty, in which households do not directly observe the coefficients β_{SST} and β_δ , but instead must learn about them over time.

4.2 Bayesian Updating of Climate Expectations

In order to incorporate model uncertainty over the coefficients β_{SST} and β_δ , consider a representative rational Bayesian household who observes the climate process defined in Section 4.1. The household does not directly observe the coefficients β_{SST} and β_δ or the true means $\bar{\delta}_t$ and $\overline{\text{SST}}_t$. Instead, they observe the noisy

realizations δ_t and SST_t . Assume that all households have common priors and common signals so that it suffices to consider the beliefs of the representative household. Assume also that the household observes all other time- t states $(w_t, m_t, \varepsilon_t^m)$ and parameters $(g^w, \beta_m, \rho_{\text{SST}}, \beta_{\text{SST}}, \rho_\delta, \beta_\delta, \sigma_m^2, \sigma_{\text{SST}}^2, \sigma_\delta^2)$ of the model and has no uncertainty over the functional form of the model.

At the beginning of each period t , the household's information set

$$\mathcal{I}_t = \left\{ \widehat{\mathbf{x}}_t^{\text{SST}}, \widehat{\mathbf{P}}_t^{\text{SST}}, \widehat{\mathbf{x}}_t^\delta, \widehat{\mathbf{P}}_t^\delta \right\}$$

encodes a multivariate Normal prior on parameters β_{SST} and $\overline{\text{SST}}_t$ governing the evolution of temperatures,

$$\mathcal{N}\left(\widehat{\mathbf{x}}_t^{\text{SST}}, \widehat{\mathbf{P}}_t^{\text{SST}}\right) \quad \text{over} \quad \begin{bmatrix} \overline{\text{SST}}_t \\ \beta_{\text{SST}} \end{bmatrix}$$

and a multivariate Normal prior on parameters β_δ and $\bar{\delta}_t$ governing the evolution of disaster damages,

$$\mathcal{N}\left(\widehat{\mathbf{x}}_t^\delta, \widehat{\mathbf{P}}_t^\delta\right) \quad \text{over} \quad \begin{bmatrix} \log \bar{\delta}_t \\ \beta_\delta \end{bmatrix}.$$

I assume that these two multivariate priors are initially independent, in which case they remain independent. The learning processes for β_{SST} and β_δ are identical. For brevity, I describe only the process for β_δ . Each period t , the expected value $\bar{\delta}_t$ of disaster damages evolves and disaster damages themselves, δ_t , are realized as described in Section 4.1. Households observe δ_t and a Normally distributed news shock n_t^δ ,

$$\begin{aligned} \log \delta_t &\sim \mathcal{N}(\log \bar{\delta}_t, \sigma_\delta^d) \\ n_t^\delta &\sim \mathcal{N}(\mu^\delta, \sigma_\delta^n). \end{aligned}$$

Though they do not know the values of β_δ or $\bar{\delta}_t$, they know that $\bar{\delta}_{t+1}$ evolves according to

$$\log \bar{\delta}_{t+1} = \rho_\delta \log \bar{\delta}_t + \beta_\delta w_t m_t$$

This is a linear Gaussian state space model. Assuming rational Bayesian learning, the posterior given emissions is linear in the shock realizations and given by the Kalman filter.

In matrix form, the unobserved state evolves according to:

$$\underbrace{\begin{bmatrix} \log \bar{\delta}_{t+1} \\ \beta_\delta \end{bmatrix}}_{\mathbf{x}_{t+1}} = \underbrace{\begin{bmatrix} \rho_\delta & w_t m_t \\ 0 & 1 \end{bmatrix}}_{\mathbf{F}_t} \underbrace{\begin{bmatrix} \log \bar{\delta}_t \\ \beta_\delta \end{bmatrix}}_{\mathbf{x}_t}.$$

and shock realizations are generated according to:

$$\underbrace{\begin{bmatrix} \log \bar{\delta}_t \\ n_t \end{bmatrix}}_{\mathbf{z}_t} = \underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}}_{\mathbf{H}_t} \underbrace{\begin{bmatrix} \log \bar{\delta}_t \\ \beta_\delta \end{bmatrix}}_{\mathbf{x}_t} + \mathbf{v}_t$$

$$\mathbf{v}_t \sim \mathcal{N}(0, \Sigma_v).$$

Given the prior $\mathcal{N}(\hat{\mathbf{x}}_t^\delta, \hat{\mathbf{P}}_t^\delta)$ for \mathbf{x}_t at the beginning of period t ,

$$\mathbf{x}_t | \mathcal{I}_t \sim \mathcal{N}(\hat{\mathbf{x}}_t^\delta, \hat{\mathbf{P}}_t^\delta),$$

the rational prior for \mathbf{x}_{t+1} at the beginning of period $t + 1$ is given by the Kalman filter:

$$\begin{aligned} \hat{\mathbf{x}}_{t+1}^\delta &= \mathbf{F}_t(\mathbf{I} - \mathbf{K}_t\mathbf{H}_t)\hat{\mathbf{x}}_t^\delta + \mathbf{F}_t\mathbf{K}_t\mathbf{z}_t \\ \hat{\mathbf{P}}_{t+1}^\delta &= \mathbf{F}_t(\mathbf{I} - \mathbf{K}_t\mathbf{H}_t)\hat{\mathbf{P}}_t^\delta\mathbf{F}_t' \\ \mathbf{K}_t &= \hat{\mathbf{P}}_t^\delta\mathbf{H}_t'(\mathbf{H}_t\hat{\mathbf{P}}_t^\delta\mathbf{H}_t' + \Sigma_v)^{-1}, \end{aligned}$$

where \mathbf{K}_t is the optimal Kalman gain.

$$\log \delta_t | \mathcal{I}_t \sim \mathcal{N}(\hat{\mu}_t, \sigma_d + \hat{\sigma}_{\mu t}).$$

4.3 Calibration of Climate Process

The Intergovernmental Panel on Climate Change (IPCC) define greenhouse gas concentration scenarios which are labelled RCP x , where x is the level of radiative forcing (net energy transfer from space) per m^2 of the Earth's surface. Figure 10 shows median temperature projections for each pathway in the Coupled Model Intercomparison Project 5 (CMIP5) ensemble (Rasmussen & Kopp, 2017).

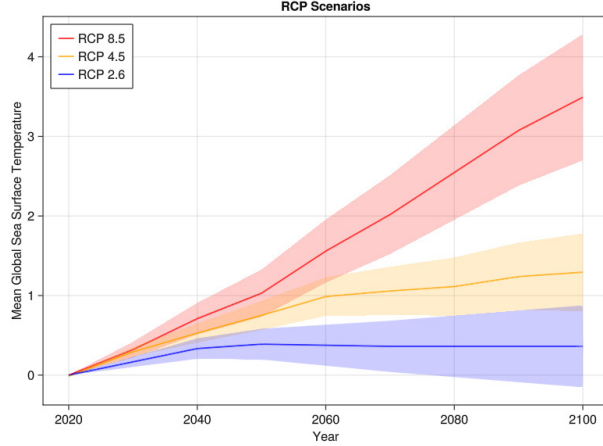


Figure 10

Representative Concentration Pathways (RCPs) from the IPCC Fifth Assessment Report (AR5). The RCPs are labelled by the level of radiative forcing (net energy transfer from space) per m^2 of the Earth’s surface.

I calibrate the free variables of the climate process as follows. I set initial global energy generation, w_0 , to global energy generation in 1900. I set the growth rate of global energy generation, g^w , to the growth rate of global energy generation between 1900 and 2020. I set initial emissions intensity, m_0 , to match average global emissions intensity in 2020. I set the parameter controlling the decline in emissions intensity, β_m , to match model-predicted median annual emissions in 2100, $w_{2100}m_{2100}$, conditional on 2020 emissions, to 2100 emissions in the RCP4.5 scenario.

I jointly set the variance of the emissions intensity shock, σ_m^2 ; the contribution of emissions to global temperatures, β_{SST} ; and the natural rate of regression of temperatures to the pre-industrial baseline, ρ_{SST} . I set them to the model-predicted first, second, and third quartiles of global temperatures in 2100, conditional on 2020 emissions, to the RCP2.0, RCP4.5, and RCP6.0 scenarios.

I set the contribution of emissions to disaster damages, β_δ , to match the increase in disaster damages between 2020 and 2050 predicted by the model of Bates et al. (2021), as reported by Wing et al. (2022). I estimate the natural rate of regression of disaster damages to the pre-industrial baseline, ρ_δ , using data on historical homeowners insurance premiums in Appendix E.

4.4 Quantification: Model Uncertainty

I use historical data on homeowners insurance premiums and losses to estimate market-relevant internal variability and model uncertainty. Homeowners insurance contracts are one-year contracts, representing a plausibly static decision on the part of the insurer. Using aggregate accounting data, I back out anticipated annual losses from storm damage, given insurers’ beliefs and information sets, in each year.

Table 3: Calibrated Parameters, Auxiliary Climate Process

Parameter	Interpretation	Value	Target or Source	Moment Value
w_0	Initial global energy generation	107 PWh	Global energy generation, 1900	107 PWh
g^w	Global energy generation growth rate	0.0155	Growth rate of global energy generation, 1900-2020	0.0155
m_0	Initial emissions intensity	0.358 t/MWh	Global emissions intensity, 2020	0.273 t/MWh
β_m	Decline in emissions intensity	0.016	RCP4.5 annual emissions, 2100 (as median)	12.3 Gt
σ_m^2	Variance of emissions intensity shock	0.40	RCP2.0/4.5/6.0 SST, 2100 (as quartiles)	
β_{SST}	Contribution of emissions to global temperatures	0.00053	""	
ρ_{SST}	Natural rate of temperature regression	0.93	""	
β_δ	Contribution of emissions to disaster damages	0.051	Increase in disaster damages 2020-2050 (Bates et al., 2021)	26.4%
ρ_δ	Natural rate of disaster damage regression	0.98	Homeowners insurance premiums	

All moments are either unitless or in units of decade⁻¹. Each parameter influences all moments; I merely report the most directly influenced moment for each parameter.

In order to calibrate uncertainty, I back out the responsiveness of homeowners insurance providers' anticipated losses (which I back out of premiums) to realized losses. Intuitively, the less certain insurers are in their climate beliefs, the more we would expect them to adjust their premiums in response to realized losses.

I model insurer's beliefs as arising from the Bayesian learning process described in Section 4.2. I then use data on premiums and losses to estimate the parameters of this process. Essentially, I am learning from the observed learning behavior of insurers. I show that the specific procedure I employ, which resembles a nested Kalman filter, can be constructed at a high level of generality.

As the procedure for estimating the Bayesian model is quite lengthy and complex and yields few qualitative insights into the climate welfare impact heterogeneity which is the primary focus of this study, I move it to Appendix E. Table 4 reports the result of the estimation.

Variable	Estimate
σ_δ^n	0.21306
σ_δ^d	0.28874
ρ	0.98
T_0	1965.87

Table 4: Results of state space estimation. Point estimates are by maximum likelihood.

5 Calibration

It remains to set the climate-unrelated parameters of the full model. Because the goal is to explore the implications of the model for the period 2020-2100, I calibrate the model's stationary steady state to be consistent with key features of the U.S. economy in the early 1990s. This is the initial steady state in which climate change unexpectedly begins, as described in Section 2.6.

I choose the early 1990s to achieve a balance between the contribution to the state of housing markets in 2020 of, on one hand, anticipatory adaptation to climate change and, on the other hand, factors which are not accounted for in the model. Calibrating the model to a more recent steady state allows for the initial conditions to be more similar to the current economy, but allows for less pre-2020 adjustment by households in anticipation of climate change.

I primarily calibrate the model using indirect inference. I match moments computed from the 5% Public Use Microdata Sample (PUMS) of the 1990 U.S. Census and the 1992 Survey of Consumer Finances (SCF).

The parameters of the model can be divided into four categories: location-level parameters, the parameters of the income process, the initial wealth distribution, and all other parameters.

Each location is calibrated to data on a single Public Use Microdata Area (PUMA), the smallest spatial unit available in the 1990 PUMS. I directly plug in PUMA-level estimates of housing prices computed from the 1990 PUMS, solve for local rents and populations in spatial equilibrium, and calibrate location-specific amenities and labor productivity shifters to match empirical location populations and mean incomes from PUMS. These estimates are independent of local housing supply elasticities. I assign these externally using the spatially-heterogeneous estimates of Saiz (2010).

The labor income process is calibrated by quantifying an annual AR(1) process on labor, welfare, and pension income from the 1989 and 1990 Panel Study of Income Dynamics (PSID), then calibrating a decadal AR(1) income process to the moments of the annual process.

I parameterize the initial wealth distribution to be log-uniform within each income quartile, yielding the five quartile boundaries (including the minimum and maximum wealth) as the free parameters. I set the minimum to \$1 and the remaining four by indirect inference, matching the mean and inner three quartiles of overall net worth in the 1992 SCF.

To quantify non-spatially-varying parameters, I use a combination of indirect inference, external estimates, and standard values. Seven parameters are calibrated by indirect inference. Abstractly, I define seven moments and find values for these seven parameters such that the simulated moments equal the empirical moments.

Concretely, I assign to each parameter a corresponding moment which the model suggests should be particularly informative about that parameter. I do this for all parameters of the model calibrated by indirect

inference: the parameters of the wealth distribution, the non-spatially-varying parameters, as well as the location-specific amenities and productivity shifters.⁴⁸ I then simply guess all parameters at once, compute simulated moments, then update each parameter guess using only the difference between the simulated and empirical values of the moment associated with that parameter. I repeat until all simulated moments converge to their empirical counterparts within a tolerance of 10^{-6} .

That is, I essentially find the parameters by gradient descent, where I set the derivative of a moment with respect to a parameter to zero unless that parameter is the “relevant” parameter for that moment. The fact that this process converges demonstrates that the chosen moments are indeed informative about their corresponding parameters, at least within the context of the model.

For brevity, in the remainder of this section it is often stated that a parameter X is calibrated to match a moment Y. This is not intended to imply that parameter X affects *only* moment Y, but that the above procedure is applied in full and that parameter X is associated with moment Y therein. All parameters calibrated by indirect inference are calibrated jointly.

5.1 Spatial Parameters

I choose the 1990 Public Use Microdata Area (PUMA) as my spatial unit. This is the smallest spatial unit reported in the 1990 PUMS. I use the $N = 1713$ PUMAs which cover the contiguous U.S. completely. A PUMA is an area with 1990 population between 100,000 and 200,000, generally following the boundaries of Metropolitan Statistical Areas (MSAs) where applicable.

For each location ℓ , I set the housing price level⁴⁹ q_ℓ to the value of a typical owner-occupied two-bedroom home in the 1990 PUMS.⁵⁰ By indirect inference, I calibrate the amenity value of a location (which enters the model as an effective multiplier on consumption spending) $\bar{\alpha}_\ell$ to match the population of location ℓ in the 1990 PUMS. By indirect inference, I calibrate the local labor productivity shifter \bar{A}_ℓ to match the mean annual labor earnings among the working-age population in location ℓ in the 1990 PUMS.

In equilibrium, these imply a value in each location for the quantity of housing H_ℓ and the rent level⁵¹ ρ_ℓ , which are untargeted. Note that the house price level q_ℓ is also an equilibrium quantity and not a parameter. In theory, q_ℓ is a moment and Π_ℓ is calibrated to match that moment, solving,

$$q_\ell = \Pi_\ell H_\ell^{\beta_\ell}.$$

In practice, I assign q_ℓ before the indirect inference step and defer the calibration of Π_ℓ until after the indirect

⁴⁸As an optimization, I also treat steady-state location-level rent levels $\rho_{\ell t}$ as parameters, with the market clearing condition as the “associated moment.” This removes the need to solve for equilibrium rents for every parameter guess.

⁴⁹As I am calibrating a steady-state version of the model, I omit time t subscripts.

⁵⁰I regress all owner-occupied home values on a PUMA fixed effect and a “number of bedrooms” fixed effect, with “two bedrooms” as the reference group and no intercept, and use the PUMA fixed effect as my house price level estimate.

⁵¹In units which correspond to the typical owner-occupied two-bedroom home.

inference step, at which point I externally assign the estimates of Saiz (2010) to set β_ℓ for each location, then set,

$$\Pi_\ell = q_\ell / H_\ell^{\beta_\ell}.$$

The estimates of Saiz (2010) are at the level of MSAs, which generally nest PUMAs. Details of the imputation procedure are in Appendix D.

5.2 Initial Wealth Endowments

By indirect inference, I set the distribution of initial wealth endowment (at age 20) to match the mean and quartiles of net worth in the SCF. A household is equally likely to be born in each quartile. Conditional on quartile, initial wealth endowment is log-uniform. The minimum initial wealth endowment is \$1. I set the maximum initial wealth endowment to match the empirical mean of net worth.

5.3 Non-Spatial Parameters

The non-spatial parameters of the model determine preferences, the lifecycle income process, initial wealth endowments, housing investment and maintenance costs, and migration costs. In addition to Census data, I also use data from the 1992 Survey of Consumer Finances (SCF).

5.3.1 Lifecycle

One model period represents ten years. Households are born at age 20 and die at age 80.

5.3.2 Preferences

By indirect inference: I set the weight of housing in consumption, γ , to match the median rent-to-income ratio among renters in the 1990 Census; I set the strength of the bequest motive, Q , to match the ratio between overall mean wealth and mean wealth of households with a head between 70 and 80; and I set the elasticity of substitution between housing and goods consumption, $1/\sigma$, to match the reduced form relationship β_{RTI} , across locations, between the housing asset price level q_ℓ and median rent-to-income ratio:⁵²

$$\text{RTI}_\ell = \alpha_{\text{RTI}} + \beta_{\text{RTI}} \log q_\ell + \varepsilon_\ell.$$

where RTI_ℓ is the median rent-to-income ratio in location ℓ .

Externally, I set the elasticity of intertemporal substitution to $1/\eta = 2.0$ as in Kaplan et al. (2020).

⁵²Ideally, I would regress on the housing rent level, but I do not observe this in the data.

5.3.3 Housing

By indirect inference, I calibrate the interest rate on mortgage borrowing, r_m , to match the ratio between the mean wealth of homeowners to overall mean wealth. I calibrate the baseline cost of maintaining a rental property to match the homeownership rate. I calibrate the excess maintenance and administrative cost on rental properties, $\chi^{\text{let}} - \chi^{\text{live}}$, to match the homeownership rate.

Externally, I set the transaction cost for housing assets, ϕ , to 7% as in Kaplan et al. (2020). I set the maximum loan-to-value ratio of a mortgage, κ , to 80%. I set χ^{live} , the baseline cost of maintaining an owner-occupied home, to \$1,000 annually.

5.3.4 Migration

I parameterize the utility migration cost,⁵³ $F^u(\ell', \ell)$, as a multiple of the distance $d(\ell', \ell)$ between the source location ℓ and destination location ℓ' :

$$F^u(\ell', \ell) = \tau d(\ell', \ell).$$

The calibration of migration costs is complicated by the fact that the fixed costs of migration, F^m , and the location preference shock dispersion, σ_ν , both affect the frequency of migration in steady state. Challengingly, different combinations of F^m and σ_ν can yield the same steady-state migration frequencies but different migration responses to spatial shocks. I externally assign F^m to the value estimated by Giannone et al. (2023) and calibrate σ_ν to match the share of households who move across PUMAs each decade.⁵⁴ I calibrate the per-distance utility cost of moving τ to match the share of households who move across states each decade.

While it is not necessarily true in reality that fixed costs of moving are monetary and distance-dependent costs of moving are preferences, this setup is necessitated by a feature of the numerical solution method that does not allow for variable monetary moving costs. However, this is not a model whose goal is to measure monetary versus utility moving costs and I argue that this assumption is generally not too violent. Monetary costs of moving do exist and it is believable that households prefer to stay close to their current location even if they are moving.

5.4 Calibration Result

Table 5 lists the units that all quantities are expressed in. The model is just-identified and all targeted moments are matched exactly. Table 6 lists the non-spatially-varying parameters, their calibrated values,

⁵³No aspect of the model solution depends on any particular functional form for F^u , however. In a future version, I plan to flexibly calibrate migration costs for each pair of locations to match data on bilateral migration flows.

⁵⁴A possible future direction is to calibrate both values to jointly match steady-state migration frequencies and empirical migration responses to some spatial shock. Identifying such a shock brings its own challenges, however, and using these empirical responses for calibration would preclude their use for model testing.

Table 5: Units

Quantity	Unit
Location	1990 Public Use Microdata Area (PUMA)
Time	10 years
Housing	Typical two-bedroom owner-occupied home
Numeraire Good	Thousand 2020 USD (\$1,000)
Distance	Kilometers (km)
Temperature	Degrees Celsius ($^{\circ}\text{C}$)

Units used throughout the quantification, unless otherwise specified. “Monetary” values such as income and monetary moving costs are paid in the form of the numeraire good, which are expressed in units of one thousand 2020 USD.

and their associated moments. To visualize spatially-varying moments, Figure 24 describes the estimated local amenity values across the U.S. Figure 25 describes the estimated local productivity shifters across the U.S. Figure 26 describes estimated local amenity values for six cities, and Figure 27 describes local productivity shifters around New York City at various scales.

