Title
Adaptation and Mitigation: Essays on Climate Economics

Author
Howden, Wesley

Publication Date
2021

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

Adaptation and Mitigation: Essays on Climate Economics

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy

in

Economics

by

Wesley Howden

Committee in charge:

Professor Mark Jacobsen, Chair
Professor Judson Boomhower
Professor Jennifer Burney
Professor Richard Carson
Professor Josh Graff Zivin

2021

The dissertation of Wesley Howden is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2021

DEDICATION

I dedicate this dissertation to my family—my parents, Susan and Stephan; my sister, Carly; and my dog, Kamala.

EPIGRAPH

I believe that the most striking feature of the economics of climate change is that its extreme downside is nonnegligible. Deep structural uncertainty about the unknown unknowns of what might go very wrong is coupled with essentially unlimited downside liability on possible planetary damages.

Marty Weitzman

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

Thank you first to my academic advisors Mark, Judd, Richard, Josh, and Jenn. Thanks to my co-author Remy; to my office mates Becky, Dani, and Sam (cut off far too abruptly); to all of my economics colleagues and the impromptu visits to Art of Espresso and talks over lunch or otherwise; everyone at the environmental lunch and various workshops (both online and off); my interdisciplinary colleagues at IFER; the bar trivia teams; and the entering cohort of 2015.

I owe a huge thanks to all of my friends—the college friends that made academia exciting; camping and hiking mates; and especially to Erica, Chris, Masha, Eric, Jasha, Margaret, Robert, and Jan. Finally, to my family (including dog!) who kept me sane—I couldn't have made it where I am without your support.

Chapter 1 is currently being prepared for submission for publication. The dissertation author is the sole author on this chapter.

Chapter 2 is currently being prepared for submission for publication. The dissertation author is the sole author on this chapter.

Chapter 3 is co-authored with Remy Levin and is currently being prepared for submission for publication. The dissertation author is a principle researcher on this chapter.

All errors are my own.

VITA

2015	S.B. in Mathematics, The University of Chicago
2015	A.B. in Economics, The University of Chicago
2015	A.B. in Political Science, The University of Chicago
2017	M.A. in Economics, University of California San Diego
2021	Doctor of Philosophy, University of California San Diego

ABSTRACT OF THE DISSERTATION

Adaptation and Mitigation: Essays on Climate Economics

by

Wesley Howden

Doctor of Philosophy in Economics

University of California San Diego, 2021

Professor Mark Jacobsen, Chair

I study the impacts of climate and environmental change and subsequent consequences for adaptation and mitigation. In Chapter 1, I use a difference in differences design to show that California households exposed to a severe heat wave are differentially more likely to adopt central air conditioning units than those less exposed, controlling for historical climate. Using this induced adoption to predict take-up, I show that induced adopters have a significant increase in their summer energy demand 3 years following the heat wave, with insignificant effects on their winter electricity demand. In addition, I present a theoretical framework where household belief-updating about the climate rationalizes heterogeneity in household learning about the climate that cannot be explained by

myopia or alternative channels.

In Chapter 2, I measure the impact of land-use change on flood risk. This study examines this by quantifying the effects of land-use change on flood damages in the state of Texas. I link claims data from the National Flood Insurance Program to a series of land-use changes to construct a tract-by-month panel, and use exogenous variation in precipitation across tract-months to estimate the effect of changes in land use on the frequency and magnitude of new flood insurance claims. I find that increases in impervious surface development within a tract increase flood insurance claims, while increases in wetland and water cover decrease these claims. In addition, using variation in tract-level elevation, I show that land-use change in neighboring geographies affects own-tract flood risk. Overall, these results suggest existence of spatial spillovers from land use and imply returns to coordination in land-use policy.

In Chapter 3, co-authored with Remy Levin, we show that individuals in a panel survey in Indonesia and Mexico exhibit changes in observed choices over a risky lottery as a result of changes in experienced temperature and precipitation levels and volatility. We use a counterfactual measure of risk under the assumption of no response to climate variables to show that total social welfare is higher under the observed distribution of risk. We interpret this as risk adaptation to climate.

Introduction

This dissertation consists of three studies that collectively explore the impacts and mechanisms of adaptation and mitigation to a changing climate. In Chapter 1, I show that an episodic heat wave has a differential effect on individuals based on prior climate change beliefs, wherein those who believe the climate is changing are more likely to update their beliefs based on these heat-wave events. Subsequently, these individuals are more likely to invest in air-conditioning than observably similar non-believers. I link this contemporaneous shock to a follow-on increase in energy use during the cooling season several years later, connecting short-run weather shocks to long run outcomes. This novel result shows an important relationship between tail climate events and long-run patterns of adaptation and elucidates the importance of studying higher moments of the climate distribution. More generally, I point to the importance of studying these effects as episodic weather events increase with climate change,

In Chapter 2 I find that flood insurance claims are affected by aggregate land use changes within a development jurisdiction, as expected, but that aggregate changes also affect claims in neighboring jurisdictions. This is particularly the case for land-use changes in neighboring jurisdictions with higher average elevation, even controlling for correlated land use changes across jurisdictions. It suggests that the current, largely parochial, approach to land-management in the US is suboptimal for controlling flood risk. These spillovers in the NFIP market add to existing frictions in this space.

Finally, Chapter 3 shows adaptation of individual risk preferences based on changes in experienced climate means and volatility. We show changes in background risk, climate, affects individuals' observed risk attitudes in Indonesia and Mexico. We link these changes in observed fundamentals to downstream behaviors, and show suggestive evidence that predicted increases in risk aversion from climate change correlates with decreases in risky behaviors including internal rates of migration and smoking. Building on recent advances in welfare economics, we develop a new method for estimating whether observed risk preference changes are, in fact, adaptive. Using our method, we find that in our sample, climate-change-induced risk preference changes are marginally welfare-improving.

In summary, in this dissertation, I present a collection of novel results that highlight the impacts of climate and environmental change and analyze the resulting patterns of adaptation and mitigation. Beyond presenting these impacts and mechanisms, this suggests implications for public policy meant to mitigate adverse effects of climate change.

Chapter 1

Adaptation to Weather Shocks and Household Beliefs on Climate: Evidence from California

1.1 Introduction

Household demand for energy responds to local climate and weather characteristics. Longer-run decisions, such as location choice or portfolio of appliances (or other energy-using durables), may reflect the state of the long-run climate for a location. Typically, we think that short-run changes to energy demand respond to short-run variations in the weather. For instance, on hotter days, households that own an air conditioner can increase their energy demand at the intensive margin by cooling their home. Using a traditional framework of household investment, extensive entry into air conditioning adoption seems unlikely to respond to a short-run increase in the number of hot days.

In this paper, I show novel evidence that households adjust their medium- to long-run capital holdings (in the form of air conditioner ownership) in response to a short-run weather shock, particularly, a severe heat wave. During the summer of 2006, a series of extreme heat waves affected regions across the Pacific and Southwest. According to reports by the California Department of Public Health, emergency room visits and hospitalizations related to heat increased significantly, and there were more than 140 heat-related deaths (Knowlton et al., 2009).

Using this heat wave-induced adoption of central air conditioning, I link this to longer-run impacts on household-level summer energy demand. Using plausibly exogenous variation in exposure to this heat wave, I show that 100 extra cooling degree days (CDDs) relative to the historical average increases the propensity for a household to own central air conditioning by about 1 percentage point. Then, estimating the reduced form impact of these weather shocks on monthly summer electricity demand 3 years following this heat wave, I show that 100 extra CDDs is associated with an average increase in energy demand by 6 kWh per household during the month of July.

This extensive entry into air conditioning adoption as a response to hot weather (as opposed to average weather, or climate, at a geographic level) is consistent with previous studies of similar phenomena. Auffhammer (2014) shows evidence for extensive entry into air-conditioning adoption in China induced by preceding hot years. One novel contribution of this paper is the link between short-run weather and long-run energy demand, as well as a qualitative replication of these Chinese results using weather data with both higher spatial and time resolution.

In addition to the link to longer-run energy demand, I explore the potential mechanisms by which households are induced to adopt an air conditioning unit by contemporaneous hot weather (relative to an average year). I use county-level measures of belief in climate change and precinct-level general elections returns to construct a proxy measure of household belief in climate change. I use this measure and introduce a third difference to the baseline specification. In this specification, I show that households that are more likely to believe that the climate is changing are also more likely to exhibit heat-wave induced adoption of air conditioning. This heterogeneity by belief in climate change is consistent with other observed empirical patterns in belief in climate change, including diverging beliefs in climate change for higher-educated Democrats (more likely to believe in climate change than lower-educated Democrats) and Republicans (less likely to believe in climate change than lower-educated Republicans).

Consistent with these patterns of heterogeneity in induced air conditioning adoption, I propose a simple framework that rationalizes these observations: households that do not believe in climate change take short-run weather anomalies (a contemporaneous heat wave) as a draw from a fixed climate distribution. Conversely, households that believe that their local climate is changing take the same weather anomaly as being informative of the future path of climate for their local ZIP code. Alternate mechanisms cannot explain this heterogeneity. This does not rule out alternative channels, such consumer myopia or contemporaneous disutility from heat, but it does suggest a role for households learning about the path of local climate.