5.5 Model Fit: Price-to-Rent Ratios

Local rent levels and housing stocks are not targeted in the calibration, but arise in the model equilibrium. Figure 28 shows that rent levels vary with housing prices, as both are driven by local amenity values, productivities, and housing supply elasticities. Furthermore, Figure 29 reveals that the model reproduces a key feature of the data: the price-to-rent ratio is higher in locations with higher housing prices. This is primarily driven by the fact that a greater concentration of wealth in more expensive locations creates greater demand for rental real estate assets, decreasing equilibrium rents. Second, a small proportion of this effect is due to a particular modeling choice: the wedge between the maintenance cost of rental housing and owner-occupied housing is a fixed cost per unit of housing, so it is relatively lower in locations with higher housing prices. However, in the calibrated model, maintenance costs account for only between 5% and 25% of rent and cannot account for the large variation in price-to-rent ratios.

Table 6: Non-Spatially-Varying Parameters

Parameter	Interpretation	Value	Target or Source	Moment Value
Lifecycle:				
a_{\max}	Length of adult life	6		
Preferences:				
γ	Taste for housing	0.2	Median (renter) rent-to-income ratio	0.23
$1/\sigma$	Elasticity of Substitution (goods vs. housing)	0.87	Rent-to-income sens. to housing price	0.011
$1/\eta$	Elasticity of intertemporal substitution	0.5		
Q	Bequest motive	4.3	Mean wealth ratio oldest-to-overall	1.62
Housing:				
r_m	Mortgage interest rate	0.80	Mean wealth ratio homeowners-to-overall	1.52
ϕ	Transaction cost	0.07	Kaplan et al. (2020)	
χ^{let}	Rental baseline maintenance cost	21.3	Homeownership rate	0.68
κ	Maximum loan-to-value ratio	0.8		
χ^{live}	Owner-occupied baseline maintenance cost	10.0		
Migration:				
σ_ν	Location preference shock dispersion	0.2	Share of households moving across PUMAs	0.40
τ	Utility cost of migration per km	0.0011	Share of households moving across states	0.22
F^m	Monetary fixed cost of migration	13.3	Giannone et al. (2023)	

All moments are either unitless or in units of decade⁻¹. Each parameter influences all moments; I merely report the most directly influenced moment for each parameter. Goods consumption is scaled in utility so that one unit of goods consumption costs \$50,000.

6 Welfare Impacts of Climate Change: Projections, Margins, and Mechanisms

In this section, I examine the projected welfare impacts of climate change in the calibrated model. The primary experiment is the unexpected onset and progression of stochastic climate change from an initial stationary steady state representing the contiguous U.S. in 1990. Immediately following the onset of climate change, households all understand that the global climate state will evolve according to the stochastic process described in Section 2.4.2.

There are two types of aggregate shock: the initial “onset shock”⁵⁵ and subsequent “climate surprises” which are drawn each period. Both are news shocks, only affecting the future trajectory of the global climate state and not its physical state at the time of the shock. Figure 11a illustrates temperature paths for a sample of draws of the stochastic climate process. Figure 11b illustrates a sample of climate draws conditional on one realization of climate surprises up to 2050.

The richness of the model allows for multiple margins and mechanisms to be captured. Furthermore, because these margins and mechanisms operate together within a single model, we can also study the implications of interactions between them. The key model features are household heterogeneity, housing wealth, uncertainty, and (stochastic) transition dynamics.

With household heterogeneity, households vary not only by location but also in their wealth, income, housing wealth, and age. Each of these margins of heterogeneity shapes a household’s climate exposure, leading to substantial variation in welfare impacts that go beyond spatial heterogeneity.

With housing wealth, households may own housing but those who do so are exposed to climate risk through housing prices, whereby news about future climate change has immediate, heterogeneous wealth and welfare impacts in the present. Housing prices are location-specific and arise endogenously in local real estate and housing rental markets.

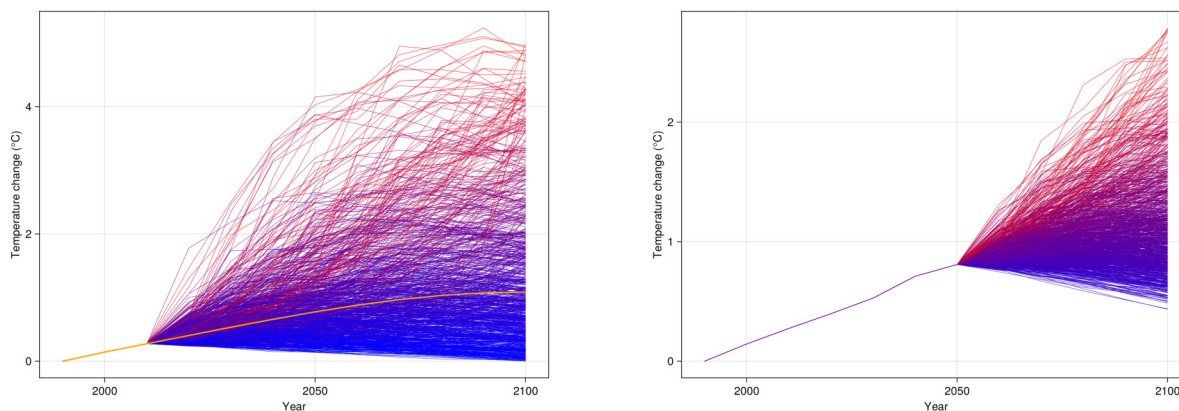
With uncertainty, households do not have perfect foresight over the future of the climate process, instead receiving news shocks over time about the future trajectory of the global climate. In particular, households are exposed to uninsurable and undiversifiable house price risk from aggregate climate uncertainty. Uncertainty leads households to be less willing to hold risky housing assets at a given price. In equilibrium, this is capitalized into endogenously lower average housing prices and higher rents.

With transition dynamics, migration frictions, financial frictions, and household wealth dynamics lead to slow adjustment over time to climate shocks, slowing equilibrium responses to climate change⁵⁶ but creating persistence in welfare impacts across households and locations.

⁵⁵The unexpected onset of the stochastic climate process in steady state.

⁵⁶Relative to an immediate adjustment to a new steady state. This is even true when the physical climate change occurs all-at-once, as I discuss in Section 6.2.

Section 6.1 reports and decomposes the projected welfare impacts of the unexpected onset of climate change by spatial margin (across- vs. within-locations), channel (housing wealth vs. other impacts), and selected margins of household heterogeneity. Section 6.2 describes the immediate impact of the climate news shock and the long-term impact of climate change on equilibrium populations and housing prices, stocks, and rents. In order to analyze the isolated impacts of expectations and physical climate changes, I also explore equilibrium responses to small, unexpected, permanent climate shocks. Section 6.3 explores how overall household and spatial inequality is impacted immediately and over time by the onset and evolution of climate change. Section 6.4 explores the role of uncertainty, highlighting the impact of climate uncertainty on housing asset demand and the resulting equilibrium fall in housing prices and relative rise in rents. Section 6.5 explores the distributional impacts of climate surprises⁵⁷ during the stochastic climate transition. From the perspective of an agent with perfect foresight, this is equivalent to the impact of news counteracting widespread household underestimation (or overestimation) in the severity of climate change.



(a) 1,000 draws of the stochastic climate process. For each draw, global mean temperature deviations from steady state are shown. The “median realization” is plotted in orange, representing the specific realization in which each draw of the climate news shock during the transition is equal to zero, $\varepsilon_t^m = 1\forall t$.

(b) 1,000 draws of the stochastic climate process, conditional on a specific physical realization at 2050. Because climate news shocks do not immediately affect the physical climate state, this is conditional on information available in 2040.

Figure 11

6.1 Overall and Decomposed Welfare Impacts of Climate Change

Table 7 describes the initial welfare impacts of the “onset shock” relative to the continuation of the steady state.⁵⁸ Aggregate welfare losses are equivalent to a loss of 0.066% of lifetime consumption with a standard deviation of 0.187%. Within-location variance in initial welfare impacts of the onset shock accounts for

⁵⁷That is, the realization of climate shocks below or above the prior median, not wholly unexpected shocks.

⁵⁸For each household, I compare the value function at the beginning of the first period of the climate transition with the steady state value function. Throughout this section, welfare is converted into equivalent lifetime consumption for a renter in a “typical location” in steady state.

36.9% of total variance.

Housing wealth impacts represent one channel of welfare impacts. Households who own housing during the onset shock experience an immediate wealth shock through the equilibrium adjustment of house prices. To quantify the housing wealth channel of welfare impacts, I consider for each household *individually* the following counterfactual: immediately following the price adjustment to climate change, the household is given a wealth transfer (in the form of the tradable good) exactly negating the house price shock.⁵⁹ The reported welfare *cost* of the housing wealth channel is the hypothetical welfare *benefit* of this compensating transfer. Within-location variance in initial welfare impacts through the housing wealth channel accounts for 84% of the variance of welfare impacts occurring through this channel. Overall, the housing price response to the onset shock is equivalent to an immediate effective loss of housing value of \$16.8bn and a transfer of \$40.7bn. Over the following century, \$186bn of aggregate housing value is lost and \$507bn is effectively transferred across space.

Figure 12 illustrates the decomposition of initial welfare impact heterogeneity by spatial margin (across- vs. within-locations), channel (housing wealth vs. other impacts), and selected margins of household heterogeneity. Homeowners are on average affected less severely than renters—however, this is not because they are homeowners but due to other factors correlated with housing ownership: when isolating the housing wealth channel of welfare impacts, households who own more housing are more severely affected.

Table 7: Initial Welfare Impact: Variance Decomposition

	All Channels	Housing Wealth Channel	Other Channels
Mean Welfare Impact	-6.56e-02	-3.04e-03	-6.26e-02
Variance	3.50e-02	1.98e-03	3.06e-02
st.d.	1.87e-01	4.45e-02	1.75e-01
Variance (Spatial Component)	2.21e-02	3.14e-04	1.77e-02
st.d. (Spatial Component)	1.49e-01	1.77e-02	1.33e-01
Variance (Within-Location Component)	1.29e-02	1.67e-03	1.29e-02
st.d. (Within-Location Component)	1.14e-01	4.08e-02	1.13e-01

Initial distribution of welfare impact of unexpected onset of climate change. Aggregate effects and variance decomposition.

Each reported variance component is a the result of variance decomposition of the overall variance for the given column. Reported standard deviations are the square roots of the variances. Welfare impacts are expressed as an equivalent percentage of lifetime consumption lost or gained.

⁵⁹Note that I do not consider a widespread policy of compensating transfers, only the impact on a single household, leaving equilibria unaffected. Equivalently, this can be thought of as the welfare benefit to a household of avoiding the price shock by selling their housing immediately (without paying the realtor’s fee) before the climate news shock and immediately re-buying it after the price drops.

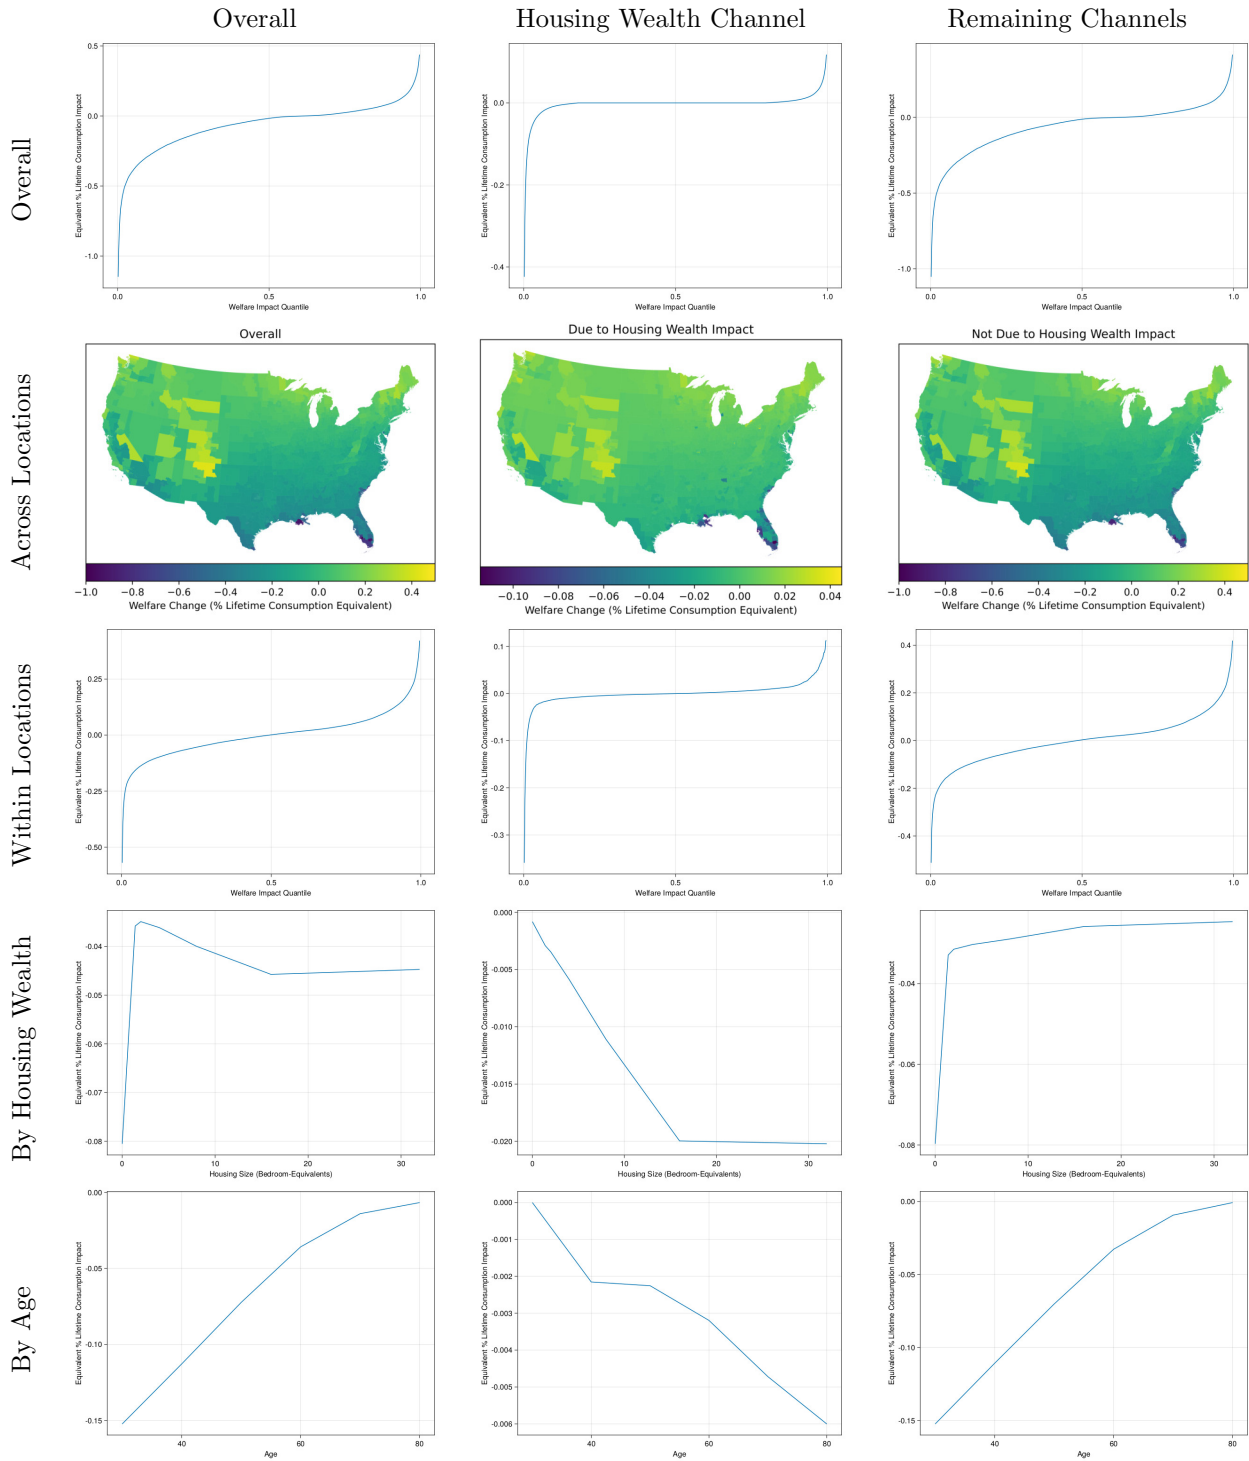


Figure 12

Distribution of welfare impacts of the unexpected onset of climate change (“onset shock”). Columns correspond to channels and rows correspond to margins of variation. Column 1 incorporates all channels. Column 2 isolates the welfare impact of the housing price response. Column 3 reports the remaining welfare impact after subtracting the contribution of the housing wealth channel. Rows decompose variation in welfare impacts into spatial, within-location, housing wealth, and age margins. Welfare impacts are expressed as an equivalent percentage of lifetime consumption lost or gained.

6.2 Equilibrium Impacts of Climate Change

Figures 30 and 31 display spatial equilibrium responses to the onset shock, as well as the 2100 equilibrium following the median realization of the climate process. Figure 30 covers the entire country and Figure 31 zooms in on Florida.

Conceptually, with regard to the immediate equilibrium response, the onset shock is a news shock taking a population initially in steady state from widespread climate denial to widespread climate “acceptance,” if “acceptance” is defined as believing that the calibrated climate process will proceed according to the stochastic process described in Section 2.4.2.

Note that the equilibrium in which these responses are determined is complex: the equilibrium in each location exists not only in spatial and intertemporal equilibrium with every other current and future location but, with climate uncertainty, with an entire distribution of future locations. Furthermore, three factors contribute to the slow evolution of variables across time: transition dynamics due to frictions, the gradual evolution of the global climate state, and the slow resolution of uncertainty.

In order to isolate transition dynamics in equilibrium responses, I also consider a small perturbation: a small, unexpected and permanent increase in global mean temperatures, by 0.1°C, to an economy initially in steady state. Impulse responses to this small perturbation are reported in Appendix J.

6.3 Climate Change and Overall Inequality

Table 8 describes overall consumption inequality in the economy before the climate news shock, immediately after the climate news shock, and one century following the beginning of physical climate change⁶⁰ under the median realization of the climate process. Although climate change has unequal welfare impacts, the immediate impact is progressive. Spatial inequality increases following the onset shock and increases over time as climate change progresses.

Table 8: Overall Welfare Inequality: Transition

	Overall Inequality	Spatial Inequality	Within-Location Inequality
All Ages			
Before Onset	293.341	71.4913	284.496
After News Shock	293.287	71.549	284.425
In 2100 (Median Realization)	292.894	72.5807	283.759
At Birth			
Before Onset	197.871	58.536	189.015
After News Shock	197.697	58.5709	188.821
In 2100 (Median Realization)	197.671	59.2204	188.592

Overall consumption inequality, before and after the onset of climate change. Consumption inequality is measured as the standard deviation of equivalent decadal consumption for a renter in a “typical location” with amenity value $\alpha = 1$ and average housing rents in steady state.

⁶⁰110 years following the climate news shock.

6.4 Role of Uncertainty

We can understand the distributional effect of climate uncertainty on wealth inequality by comparing the global solution of the model with aggregate climate uncertainty (the “Stochastic Model”) with a model of perfect foresight in which the transition is effectively deterministic (the “Deterministic Model”).⁶¹ In particular, consider the median path illustrated in Figure 15a. This arises as one possible realization of the full stochastic model. Consider now a deterministic model⁶² in which households expect this median path to occur with certainty.

Each model begins in a stationary steady state representing the year 1990, then is hit by the news shock that a (stochastic or deterministic) climate transition is beginning. Figure 13 illustrates how average welfare varies by wealth percentile in the stochastic and deterministic models 30 years after the beginning of the climate transition.

The difference between the two models is climate uncertainty. If we shut down uncertainty (the “Deterministic Model”), then the welfare of households in the lower half of the wealth distribution is higher than in the model with uncertainty (the “Stochastic Model”), though the welfare of households in the top half is relatively unchanged. This regressive distributional consequence of climate uncertainty is equivalent to a \$94bn wealth transfer from the bottom half of the wealth distribution to the top.⁶³

⁶¹The distinction between these two models is one of household expectations, not that nature is somehow “more stochastic” in the global solution.

⁶²The same model generating the solid lines in Figure 15b.

⁶³I choose to focus on the relationship between welfare and wealth quantile as my measure of wealth inequality. The alternative would be to construct a price index for each simulated period, but this comes with the challenge that relative housing and goods prices vary across time as well as space and affect renters and homeowners differently.

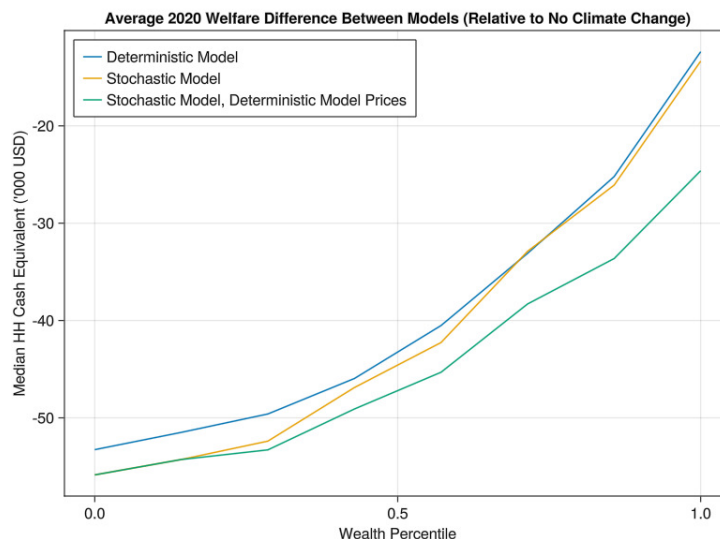


Figure 13: Welfare, relative to the no-climate-change stationary steady state, along the wealth distribution, in 2020, given the observed 1990-2020 climate realizations. The full stochastic model is shown, as well as a “deterministic model” under which the climate evolves according to the median path of the climate process with certainty, and a “stochastic model with deterministic prices” under which the climate evolves according to the full climate process, but prices are deterministically set to their path under the deterministic model (ignoring market clearing).

To understand what is driving this regressive distributional effect of climate uncertainty, I shut down the exposure of house prices to climate uncertainty. In order to do this, I define a model in which households still face uncertainty over climate impacts on productivity, amenities, energy costs, and disaster risk, but house prices are exogenously constrained, no matter the realization of climate shocks, to equal the value they would have taken in the deterministic model.⁶⁴ Figure 13 displays wealth inequality in this “Stochastic Model with Deterministic Model Prices,” compared to the Stochastic Model and the Deterministic Model. Wealthier households are worse off than they were in the stochastic model, but poorer households are unaffected on average, indicating that the uncertainty in housing prices in the Stochastic Model actually benefits richer households. This is due to an endogenous risk premium on housing assets that arises under aggregate climate uncertainty.

Figure 14 displays the paths of average house prices, across all locations, under the deterministic model and under the stochastic model. In the deterministic model, housing prices are an average of 1.16% higher in 2020. Note that the difference narrows over time, as uncertainty resolves. This also hints at the potential role of overconfidence in shaping economic climate outcomes. When markets are overconfident in their climate projections, risk premia will be smaller. Prices will be closer to the deterministic case, and the impacts of

⁶⁴As before, market clearing conditions will not hold in this scenario. This no-price-change scenario can be understood as the result of a government policy maintaining house prices at the previously anticipated levels through constructing and selling homes below cost (if market prices are too high) or purchasing and demolishing homes (if market prices are too low).

surprises will be closer to those described in Section 6.5.




Figure 14: Average house prices across all locations along the median path, under both the stochastic and deterministic models. In the stochastic model, expectations are that the climate will follow the process described in Section 2.4.2. In the deterministic model, expectations are that the climate will deterministically follow the median path.

6.5 Impact of Surprises

If households underestimate the severity of future climate change, then the news that climate change will be more severe than expected will eventually arrive as a negative climate news surprise. This reduces the average value of housing assets, which are concentrated in the portfolios of wealthier households, reducing overall wealth inequality. However, because some locations experience an increase in housing values, a subset of wealthy households benefit substantially from a negative climate news surprise.

Figure 15a describes an illustrative scenario. I consider three paths from the distribution of climate paths defined by the stochastic process in Section 2.4.2. The first is the median or “no surprise” path in which the realized value of the climate shock drawn each period is equal to 0. This path approximately corresponds to the RCP4.5 scenario.

$$\varepsilon_t^{m,\text{median}} = 0 \quad \forall t.$$

The second is the “negative surprise in 2050” path, in which the realized value of the climate shock drawn each period is zero, except for the shock drawn in 2050, ε_{2050}^m , which is equal to the 95th percentile of the shock distribution (implying higher future temperatures). The third path is the “positive surprise in 2050” path, in which all shocks are zero, except that ε_{2050}^m is equal to the 5th percentile of the shock distribution (implying lower future temperatures).

The welfare of wealthier households is more sensitive to climate surprises than the welfare of less wealthy households. To understand why, I shut down one channel at a time: first, uncertainty; second, exposure of house prices to climate surprises. After shutting down exposure of house prices to climate surprises, the correlation between wealth and climate sensitivity flips: climate sensitivity is then decreasing in wealth.


To shut down uncertainty, consider a model in which households believe that they have perfect foresight about the climate process and that the transition will effectively be deterministic along the median path. The positive or negative surprise in 2050, if it occurs, is fully unexpected. Following the surprise, households believe that the new path is deterministic. This model from the same initial 1990 steady state as the stochastic model and the “onset shock” is still unexpected. However, instead of a stochastic climate process beginning, the climate process that unexpectedly begins is deterministic.

Figure 15b plots, in solid lines, the welfare impact of the climate surprise illustrated in Figure 15a, as a function of wealth, shutting down uncertainty. That is, until 2050, households believe that they are on a deterministic transition along the median path.

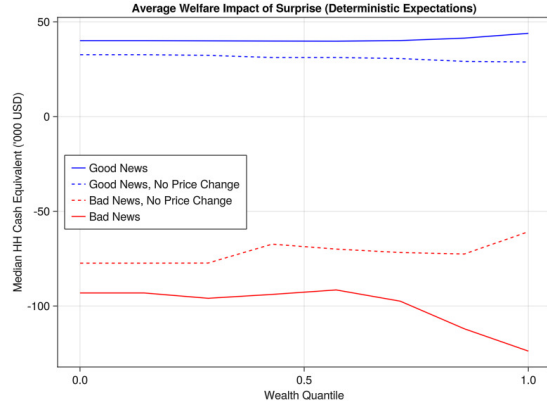
For each wealth group, I plot the value at the beginning of 2050 under the positive (blue) and negative (red) surprises, relative to the no-surprise scenario. Wealthier households benefit more from the positive surprise and are more harmed by the negative surprise. The relationship between wealth and climate sensitivity is positive.

To quantify the contribution of housing wealth to this positive wealth-to-climate-sensitivity relationship, we can consider the case in which housing prices are unaffected by the surprise. That is, we ignore market clearing conditions, and simply impose exogenously that house prices along the new path are equal to what equilibrium house prices would have been along the old path. The climate surprise still affects local labor productivity, amenities, energy costs, and disaster risk.

The resulting welfare effects, along the wealth distribution, are plotted in Figure 15b in dashed lines. Shutting down house price exposure to the climate surprise, the wealth-to-climate-sensitivity relationship becomes negative, illustrating that the higher climate sensitivity of wealthier households is due to their exposure to housing assets.



(a) Three realizations of climate paths defined by the stochastic process in Section 2.4.2. The first is the median or “no surprise” path in which the realized value of the climate shock drawn each period is equal to 0. The high and low paths are the result of a single high or low draw of the news shock ε_t^m , followed by a sequence of zero draws.



(b) Average welfare impacts of each surprise, as defined in the left figure. Solid lines represent welfare impacts if prices adjust to the new path. Dashed lines represent welfare impacts if prices are fixed to their values along the initial middle path (in violation of market clearing).

Figure 15

7 Conclusion

How large is heterogeneity in the welfare effects of climate change across households? I study how homeownership shapes the answer to this question by building and quantifying a dynamic spatial equilibrium model of the contiguous U.S. with 1713 locations featuring rich household heterogeneity, housing wealth at the household level, and stochastic transition dynamics with aggregate climate uncertainty. In the model, heterogeneous households make forward-looking consumption-savings, migration, and real estate investment decisions which endogenously determine spatially heterogeneous housing prices. Crucially, housing prices both reflect expectations about future climate change and determine an empirically realistic proportion of household wealth.

With homeownership, inequality in the welfare impacts of climate change is large and uneven across space, wealth, housing wealth, and age. In response to a change in expectations from widespread climate denial to widespread climate acceptance, the spatial adjustment of housing prices to the climate news shock is equivalent to an effective transfer of housing value across space of \$41bn immediately and \$507bn over the following century.

With anticipation and housing wealth, uncertainty introduces regressive effective wealth transfers through higher rent. Climate risk causes households to be less willing to own rental housing, raising rent and expected returns on housing. Among households who do own housing, wealthier households are more willing to bear this risk, benefiting from the higher risk premia required by marginal, less-wealthy landlords. Less-wealthy

renters are hurt by higher rents. Altogether, the regressive welfare impact of uncertainty is equivalent to a transfer of \$94bn from the bottom half of the wealth distribution to the top half, relative to a model with perfect foresight, for the cohort reaching adulthood in 2050. The analysis of climate uncertainty depends crucially on having found the global solution of the model. This is made possible by the development of a deep learning method for solving recursive economic models under aggregate uncertainty.

Several factors not incorporated in this model are likely to further complicate the role of housing wealth in shaping inequalities in climate damages. Three such factors are belief heterogeneity, agglomeration, and growth. In reality, beliefs about climate change vary significantly across individuals. Thus households who more accurately predict the consequences of climate change might be more likely to gain from forward-looking decisions informed by those beliefs. In reality, productivity is shaped by agglomeration forces as well as exogenous geographic advantages. If these amplify the spatial reallocation of economic activity due to climate change, they may also amplify the impacts of climate change on housing prices. Economic and population growth are also significant factors shaping housing prices. Especially because climate change occurs on the scale of centuries, the interaction of growth with climate change may have the potential to dramatically shape house price and welfare impacts.

Finally, this model develops methods with potentially interesting applications beyond the climate literature. In reality, all housing wealth depends on housing prices which reflect uncertain expectations. Expectations and uncertainty over many long-term developments thus shape housing prices in the present. Dynamic spatial models with heterogeneous households, homeownership, and aggregate uncertainty—together with the tools I develop to solve them—could be used to study the impacts of expectations regarding future zoning policy, population growth, or future agglomeration forces on housing markets.

References

- Albouy, D., Graf, W., Kellogg, R., & Wolff, H. (2016). Climate Amenities, Climate Change, and American Quality of Life [Publisher: The University of Chicago Press]. *Journal of the Association of Environmental and Resource Economists*, 3(1), 205–246. <https://doi.org/10.1086/684573>
- Azinovic, M., Gaegauf, L., & Scheidegger, S. (2022). Deep Equilibrium Nets [eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12575>]. *International Economic Review*, 63(4), 1471–1525. <https://doi.org/10.1111/iere.12575>
- Balboni, C. (2021). In Harm’s Way? Infrastructure Investments and the Persistence of Coastal Cities.
- Bates, P. D., Quinn, N., Sampson, C., Smith, A., Wing, O., Sosa, J., Savage, J., Olcese, G., Neal, J., Schumann, G., Giustarini, L., Coxon, G., Porter, J. R., Amodeo, M. F., Chu, Z., Lewis-Gruss, S., Freeman, N. B., Houser, T., Delgado, M., ... Krajewski, W. F. (2021). Combined Modeling of US Fluvial, Pluvial, and Coastal Flood Hazard Under Current and Future Climates [eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2020WR028673>]. *Water Resources Research*, 57(2), e2020WR028673. <https://doi.org/10.1029/2020WR028673>
- Bilal, A., & Rossi-Hansberg, E. (2023, June). Anticipating Climate Change Across the United States. <https://doi.org/10.3386/w31323>
- Borusyak, K., Dix-Carneiro, R., & Kovak, B. (2022). Understanding Migration Responses to Local Shocks.
- Caliendo, L., Dvorkin, M., & Parro, F. (2019). Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica*, 87(3), 741–835. <https://doi.org/10.3982/ECTA13758>
- Crews, L. (2023). A Dynamic Spatial Knowledge Economy.
- Deryugina, T., & Hsiang, S. (2017, November). *The Marginal Product of Climate* (tech. rep. No. w24072). National Bureau of Economic Research. Cambridge, MA. <https://doi.org/10.3386/w24072>
- Desmet, K., Kopp, R. E., Kulp, S. A., Nagy, D. K., Oppenheimer, M., Rossi-Hansberg, E., & Strauss, B. H. (2021). Evaluating the Economic Cost of Coastal Flooding. *American Economic Journal: Macroeconomics*, 13(2), 444–486. <https://doi.org/10.1257/mac.20180366>
- Favilukis, J., Ludvigson, S. C., & Van Nieuwerburgh, S. (2017). The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk Sharing in General Equilibrium [Publisher: The University of Chicago Press]. *Journal of Political Economy*, 125(1), 140–223. <https://doi.org/10.1086/689606>
- Fernandez-Villaverde, J., Nuno, G., Sorg-Langhans, G., & Vogler, M. (2020). Solving High-Dimensional Dynamic Programming Problems using Deep Learning, 48.
- Fried, S. (2022). Seawalls and Stilts: A Quantitative Macro Study of Climate Adaptation. *The Review of Economic Studies*, 89(6), 3303–3344. <https://doi.org/10.1093/restud/rdab099>

- Gassert, F., Cornejo, E., & Nilson, E. (2021). Making Climate Data Accessible: Methods for Producing NEX-GDDP and LOCA Downscaled Climate Indicators. *World Resources Institute*. <https://doi.org/10.46830/writn.19.00117>
- Giannone, E., Paixão, N., Pang, X., & Li, Q. (2023). A Quantitative Spatial Equilibrium Model with Wealth.
- Greaney, B. (2021). The Distributional Effects of Uneven Regional Growth, 49.
- Han, J., Yang, Y., & E, W. (2021). DeepHAM: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3990409>
- Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D. J., Muir-Wood, R., Wilson, P., Oppenheimer, M., Larsen, K., & Houser, T. (2017). Estimating economic damage from climate change in the United States. *Science*, 356(6345), 1362–1369. <https://doi.org/10.1126/science.aal4369>
- Hsiao, A. (2023). Sea Level Rise and Urban Adaptation.
- Kaplan, G., Mitman, K., & Violante, G. L. (2020). The Housing Boom and Bust: Model Meets Evidence [Publisher: The University of Chicago Press]. *Journal of Political Economy*, 000–000. <https://doi.org/10.1086/708816>
- Kleinman, B., Liu, E., & Redding, S. J. (2023). Dynamic Spatial General Equilibrium, 46.
- Knutson, T., Camargo, S. J., Chan, J. C. L., Emanuel, K., Ho, C.-H., Kossin, J., Mohapatra, M., Satoh, M., Sugi, M., Walsh, K., & Wu, L. (2020). Tropical Cyclones and Climate Change Assessment: Part II: Projected Response to Anthropogenic Warming [Publisher: American Meteorological Society Section: Bulletin of the American Meteorological Society]. *Bulletin of the American Meteorological Society*, 101(3), E303–E322. <https://doi.org/10.1175/BAMS-D-18-0194.1>
- Komissarova, K. (2022). Location Choices over the Life Cycle: The Role of Relocation for Retirement.
- Krusell, P., & Smith, A. A., Jr. (1998). Income and Wealth Heterogeneity in the Macroeconomy. *Journal of Political Economy*, 106(5), 867–896. <https://doi.org/10.1086/250034>
- Krusell, P., & Smith, A. A., Jr. (2022). Climate Change Around the World.
- Maliar, L., Maliar, S., & Winant, P. (2021). Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122, 76–101. <https://doi.org/10.1016/j.jmoneco.2021.07.004>
- Pang, X., & Sun, P. (2023). MOVING INTO RISKY FLOODPLAINS:
- Rasmussen, D., & Kopp, R. E. (2017). Probability-Weighted Ensembles Of U.S. County-Level Climate Projections For Climate Impact Modeling [Publisher: Zenodo]. <https://doi.org/10.5281/ZENODO.582327>
- Rode, A., Carleton, T., Delgado, M., Greenstone, M., Houser, T., Hsiang, S., Hultgren, A., Jina, A., Kopp, R. E., McCusker, K. E., Nath, I., Rising, J., & Yuan, J. (2021). Estimating a social cost of carbon for global energy consumption [Number: 7880 Publisher: Nature Publishing Group]. *Nature*, 598(7880), 308–314. <https://doi.org/10.1038/s41586-021-03883-8>

- Saiz, A. (2010). The Geographic Determinants of Housing Supply [Publisher: Oxford Academic]. *The Quarterly Journal of Economics*, 125(3), 1253–1296. <https://doi.org/10.1162/qjec.2010.125.3.1253>
- Sinai, T., & Souleles, N. S. (2005). Owner-Occupied Housing as a Hedge Against Rent Risk*. *The Quarterly Journal of Economics*, 120(2), 763–789. <https://doi.org/10.1093/qje/120.2.763>
- Wing, O. E. J., Lehman, W., Bates, P. D., Sampson, C. C., Quinn, N., Smith, A. M., Neal, J. C., Porter, J. R., & Kousky, C. (2022). Inequitable patterns of US flood risk in the Anthropocene [Number: 2 Publisher: Nature Publishing Group]. *Nature Climate Change*, 12(2), 156–162. <https://doi.org/10.1038/s41558-021-01265-6>
- Yao, R., & Zhang, H. H. (2005). Optimal Consumption and Portfolio Choices with Risky Housing and Borrowing Constraints. *Review of Financial Studies*, 18(1), 197–239. <https://doi.org/10.1093/rfs/hhh007>

A Computational Method

For a large class of discrete-time heterogeneous-agent models with idiosyncratic household states $x \in \mathbf{X}$ and (possibly time-varying) aggregate state Γ , the global solution can be expressed by two systems of equations. First, the beginning-of-period value function is equal to the maximum obtainable combination of utility and expected time-discounted end-of-period value,

$$\begin{aligned}
 V^{\text{start}}(x; \Gamma) &= \max_a u(a; \Gamma) + \mathbb{E}_{x'} [\beta V^{\text{end}}(x'; \Gamma) \mid x, a, \Gamma] \\
 \text{s.t. } a &\in A(x; \Gamma),
 \end{aligned}
 \tag{8}$$

where the household can influence utility and its end-of-period state x' through the use of a (possibly vector-valued) choice variable a chosen from a feasible set $A(x; \Gamma)$.

Second, the end-of-period value function, or “continuation value” function, is equal to the expected value of the start-of-period value function in the following period, after the realization of the aggregate shock:

$$V^{\text{end}}(x; \Gamma) = \mathbb{E}_{\Gamma'} [V^{\text{start}}(x; \Gamma') \mid \Gamma].
 \tag{9}$$

I refer to (8) as the “intra-period household’s problem” and (9) as the “continuation value equation.” I divide the problem of finding the global solution into two parts corresponding to these two systems of equations. While the intra-period household’s problem is a solved problem in the sense that various methods exist to solve it, I both highly optimize and modularize the solution method by breaking a single period into a sequence of “stages” and solving each stage separately. This provides high-performance, modular code which can be easily looped over or recombined to solve a wide variety of models. I do not use neural networks at all to solve the intra-period household’s problem, obviating the need to use neural networks to approximate policy functions.

To solve the continuation value equation, I use a neural network to approximate the continuation value function and train it using a simple supervised learning algorithm. The deep learning algorithm takes the solution to the intra-period household’s problem as given, eliminating the need to specialize the deep learning algorithm to the details of the household’s problem.

By dividing the problem into two parts in this way, I can develop algorithms for each part of the problem that are both highly optimized and general. Furthermore, I can use neural networks as little as possible, only where they are needed and well-suited. This reduces room for error in the model solution and allows for greater generality in the household’s problem.

Section A.1 describes the solution method in a stationary steady state. Section A.2 describes the solution method when aggregate variables might be changing, but do so in a perfectly-foreseen way, without aggregate

uncertainty. Section A.3 describes how to use neural networks to incorporate aggregate uncertainty. Section A.4 describes a proposal for the general use of these techniques in more challenging models where aggregate uncertainty persists forever and there is a nontrivial ergodic distribution.

A.1 Solving Discrete-Time Heterogeneous-Agent Models in Stationary Steady State

The model I describe in my job market paper has a relatively large number of model features (frictions, household choices, etc.). The majority of this complexity is in the household problem. In this section, let us consider the stationary steady state of the full model, in which emissions are equal to zero, $e_t = 0 \forall t$, and the economy has converged to stationary steady state.

In the stationary steady state, equilibrium quantities (housing prices, rents, and stocks) are fixed over time. The stationary steady state can be solved by first solving the household’s problem given equilibrium quantities, then, in an outer loop, solving for equilibrium quantities to satisfy market-clearing conditions.

Algorithm 1 Stationary Steady State Outer Loop

1. Guess housing prices and rents for each location
 2. Solve household problem for all households
 3. Compute market clearing conditions
 4. Update housing price and rent guesses and repeat until convergence
-

A.1.1 Solving the Household’s Problem

In the stationary steady state, while household value functions do not depend on time, the household’s problem is still quite complex. The household’s idiosyncratic transition each period has many elements: choosing location, choosing housing investment, paying realtor’s fees, etc. I break this complexity into manageable pieces by treat these elements as occurring one at a time and solve for each one independently. That is, I treat a period of the household’s problem as consisting of several *stages*, where the idiosyncratic transition happening in each stage is simple and manageable.

This leverages a useful property of discrete-time models. While, in a continuous-time model, all stages (elements of the household’s idiosyncratic transition) are happening “simultaneously,” in a discrete time model they can happen one at a time.