California is a unique case for studying these events within the United States, given the high propensity of mild climates. Because of this, there is likely a large margin of households that exist near a threshold for adoption that does not exist in other United States settings. While the initial estimates may seem large, 100 extra CDDs could mean 10 nights of 80 degree Fahrenheit nights instead of 70 degree nights, contextualizing the potential impact of a marginal heat wave. And as other developed European countries are exposed to heat waves, it will become more policy relevant to think about how households make investment decisions in energy-using durable goods.[1] Further, as incomes grow in the developed world and households broach the adoption margin, these dynamics may become more relevant, with implications for the dynamic path of aggregate energy demand.

The rest of the paper proceeds as follows: Section 1.2 reviews past literature on climate change and energy consumption and discusses why this particular study is novel. Section 1.3 discusses the data and empirical strategy. Section 1.4 presents the models used to estimate the effect of weather on AC adoption and energy demand. Section 1.5 contains the model estimates. In Section 1.6, I discuss these results and their implications

[1]See the 2019 European heat wave: https://www.nytimes.com/2019/07/25/world/europe/heatwave-record-temperatures.html.

for household behavior. Finally, Section 1.7 concludes.

1.2 Literature

This paper explores the link between short-run weather shocks and longer-run outcomes, specifically through induced adoption of air conditioning and the implications for energy demand. In this section, I first summarize the literature that links energy demand to local climate impacts, as well as the literature that specifically focuses on air conditioning adoption. In addition, I discuss the literature on household and market-wide beliefs in climate change. I contribute to a new but growing literature that finds heterogeneity in household investment decisions based on beliefs about the climate.

1.2.1 Energy demand and climate impacts

The first major contribution of this paper is to explore the link between short-run weather shocks on longer-run outcomes for energy demand. Auffhammer and Mansur (2014) review the literature on energy consumption and climate trends and delineate between two general methods of estimating this relationship. First, panel methods focus largely on local weather variation and estimate energy demand response at the intensive margin (Deschênes and Greenstone, 2011; Auffhammer and Aroonruengsawat, 2011). The disadvantage of using this intensive-margin relationship to estimate long-run projections of energy demand is the inability to account for adaptation over time. In the residential setting, fixing a household's portfolio of energy-using goods could lead to underestimates (if they buy an air conditioner) or overestimates (if they install rooftop solar) as temperatures increase.

Second, cross-sectional or time-series methods use wide spatial variation or long differences in climate to estimate the impact of long-run changes in climate. The advantage of these methods is that over large geographic or temporal dimensions, the extensive margin effects can be captured (assuming that individuals have "re-optimized" to their long-run equilibrium preferences). Albouy et al. (2016) use cross-sectional variation in the climate and estimate American's willingness to pay for climate amenities. Aside from concerns about omitted variables bias (OVB) in these methods, they have largely been unable to address shorter run fluctuations in the weather. I contribute to the synthesis of these by using short-run weather shocks to estimate adaptation at the extensive margin, and link this to longer-run implications for energy demand.

Studies focused on the extensive margin of air conditioning adoption often highlight the developing-country context, as non-linearities in the income-adoption curve imply large changes in future energy demand as incomes grow (Wolfram, Shelef and Gertler, 2012). Auffhammer (2014) uses monthly variation in temperature over a panel of provinces in China to measure the extensive effect of temperature on air conditioning adoption, and shows strong evidence that years following a hot summer see larger increases in adoption, but does not link this to realized energy demand. In a similar (but shorter) setting, Asadoorian, Eckaus and Schlosser (2008) use monthly variation in temperature over a panel of provinces in China to measure both intensive and extensive effects on energy demand through air conditioning. While they find that air conditioning adoption is highly sensitive to energy prices, they find no significant effects of monthly temperature on air conditioning in both urban and rural settings. There are two potential explanations for the different temperature/adoption relationship that I observe. First, the results from the California setting may not be externally valid to the China or other developing setting. Second, both of these studies focus on monthly variation. Instead, I focus on daily events that capture more information about the tails of the temperature distribution. If particularly severe events are more important to the adoption decision than mean changes, then these monthly panels may not adequately reflect the underlying temperature/adoption relationship.[2]

[2] This could also explain differences between Auffhammer (2014) and Asadoorian, Eckaus and Schlosser (2008), the former of which uses a longer province panel (1995–2009, compared to 1995–2000)

My findings suggest that in the California setting, severe shocks can increase adoption, even when changes in longer-run measures (such as monthly means) are modest. However, focusing on lower temporal frequencies is common in this space. Biddle (2008) shows that differences in long-run measures of climate can explain most of the differences in air conditioning penetration at the Metropolitan Statistical Area (MSA) level, with most of the residual difference explained by household income. In an engineering paper, McNeil and Letschert (2010) document correlations between climate and air conditioning adoption to define a measure they call "saturation" for air conditioning, and suggest this as a statistic in energy demand forecasts.[3] However, I document that, even in the absence of changing climate, a one-off shock can induce significant increases in medium- to long-run energy demand through air conditioning adoption, highlighting the importance of tail events.

1.2.2 Beliefs

What is the mechanism that links short-run weather shocks to long-run air conditioning holdings and energy consumption? I provide suggestive evidence for a household learning model. Using proxies for household belief in climate change, I show that households that believe in climate change are significantly more likely to be induced into adopting air conditioning following a 2006 heat wave. This is in line with a model of households updating their beliefs about their local climate when a weather shock (heat wave) provides them with new information about their local climate.

In some market settings, equilibria may be more likely to appear as fully incorporating information about the climate. For instance, Schlenker and Taylor (2019) show that aggregate financial markets for weather futures in the United States reflect consensus

[3]Though non-causal, this points to an important feature of the adoption margin relative to the local climate that is relevant for this paper: as adoption reaches saturation for a particular climate region, extreme heat wave can mechanically induce smaller changes in the adoption margin for air conditioning. This means that average effects from heat wave in empirical models presented below include the net effect of areas further and closer to saturation, implying heterogeneity in the potential to react to a heat wave.

climate change projections from the scientific community. In other instances, there are significant frictions where we observe large deviations from the equilibrium outcomes we might expect if individuals had perfect information about the climate.

Given significant household-level heterogeneity in climate change belief, there is a growing literature that explores how this affects investment decisions. Bakkensen and Barrage (2017) show that Rhode Island households that own coastal property are systematically more likely to underestimate flood risk, leading to housing prices that exceed their fundamentals. Barrage and Furst (2019) show further that new construction starts are more likely to occur in climate-skeptic communities. I consider similar household-level heterogeneity in belief, and use this to explain diverging adoption patterns in air conditioning based on beliefs about climate change. While this heterogeneity is consistent with a model of household learning about the climate, there are other potential channels that could drive adoption in this setting.

First, salience about the risk and frequency of climate events affects how individuals mitigate or insure against such risk. For example, households in flood-prone regions are more likely to purchase flood insurance after a flood event (Gallagher, 2014; Bakkensen, Ding and Ma, 2019). Additionally, information provision of both flood and wildfire risk maps affects equilibrium housing prices, even if the fundamental risk is unchanged (Gibson and Mullins, 2020; Garnache and Guilfoos, 2019). There are plausible salience mechanisms that could explain a link between heat waves and air conditioning adoption: disutility of heat, for example.

Second, behavioral channels could rationalize weather-induced adoption of air conditioning. Busse et al. (2015) study household purchases of vehicles in the presence of idiosyncratic weather phenomena. They find that the investment decision responds to idiosyncratic weather, which is inconsistent with a fully-rational purchase decision. Instead, they provide evidence of projection bias from current weather, where future utility is a convex combination of utility based on the current idiosyncratic state and the realized state. To this extent, households may exhibit similar behavioral biases as a response to an unexpected heat wave.

Finally, a simple rationalization of weather-induced adoption of air conditioning would be highly convex costs of temperature imposed by a heat wave. Contemporaneous costs imposed by a severe heat wave could rationalize contemporaneous adoption of air conditioning regardless of the net present value of ownership in future periods. High contemporaneous costs are consistent with results from Albouy et al. (2016), where households are willing to pay significantly more to avoid extreme high temperatures than to avoid similarly extreme low temperatures in a cross-sectional hedonic analysis. This follows the same pattern as impacts on crop yields in Schlenker and Roberts (2009), suggesting a similar physiological aversion to extreme temperatures.

However, in each of these alternative channels, unless the mechanism is systematically correlated with belief in climate change, it cannot fully explain the patterns in air conditioning adoption in this setting. Additionally, the significance of the belief channel suggests that it is not negligible relative to alternative mechanisms. I discuss the contribution of alternative mechanisms further in Section 1.6.

1.3 Data

1.3.1 Household appliance and energy data

The primary data I use in the empirical analysis contain information on household appliance ownership and one year of monthly energy use from the Residential Appliance Saturation Survey (RASS), commissioned by the California Energy Commission in order to project future energy demand. This cross-sectional survey includes 21,920 and 24,464 California households in 2003 and 2009, respectively. These households were randomly drawn from the service areas of three primary independently-owned utilities (IOUs)—Pacific Gas and Electric (PG&E), Southern California Edison (SCE), and San Diego Gas and Electric (SDG&E)—and the largest publicly-owned utility, Los Angeles Department of Water and Power, which collectively serve 87 percent of California's customers.[4] Households are identified geographically at the ZIP-code level.

The primary outcome of interest is household ownership of central or room air conditioning units and monthly electricity demand. In addition, the RASS includes other household characteristics including household size, home age, income, and education of the head of household. In addition, I use household-reported installations of other appliances including dishwashers, standalone freezers, and CFL lighting as a placebo test for the baseline estimation strategy.

Table 1.1 breaks down the level of central and room air conditioning each year by California utility. Percent of households owning central air conditioning increased in each utility's jurisdiction. In aggregate, central air ownership increased 8 percentage points between 2003 and 2009. In all jurisdictions aside from SDG&E, ownership of room air conditioning units (that is, window units or standalone units) decreased, suggesting substitution towards central air conditioning.

Table 1.1. Air conditioning saturation by utility from RASS

	Survey wave					
	2003			2009		
	Central Air	Room Air	N	Central Air	Room Air	N
PG&E	.39	.14	6,265	.44	.11	6,458
SDG&E	.35	.09	5,445	.43	.13	5,970
SCE	.48	.20	6,102	.58	.18	6,444
LADWP	.29	.25	4,071	.41	.24	5,538

Note: summaries for proportion of installations of central or room-style air conditioning by utility. RASS covers the three largest IOUs and the largest POU, LADWP to represent greater than 80 percent of California households.

Installation of central air conditioning represents a sizable investment for a household. The 2019 national average cost of installation is reported to be typically between $4,000 and $7,000, with potentially higher costs depending on idiosyncratic home char-

[4]As reported for 2010 by the State of California Energy Commission.

acteristics.[5] Conversely, portable room units can be purchased for only a few hundred dollars, and because of relative portability, have an active secondary market.[6] Central units are tied to the structure, and represent a longer-term investment decision for the house. Because of this, I focus the primary analysis on central air-conditioning units, but report the robustness of the primary results using room air-conditioning units, and the combination of all units in Section A.1.

1.3.2 Temperature data

I obtain local weather data at the ZIP-code level from the Parameter elevation Regression on Independent Slopes Model (PRISM), which uses meteorological models and weather station data to interpolate daily temperatures at a four kilometer resolution (PRISM Climate Group, 2021). For each ZIP code, I take a simple mean of pixels that are bounded within a ZIP code for a daily observation; or in the case that no pixel falls within a ZIP code, I take the closest pixel observation. I winsorize the ZIP-code average daily temperatures at the top and bottom one percent, and match these to the household appliance and energy use data (identified geographically at the ZIP-code level).

1.3.3 Construction of historical climate and temperature anomalies

In order to relate contemporaneous weather observations to the climate of a locality, I construct a measure of local historical climate and define yearly anomalies relative to this historical climate. For a ZIP code z, I count the number of cooling degree days (CDD) in a year t, where a CDD occurs when the mean daily temperature is above 65° Fahrenheit, as given by Equation 1.1. I define the historical climate for ZIP code z to be the mean number of CDDs per year from 1981 to 2005 (Equation 1.2). I use 2005 as the upper

[5] See https://www.homeadvisor.com/cost/heating-and-cooling/install-an-ac-unit/.

[6] For example, Figure A.2 depicts results from an August 13, 2019 search for air conditioning on an online resale website in San Diego.

cutoff for this historical climate and focus on the plausibly exogenous temperature shock during the summer of 2006—a particularly hot year for California.

$$CDD_{z,t} = \sum_{\text{days } \in t} (\text{mean temp} - 65°F) \times \mathbb{1}(\text{mean temp} > 65°F) \tag{1.1}$$

$$Climate_z \equiv \overline{CDD}_{z,t\in\{1981,...,2005\}} \tag{1.2}$$

Figure 1.1 depicts the distribution of relative anomalies by ZIP code, where the large mass of this distribution lies to the right of zero. This can be interpreted as saying that most California ZIP codes (covered by the RASS) experienced more CDDs during 2006 than during a typical year leading up to that point. I use this 2006 heat wave as an event between the two RASS survey waves (2003 and 2009), and consider the ZIP-code level heterogeneity in exposure to this heat wave as a source of identifying variation. I define the *CDD anomaly* as the difference between the number of CDDs in 2006 and the historical climate, enumerated by Equation 1.3.

$$CDD\,anomaly_z \equiv CDD_{z,2006} - Climate_z \tag{1.3}$$

One alternative formulation would be to define the CDD anomaly as the number of extra CDDs compared to an average for all years between the two survey waves, that is, construct the CDD anomaly for 2004 through 2008. I report the results of this exercise in Section A.1 and estimate qualitatively similar estimates for the baseline empirical specification, but with less precision.

The immediate concern of using the CDD anomaly as the identifying variation is whether this anomaly can be taken as plausibly exogenous. For example, one may think that the historical climate of a ZIP code as measured by average CDDs could be correlated with the 2006 anomaly. That is, if historically hot ZIP codes are more likely to experience

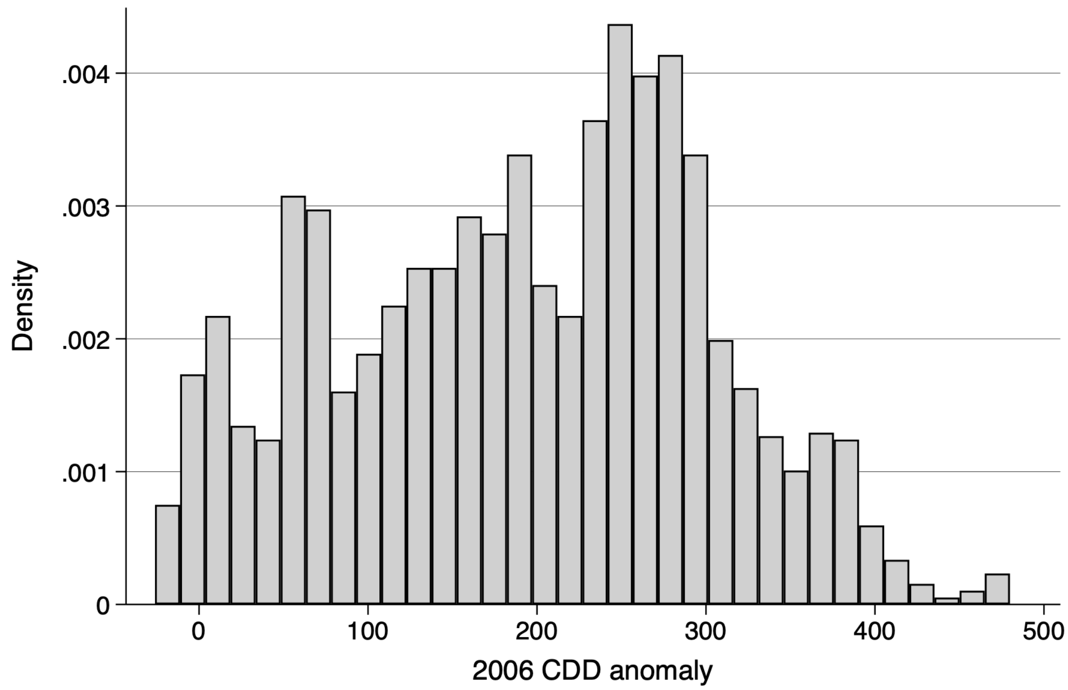

Figure 1.1. Heat wave anomaly for 2006

Note: this histogram depicts the number of extra CDDs in 2006 by ZIP code relative to the historical number of CDDs for the ZIP code.

larger anomalies, this would be a concern for the identifying assumption that this CDD anomaly is orthogonal to the historical climate. Figure 1.2 shows that the residual climate anomaly (after netting out city fixed effects) is not predicted by the historical number of CDDs for a ZIP code.

Similarly, one may be concerned that ZIP codes that are more severely affected by a heat wave in 2006 may typically experience more variance in year-to-year daily temperatures. That is, if a household living in a particular ZIP code experiences severe temperature anomalies in 2006, but they were also likely to be more exposed in previous years, the 2006 anomaly from historical *mean* years would not capture this fact. However, I find that the 2006 anomaly is not predictive of past daily temperature variance. Figure 1.3 plots the 2006 CDD anomaly against the historical CDD variance from 1981 to 2005 after netting out city fixed effects, and shows that the extreme heat wave in 2006 is not predictive of past variability in average number of CDDs.