Intuitively, each stage of the household’s problem consists of one thing happening to a household or one choice made by the household. We can then backwards induct through each stage in a period to solve

the household’s problem. Crucially, because these stages happen sequentially, the computation time needed to solve the household’s problem is *linear* in the number of stages (but still exponential in the number of idiosyncratic state variables). In principle, this allows us to solve household’s problems with arbitrarily many elements in the household’s transition, as long as the number of idiosyncratic state variables remains manageable. I provide self-contained code implementing each stage, so that stages can be added, removed, and reordered at virtually zero cost of programmer time.⁶⁵

I now formally define a “stage” and the way in which stages can be computationally approximated in a uniform, composable way. The formalism is a little abstract, but the basic idea is simple: each stage represents a simple household transition. For each stage, a Bellman equation implies a backwards induction function taking continuation values to beginning-of-stage values. The purpose of the abstraction is to show that every element of the household’s problem can be represented in a unified way by the same general abstract structure that is amenable to computation.⁶⁶

A.1.2 Decomposing a Complex Idiosyncratic Decomposition into Stages

Consider one element of the household’s idiosyncratic transition, such as the arrival of income, the consumption-savings decision, or the application of a borrowing constraint. Assume that the

Definition 1. *Stage.* A stage $S = (\mathbf{X}^{\text{pre}}, \mathbf{X}^{\text{post}}, \Xi, \Pi)$ consists of a beginning-of-stage household state space \mathbf{X}^{pre} , an end-of-stage household state space \mathbf{X}^{post} , a “backward induction operator” taking an end-of-period continuation value function V^{post} to a beginning-of-period value function V^{pre} ,

$$\begin{aligned} \Xi : \mathbb{R}^{\mathbf{X}^{\text{post}}} &\longrightarrow \mathbb{R}^{\mathbf{X}^{\text{pre}}} \\ &: V^{\text{post}} \mapsto V^{\text{pre}} \\ &: (V^{\text{post}} : \mathbf{X}^{\text{post}} \rightarrow \mathbb{R}) \mapsto (V^{\text{pre}} : \mathbf{X}^{\text{pre}} \rightarrow \mathbb{R}), \end{aligned}$$

and a “forward simulation operator” taking a beginning-of-stage household distribution λ^{pre} and an end-of-stage value function V^{post} to an end-of-stage household distribution λ^{post} ,

$$\begin{aligned} \Pi : \Lambda(\mathbf{X}^{\text{pre}}) \times \mathbb{R}^{\mathbf{X}^{\text{post}}} &\longrightarrow \Lambda(\mathbf{X}^{\text{post}}) \\ &: (\lambda^{\text{pre}}, V^{\text{post}}) \mapsto \lambda^{\text{post}} \\ &: ((\lambda^{\text{pre}} : \mathcal{P}(\mathbf{X}^{\text{pre}}) \rightarrow \mathbb{R}^+), (V^{\text{post}} : \mathbf{X}^{\text{post}} \rightarrow \mathbb{R})) \mapsto (\lambda^{\text{post}} : \mathcal{P}(\mathbf{X}^{\text{post}}) \rightarrow \mathbb{R}^+), \end{aligned}$$

⁶⁵These “stage functions” could even be written by different researchers and combined to easy program a large variety of models.

⁶⁶This is true in my model but certainly not over the space of all possible models, and in fact relies on the property that household decisions do not interact with each other except through prices. There is some analogy here to the increasingly-popular mean-field-game approach.

where $\Lambda(\mathbf{X})$ is the space of measures $\lambda : \mathcal{P}(\mathbf{X}) \rightarrow \mathbb{R}^+$ on \mathbf{X} .

Alternatively, we can think of a stage as a single operator iterating V^{post} backwards and λ^{pre} forwards:

$$\begin{aligned} S : \Lambda(\mathbf{X}^{\text{pre}}) \times \mathbb{R}^{\mathbf{X}^{\text{post}}} &\longrightarrow \Lambda(\mathbf{X}^{\text{post}}) \times \mathbb{R}^{\mathbf{X}^{\text{pre}}} \\ : (\lambda^{\text{pre}}, V^{\text{post}}) &\mapsto (\lambda^{\text{post}}, V^{\text{pre}}). \end{aligned}$$

Furthermore, we can compose two stages to form a composite stage.⁶⁷

Definition 2. *Composition of Stages.* Let $S_1 = (\mathbf{X}_1^{\text{pre}}, \mathbf{X}_1^{\text{post}}, \Xi_1, \Pi_1)$ and $S_2 = (\mathbf{X}_2^{\text{pre}}, \mathbf{X}_2^{\text{post}}, \Xi_2, \Pi_2)$ be two stages. Their composition $S_{12} = S_1 \circ S_2 = (\mathbf{X}_1^{\text{pre}}, \mathbf{X}_2^{\text{post}}, \Xi_{12}, \Pi_{12})$ is a stage given by,

$$\begin{aligned} \Xi_{12} &\equiv \Xi_1 \circ \Xi_2 : \mathbb{R}^{\mathbf{X}_2^{\text{post}}} \longrightarrow \mathbb{R}^{\mathbf{X}_1^{\text{pre}}} \\ \Pi_{12} &: \Lambda(\mathbf{X}_1^{\text{pre}}) \times \mathbb{R}^{\mathbf{X}_2^{\text{post}}} \longrightarrow \Lambda(\mathbf{X}_2^{\text{post}}) \\ \Pi_{12}(\lambda_1^{\text{pre}}, V_2^{\text{post}}) &= \Pi_2(\Pi_1(\lambda_1^{\text{pre}}, \Xi_2(V_2^{\text{post}})), V_2^{\text{post}}). \end{aligned}$$

Thus, an entire period can be thought of as a single stage. The purpose of this, however, is to decompose periods into the simplest possible stages.

This level of abstraction might seem unnecessary. but the essential point is that, for each stage, all we need to do is map a value function backwards and map a household distribution forwards. In particular, we can make these stages individually very simple and then compose stages to define complex models. We can solve those models by composing the numerical approximations of each stage.

To make things more concrete, I describe how one period of the household problem of a simple Aiyagari model can be subdivided into five stages: receiving income, the consumption-savings decision, the application of a borrowing constraint, an idiosyncratic income shock, and the passage of time.

For each example stage, let an individual household's beginning-of-stage state be given by

$$x = (b, z) \in \mathbf{X} = \mathbf{X}^{\text{pre}} = \mathbf{X}^{\text{post}},$$

where b represents the household's bondholdings and z represents the household's income. The beginning-of-period state space \mathbf{X}^{pre} and end-of-period state space \mathbf{X}^{post} are equal. Let the end-of-state continuation value function be given by $V^{\text{post},s}(b, z)$ for $s \in \{\text{income, consume, constraint, shock, time}\}$.⁶⁸ To save on notation, I omit the stage-specific superscript within each example.

⁶⁷This composition rule is associative and there is a natural "identity" stage on any state space \mathbf{X} which changes neither V nor λ . It might be of some conceptual interest to note that we can therefore define a mathematical category **Stage** with the stages as the morphisms and the state spaces as the objects.

⁶⁸I abuse notation somewhat in not writing $V^{\text{post},s}((b, z))$.

Example 1. Receiving Income.

Suppose that after receiving income, a household with initial state (b, z) has continuation state $(Rb + z, z)$. The household's backward induction operator is given by,

$$\Xi^{\text{income}} V^{\text{post}}(Rb, z) = V^{\text{post}}(Rb + z, z),$$

and the forward simulation operator is given by,

$$\Pi^{\text{income}} \lambda^{\text{pre}}(T \subseteq \mathbf{X}) = \lambda^{\text{pre}}(\{(Rb, z) \mid (Rb + z, z) \in T\}).$$

A similar thing is true for any deterministic household transition $f : \mathbf{X}^{\text{pre}} \rightarrow \mathbf{X}^{\text{post}}$.⁶⁹

Example 2. Consumption-Savings Decision.

Suppose that a household with initial state (b, z) chooses continuation state (b', z) maximizing

$$\Xi^{\text{consumption}} V^{\text{post}}(b, z) = u(b'^* - b) + V^{\text{post}}(b'^*, z)$$

$$\text{where } b'^*(b, z; V^{\text{post}}) \equiv \arg \max_{b'} u(b' - b) + V^{\text{post}}(b', z)$$

This defines the backward induction operator of a consumption-savings decision stage.⁷⁰ The forward simulation operator is given by

$$\Pi^{\text{consumption}} \lambda^{\text{pre}}(T \subseteq \mathbf{X}; V^{\text{post}}) = \lambda^{\text{pre}}(\{(b, z) \mid (b'^*(b, z; V^{\text{post}}), z) \in T\}).$$

Example 3. Borrowing Constraint.

Suppose that a household is subject to the borrowing constraint $b \geq 0$. We can then define a borrowing-constraint stage by,

$$\Xi^{\text{constraint}} V^{\text{post}}(b, z) = \begin{cases} V^{\text{post}}(b, z) & b \geq 0 \\ -\infty & b < 0 \end{cases}$$

$$\Pi^{\text{constraint}} \lambda^{\text{pre}} = \lambda^{\text{pre}}.$$

Example 4. Idiosyncratic Income Shock.

⁶⁹Specifically, for any such f , the backward induction operator is just the pullback or precomposition of V^{post} by f and the forward simulation operator is the pullback of λ^{pre} by the inverse image map $f^{-1} : \mathcal{P}(\mathbf{X}^{\text{post}}) \rightarrow \mathcal{P}(\mathbf{X}^{\text{pre}})$ associated with f .

⁷⁰Note that b' is fully unconstrained here, and there is no time-discounting; these will occur in different stages.

Suppose that a household's income z evolves according to

$$z' \sim \mathcal{D}(z)$$

for some distribution \mathcal{D} . Then the idiosyncratic income shock stage is given by,

$$\begin{aligned} \Xi^{\text{shock}} V^{\text{post}}(b, z) &= \mathbb{E}_{z'} [V^{\text{post}}(b, z') \mid z] \\ \Pi^{\text{shock}} \lambda^{\text{pre}}(T; V^{\text{post}}) &= \int_{(b, z)} P((b, z') \in T \mid z) d\lambda^{\text{pre}} \end{aligned}$$

where $T \subseteq \mathbf{X}$.

Example 5. Passage of Time.

Suppose that time passes between the beginning and end of the stage, with discount factor β . Then, the passage-of-time phase is given by,

$$\begin{aligned} \Xi^{\text{time}} V^{\text{post}}(b, z) &= \beta V^{\text{post}}(b, z) \\ \Pi^{\text{time}} \lambda^{\text{pre}} &= \lambda^{\text{pre}}. \end{aligned}$$

At this point, this might well seem like a great deal of formalism with no interesting content. The purpose of this all, however, is just to show that solving a model with a complex household transition is not especially difficult, as long as we can break the period into stages and numerically approximate each stage. For example, the full list of stages in my full model are:

- Apply impact of house price changes on household net worth.
- Decide whether or not to move.
- Choose location if moving.
- Choose whether to sell owner-occupied housing.
- Choose whether to buy owner-occupied housing, and how much.
- Choose whether to sell rental real estate.
- Choose whether to buy rental real estate, and how much.
- Receive income and pay expenses.
- Make consumption-saving decision.

- Receive idiosyncratic income shock.
- Passage of time.
- End-of-life and bequest (if at maximum age).

A.1.3 Numerically Solving Each Stage

The first step in numerically solving a stage is to discretize the state space. My full model has six idiosyncratic state variables: bondholdings, income type, owner-occupied housing, rental real estate, location, and age. I use a separate array for each age group. For the other five variables, I define a five-dimensional array which is the product of portfolio value, income type, owner-occupied housing, rental housing, and location grids. I use portfolio value (bondholdings + real estate value) instead of bondholdings for the first dimension so that buying or selling real estate does not move households around the grid too dramatically. While I use the full product grid, significant savings might be possible with a dynamic sparse grid approach, such as Brumm et al. (2022) use for continuous-time models.

I then represent both value functions V and household state distributions λ as arrays A_V and A_λ of this form, and construct a pair of arrays for each (more or less) stage and age group. A_V represents the interpolation nodes of the function V on the grid and A_λ represents a grid of point masses on those same grid points. I use the same grid for every stage to improve the reorderability and composability of the stages. I then implement each stage as a pair of functions taking arrays to arrays: one taking end-of-stage V^{post} to beginning-of-stage V^{pre} and one taking end-of-stage V^{post} and beginning-of-stage λ^{pre} to end-of-stage λ^{post} .

Four challenges arise in doing this in a performant manner. First, reinterpolating V^{pre} onto the grid points after each backwards induction step. Second, reinterpolating λ^{post} back onto the grid points after each forward simulation step. Third, choosing optimal consumption for each grid point at the consumption-savings step. Fourth, simulating implementing the location choice stage, which requires a household at each grid point to choose between 1713 locations, with over 200 million gridpoints. I provide optimized algorithms implemented in Julia for overcoming each of these bottlenecks.

Because my model is a lifecycle model, I can begin by computing the bequest value at each terminal gridpoint, then solve the entire household’s problem backwards from there. For age group a , the value function V_a^{end} at the end of each period is simply equal to the value function at the start of the following period for the next age group,

$$V_a^{\text{end}} = V_{a+1}^{\text{start}}.$$

To solve for prices, I apply the outer loop algorithm I describe above.

A.2 Transition Dynamics

Having solved for stationary steady state, we can now extend to transition dynamics. Consider a version of my main model with perfect foresight after the onset shock. That is, while the onset of climate change is unexpected, after the onset shock all households understand that climate climate will continue deterministically along the median path, where $\varepsilon_t^m = 0 \forall t$. To solve this, we can simply pick an end date far in the future (I choose 2300) at which we assume that the economy has returned to its stationary steady state. Then, given prices, we can simply use backwards induction to solve value functions from 2300 to 1990, then simulate the household state distribution forward from 1990 to 2300. We can then solve for prices in an outer loop.

Algorithm 2 Deterministic Transition Outer Loop

1. Solve 2300 stationary steady state.
 2. Guess housing prices and rents for each location and decade.
 3. Solve household problem for all households across entire transition path, from 1990 onset shock to 2300 new steady state.
 4. Compute market clearing conditions.
 5. Update housing price and rent guesses and repeat from Step 3 until convergence.
-

For age group a , the value function V_{at}^{end} at the end of period t , is simply equal to the value function for the following age group at the start of the following period,

$$V_{at}^{\text{end}} = V_{a+1,t+1}^{\text{start}}.$$

A.3 Aggregate Risk

To solve the model with aggregate risk, I apply a similar strategy. The difference is that value function now depend on the aggregate state Γ , not only on the idiosyncratic state and a time index. The aggregate state changes between the end of the household's problem in one period and the beginning of the household's problem in the following period,

$$V_{at}^{\text{end}}(x; \Gamma) = \mathbb{E}_{\Gamma'} [V_{a+1,t+1}^{\text{start}}(x; \Gamma') \mid \Gamma].$$

Note that there is no need to solve the policy function. The policy function is calculated on-the-fly during

the simulation of the model within each period. Because we do not need to solve for the policy function, we only need to use neural networks to approximate the value function (and, as an optimization, the price function). This is essentially a form of supervised learning, which is generally much easier and more stable than using a neural network to solve for a policy function.

To use a neural network to approximate $V_{at}^{\text{end}}(x; \Gamma)$ requires two steps. First, I must reduce the dimension of Γ somehow. Second, I must train the neural network.

Dimension Reduction There are many ways to reduce the dimension of Γ . One fairly robust way to do it is introduced in Han, Yang, and E (2023) with the DeepHAM method. However, I simply exploit the fact that the aggregate state in my model is simply a function of initial stationary steady state (which does not change and thus does not need to be inputted into the neural network) and the history of aggregate shocks ε_t^m following the onset shock.⁷¹ Thus, the neural net approximator for $V_{at}^{\text{end}}(x; \Gamma)$ requires only $5 + t$ inputs, where t is the number of periods since the onset shock. Let $\mathcal{N}_{at}^{\text{end}}$ represent this neural net approximator. Then,

$$V_{at}^{\text{end}}(x; \Gamma) = \mathcal{N}_{at}^{\text{end}}(x; v).$$

I simply define a separate neural network approximator $\mathcal{N}_{at}^{\text{end}}$ for each age a and period t .

Training To train this neural network, I first discretize the aggregate shock into 5 possible states for each ε_t^m , representing the median and each one- or two-standard deviation surprise. Let $\{\eta_k\}_{k=1}^5$ represent each of these five realizations. Then the following holds with equality:

$$V_{at}^{\text{end}}(x; \Gamma) = \sum_{k=1}^5 V_{a+1, t+1}^{\text{start}}(x; \Gamma'(\Gamma, k)) P(\varepsilon_t^m = \eta_k). \quad (10)$$

I can then train the neural network by the following algorithm:

⁷¹I thank Robert Wagner for this idea.

Algorithm 3 Neural Network Training Loop

1. Solve 2300 stationary steady state.
 2. Initialize a neural network $\mathcal{N}_{at}^{\text{end}}$ for each age a and time t .
 3. Draw a path of aggregate shocks $\{\varepsilon_t^m\}_{t=1}^T$ and compute the resulting temperature path $\{\text{SST}_t^m\}_{t=1}^T$.
 4. Working forwards, for each age a and time τ :
 - (a) Guess period- t prices $\{\hat{\rho}_{it}(\varepsilon_1^m, \dots, \varepsilon_t^m), \hat{q}_{it}(\varepsilon_1^m, \dots, \varepsilon_t^m)\}_{i \in I}$.
 - (b) Compute from neural network $\hat{V}_{at}^{\text{end}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m) = \mathcal{N}_{at}^{\text{end}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m)$ at each grid point x .
 - (c) Simulate $\lambda_{at}^{\text{start}}$ forward based on this $\hat{V}_{at}^{\text{end}}$.
 - (d) Compute market clearing conditions within period t .
 - (e) Repeat from Step 4a until market clearing conditions satisfied within tolerance.
 5. Working backwards, for each age a and time τ :
 - (a) For each $k \in \{1, \dots, 5\}$:
 - i. Guess period- $t+1$ prices $\{\rho_{it+1}(\varepsilon_1^m, \dots, \varepsilon_t^m, \eta_k), q_{it+1}(\varepsilon_1^m, \dots, \varepsilon_t^m, \eta_k)\}_{i \in I}$ for each location i .
 - ii. Compute $\hat{V}_{at+1}^{\text{end}}(x; \Gamma, \varepsilon_1^m, \dots, \varepsilon_t^m, \eta_k) = \mathcal{N}_{at}^{\text{end}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m, \eta_k)$.
 - iii. Simulate $\lambda_{at+1}^{\text{start}}$ forward based on this $\hat{V}_{at+1}^{\text{end}}$.
 - iv. Compute market clearing conditions within period $t+1$.
 - v. Repeat from step 5a(i) until market clearing conditions satisfied within tolerance.
 - vi. Iterate backwards to obtain $\hat{V}_{at+1}^{\text{start}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m, \eta_k)$.
 - (b) Compute
$$\tilde{V}_{at}^{\text{end}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m) = \sum_{k=1}^5 \hat{V}_{at+1}^{\text{start}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m, \eta_k) P(\varepsilon_t^m = \eta_k)$$
 - (c) Update $\mathcal{N}_{at}^{\text{end}}$ with error target $\mathcal{N}_{at}^{\text{end}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m) - \tilde{V}_{at}^{\text{end}}(x; \varepsilon_1^m, \dots, \varepsilon_t^m)$.
-

In practice, I move the determination of prices to the outer loop and use a neural network to guess prices once for each draw of the aggregate shock path. I then update the price network using excess supply of rental real estate and rental properties for each location: the errors in the market clearing conditions. This is the approach of Azinovic, et al. (2023)

A.4 Greater Generality

A.5 Optimized Algorithms for Interpolations, Choice Probabilities, and Optimal Choice

The primary numerical methods contribution in this paper is the use of deep learning (neural networks) to solve for prices and household conditional continuation values, as a function of aggregate uncertainty, given the aggregate state.

The key insight of the deep learning method is that, if a household knows the current utility and continuation values resulting from any action, then conventional methods can be used to solve the policy function. With one or two dimensions of heterogeneity, this is generally no problem. But because the model has so many individual states (213 million), it was necessary to develop a few custom algorithms for performance.

These supplementary algorithms overcome some computational bottlenecks when scaling up discrete-time heterogeneous-agent models to large numbers of individual states and locations. In the hope that someone finds them useful, I have made a separate Github repository for them, in highly-optimized parallelized Julia code.

A.5.1 All-at-once interpolating functions and distributions on endogenous gridpoints

Speedup factor: $\log n$ ($n = \#(\text{gridpoints})$, relative to individual lookups)

When solving value functions using backwards induction or simulating distributions forward, if the gridpoints are defined in terms of endogenous quantities (e.g. wealth), you often have to reinterpolate the function or distribution back onto the grid. Existing interpolation libraries look up each gridpoint independently. With n gridpoints, each query is $O(\log n)$ and the whole reinterpolation is $O(n \log n)$. But if the grid is monotonic, you only need to iterate over each grid once, which you can do in $O(n)$ time.

Reinterpolating is slightly different for functions and distributions, since distribution approximations depend on the density of gridpoints.

A.5.2 Computing choice probabilities for heterogeneous agents with a single matrix multiplication

Speedup factor: up to $n^{0.62}$ or so, with big improvement on constant (if $n = \#(\text{locations})$, relative to naive approach)

With n locations and m possible within-location agent states, if agents receive Gumbel-distributed location preferences, then to solve for the value function and choice probabilities, you have to solve for $n*n*m$ possible agent-choice pairs. Luckily, this can be constructed as a single matrix multiplication, which can even be done non-allocatingly if set up properly. Highly-optimized linear algebra libraries can then be used.

A.5.3 All-at-once solving monotone decision function for heterogenous agents

Speedup factor: n ($n = \#(\text{gridpoints})$, relative to naive individual lookups. Speedup factor $\log n$ over optimized individual lookups)

In many models, agents must make consumption savings decisions. Assuming a discretized wealth grid, the most flexible solution is to maximize over all possible continuation values for each agent, but this requires $O(n^2)$ total computations. If both continuation value and optimal saving are increasing in wealth, each agent needs only consider possible choices that lie between the choices of adjacent agents. By using binary search for individual lookups and a form of modified binary search to constrain the search space for individual agents, we can compute the optimal decisions of n agents who differ only by wealth using only $O(n)$ computations.

B Full Household's Problem

In sum, the household's problem is given by

$$\begin{aligned}
& V_t(b, z, h_0^{\text{live}}, h_0^{\text{let}}, \ell_0, a, \{\varepsilon_l\}_l) \tag{11} \\
& = \max_{h^{\text{live}}, h^{\text{let}}, \ell, g, h} \varepsilon_l - \tau d(\ell, \ell_0) + u(g, h, \alpha_{\ell t}) + \beta \mathbb{E} V_{t+1}(b', z', h^{\text{live}'}, h^{\text{let}'}, \ell, a + 1, \{\varepsilon'_l\}_l) \\
\text{s.t. } & u(g, h, \alpha_{\ell t}) = \alpha_{\ell} \frac{(g^{1-\sigma} + \gamma h^{1-\sigma})^{\frac{1-\eta}{1-\sigma}} - 1}{1-\eta} \\
& b' = R(\tilde{b}) + A_{\ell t} e^z + \rho_{\ell t} h^{\text{let}} - g - \rho_{\ell t} h^{\text{rent}} - \chi_{\ell t}^{\text{live}} h^{\text{live}} - \chi_{\ell t}^{\text{let}} h^{\text{let}} \\
& \tilde{b} = b - F^m D^{\text{move}} + \sum_{s \in \{\text{live}, \text{let}\}} D^{\text{sell}, s} [(1 - \phi) q_{\ell_0 t} h_0^s - q_{\ell t} h^s] \\
& R(\tilde{b}) = \begin{cases} (1 + r^f) \tilde{b} & \tilde{b} \geq 0 \\ (1 + r^m) \tilde{b} & \tilde{b} < 0 \end{cases} \\
& b' \geq -\kappa q_{\ell t} (h^{\text{live}} + h^{\text{let}}) \\
& h = \begin{cases} h^{\text{live}} & h^{\text{live}} > 0 \\ h^{\text{rent}} & \text{otherwise} \end{cases} \\
& D^{\text{move}} = \mathbb{1}[\ell \neq \ell_0] \\
& D^{\text{sell}, s} = D^{\text{move}} \vee \mathbb{1}[h_0^s \neq h^s] \quad \forall s \in \{\text{live}, \text{let}\}
\end{aligned}$$

Expected continuation utility is defined as

$$\mathbb{E} V_{t+1}(b', z', h^{\text{live}'}, h^{\text{let}'}, \ell, a + 1) \equiv \mathbb{E}_{z', h^{\text{live}'}, h^{\text{let}'}, \{\varepsilon'_l\}_l} [V_{t+1}(b', z', h^{\text{live}'}, h^{\text{let}'}, \ell, a + 1, \{\varepsilon'_l\}_l) \mid z, h^{\text{live}}, h^{\text{let}}]$$

with stochastic transitions,

$$\begin{aligned}
\text{s.t. } z_{it+1} &= pz_{it} + \zeta_{a_{it}} + \nu_{it} \\
\nu &\sim \mathcal{N}(0, \sigma_\nu) \\
\varepsilon_{lit} &\sim \text{Gumbel}(0, 1)
\end{aligned} \tag{12}$$

The properties of the Gumbel distribution imply that households choose locations in proportion to the (exponentiated) value of living in that location (see Appendix B.1).

B.0.1 Household Distribution Transition Function

The individual policy function solving (11) maps beginning-of-period state $x \in \mathbf{X}$ to chosen end-of-period (pre-shock) state $x^{\text{choice}} \in \mathbf{X}^{\text{choice}}$,

$$\begin{aligned}
\text{Policy}_t(x, \{\varepsilon_\ell\}_\ell) &= x^{\text{choice}} \in \mathbf{X}^{\text{choice}} \\
x^{\text{choice}} &= (b', z, h^{\text{live}}, h^{\text{let}}, \ell, a).
\end{aligned}$$

For $\mathbf{X}_0^{\text{choice}} \subseteq \mathbf{X}^{\text{choice}}$, taking $\{\varepsilon_\ell\}_\ell$ as unknown, the structure of the Gumbel distribution gives,

$$\Pr_\varepsilon(\ell \mid b, z, h^{\text{live}}, h^{\text{let}}, \ell_0, a + 1; \Gamma_t) = \frac{\exp\left(\psi \tilde{V}_\ell(b, z, h_0^{\text{live}}, h_0^{\text{let}}, \ell_0, a)\right)}{\sum_l \exp\left(\psi \tilde{V}_l(b, z, h_0^{\text{live}}, h_0^{\text{let}}, \ell_0, a)\right)}$$

where

$$\tilde{V}(b, z, h_0^{\text{live}}, h_0^{\text{let}}, \ell_0, a) = \max_{h^{\text{live}}, h^{\text{let}}, g, h} \tau d(\ell, \ell_0) + u(g, h, \alpha_{\ell t}) + \beta \mathbb{E}V_{t+1}(b', z', h^{\text{live}'}, h^{\text{let}'}, \ell, a + 1, \{\varepsilon'_l\}_l),$$

subject to the same constraints as in (11). This naturally defines idiosyncratic and distributional transition functions from \mathbf{X} to $\mathbf{X}^{\text{choice}}$. For $\mathbf{X}_0^{\text{choice}} \subseteq \mathbf{X}^{\text{choice}}$,

$$\begin{aligned}
P_t(\mathbf{X}_0^{\text{choice}} \mid x) &= \Pr_\varepsilon(\text{Policy}_t(x, \{\varepsilon_\ell\}_\ell) \in \mathbf{X}_0^{\text{choice}}) \\
\lambda_t^{\text{choice}}(\mathbf{X}_0^{\text{choice}}) &= \int_{x \in \mathbf{X}} P_t(\mathbf{X}_0^{\text{choice}} \mid x) d\lambda_t(x).
\end{aligned}$$

The idiosyncratic transition function (12) defines idiosyncratic and distributional transition functions

from $\mathbf{X}^{\text{choice}}$ to \mathbf{X} . For $\mathbf{X}'_0 \subseteq \mathbf{X}$,

$$T_t(\mathbf{X}'_0 | x^{\text{choice}}) = \Pr(x' \in \mathbf{X}'_0 | x^{\text{choice}})$$

$$\lambda_{t+1}(\mathbf{X}'_0) = \int_{x \in \mathbf{X}} T_t(\mathbf{X}'_0 | x^{\text{choice}}) d\lambda_t^{\text{choice}}(x).$$

B.1 Extreme-Value Location Preference Shocks

Integrating over the location preference shocks $\{\varepsilon'_l\}_l$ gives

$$\mathbb{E}V_{t+1}(b', z', h^{\text{live}'}, h^{\text{let}'}, \ell, a + 1) = \psi^{-1} \log \left(\sum_{\ell} \exp \left(\psi \tilde{V}_{\ell}(b, z, h_0^{\text{live}}, h_0^{\text{let}}, \ell_0, a) \right) \right)$$

where

$$\tilde{V}_{\ell}(b, z, h_0^{\text{live}}, h_0^{\text{let}}, \ell_0, a) = \max_{h^{\text{live}}, h^{\text{let}}, g, h} \tau d(\ell, \ell_0) + u(g, h, \alpha_{\ell t}) + \beta \mathbb{E}V_{t+1}(b', z', h^{\text{live}'}, h^{\text{let}'}, \ell, a + 1, \{\varepsilon'_l\}_l).$$

C Harmonization of Climate Damage Estimates

C.1 Labor Productivity

I take estimates of the effect of temperature on wages from Deryugina and Hsiang, 2017. They use a panel of U.S. counties from 1969 to 2014 to estimate the effect of temperature on wages, using idiosyncratic weather variation. Although they do not study the effect of long-term adaptation, they argue that, in equilibrium, the marginal response of income to short-term weather realizations is equal to the marginal effect of income to long-term climate variation. They use an envelope-theorem argument: in the neighborhood of equilibrium temperatures, the benefit of long-term climate adaptations must be zero.

They thus estimate the effect of daily temperatures on log annual total income per capita. Their Figure 4a (my Figure 16) shows their estimates. They choose 55°F as the reference temperature, because it is the temperature with the greatest positive effect on productivity in their estimation. I recenter at 65°F. Upon recentering, the marginal effect of temperatures above 65°F remains fairly constant. The overall effect of temperatures below 65°F is negative but small, with an average effect of 6.6×10^{-6} per °C-day between 15°F and 65°F, an order of magnitude smaller than the average effect of -4.5×10^{-5} per °C-day between 65°F and 85°F. It is true, however, that the marginal effect of temperatures below 65°F is not constant, and it would be ideal to use projections of HDD and CDD centered at 55°F. Table 1 shows these extracted estimates.

Using climate projections on HDD and CDD changes per °C global warming, as in Section 3.3.2, I compute income damages for each county per °C global warming. Figure 2 displays the resulting projections

of income reductions for each PUMA.

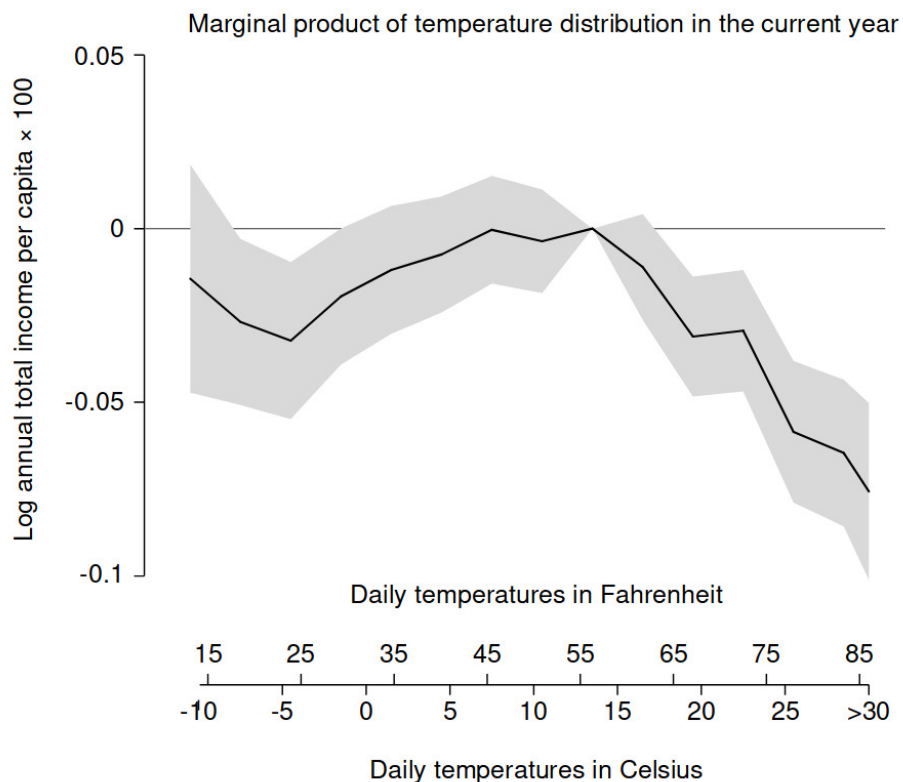


Figure 16: Impact of a single day at a certain temperature on log annual total income. Reprinted from Deryugina and Hsiang, 2017, Figure 4, Panel 1.

C.2 Disaster Projections

I use projections of climate impacts on total combined damages from flooding from Bates et al. (2021), as reported in Wing et al. (2022).

I argue that their projections, together with the estimates I use on heat and cold amenities, capture the great majority of climate impacts on natural disasters. Of the 18 FEMA hazard categories, “Earthquake” and “Volcano” are likely not climate-sensitive. “Wildfire,” “Heat Wave,” “Cold Wave,” “Winter Weather,” “Ice Storm,” “Drought,” “Avalanche,” and “Hail” are not separately identified from the estimates on heat and cold climate amenities which I take from the literature (Section 3.3.2). This leaves primarily flooding and wind-related hazards and wind-related hazards—“Coastal Flooding,” “Landslide,” “Riverine Flooding,” “Tsunami,” “Hurricane,” “Tornado,” “Strong Wind,” “Lightning.”

The projections I use are likely to capture the great majority of our current best understanding of the climate-sensitivities of flooding and wind-related disasters. These projections include all forms of flooding, including those which happen in conjunction with storms, leaving primarily wind-related damages. Current

projections of the impact of climate change on wind-related damages *per se* are highly uncertain, with confidence intervals centered near zero.

A recent meta-analysis of projections of tropical cyclone activity under climate change (Knutson et al., 2020), find confident projections of increasing tropical cyclone precipitation, but small and highly uncertain changes in terms of wind speeds. Among the models they survey, median projections for a 2° warming scenario predict a 16% *decrease* in tropical cyclone frequencies in the North Atlantic basin, but a 3% increase in average intensity. The net result is a 10% increase in the annual number of category 4-5 cyclones, with an interquartile range of [-21%, 28%]. Cyclone intensity refers to wind speed only, and this uncertainty does not extend to precipitation and flooding. The median projected increase in precipitation under the 2°C warming scenario is 16% with interquartile range [8%, 21%]. Flooding due to storm surges is also projected to increase, due to rising sea levels (Little et al. 2015). Indeed, the greatest impact of climate change on property damages is expected to take the form of these “compound events,” (AghaKouchak et al., 2020) and compound events which include flooding are encompassed the projections I use.

As a test of this coverage, I compare the mean annual residential flooding damages, or Average Annual Losses (AAL), from all contiguous U.S. flooding, under 2020 climate conditions, predicted by the model of Bates et al. (2021). Their prediction is \$30bn, which actually exceeds the mean annual damages in SHELDUS between 2011 and 2020 in the “Coastal Flooding,” “Hurricane,” “Riverine Flooding,” “Severe Storm,” and “Tsunami” categories. These realized damages are \$21.39bn (2021 USD) and account for 71% of all realized property damage from disasters in that period.

Due to data limitations, I have only two variables per county from the analysis of Bates et al. (2021). The first is Average Annual Losses (AAL), per county, given 2020 climate conditions. The second is the percentage change in AAL per county between 2020 and 2050 under RCP 4.5. In order to harmonize this with the other estimates, I divide the AAL change by 0.63°C, the median global warming under RCP 4.5 in that time period (Rasmussen and Kopp, 2017). This gives me a linearized measure of the projected increase in AAL per °C global warming:

$$\frac{d}{dT} \widehat{\text{AAL}}_{\ell}(T) = \frac{\text{AAL}_{\ell,2050} - \text{AAL}_{\ell,2020}}{0.63 \text{ }^{\circ}\text{C}}$$

C.3 Heat and Cold Amenities

I take estimates of heat and cold amenities from Albouy et al., 2016, who use a Rosen-Roback style model and the IPUMS 5% sample of the 2000 Census to impute amenities for each U.S. PUMA as the residual of real wage differentials. Thus, I am not double-counting the impact of temperatures on wages. They factor energy costs into real wages, so that I am not double-counting the impact of energy costs. They estimate a range of specifications.

I pick the richest model that is consistent with my model, reported in their Figure 5F (my Figure 17b). The dependent variable is willingness-to-pay (WTP) to avoid a single day at a certain temperature, and experience a day at 65°F instead. They fit WTP to a four-piece linear spline with geographic, non-temperature weather, and demographic controls, and state fixed effects. All other specifications use fewer controls, except for their Figure 5E (my Figure 17a), which fits WTP to a 7th-degree cubic spline.

The cubic spline specification is very close to linear for temperatures less than 65°F, and close to linear between 65°F and 85°F, but constant above 85°F, possibly due to lack of data. The linear specification separately estimates slopes below 45°F, between 45°F and 65°F, between 65°F and 80°F, and above 80°F. They find similar slopes for both segments below 65°F and both segments above 65°F. I therefore believe that it is not unreasonable to assume constant, but different, marginal effects for days below and above 65°F, respectively. I take their estimates for WTP at each extreme, 0°F and 110°F, and assume a constant marginal value of each degree-day below and above 65°F.⁷²

Assuming constant marginal effects of each degree-day below and above 65°F allows me to impute cold disamenities as a function of annual heating-degree-days (HDD), the sum over each day t of the year

$$\sum_t \max(65^\circ\text{F} - T_t, 0)$$

of the number of degrees by which that day's temperature T_t falls below 65°F. I similarly impute heat disamenities as a function of annual 65°F-baseline cooling-degree-days (CDD). I report these WTP per degree-day estimates in Table 2.

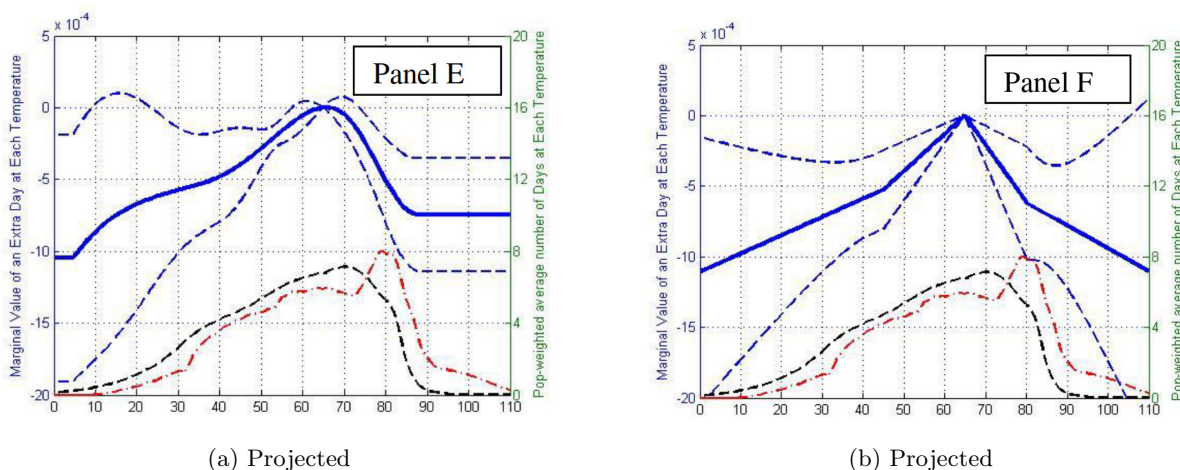


Figure 17: Marginal value of a day at a given temperature, relative to a day at 65°F, as a percentage of annual income. Reprinted from Albouy et al., 2016, Figure 9.

⁷²Equivalently, I take the arithmetic mean of each pair of segments, weighted by segment length.

C.4 Amenities by Margin

D Imputation Procedures

E State Space Model for Insurance Estimation

E.1 Data Sources

I use annual data for the U.S. between 2003 and 2020, compiled by the Insurance Information Institute (III), on average homeowners insurance premiums, average losses per policy-year, sources of losses, and sources of expenses. Expenses are divided first into incurred losses (i.e. payouts to policyholders), adjustment expenses, and underwriting and other expenses. Incurred losses are categorized by hazard (e.g. water damage, theft, etc.). These data are originally reported by the National Association of Insurance Commissioners and by Insurance Services Office, Inc. I use data on global annual greenhouse gas emissions from the World Resources Institute.

E.2 Estimating Anticipated Losses

Accounting variables such as insurance losses and expenses are reported as the ratio between that variable and total earned premiums. “Incurred losses” are total claims paid to policyholders for a year, and “loss adjustment expenses” are additional expenses related to investigating and settling claims. I group all other expenses into “underwriting expenses,” or expenses related to originating or maintaining insurance policies rather than handling claims.