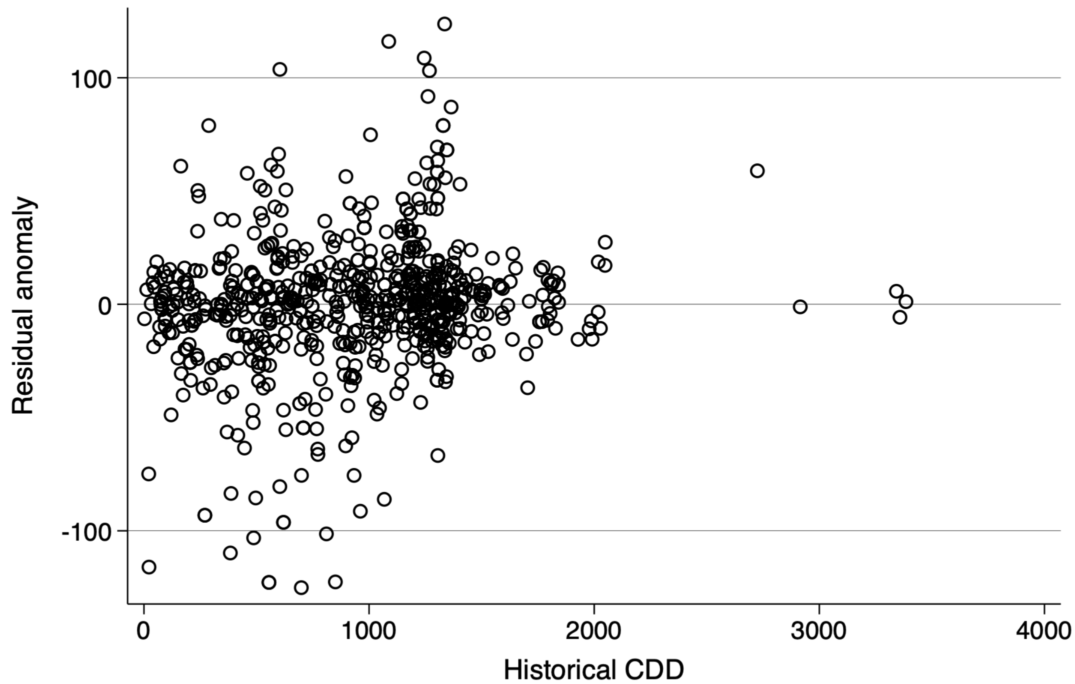

Figure 1.2. Orthogonality of residual anomaly

Note: the plot shows the residual climate anomaly after netting out city fixed effects. This depicts the fact that historical number of CDD is not predictive of the residual 2006 CDD anomaly.

1.3.4 Beliefs about climate

Finally, I match these household data identified at the ZIP-code level to two different measures of beliefs in climate change. First, I use the 2018 Yale Climate Survey that reports county-level measures of belief in climate change. The specific series I use is any belief in climate change, regardless of belief in its severity or cause.[7] Since this is measured at the county level, using this as a measure for household-level beliefs will introduce significant measurement error.

[7]The specific question is: "Do you think that global warming is happening?" Howe et al. (2015)

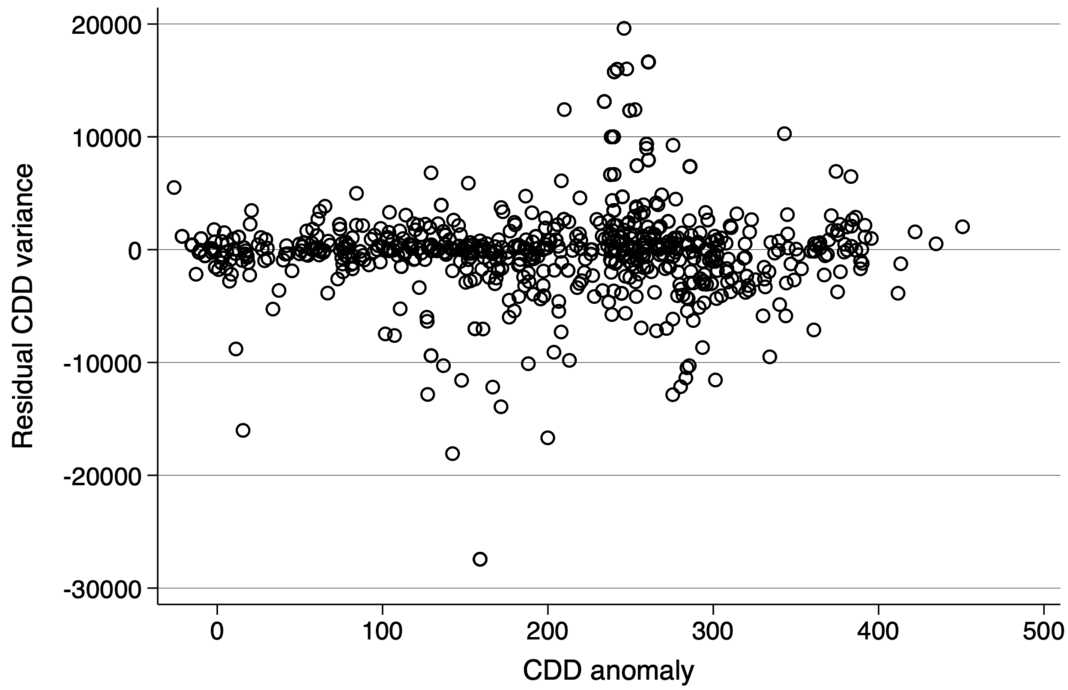

Figure 1.3. Contemporaneous anomaly and historical climate variance

Note: the plot shows the residual CDD variance after netting out city fixed effects. The variance of CDD for ZIP code z is defined as the variance of the yearly number of CDD for a ZIP code from 1981 to 2005. The CDD anomaly in 2006, on the horizontal axis, is not predictive of past variability in the number of CDD in a given year.

To supplement this county level measure, I also use precinct-level elections returns data for the 2004 presidential election from the California Secretary of State. Partisanship is highly predictive of belief in climate change and its severity (McCright and Dunlap, 2011). I spatially match households in a particular ZIP code with percent Democratic share in the nearest election precinct, using both ZIP code and precinct centroids for geographic distance. While I use both of these measures as a very imperfect proxy for household level belief in climate change, I also use education and precinct Democratic share to instrument for belief defined at the county level.

In addition, I use a fact documented by a 2015 Gallup poll about belief heterogeneity in climate change by level of education.[8] For Democrats, belief in climate change is increasing in education. However, for Republicans, belief is decreasing in education. I exploit this empirical fact in my discussion of mechanisms for induced air conditioning adoption.

1.4 Empirical Strategy

In this section, I discuss the identification strategy that I use in order to measure the effect of short-run temperature anomalies on household ownership of air conditioning and longer-run energy demand.

In an ideal laboratory experiment, I would randomly expose ZIP codes to an extreme heat wave indicator and observe air conditioning penetration pre- and post-treatment. Assuming households are geographically fixed, a regression of post-treatment air conditioning penetration on the randomized heat wave and pre-treatment penetration will identify the average take-up in air conditioning induced by the heat wave. Assuming that penetration is diminishing in average temperature, it is important to note that ZIP codes with penetration near saturation levels will minimize potential contribution through this channel, and that the specific parameter identified will be the average effect over the distribution of households.

In reality, heat waves are not binary, and heat wave severity will be correlated with mean climate characteristics, so I use the CDD anomaly measure as defined in the previous section by Equation 1.3 as the source of identifying variation. Above, I discussed some of the potential threats to identification and why these would not be a problem for the estimation. However, another concern would be if CDD anomaly were correlated with observed or unobservable household characteristics. To this effect, I compare observables for households located in the fourth quarter of the CDD anomaly to a random sample of 100 ZIP codes covered by the RASS and report the summary statistics in Table 1.2.

[8] https://news.gallup.com/poll/182159/college-educated-republicans-skeptical-global-warming.aspx

Households particularly affected by the CDD anomaly do not appear noticeably different from randomly chosen households, which is reassuring for the estimation.

Table 1.2. Testing the orthogonality of the heat wave using simulated weather shocks

	(1)		(2)	
	Q4 anomaly		Random sample	
	mean	sd	mean	sd
Climate (avg yearly CDDs)	1001.64	410.90	901.98	41.84
Central air installation	0.42	0.49	0.44	0.03
Dem presidential vote share 2004	0.55	0.22	0.54	0.02
College educated	0.41	0.48	0.44	0.02
HH income ($1000s of 2009 USD)	52.53	40.55	60.21	2.19
Owner occupied	0.58	0.49	0.62	0.02
Number of bedrooms	2.52	1.14	2.61	0.06
Home age (years)	35.79	17.84	33.24	0.84
Observations	5713		100	

Note: this reports a test of selecting a random sample of ZIP codes to the fourth quartile of 100 ZIP codes as affected by the heat wave anomaly. Households are sampled with equal probabilities from the ZIP codes covered by the four utilities in the RASS.

1.4.1 Baseline specification

The baseline model is a difference in differences (DD) linear probability model in air conditioning ownership, where I compare the probability that a household in a ZIP code in 2009 owns an air conditioning unit compared to a household in the same ZIP code in 2003. This is differenced by CDD anomaly for that specific ZIP code. City fixed effects imply that the identifying variation is differential ZIP-code exposure to the CDD anomaly within a city. I use the same estimation strategy for measuring household-level summer electricity demand, and interpret this as the longer-run effect of the CDD anomaly on monthly electricity demand in kilowatt hours (kWh).

Since the RASS reports two cross sections, I am unable to test for parallel trends in ZIP-code level air conditioning penetration. However, due to the nature of the CDD anomaly and its credible exogeneity (discussed above), this does not pose the same identification problem that may exist in other DD designs.

The baseline estimating equation is:

$$y_{zit} = \beta_0 \left(CDD\, anomaly\right)_z \times \mathbb{1}\{2009\}_t + \beta_1 \mathbb{1}\{2009\}_t + \Theta X_i + \gamma_k + \delta_t + \varepsilon_{zit}, \qquad (1.4)$$

where y_{zit} is the outcome for a household i in ZIP code z in year t. X_i is a vector of housing and homeowner characteristics including a dummy for college education for head of household, household income, home age, and number of residents. γ_k and δ_t are spatial and year fixed effects.