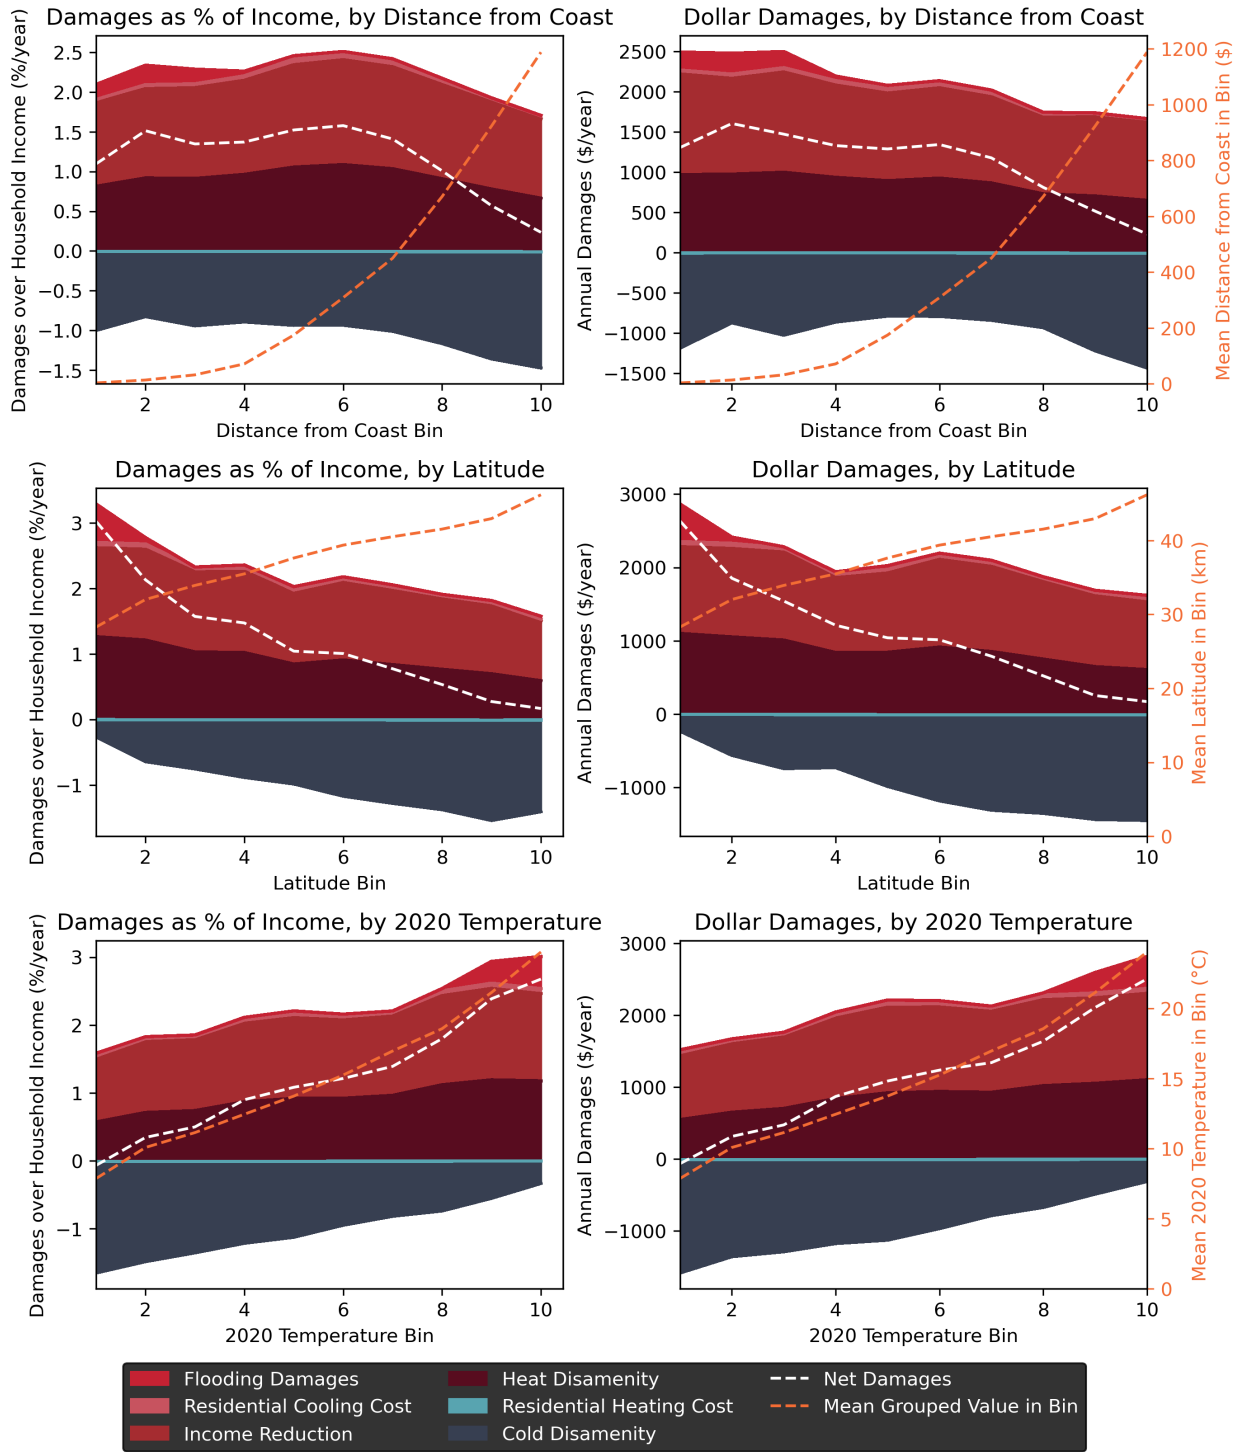



Figure 18

Mean damages by household income, distance from coast, and urban population share of household location (PUMA). Estimates are obtained by summing over damages to labor income, location amenities, residential energy costs, and residential disaster damages, as computed in this section, for each household in the 2020 American Community Survey.



(a) Breakdown of expenses experienced by insurers from homeowners insurance business (ratio with premiums earned). The combined ratio is the sum of the loss ratio (incurred losses divided by premiums earned) and the expense ratio (expenses divided by premiums earned). When the combined ratio is greater than 100, the insurer is losing money. Data are from the Insurance Information Institute (III).

(b) Breakdown of expenses experienced by insurers from homeowners insurance business (2021 dollars). Data are the same as the left figure, but values are given in dollar terms instead of as a ratio with premiums earned.

Figure 19

Figures 19a and 19b show the breakdown of expenses experienced by homeowners insurance companies. The majority of expenses are incurred losses, with loss adjustment expenses and underwriting expenses each accounting for about a quarter of total expenses.

The solid line indicates the “combined ratio,” a common metric used to assess the profitability of an insurance company. I assume that insurers are zero-profit, so that the anticipated combined ratio is equal to 100. I use this assumption to back out the expected level of incurred losses that insurers anticipate paying each year.⁷³

Figure 20 shows the average homeowners insurance premium and estimated anticipated losses, 2003-2020, in 2021 USD. The estimated anticipated losses appear to track the mean of realized losses well.

⁷³Additional details of this imputation procedure are in Appendix F. Expenses consist of both fixed and variable expenses that depend on realized losses, which I must invert.

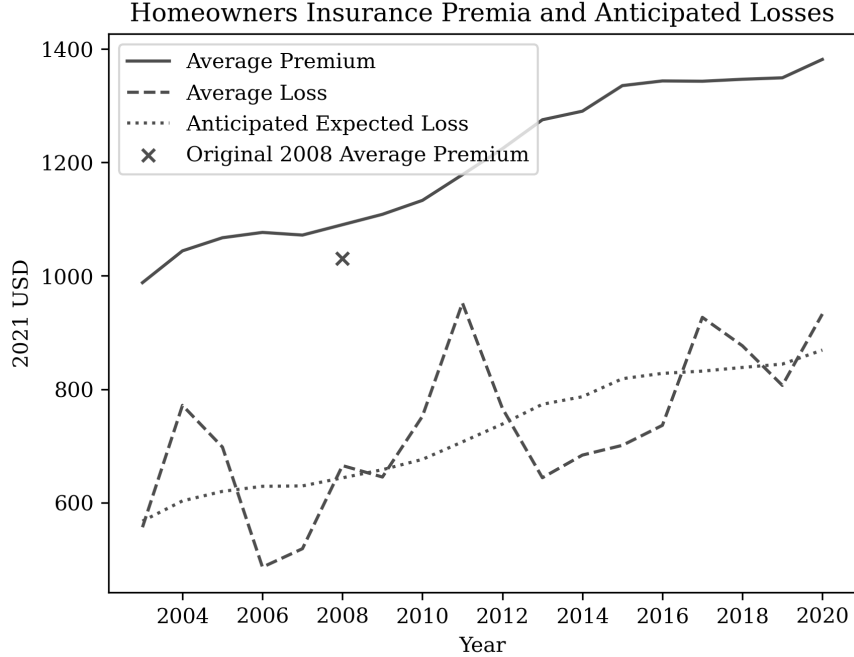


Figure 20: Average homeowners insurance premium and estimated anticipated losses, 2003-2020, in 2021 USD. The only significant decrease in average premium, in the original data, occurs in 2008, and is immediately reversed in 2009. I drop this observation on the grounds that it is likely a transient change due to the financial crisis rather than a persistent change in climate beliefs. Dollar losses (expenses) are imputed by multiplying the loss (expense) ratio by the average premium. Prior to 2010, dollar losses are imputed by multiplying claim frequencies (claims per 100 policy-years) by claim severity (average amount paid per claim). Nominal values are deflated by CPI.

E.2.1 Linear Regression

I show here that reduced-form evidence suggests that homeowners insurance premiums are responsive to disaster realizations in such a way that suggests that insurers are uncertain about the true distribution from which these realizations are drawn. I use this to inform a Bayesian estimation of the insurer’s uncertainty over the contribution of emissions to disaster risk in Appendix E.

I begin by defining the simple estimating equation,

$$m_{t+1} - m_t = \alpha + \frac{\tilde{\sigma}^2}{\sigma_\varepsilon^2}(\varepsilon_t - m_t) + \frac{\tilde{\sigma}^2}{\sigma_\eta^2}(\eta_t - m_t). \tag{13}$$

where m_t is the insurer’s prior mean of losses in year t , which I impute from the insurance premium (with details in Appendix F). ε_t is the insurer’s realized losses in year t , and η_t is a news shock which explains residual variation in premiums.


That is, I simply regress changes in imputed prior means on observed (log) loss surprises ε_t , and use the model to interpret the reduced-form coefficient on ε_t . Although $m_t = L_t - (\sigma^2 + \sigma_\varepsilon^2)/2$ depends on

parameters, the effect of these parameters is small, and can be solved for by fixed point. More seriously, unless $\mu_t = m_t$, the error term $\frac{\tau_\eta}{\tau}(\eta_t - m_t)$ will be biased and correlated with the regressor though the insurer's prior bias $\mu_t - m_t$,

$$\mathbb{E}[\eta_t - m_t] = \mu_t - m_t$$

$$\text{Cov}(\varepsilon_t - m_t, \eta_t - m_t) = (\mu_t - m_t)^2 - 2(\mu_t - m_t),$$

so this should be considered a descriptive exercise. Figure 21b shows the data used in this regression.



(a) Actual incurred losses ε_t , and anticipated incurred losses m_t imputed from the state space model.

(b) The regression data: $\Delta m \equiv m_{t-1} - m_t$ and $\varepsilon_t - m_t$.

I estimate Equation 13 by OLS, and report the results in Table 9. In columns 5-8, I fix the intercept α at 0.77% to match the forecast of Wing et al. (2022) of a 26% increase in flooding damages over 30 years. In columns 1,2,5, and 6, I estimate the slightly modified Equation 14 in which the error is reduced-form and not modeled as a signal. In columns 2,4,6, and 8, I drop the first observation because it is by far the largest value of the regressor. In all specifications there is a significant effect of loss surprises on premiums, although the significance is lower when the first observation is included.

Table 9: Parameter estimates.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
α	0.046*** (0.012)	0.054*** (0.010)	0.046*** (0.012)	0.054*** (0.010)				
$\varepsilon - m_t$	0.083* (0.043)	0.076*** (0.022)	0.083* (0.043)	0.076*** (0.022)	0.166** (0.058)	0.095*** (0.024)	0.163** (0.059)	0.095*** (0.025)
N	18	17	18	17	18	17	18	17
R^2	0.190	0.453	0.190	0.453	0.326	0.485	0.313	0.483

$$m_{t+1} - m_t = \alpha + \frac{1}{1 + \sigma_\varepsilon^2/\sigma^2}(\varepsilon_t - m_t) + \zeta_t. \quad (14)$$

E.3 Model

I attribute to a representative home insurer the same climate model as in the main model. Insurer storm losses $\exp(d_t)$ are generated by a persistent-transitory process. Log-losses d_t are drawn from a Normal distribution with unobserved mean μ_t , which follows a Gaussian random walk that is affected by emissions e_t . Each year, the insurer also observes a news shock n_t , unobserved to the econometrician, whose mean is a linear function of μ_t . Without loss of generality, η_t has mean μ_t . Formally,

$$\begin{aligned} \mu_{t+1} &= \rho\mu_t + \beta e_t \\ n_t &\sim \mathcal{N}(\mu_t, \sigma_n^2) \\ d_t &\sim \mathcal{N}(\mu_t, \sigma_d^2) \end{aligned}$$

The parameters $\{\sigma_n^2, \sigma_d^2, \rho\}$ are known to the insurer. The values $\{n_t, d_t, e_t\}$ are observed each period. The parameters β and μ_t are unobserved and learned about over time.

The timing of the model is as follows. Each period:

1. The insurer sets a premium π_t .
2. The shocks n_t and d_t are realized and observed.
3. The insurer updates their prior over μ_t and β .
4. The true mean μ_t is updated according to the AR(1) process.

If the insurer has an initial multivariate normal prior over μ_0 and β , then, if fully rational Bayesian learners, they update their beliefs according to the Kalman filter. Rewriting the model in terms of matrices:

$$\begin{aligned} \underbrace{\begin{bmatrix} d_t \\ n_t \end{bmatrix}}_{\mathbf{z}_t} &= \underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}}_{\mathbf{H}_t} \underbrace{\begin{bmatrix} \mu_t \\ \beta \end{bmatrix}}_{\mathbf{x}_t} + \mathbf{v}_t, \quad \mathbf{v}_t \sim \mathcal{N}(0, \Sigma_v) \\ \underbrace{\begin{bmatrix} \mu_{t+1} \\ \beta \end{bmatrix}}_{\mathbf{x}_{t+1}} &= \underbrace{\begin{bmatrix} \rho & e_t \\ 0 & 1 \end{bmatrix}}_{\mathbf{F}_t} \underbrace{\begin{bmatrix} \mu_t \\ \beta \end{bmatrix}}_{\mathbf{x}_t} \end{aligned}$$

If the insurer's prior for \mathbf{x}_t at the end of period $t - 1$ is

$$\mathbf{x}_t \mid \mathcal{I}_t \sim \mathcal{N} \left(\hat{\mathbf{x}}_t \equiv \begin{bmatrix} \hat{\mu}_t \\ \hat{\beta}_t \end{bmatrix}, \hat{\mathbf{P}}_t \right),$$

then their prior for \mathbf{x}_{t+1} at the end of period t is

$$\begin{aligned} \hat{\mathbf{x}}_{t+1} &= \mathbf{F}_t(\mathbf{I} - \mathbf{K}_t\mathbf{H}_t)\hat{\mathbf{x}}_t + \mathbf{F}_t\mathbf{K}_t\mathbf{z}_t \\ \hat{\mathbf{P}}_{t+1} &= \mathbf{F}_t(\mathbf{I} - \mathbf{K}_t\mathbf{H}_t)\hat{\mathbf{P}}_t\mathbf{F}_t' \\ \mathbf{K}_t &= \hat{\mathbf{P}}_t\mathbf{H}_t'(\mathbf{H}_t\hat{\mathbf{P}}_t\mathbf{H}_t' + \Sigma_v)^{-1} \end{aligned}$$

where \mathbf{K}_t is the optimal Kalman gain.

E.3.1 Nesting Step

I observe $\{d_t, e_t, \hat{x}_t\}$ and wish to estimate $\{\sigma_n, \sigma_d, \rho, \hat{\beta}_0\}$. First, I simply rewrite the above equations in terms of block matrices:

$$\begin{aligned} \begin{bmatrix} \mathbf{x}_{t+2} \\ \mathbf{z}_{t+1} \\ \hat{\mathbf{x}}_{t+1} \end{bmatrix} &= \begin{bmatrix} \mathbf{F}_{t+1} & 0 & 0 \\ \mathbf{H}_{t+1} & 0 & 0 \\ 0 & \mathbf{F}_t\mathbf{K}_t & \mathbf{F}_t(\mathbf{I} - \mathbf{K}_t\mathbf{H}_t) \end{bmatrix} \begin{bmatrix} \mathbf{x}_{t+1} \\ \mathbf{z}_t \\ \hat{\mathbf{x}}_t \end{bmatrix} + \begin{bmatrix} 0 \\ \mathbf{v}_t \\ 0 \end{bmatrix} \\ \begin{bmatrix} 0 \\ \mathbf{v}_t \\ 0 \end{bmatrix} &\sim \mathcal{N} \left(0, \begin{bmatrix} 0 & 0 & 0 \\ 0 & \mathbf{R}_t & 0 \\ 0 & 0 & 0 \end{bmatrix} \right) \\ \begin{bmatrix} d_t \\ \hat{\mu}_t \end{bmatrix} &= \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{x}_{t+2} \\ \mathbf{z}_{t+1} \\ \hat{\mathbf{x}}_{t+1} \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mu_t \\ \beta \\ n_t \\ d_t \\ \hat{\mu}_t \\ \hat{\beta}_t \end{bmatrix} \end{aligned}$$

Observe that, for any guess of the parameters $\{\sigma_n, \sigma_d, \rho, \hat{\beta}_0\}$, all of the matrices in this system are known to us. That is, we can treat the above as a linear Gaussian state space model and apply the Kalman filter to it. This is the nesting step.

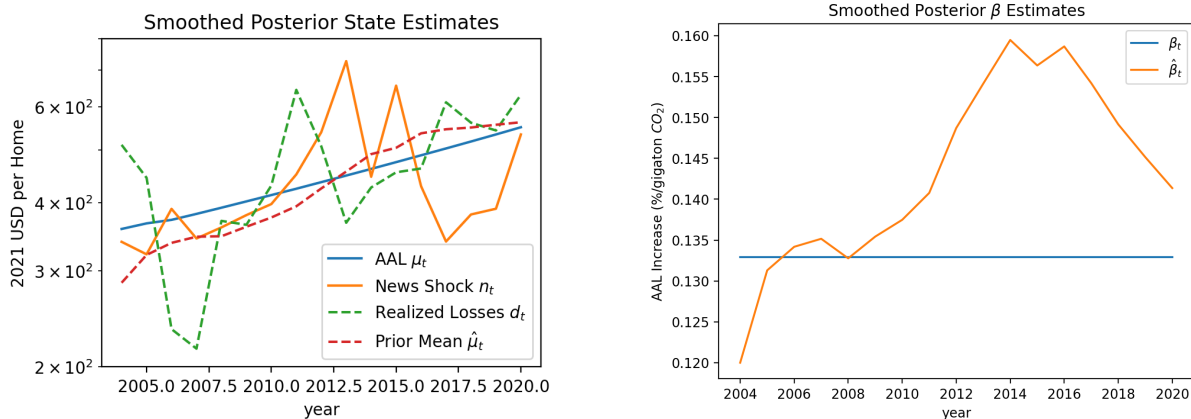
E.3.2 State Space Model Estimation

For any prior and given set of parameters, the Kalman filter yields a likelihood of the observed data, in our case $\{d_t, \hat{\mu}_t\}$. I thus estimate the parameters by maximum likelihood. Table 10 displays the result.

After estimating parameters, the Kalman smoother provides estimates of the unobserved state. Figure 22 shows the results. Given the parameter estimates, variation in $\hat{\mu}_t$ is almost entirely due to variation in beliefs $\hat{\beta}_t$ about β , as opposed to being driven by beliefs that the unobserved mean μ_t has changed over time.

Variable	Estimate
σ_n^{err}	0.21306
σ_d^{err}	0.28874
ρ	0.98
T_0	1965.87

Table 10: Results of state space estimation. Point estimates are by maximum likelihood.



(a) Estimates of unobserved μ_t and n_t . The dashed lines $\hat{\mu}_t$ and d_t are data.

(b) Estimates of β and $\hat{\beta}_t$.

Figure 22

F Homeowners Insurance Anticipated Incurred Losses

In order to account for the fact that expenses other than the incurred expenses (henceforth “non-incurred expenses”) have a time trend and covary with incurred expenses, I regress non-incurred expenses on incurred losses and a time trend. That is, I decompose it into fixed and variable costs, where I allow for a time trend in fixed costs. Though the number of observations is small, the fit is very good. I regress non-incurred expenses on incurred losses and a time trend, and find a small time trend and . That is, I decompose it into

fixed and variable costs, where I allow for a time trend in fixed costs.

$$\text{Expenses}_t = \alpha + \beta(\text{Incurred Losses})_t + \gamma\text{Year}_t + \varepsilon_t$$

Table 11: Expenses on incurred losses.

	Expense Ratio
Intercept	731.743*** (42.405)
Incurred Loss Ratio	0.079*** (0.009)
Year	-0.346*** (0.021)
R-squared	0.976
R-squared Adj.	0.970
N	12

I then assume risk-neutrality of insurers and apply the zero-profit condition that insurers are targeting a constant combined ratio of 100. I can thus estimate anticipated expected annual losses by setting anticipated expected combined losses equal to premiums earned. By “anticipated expected” (henceforth simply “anticipated”) losses, I mean

$$\mathbb{E}[\text{Losses} \mid \mathcal{I}_t],$$

the expected value of anticipated losses given the insurer’s information set at time t . The mean combined ratio over the sample period was 101.425, suggesting that it is reasonable to assume that a value of 100 is targeted. Note that we would expect the combined ratio to be greater than 100 if insurers experience greater-than-expected losses over the sample period.