When household central air ownership is the dependent variable, β_0 is interpreted as the effect of one extra CDD in 2006 relative to the historical local climate on a household's propensity to own a central air conditioner. The identifying variation is differential exposure to the 2006 temperature anomaly within a city. The standard errors are clustered at the ZIP code, which is the level of the treatment (CDD anomaly in 2006).[9] In different specifications of the linear probability model, I use household ownership of other types of appliances—dishwashers, standalone freezers, and CFLs—as a placebo test for the baseline model. The implicit assumption is that ownership of these items is unlikely to covary with the weather.

When July electricity demand is the dependent variable, β_0 is interpreted as the effect of one extra CDD in 2006 relative to the historical ZIP code climate on household electricity demand in 2009. I interpret this reduced-form estimate as the combined effect of induced air conditioning adoption and choosing to run the air conditioner during July 2009. The analogous placebo test to the linear probability model above is electricity demand in the winter (specifically, February).

[9]When clustering the standard errors at a higher geographic level, such as city or climate zone, the point estimates are still measured precisely. See Section A.1

1.4.2 Heterogeneity and mechanism analysis

In order to explore the mechanism through which households are induced into acquiring air conditioners by temperature anomalies, I interact specific variables with the DD in the baseline specification.

First, I consider a triple difference model, where I difference the baseline specification against household-level belief in climate change. The estimating equation is:

$$AC_i = \beta_0\, belief_k \times (CDDanomaly)_z \times (2009)_i + \sum_j \beta_j [\text{two-way interactions}] + \Theta X_i + \gamma_k + \delta_t + \varepsilon_i, \quad (1.5)$$

where the right-hand side variables are as defined above, and the two-way interactions are all combinations of the three variables in the triple difference. When using belief as defined at the county level, I drop the city fixed effects that would otherwise absorb county-level belief in climate change.

Because the estimate of β_0 will be attenuated by mis-measurement of household-level belief in climate change, I also use precinct-level Democratic share in the 2004 presidential election as a proxy for household belief in climate change. In addition, I use a measure of predicted belief in Equation 1.5, where the first stage regresses county belief on household education and precinct Democratic share, as well as the interaction of the two.

In addition to measures of belief, I also consider a triple-differenced model with education of head-of-household as follows:

$$AC_i = \beta_0\, College_i \times (CDDanomaly)_z \times (2009)_i + \sum_j \beta_j [\text{two-way interactions}] + \Theta X_i + \gamma_k + \delta_t + \varepsilon_i. \quad (1.6)$$

I also consider a fourth-differenced model, where we can difference across precinct Democratic share:

$$AC_i = \beta_0 College_i \times Demshare_p \times (CDDanomaly)_z \times (2009)_i + \sum_j \beta_j [\text{two-way interactions}] + \sum_l \beta_l [\text{three-way interactions}] + \Theta X_i + \gamma_k + \delta_t + \varepsilon_i. \quad (1.7)$$

A recent Gallup poll of Americans indicates that belief in climate change is not strictly increasing in education.[10] Instead, belief differs by partisanship. Conditional on identifying as a Democrat, belief in climate change is increasing in education. Conversely, conditional on Republican identification, belief in climate change decreases with education.

Because of this, the *ex ante* expectation of the sign on the estimate of β_0 in Equation 1.6 is unclear. If one of the mechanisms for induced adoption is belief that the climate is changing, higher levels of education could be associated with a higher or lower propensity of belief depending on partisan identification. However, when differencing this again by a proxy for partisanship, as in Equation 1.7, β_0 is interpreted as the propensity for college-educated households in a more Democratically-leaning ZIP code to have been induced into adopting an air conditioning unit by the 2006 CDD anomaly.

In addition to informing these third- and fourth-differenced models, this non-monotonic relationship between education and belief in climate change implies the importance of the interaction with education for the predicted household belief, which I use for the preferred estimate for the model defined by Equation 1.5 when using a predicted measure for household belief.

[10] See Section A.2. Source: Gallup poll available here.

1.5 Results

In this section, I report the main results of the estimated models defined in the previous section.

1.5.1 Baseline results of weather on air conditioning adoption

Table 1.3 reports the model estimates for the baseline models of temperature-anomaly induced adoption of central air conditioning (columns 1 and 2) and long-run household electricity demand in kWh. Column 1 is the preferred specification, and the coefficient on "Anomaly $\times$ 2009" can be interpreted as follows: the mean effect of an extra 100 CDDs in 2006 relative to the historical average from 1981–2005 increases the propensity for a household to own central air conditioning by one percentage point. The identifying variation is differential ZIP code exposure to the 2006 heat wave within a city. Utility fixed effects and year fixed effects absorb variation induced by differential utility structures and variables common to households within a year respectively.

In addition to the average effect across households, column 2 in Table 1.3 separates this effect across quartiles of the historical climate distribution. That is, "Q1" refers to ZIP codes where the historical climate lies in the first quartile of the California distribution (defined by average number of CDDs in a year from 1981–2005). Similarly, "Q4" refers to the quartile of ZIP codes where households historically experience the highest number of CDDs within a given year (based on the daily temperature data from 1981–2005). Households in the third quartile of this climate distribution have the largest and most precise point estimate—for every 100 CDDs of anomaly, households are 2 percentage points more likely to be induced into adopting central air conditioning. Column 2 also provides suggestive evidence that households in the second quartile of California climate are more likely to be induced into air conditioning adoption than either the top or bottom quartiles.

Table 1.4 reports the results of three separate placebo tests following Equation 1.4,

Table 1.3. Baseline model

	(1) Central air b/se	(2) Central air b/se	(3) Electricity b/se	(4) Electricity winter b/se
Anomaly X 2009	0.0001***		0.06023***	0.00867
	(0.00005)		(0.010253)	(0.006725)
Q1 interaction		-0.0000		
		(0.00018)		
Q2 interaction		0.0001		
		(0.00008)		
Q3 interaction		0.0002***		
		(0.00006)		
Q4 interaction		-0.0000		
		(0.00006)		
Controls	X	X	X	X
UtilityFE	X	X	X	X
CityFE	X	X	X	X
N	38581	38581	35734	33503

Note: standard errors clustered at the ZIP-code level. In column 3, electricity corresponds to household electricity demand for the billing cycle covering mostly July in kWh. In column 4, this corresponds to electric billing cycle covering most of February. "Most coverage" is necessary because of the staggered billing cycles across households. *p< 0.1, **p< 0.05, ***p< 0.01.

Table 1.4. Placebo tests

	(1) Dishwasher b/se	(2) Freezer b/se	(3) CFLs b/se
Anomaly × 2009	-0.0000	0.0001	-0.0000
	(0.00004)	(0.00004)	(0.00004)
Controls	X	X	X
Utility FE	X	X	X
City FE	X	X	X
N	38581	37209	36756

Note: standard errors clustered at the ZIP-code level. In column 2, freezer refers to household ownership of a standalone freezer. CFLs refer to household ownership of compact fluorescent lamps, the energy efficient lightbulb at the time of the RASS. N varies by specification due to different household response rates to ownership of the various appliances. N falls for the electricity demand models as billing data is not fully populated in the RASS. *p< 0.1, **p< 0.05, ***p< 0.01.

where the outcome variable is household ownership of a dishwasher (column 1), a standalone freezer (column 2), or installation of fluorescent light bulbs (column 3). These precisely estimated zeroes for each of these specifications imply 95 percent confidence intervals that do not include any response larger than one one-hundredth of a percent in order of magnitude. Though this response is unsurprising, this lends credibility to the assumption that the model only captures household investment decisions that should be directly affected by temperature anomalies. One could imagine that a negative and significant coefficient on any of these appliances could arise if households substitute purchases of one appliance for another, but this does not seem to be the case.

1.5.2 Baseline results of weather on energy demand

Column 3 of Table 1.3 reports the baseline results of household-level electricity demand for July following Equation 1.4. The point estimate is interpreted as the following: for every 100 CDDs of anomaly, households on average increase their electricity demand by approximately 6 kWh.

Finally, column 4 of Table 1.3 shows the similar placebo results for household energy demand during the month of February. The coefficient can be interpreted as a precisely-estimated zero estimate. For every 100 CDDs of anomaly, I can rule out an increase in household electricity demand by more than 2 kWh at the 95 percent confidence interval.[11] Of course, there is no *ex ante* reason to think that induced adoption of central air conditioning does not affect winter energy demand. For instance, HVAC installation may correlate with installation of heating units, insulation, or other homeowner investments.

1.5.3 Heterogeneity in climate change belief and education

Table 1.5 reports the estimation results following Equation 1.5, the linear model differencing the baseline specification across different proxies for household-level belief in climate change. Column 1 of Table 1.5 uses county level belief in climate change. Since "belief" is the county share of adults that believe in climate change, there exists a large amount of household level measurement error for belief in climate change. The lowest proportion of county-level belief in California is 61 percent of adults, with the highest proportion being 79 percent of adults believing in climate change. The estimate in column 1 implies that the induced adoption effect is stronger in counties where people are more likely to believe in climate change.

Column 2 of Table 1.5 instead uses the Democratic share from the 2004 presidential general election from the nearest precinct closest to a ZIP code to difference across the baseline specification. Again, this will measure household-level belief in climate change with a large degree of error, but the point estimate may still be interpreted as saying that households that are more likely to be Democratic identifying (more likely to believe in climate change) are more likely to have been induced into adopting air conditioning.

[11] Note that the differing number of household observations in each specification is due differential response to the RASS.

Table 1.5. Differencing over belief in climate change

	(1) Central Air b/se	(2) Central Air b/se	(3) Central Air b/se
Belief × CDD anomaly × 2009	0.0001***		
	(0.00002)		
Dem share × CDD anomaly × 2009		0.0004***	
		(0.00013)	
$\hat{\text{Belief}}$ × CDD anomaly × 2009			0.00064***
			(0.000194)
Two-way interactions	X	X	X
Controls	X	X	X
Utility FE	X	X	X
City FE		X	X
N	38674	38581	41491

Note: standard errors clustered at the ZIP-code level. Education is an indicator for college education for head of household. Democratic share is precinct-level Democratic share in the 2004 presidential election. *p< 0.1, **p< 0.05, ***p< 0.01.

In order to try to correct for this measurement error, I use head of household education and closest precinct Democratic share to predict household-level belief in climate change. Recalling the non-monotonicity of belief in education, discussed above, I also include the interaction in the preferred specification. Table 1.6 reports the results of this first stage, where both columns 1 and 2 have R-squared of about 0.3. College education, Democratic precinct share in the 2004 presidential election, and being both college educated and living in a more Democratic ZIP code are all positively correlated with belief in climate change in this first stage.

When using this predicted measure for belief, column 3 of Table 1.5 reports my preferred specification of the triple difference. Since county level of belief ranges from 61 to 79 percent belief, moving from households that are least likely to most likely to believe in climate change increases the induced propensity by about one half of one percentage point per 10 CDD anomaly.

Table 1.7 reports the estimation results following the models specified in Equation 1.6

Table 1.6. First stage for belief in climate change

	(1) Belief b/se	(2) Belief b/se
Education	0.0089***	0.0029***
	(0.00027)	(0.00077)
Democratic share	0.0879***	0.0823***
	(0.00068)	(0.00095)
Education × democratic share		0.0114***
		(0.00136)
Controls		
Utility FE		
City FE		
N	42312	42312

Note: standard errors clustered at the ZIP-code level. Belief is a county-level report of the percentage of adults that believe that the climate is changing. Twoways refers to the two-way interactions necessary for the triple difference model. *p< 0.1, **p< 0.05, ***p< 0.01.

and Equation 1.7. Column 1 differences the main specification against an indicator for college education for the head of household. Though not significant at the 5 percent level, the coefficient provides some suggestive evidence that college educated households are more likely to be induced into adopting air conditioning.

Column 2 of Table 1.7 indicates that college educated household in the most Democratic ZIP codes are significantly more likely to be induced into air conditioning adoption than college educated households in the least Democratic ZIP codes. The negative (but insignificant) coefficient on the first row is in line with a small negative or zero effect for households located in comparatively non-Democratic ZIP codes.

Because households may choose to endogenously relocate based on weather shocks, I run all of these specifications restricted to homeowners that have lived in their homes since prior to the 2006 heat wave, and I report these results in Section A.1. Here, I also report the results of all models run using room air conditioning or any air conditioning as

Table 1.7. Heterogeneity of effects by education

	(1) Central Air b/se	(2) Central Air b/se
College × CDD anomaly × 2009	0.0001	-0.0002
	(0.00005)	(0.00010)
College × Dem share × CDD anomaly × 2009		0.0005**
		(0.00019)
Two-way interactions	X	X
Three-way interactions		X
Controls	X	X
Utility FE	X	X
City FE	X	X
N	38581	38581

Note: standard errors clustered at the ZIP-code level. Twoways (threeways) refers to the two-way (three-way) interactions necessary for the triple difference model. *p< 0.1, **p< 0.05, ***p< 0.01.

the outcome variable, with qualitatively similar results.

1.6 Discussion

In this section, I discuss the empirical results and some of the implied mechanisms for induced air conditioning adoption. I first summarize the baseline results and their implications, and move towards the suggestive evidence about the role of household beliefs in climate change.

I provide strong causal evidence that households respond to a CDD anomaly in 2006 by increasing their propensity to own air conditioning units. Following the results reported in Table 1.3, we can observe that the households most likely to be induced to adopt air conditioning live in the third quarter of the climate distribution in California, followed by the second quartile. This is unsurprising if ZIP codes in the top quartile of the climate distribution are fully saturated with air conditioning, and if if households in the first quartile never need air conditioning.

Using the preferred estimate from the baseline model, and using the median CDD anomaly of about 200 CDDs, the model can explain about 2 percentage points of increase in the propensity for these California households to own an air conditioner. From Table 1.1, this means that response to this severe heat wave can explain about one quarter of the overall increase in central air conditioning ownership from 2003 to 2009 (an increase of about 8 percent). Using the same baseline model and median CDDs, this translates to about a 12 kWh increase in monthly demand for all California households covered by the survey.

The total increase in energy demand depends on the behavior of counterfactual households in the absence of the 2006 heat wave. If these households were to eventually adopt air conditioners, but waited for a particularly hot year, this might amount to energy "pull-forward." In this case, the net effect on energy demand depends on the timing of a counterfactual shock—from earlier to the shock of interest to far in the future. However, if the 2006 heat wave actually permanently changed the stock of installed air conditioners, the total effect on energy demand would be significantly larger.

1.6.1 Alternative mechanisms

Prior to discussing the climate-change belief channel for air conditioning adoption, I discuss several other channels for this induced adoption. For context, consider an extremely stylized model of a household making the decision about air conditioning ownership. Formally, suppose an air conditioner lasts for T periods, weather in period t is given by ω_t from a climate distribution Ω, and household preferences over the AC decision are given by u, then a household invests in air conditioning in year t if $\mathrm{E}\sum_{\tau=t}^{t+T} u(AC, \omega_\tau) > 0$. That is, they purchase an air conditioner if the expected utility over the lifetime of the air conditioner is positive (net of all costs—fixed and flow costs, and any other associated utility costs).

First, if the contemporaneous heat wave is bad enough to cause extreme household disutility, it could be rational to install an air conditioner to assuage these high short-term costs, even if the expected net benefit for the following periods were negative. In the context of the simple model of household investment in air conditioning, this could be the case if $u(AC, \omega_t) > 0$ even if the expected utility from ownership $\mathrm{E}\sum_{\tau=t+1}^{t+T} u(AC, \omega_\tau) < 0$. However, qualitative reports of HVAC installation imply that the time horizon for central air installation would make it difficult to believe that households are responding simply because of contemporaneous disutility.

A second mechanism discussed above would simply be a shift in timing. If a weather realization in t changed the timing decision for a household acquiring air conditioning based on static expectations about the weather, this would amount to pull-forward in energy demand for some number of periods. If this were the case, we may expect to see smaller effects for induced adoption with increasing duration of home ownership. Section A.1 shows that effects are not decreasing when restricting to homeowners with substantial tenure.

Other alternative mechanisms for adoption could include behavioral channels. One potential behavioral channel would be myopia or projection bias on the part of the household. This would be the case if the future utility of air conditioning ownership is a convex combination of ownership today and the actual realized utility of air conditioning ownership. Busse et al. (2015) provide strong evidence for projection bias in the automobile purchase decision, where consumers buy more convertibles on sunny days, the analogous phenomenon in this setting would be buying more air conditioning units during hot years.

1.6.2 Beliefs about the climate

While I cannot rule out the contribution of alternative mechanisms for weather-induced adoption of air conditioning, I propose a simple framework that rationalizes the heterogeneity in the induced adoption observed in patterns of adoption by belief in climate change, education, and partisanship.

Consider again a household that is posed with the choice to buy an air conditioner if $\mathrm{E}\sum_{\tau=t}^{t+T} u(AC, \omega_\tau) > 0$. While the above mechanisms focused on the expected utility calculation given a series of weather draws ω_t from a fixed climate distribution Ω, consider instead two types of households: those that believe that the climate distribution Ω is fixed, and those that believe that the climate distribution is actually changing—"updaters." In this setting, the updaters observe the weather today, take this as a signal of the future path of weather, and update their beliefs of the climate based on the contemporaneous weather. If nothing in the consumer's choice changes except for expectations over the climate, this can rationalize investment in central air conditioning.

This framework is consistent with the heterogeneity results discussed in Subsection 1.5.3. We can observe that when interacted with the baseline model, both county-level survey data about belief in climate change and a proxy using precinct-level election returns suggest that households that are more likely to believe in climate change (measured with large error) are more likely to be induced adopters. And when using the predicted belief using the first stage reported in column 2 of Table 1.6, the coefficient of interest indicating induced adoption is larger, since this mitigates attenuation caused by mismeasurement in columns 1 and 2.

Further, the results in Table 1.7 qualitatively follow what would be predicted in this framework. Since higher education contributes to divergent beliefs about climate by partisanship, it is not immediately clear what the expected estimated coefficient of the triple difference following Equation 1.6 would be. Rather, it should depend on the relative partisan share of the population. For instance, if all households identified as Republicans, I would expect the difference across education to be negative if households are changing their investment behavior after updating their beliefs about climate. Again, this follows the empirical facts about belief in climate change and partisanship discussed in Subsection 1.3.4. However, when I introduce a *fourth* difference by Democratic share, the coefficient of interest (differential induced adoption by educated Democratic households) is positive and significant. This is consistent with the fact that college educated Democrats are the most likely to believe in climate change and are induced into adoption differentially by the 2006 heat wave.