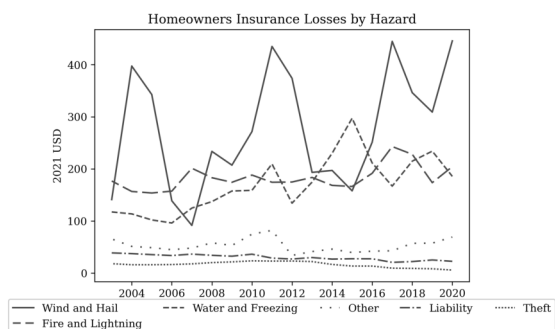
Let $\mathbb{E}[x \mid \mathcal{I}_t]$ denote the anticipated value of x conditional on the insurer’s information set at time t . For each year t , let L_t denote realized incurred loss ratio, N_t denote non-incurred expense ratio, and C_t denote

the combined ratio. Then,

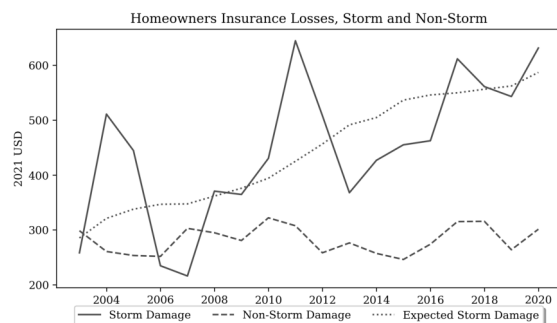
$$\begin{aligned}
 \mathbb{E}[C_t | \mathcal{I}] &= 100 \\
 &= \mathbb{E}[L_t | \mathcal{I}] + \mathbb{E}[N_t | \mathcal{I}] \\
 &= \mathbb{E}[L_t | \mathcal{I}] + \alpha + \beta \mathbb{E}[L_t | \mathcal{I}] + \gamma \text{Year}_t \\
 \mathbb{E}[L_t | \mathcal{I}] &= (\mathbb{E}[C_t | \mathcal{I}] - \alpha - \gamma \text{Year}_t) / (1 + \beta)
 \end{aligned}$$

F.1 Storm Component of Losses

Figure 23a shows the breakdown of losses by hazard. Note that there is no clear trend in any category except “Wind and Hail” and “Water and Freezing.” In Figure, I group “Wind and Hail” and “Water and Freezing” losses into “Storm Losses” and all other incurred losses into “Non-Storm Losses.” The time trend of non-storm losses is vanishingly small, \$0.72/year. I thus assume that the time trend in anticipated non-storm losses is zero, and that all variation in anticipated losses is due to variation in anticipated storm losses. For each year, I simply subtract the sample mean of non-storm losses from total anticipated incurred losses to obtain my estimate of anticipated storm losses.



(a) Annual average homeowners insurance losses per policy-year, by hazard. Losses are reported in data as shares of total incurred losses, and are here imputed by multiplying by imputed total incurred losses.



(b) Annual average homeowners insurance losses per policy-year, by storm category. Losses are reported in data as shares of total incurred losses, and are here imputed by multiplying by imputed total incurred losses.

G Calibrated Spatial Parameters

Calibrated Amenities




Figure 24: Estimated Amenity Values

Productivity Shifter

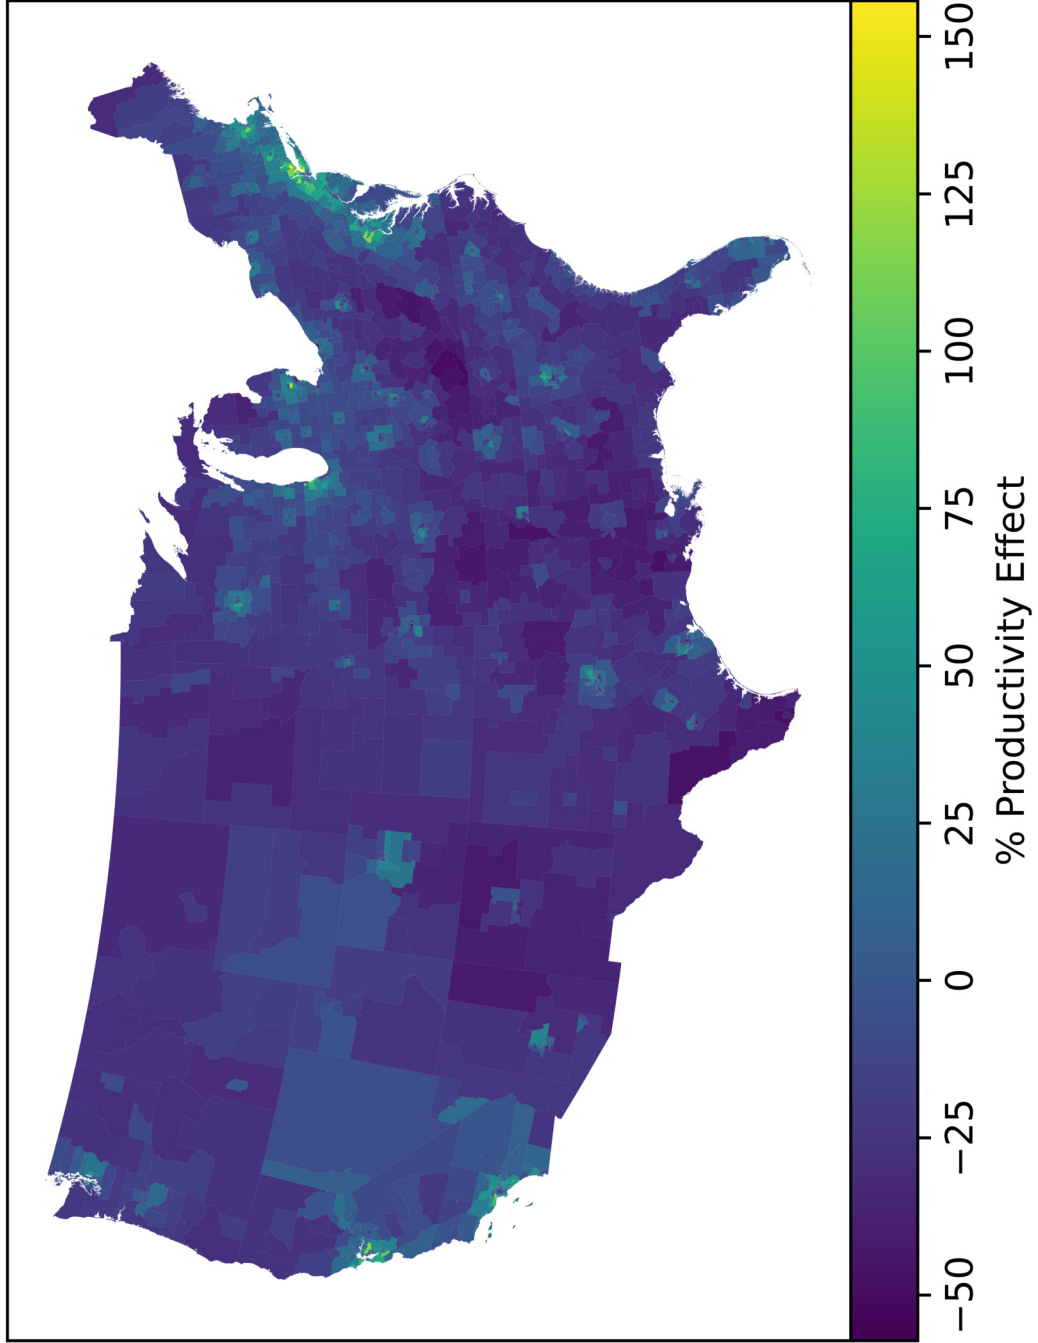


Figure 25: Estimated Productivity Shifters

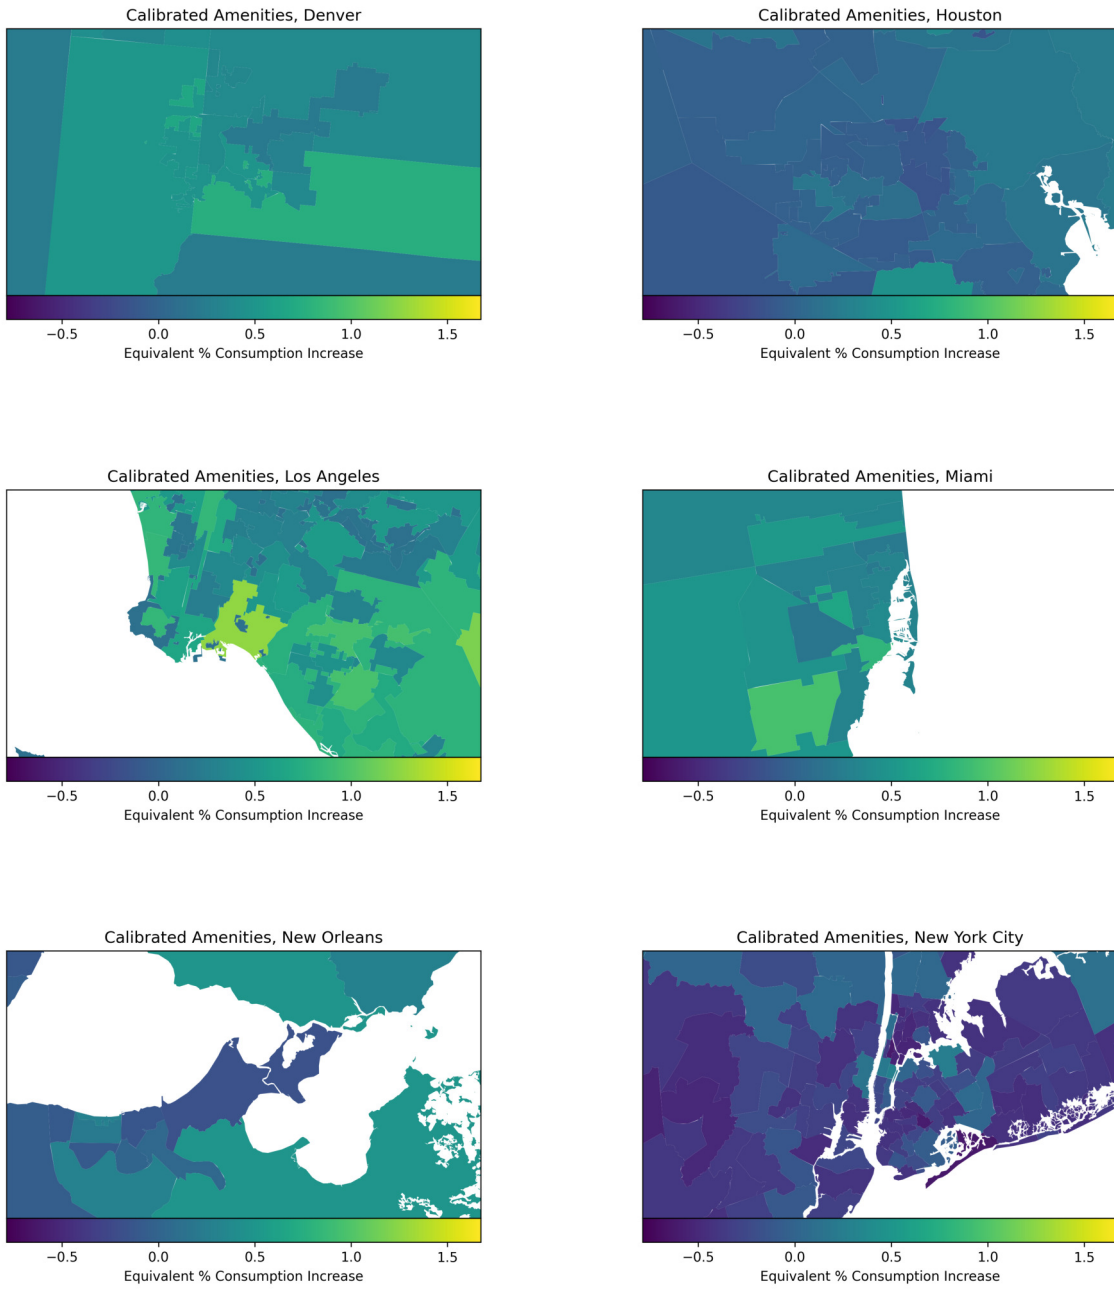


Figure 26: Calibrated Amenities by City







Figure 27: Calibrated Amenities Around New York City

Rent Level




Figure 28: Equilibrium Rent Level

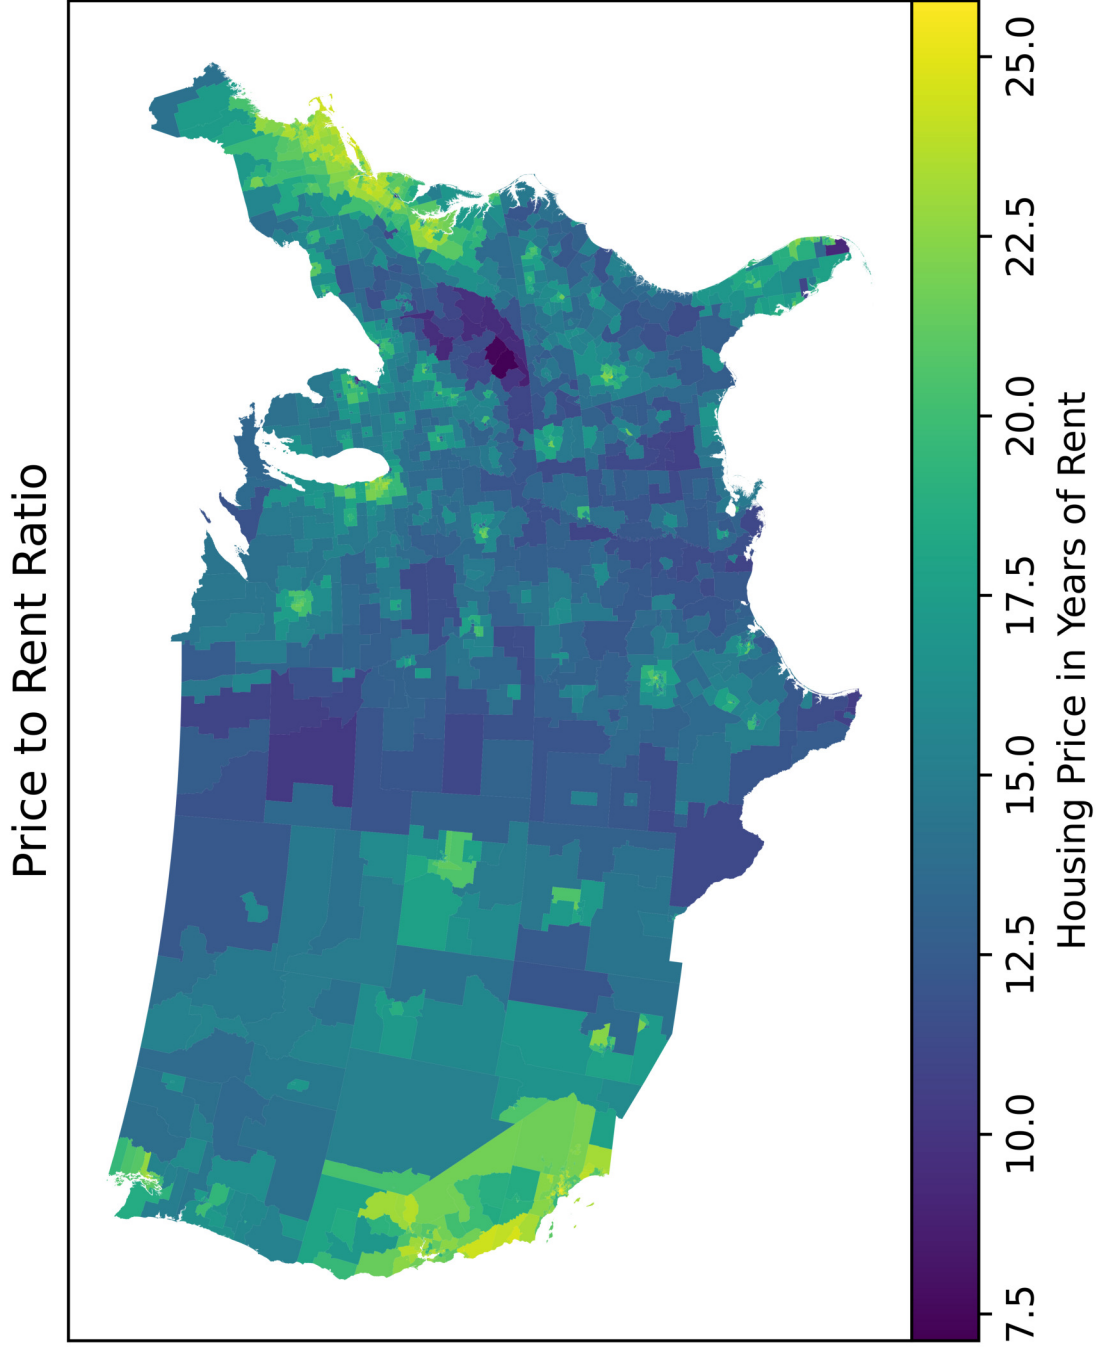


Figure 29: Equilibrium Price to Rent Ratio

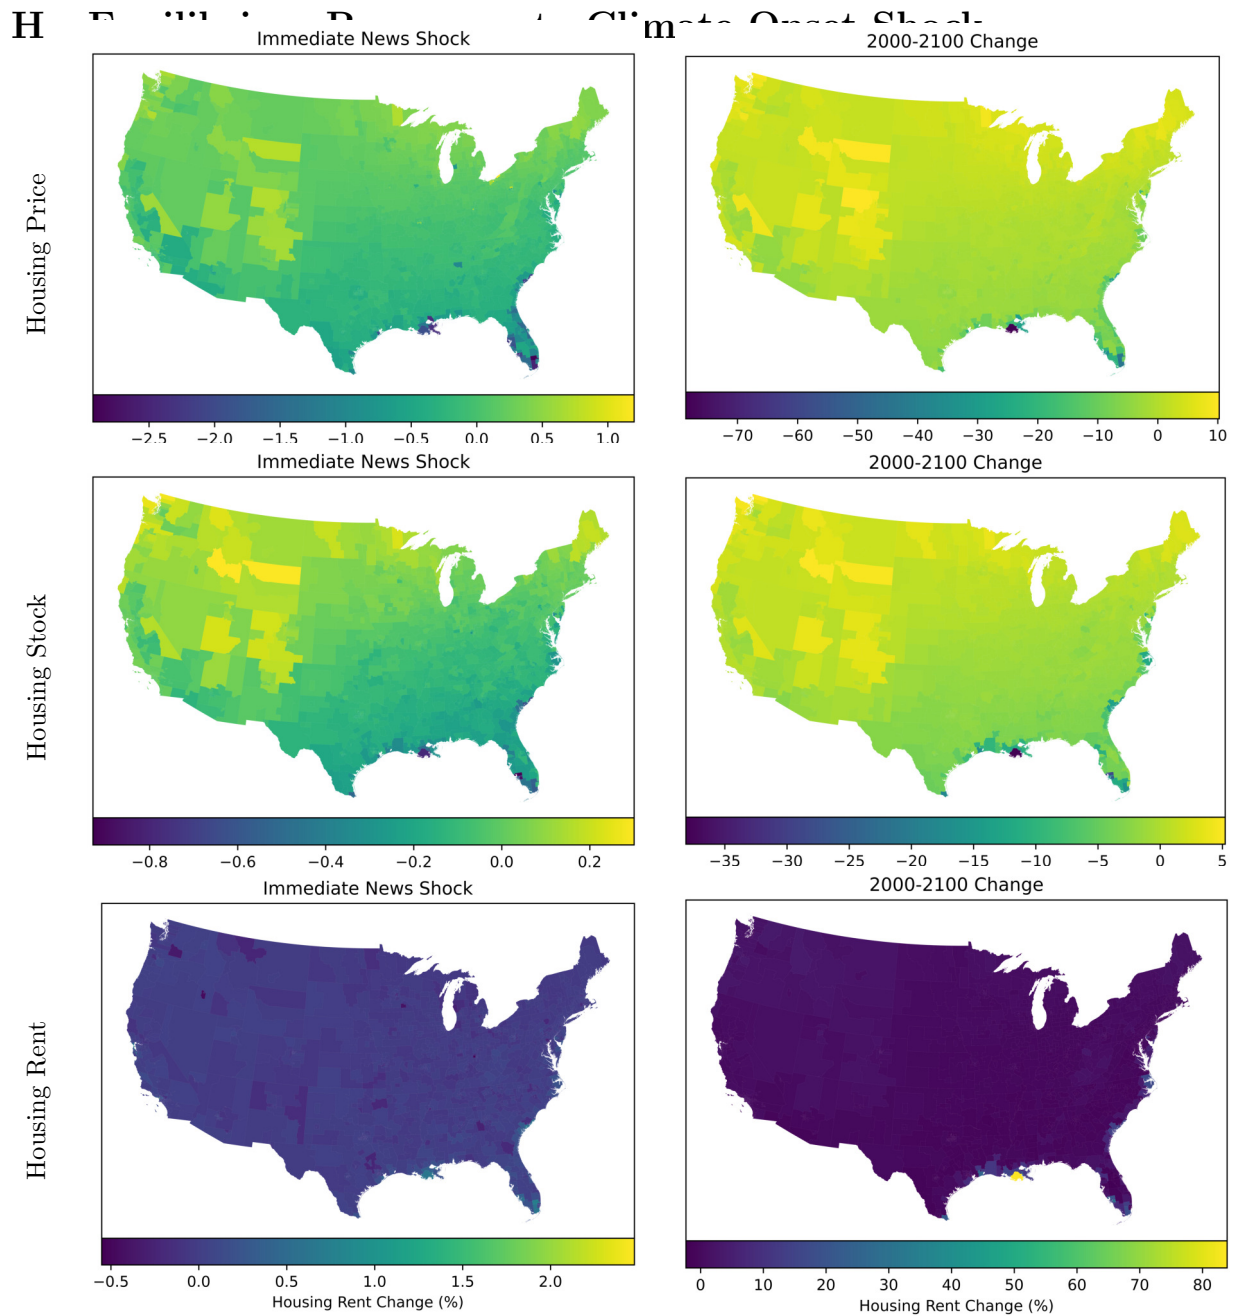


Figure 30

Equilibrium responses to climate change across the U.S. Column 1 describes the immediate equilibrium responses to the onset shock, before the climate physically changes. Column 2 describes the spatial equilibrium in 2100 under the median realization of the climate process (corresponding to RCP4.5).