Though most of these heterogeneity analyses are identified imprecisely with imperfect measures of household-level belief in climate change, I find the preponderance of evidence consistent with this framework to be highly suggestive that there is some role for a belief-updating channel with respect to the climate. That is, this induced adoption of air conditioning cannot be fully rationalized by alternative channels. I take this as novel evidence that households change their investment behavior as a response to updating beliefs about the climate when it comes to household investment decisions. In this context, this has implications for the dynamics of air conditioning ownership and the path of energy demand over time. More expansively, it is likely similar mechanisms may affect a variety of dynamic consumer problems that are related to a changing climate.

1.7 Conclusion

In this paper, I provide causal evidence that households respond to short-run weather shocks by making investment decisions with long-run implications. Particularly, California households differentially exposed to a 2006 heat wave increased their propensity to own a central air conditioning unit, and this event can explain nearly 2 percentage points of the increase in central air conditioner ownership—or about one quarter of the total increase—from 2003 to 2009. Through this induced adoption channel, households also increased their July energy demand three years following the heat wave. Exploring this link between short-run weather and long-run energy demand is immediately important for forecasting exercises, but also exposes an important mechanism by which households make investment decisions.

This has direct implications for forecasting air conditioning ownership and long-run energy demand as households are exposed to extreme weather. Previous studies define air conditioning penetration as a function of a fixed climate and other state variables (Deschênes and Greenstone, 2011). However, here I provide evidence of the dynamic adoption of air conditioning depending at least in part on tail events (heterogeneity in exposure to a severe heat wave). The total effect on energy demand is still an open question, as I cannot identify whether, on the low end, if this is simply a timing decision where the effect on energy demand would be some amount of "pull-forward," or, on the high end, whether this short-run weather realization induces adoption for a household that otherwise would never purchase air conditioning absent of the heat wave.

In addition to these baseline results, I also explore the heterogeneity in response to the 2006 heat wave and provide suggestive evidence for a household belief-updating framework. When differencing the baseline DD model by measures of household belief in climate change, I show that households that are more likely to believe that the climate is changing are also more likely to be induced into adopting air conditioning in response to the

2006 heat wave. This effect is strong enough to be detected when using even very imperfect measures of household-level belief in climate change. This behavior can be rationalized by a belief-updating framework, where households that believe in climate change take contemporaneous weather as a signal of the future path of their local climate, which changes the household decision for purchasing air conditioning. Conversely, a household that does *not* believe that the climate is changing may not update their expected upside of adopting air conditioning.

Finally, I exploit the non-monotonicity of partisan belief in climate change with respect to education to test this belief-updating framework. College-educated individuals are not necessarily more likely to believe in climate change; but, conditional on partisanship, college education increases belief for Democrats and *decreases* belief for Republicans. In a triple- and fourth- difference specification, I show that the induced-adopter effect differentially applies to college educated Democrats, who are most most likely to believe in climate change.

This evidence that long-run household investment decisions responds to short-run weather phenomena is important, since the prevailing literature links such decisions to long-run dynamics. In this paper, I provide novel, suggestive evidence that households take contemporaneous weather events and form beliefs over longer-run state variables in the climate. This suggests implications for household decision-making both within the environmental setting and more broadly.

1.8 Acknowledgements

Chapter 1 is currently being prepared for submission for publication. The dissertation author is the sole author on this chapter.

Chapter 2

Inundated by Change: The Effects of Land Use on Flood Damages

2.1 Introduction

Climate change is expected to increase the frequency and severity of storms and tail rainfall events (Cleetus, 2013; Donat et al., 2016). Proper urban planning around the built environment can help mitigate against these increased risks, and mitigation can also follow through management of wetlands and other terrestrial ecosystems (Woodruff, Irish and Camargo, 2013). In the past decades, urban planning, growth, land management and degradation have contributed to large changes in land use and land cover, with more than eight percent of the United States' land cover changing at least once between 1973 and 2000 (Sleeter et al., 2013). This encompasses changes in development, shifting agricultural patterns, loss (and restoration) of forest cover, and large losses of coastal and inland wetlands. Per a 2019 Intergovernmental Panel on Climate Change (IPCC) special report, changes in the climate are expected to contribute to and accelerate these changes (Shukla et al., 2019).

In this paper, I measure the effect of land-use change on flood risk through its impact on flood insurance claims. I combine fine temporal and spatial variation in land use in a Texas panel from 2010–2016 with monthly anomalous rainfall to estimate a difference in differences (DD) model that identifies the impact of within-tract changes in land use on flood insurance claims. In the baseline specification, I find that increases in impervious surface coverage and decreases in wetlands and water coverage increase the propensity of flood claims, conditional on a rainfall shock. Additionally, I explore whether these effects are symmetric for positive and negative changes in particular land cover. Importantly, I find that both positive changes in wetlands coverage (restoration) decrease the the numeracy of flood claims, and corresponding decreases in wetlands coverage increase flood claims.

Second, I combine this DD design with variation in elevation to show the existence of spillovers in land-use-based flood risk. Relative to a neighbor that is downhill, land-use change for an uphill neighbor is more consequential for flood risk. I show that changes in development, water, and wetland coverage cause significantly larger impacts for neighboring regions when these land-use changes happen at a relatively higher elevation. Under the assumption that neighboring land-use correlation is not systematically biased towards higher or lower elevation neighboring geographies, this design should control for omitted variables bias (OVB) due to spatial correlation in land use.

To my knowledge, these first two results are novel in their space and scope in the economics literature. There exists a broad scientific and engineering literature on the mechanism of how this changes risk. The physical mechanism of different land covers matters for surface runoff—for example, retention ponds, wetlands, and developed permeable surfaces have saturation points at which they behave like an impervious surface during tail rainfall events (Konrad, 2003). The United States Geological Survey (USGS) makes urban management recommendations that specifically try to account for these mechanisms (Cappiella et al., 2012). I remain relatively agnostic on these underlying mechanisms. Instead, I quantify this risk in the context of flood insurance claims, and show how these impacts are realized in an imperfect market setting.

In this analysis, I use the universe of flood insurance claims from the National Flood Insurance Program (NFIP) for the extent of the study. These claims data are widely used in the economics of floods literature, but I focus on a more novel scale. Many existing studies are focused on high-signal flooding events, either temporally (e.g., a hurricane or other superstorm) or spatially (e.g., geographies that are repeatedly subject to floods over time). My analysis incorporates these while also accounting for diffuse spatial and temporal flooding events, which have received less attention than the scope of their contribution to claims would suggest. For the entire universe of NFIP claims, 31 percent of a total of 2.4 million claims reported a date of loss outside of the Atlantic hurricane season (June 1 through November 30). The payouts for these claims account for $15 billion of the total $91 billion (2019 USD) in claims expenditures. This is a conservative account for attribution of non-hurricane flood insurance payouts, suggesting that even at a minimum, flood insurance claims during the non-hurricane season account for a significant portion of all claims. In Texas alone from 2011–2016, claims representing non-hurricane season losses represent 76 percent of total claims, and $1.3 billion of a total of $1.5 billion in expenditures. Flood events resulting in NFIP claims occurred in 162 of 254 of Texas counties, and 2376 of 5265 census tracts during this same period.

The findings and scope of this paper point to the importance of organized mitigation strategies for a changing climate and contribute to the literature on the impacts of urban planning and regulation in the face of natural disaster impacts. Coastal wetlands have been identified to mitigate against property damage from hurricanes along the US Gulf Coast (Sun and Carson, 2020). A large fraction of wetlands are non-coastal, and I account for the importance of these during extreme events. For context, In 2016 Texas, the average census tract in a coastal county has 4.1 percent wetlands cover, and the average census tract in a non-coastal county has 2.4 percent wetlands cover. I show that changes in wetland, developed, and water coverage at the census-tract level impact observable impacts from floods.

I contribute to the literature on climate impacts and spillovers. Baylis and Boomhower (2021) provide evidence that building codes to mitigate wildfire damage generate positive spillovers in protection for unregulated structures. Similarly, there is current work that investigates spillovers in flood mitigation from historical building codes.[1] I show that differing patterns of land use also generate externalities in flood risk. The monetary impact of these spillovers has implications for urban planning, and this quantification could feasibly incorporate into the assessment of impact fees for the finance of public mitigation infrastructure Brueckner (1997). More largely, these results serve as suggestive of potential gains in flood mitigation to coordination in land management.[2]

The rest of the paper proceeds as follows: Section 2.2 provides background on the NFIP and situates this study within the existing floods literature. Section 2.3 gives a description of the data. Section 2.4 describes the empirical strategy. Section 2.5 discusses the results and their implications. Section 2.6 presents the results of a series of robustness checks on the empirical models, and Section 2.7 concludes.

2.2 Background

Flooding and flood insurance are active topics in the economics literature. I use NFIP claims as a measure of flood risk in the presence of changing land use. The NFIP is a federally-run program established in 1968 with the goal of setting actuarially fair residential flood insurance policies (Federal Emergency Management Agency, 2002). Managed by the Federal Emergency Management Agency (FEMA), the NFIP creates and issues flood risk maps, sets premiums, and underwrites flood insurance policies for participating communities. Residential coverage tops out at \$250,000 of coverage for structures and \$100,000 in contents. While original intent was to provide actuarially fair insurance, most policies in force are priced below the expected payout (Wagner, 2020).