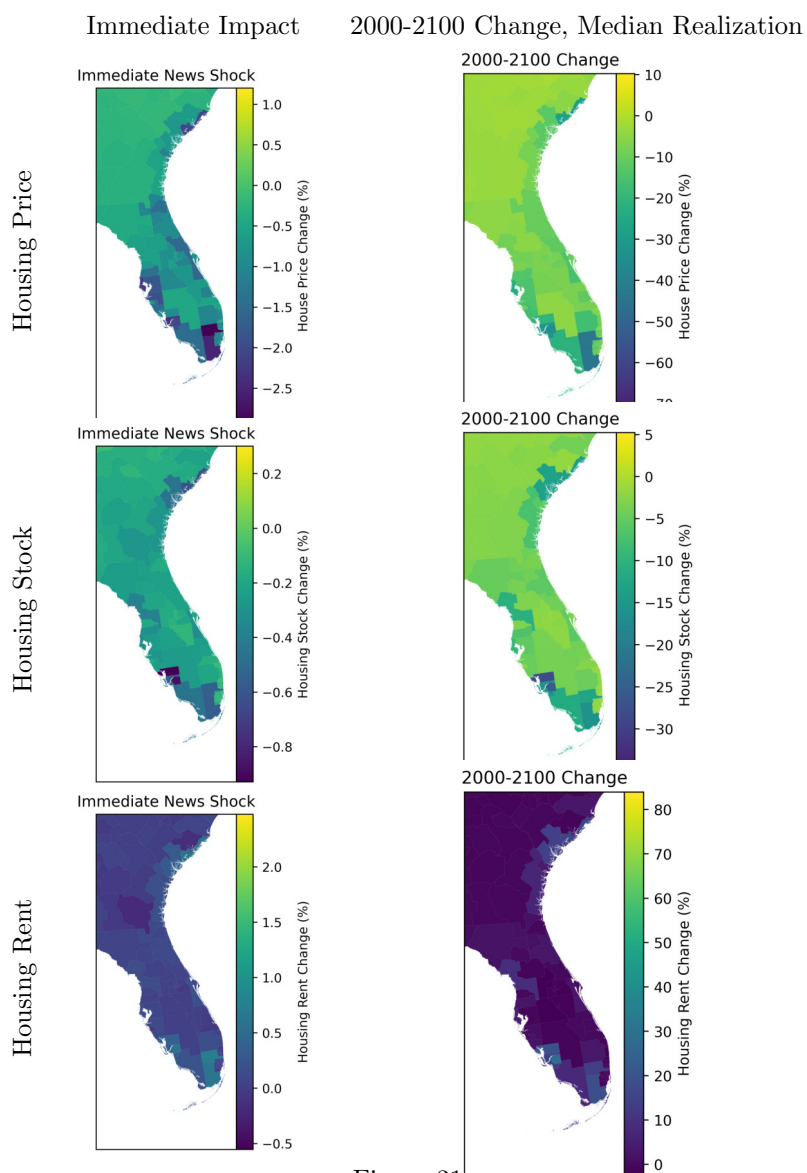


Figure 31

I Inequality Response Along Median Path

Figure 32 displays the paths of consumption inequality over time following the onset shock, along the median path.

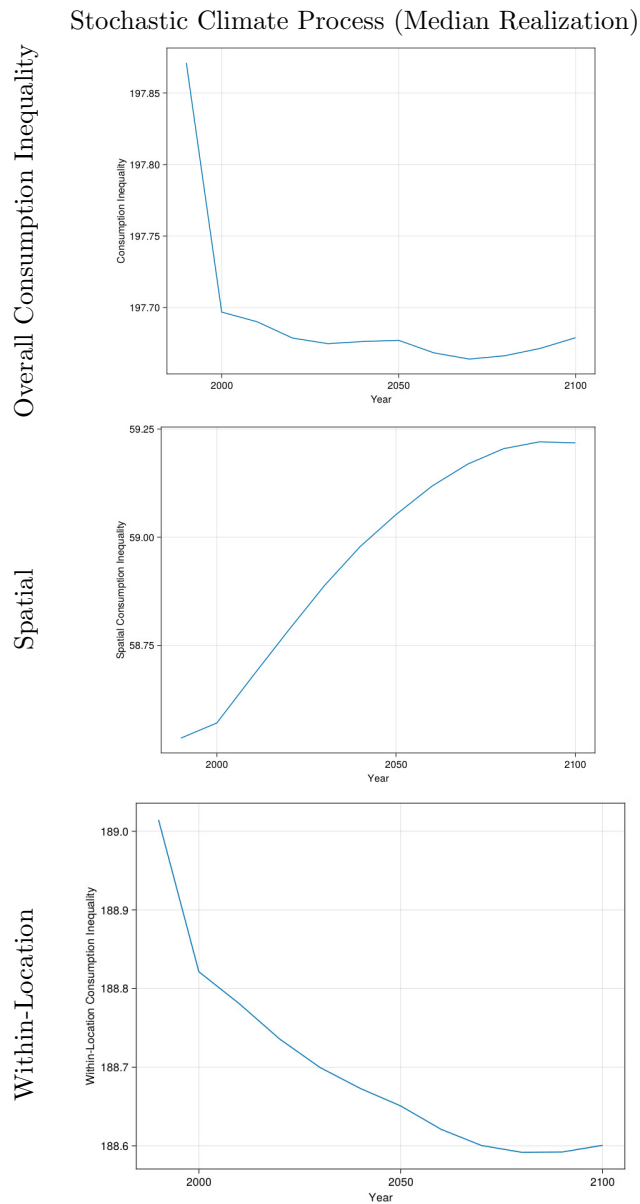


Figure 32

Paths of consumption inequality over time following the onset shock, along the median path. Consumption inequality is measured as the standard deviation of equivalent decadal consumption for a renter in a “typical location” with amenity value $\alpha = 1$ and average housing rents in steady state.

J Impulse Response Functions To Small Climate Perturbation

Figure 33 plots the impulse response functions to a 0.1°C MIT shock. In 1990, the temperature permanently and unexpectedly increases by 0.1°C.

K Results Without Migration

Table 12 reports the distribution of the initial welfare impact of the unexpected onset of climate change, in a version of the model starting from the same initial steady state but in which migration is not allowed after the onset shock. Table 13 reports the impact on overall welfare inequality following the onset of climate change and in 2100 under the median realization in this model. Figure 34 plots the decomposition of welfare impacts by channel and margin in this model.

Table 12: Initial Welfare Impact: Variance Decomposition

	All Channels	Housing Wealth Channel	Other Channels
Aggregate Welfare Impact	-0.000513662	-0.000183067	-0.000370016
Variance	0.000768325	1.62918e-5	0.000693449
Variance (Spatial Component)	0.000429854	9.04674e-7	0.000397049
Variance (Within-Location Component)	0.000338472	1.53872e-5	0.000296401

Distribution of initial welfare impact of unexpected onset of climate change, with no migration. Aggregate effects and variance decomposition.

Table 13: Overall Welfare Inequality: Transition

	Overall Inequality	Spatial Inequality	Within-Location Inequality
All Ages			
Before Onset	408.498	162.817	374.648
After News Shock	403.977	157.377	372.062
In 2100 (Median Realization)	410.283	161.868	377.003
At Birth			
Before Onset	286.599	232.677	167.334
After News Shock	281.774	227.955	165.631
In 2100 (Median Realization)	280.091	219.07	174.526

Overall welfare inequality, before and after the onset of climate change.

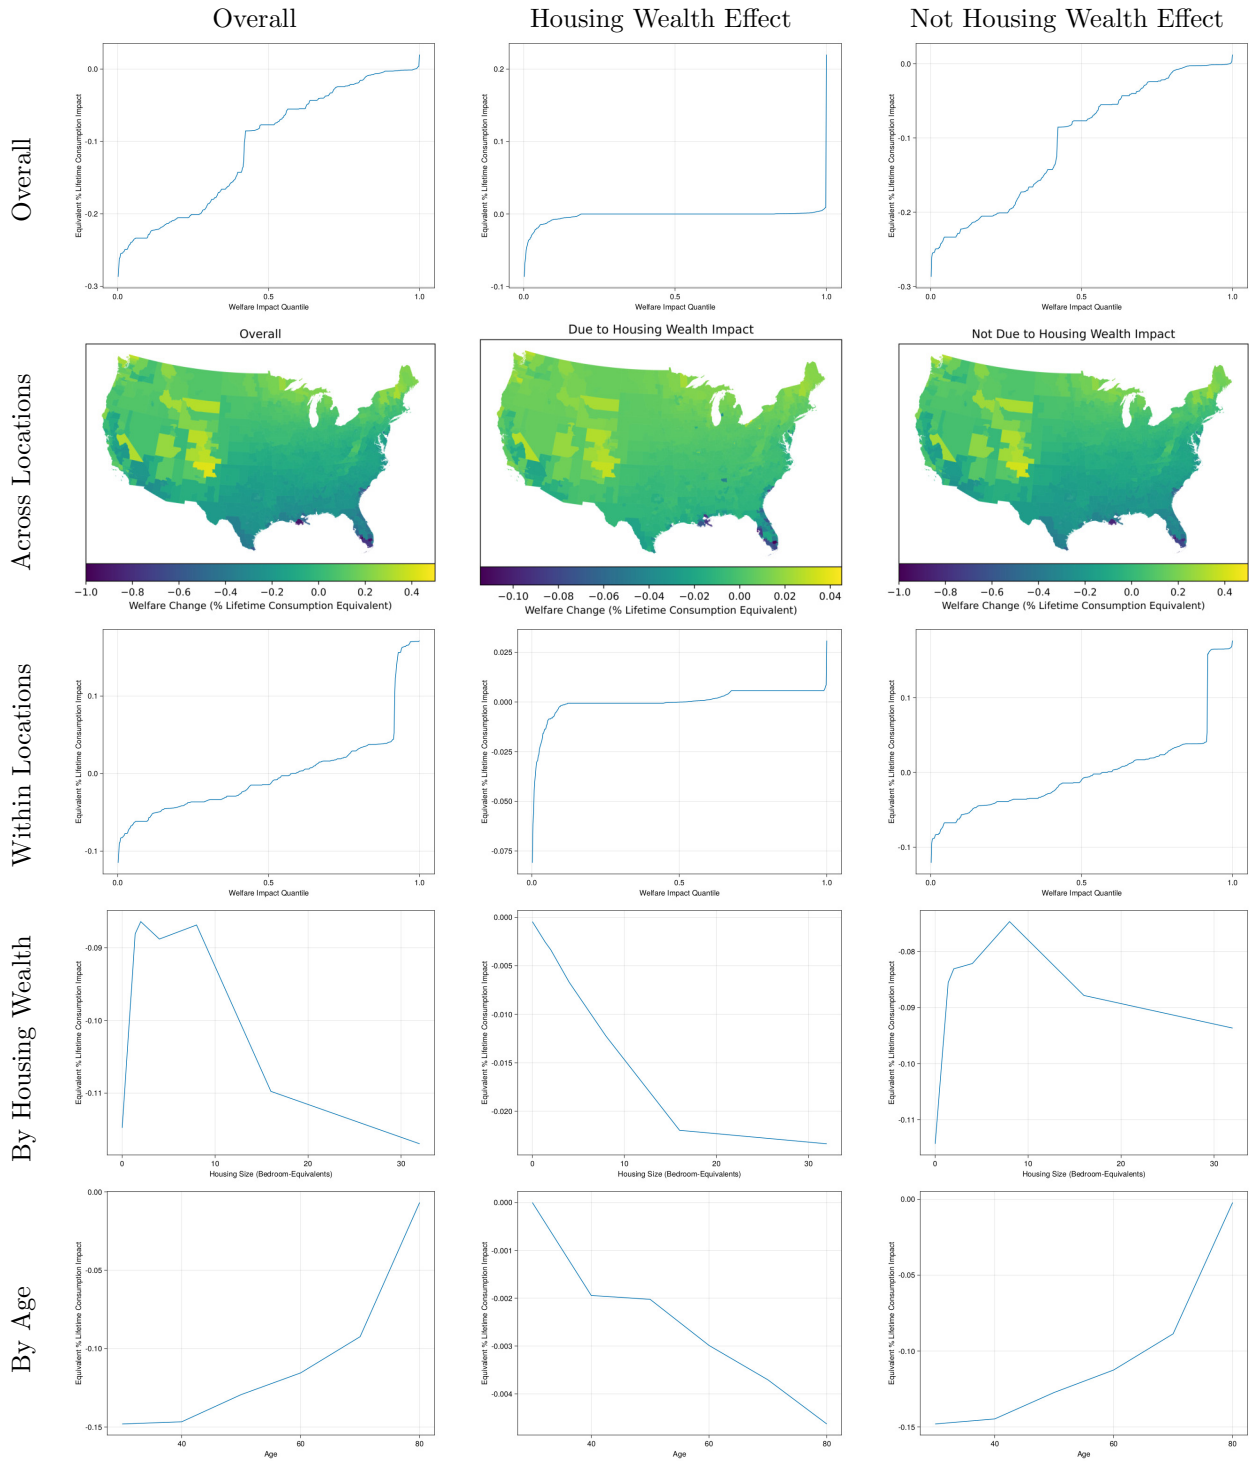


Figure 33

Initial welfare impact inequality by channel and margin: 0.1°C perturbation. Column 1 reports the welfare impact of the beginning of the climate transition, relative to continuation of the steady state. Column 2 reports the welfare impact of the housing price response, computed as (negative of) the welfare impact for each household of a post-shock transfer exactly negating the portfolio value impact of the house price change. Column 3 reports the welfare impact of the combined climate shock, house price adjustment, and compensating transfer, relative to a continuation of the steady state. Row 1 plots the cumulative distribution of the initial welfare impact across households. Row 2 plots the mean welfare impact across locations. Row 3 plots residual impacts, controlling for location-level means. Welfare impacts are expressed as an equivalent percentage of lifetime consumption lost or gained.

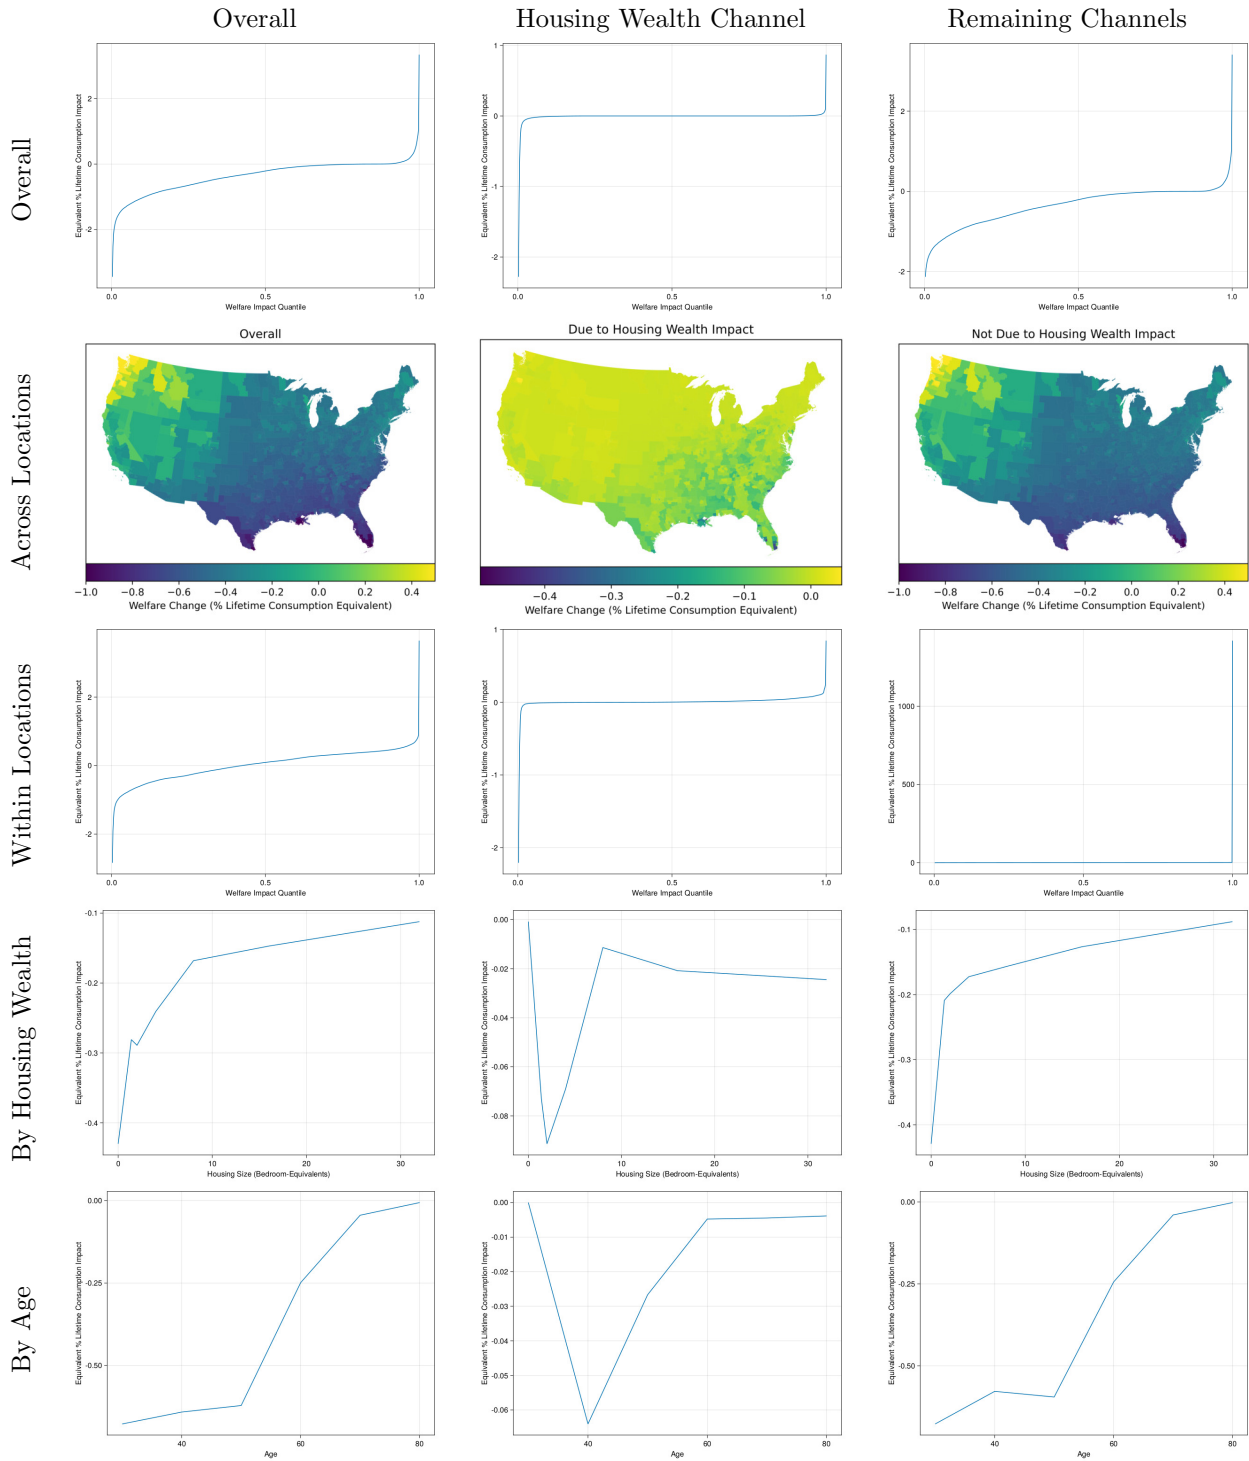


Figure 34

Distribution of welfare impacts of the unexpected onset of climate change (“onset shock”), in a variant of the model without migration. Columns correspond to channels and rows correspond to margins of variation. Column 1 incorporates all channels. Column 2 isolates the welfare impact of the housing price response. Column 3 reports the remaining welfare impact after subtracting the contribution of the housing wealth channel. Rows decompose variation in welfare impacts into spatial, within-location, housing wealth, and age margins. Welfare impacts are expressed as an equivalent percentage of lifetime consumption lost or gained.