The NFIP is an imperfect measure of flood impacts. If the NFIP were implemented

[1] This follows from discussions with Laura Bakkensen regarding unreleased work on flood mitigation spillovers in Florida following hurricanes.

[2] In 2019, the Texas state government approved a new floodplain management program, which develops flood planning regions based on river basins through the Texas Water Development Board TX (2020).

in a competitive market without frictions, using NFIP claims in the analysis would lead to a lower-bound estimate on impacts, since we do not observe damages that do not show up in the claims data or occur to structures outside of the NFIP coverage. However, existing papers demonstrate that flood insurance, and specifically the NFIP, is far from a perfect insurance market. I summarize these and relate them to the context and findings of this paper, and I discuss how these factor into the identification and interpretation of my results.

One observation from the existing literature on the NFIP is that flood insurance take-up varies widely with observable risk. Wagner (2020) shows that households might substitute observable mitigation (raised homes) for flood insurance, leading to adverse selection from low-lying houses in the insurance pool. Gallagher (2014) shows that flood insurance take-up spikes after significant flooding events (with federal emergency declarations), and slowly declines to baseline take-up in the years following. Bakkensen, Ding and Ma (2019) replicate this salience event in Florida following hurricanes. Qualitatively, my data validate this general trend, showing a secular decline in policies in force during the main study, 2010–2016. The main sample was preceded by Hurricane Ike in 2008 and followed by Hurricane Harvey in 2017, where we may expect to see another spike in NFIP take-up.

Another observation is that flood insurance take-up responds to information provision about unobservable risk. Provision of both flood and wildfire risk maps affects equilibrium housing prices, even if the fundamental underlying risk is unaffected (Gibson and Mullins, 2020; Garnache and Guilfoos, 2019). Other risk signals may be difficult or costly to observe. Keenan and Bradt (2020) show evidence of information asymmetries in risk assessment for financial institutions, indicating that local lenders may have lower frictions for obtaining risk assessment specifically for local properties. In the case of land use, if aggregate changes in risk driven by land use are not easily observable, or information frictions make it difficult to know the risk contribution of land use, then static premiums through the NFIP are further evidence of distortions in the flood insurance market. I show that this is the case.

Though these variations in take-up could be attributed to optimizing agents (through rational inattention, substitution towards other mitigation measures, etc.), this is unlikely to be the case. If households are generally risk averse, take-up of actuarially fair policies should increase expected utility, and even more so if premiums are below the actuarially fair price. As a potential explanation, Wagner (2020) raises the issue of adverse selection in flood insurance markets, and suggests welfare gains from insurance mandates. Other interventions to address aggregate trends towards "under-insurance" have been explored in the public finance literature (Kunreuther, 1996; Kriesel and Landry, 2004; Kunreuther and Michel-Kerjan, 2009). An alternative explanation for this is heterogeneous sorting of households. In a broad-based field survey, Bakkensen and Barrage (2017) show that Rhode Island households that owned coastal properties were systematically likely to underestimate flood risk. Barrage and Furst (2019) show further that new construction was more likely to occur in flood-risk zones amongst more climate- and flood-skeptic communities. In the context of this paper, these do not undermine necessary assumptions for identification, but may have implications for the interpretation of the main results. I address this in Section 2.5.

Finally, these results have potential implications for existing community-level mitigation programs within the context of NFIP. The Community Rating System (CRS) provides community-level discounts on flood insurance offered through NFIP for participants whose floodplain management includes auxiliary flood mitigation measures. These include community-level credits that fall into one of four categories including management for flood damage reduction.[3] Communities participating in the program are scored and given discounts on premiums. Frimpong et al. (2020) study the impacts of this program,

[3]The general categories include: public information provision, mapping and regulation, flood damage reduction, and flood preparedness.

and find CRS participation increases uptake and reduce damage claims for communities that score in the best (lowest) class.

To summarize, I uncover additional frictions in the market for flood insurance attributable to unobservable risk impacts from changes in land use. My analysis provides a structure to quantify this risk and suggests implications for optimal management policy.

2.3 Data

I construct a monthly panel of flood insurance claims and precipitation patterns, and connect this to longer-run land use characteristics for each Texas census tract. I match a subset of these with available data on flood insurance policies in force, housing summaries from the American Community Survey (ACS). The baseline analysis is restricted to 2010–2016. I also report the results of a broader analyses with more limited data from 2001–2016. The summary statistics for these data are reported in Table 2.1.

2.3.1 National flood insurance claims and the universe of policies

The NFIP is a federal program established in 1968 with the goal of setting actuarially fair flood insurance policies.[4] I use the universe of national flood insurance claims for the state of Texas from 2001–2016. Individual claims report the day of the flooding event identified spatially at the census-tract level.[5] I aggregate the number and value of claims (in 2019 USD) to create a panel of month-by-census tract claims data.

I also use the FEMA NFIP universe of policy data, collected through a Freedom of Information Act (FOIA) request. The set of unredacted reports are available from 2009 to the present at the census-tract level.[6] I define a policy as active during a particular month if the policy was opened prior to the 15th day of the month, and ended (or remained open)

[4]See Federal Emergency Management Agency (2002)
[5]These are reported using the 2010 set of census tracts.
[6]See Section B.4 for FOIA information on previous policy datasets.

Table 2.1. Summary statistics

	Mean	SD	Min	Max
Δ *Land use 2011–2016*:				
Water cover	0.001	0.007	-0.153	0.110
Impervious cover	0.007	0.013	-0.002	0.209
Wetland cover	-0.001	0.005	-0.140	0.085
Rainfall days/month:				
0–1 in.	7.13	4.14	0	24
1–2 in.	0.66	0.91	0	8
2–3 in.	0.17	0.42	0	5
3+ in.	0.12	0.36	0	3
NFIP Claims by tract per month:				
Claims	0.09	2.25	0	436
$ 2019 USD	5069	255430	0	93 million
Observations	287793	287793	287793	287793

Note: summary statistics reported for the baseline specification. Days/month and NFIP claims cover the universe of tract-months in Texas for 2010, 2011, 2012, 2015, and 2016. Land use reported as percentage cover of census tract-coverage.

after that 15th day. I then aggregate the total number of active policies in force (and the total cost of the premiums) at the tract level to merge with the claims data. I use the 2019 Consumer Price Index (CPI) to adjust individual premiums and claims to their real value in 2019 USD.

2.3.2 Precipitation data

I use precipitation data reported by the Parameter elevation Regression on Independent Slopes Model (PRISM) (PRISM Climate Group, 2021). PRISM reports a gridded dataset of daily precipitation in the continental US at a 4 kilometer (km) resolution. I collect PRISM daily precipitation from 2001 through 2016.

To merge this with flood insurance claims data, I match each census tract centroid to the nearest PRISM pixel. The median distance of this match is 1.7 km when matching Texas census tracts, with a maximum distance of 2.8 km.

Following Deschênes and Greenstone (2011), I construct bins for daily precipitation data. $\mathbf{R}_{imt}$ is a vector of precipitation bins, where each element of $\mathbf{R}_{imt}$ represents the number of days in month-year mt that fall in a particular bin in census tract i. In the baseline specification, this includes one-inch bins from zero to over three inches of daily rainfall.[7] This definition of daily rainfall allows for a non-parametric precipitation-flood relationship, with the implicit assumption that days of precipitation within a bin will the same effect on flooding.This also assumes that the effect of a marginal day of precipitation is the same regardless of when this occurs within a given year.

In addition to the baseline definition of this binned precipitation variable, I consider alternative specifications of this tract-by-year variable by constructing a measure of repeated daily rainfall events. In this case, $\mathbf{R}_{imt}$ measures the number of rain "spells" that occur within a census tract-year. I define a "spell" as a threshold number of consecutive days that a census tract experiences some minimum amount of precipitation within a day (2 inches).

2.3.3 Geographic data on land use and elevation

The USGS National Land Cover Dataset (NLCD) classifies land and impervious surface coverage for the continental United States at a resolution of 30 meters (Dewitz, 2019). These land cells are broadly classified into developed land, water, and various green covers. Each of these categories is further refined to more specific covers.[8] These coverage data are reported for seven periods from 2001 to 2016. I match these data to census tracts by taking the mean coverage of each class of pixels that falls within a tract for each survey year. I focus specifically on impervious developed coverage, water coverage, and wetlands coverage. I use the NLCD measure of urban imperviousness to calculate the average level of impervious surface coverage at the tract level, and omit permeable surface development

[7]That is, (0,1] inches, (1,2] inches, (2,3] inches, and 3+ inches. 0 inches is the omitted bin.

[8]These include open water, perennial ice/snow, developed open space, developed low intensity, developed medium intensity, developed high intensity, barren land (rock/sand/clay), deciduous forest, evergreen forest, mixed forest, shrubland, herbaceous land, planted and cultivated land, woody wetlands, and emergent herbaceous wetlands.

from the baseline model. Impervious surface coverage refers to urban development that restricts water from penetrating the ground.[9]

Finally, I calculate mean elevation for each census tract using the USGS National Elevation Dataset (NED) by taking the mean elevation of pixels that fall within a particular census tract. I use this tract-level elevation variation to estimate spillovers from changes in land use (USGS EROS, 1999).

2.3.4 Housing summaries

I use the decennial census and ACS 5-year estimates to construct housing variables at the census-tract level. These include full counts of housing units at the tract level for each decennial census (2000 and 2010) and estimates for each year reported by the ACS (2009–2016). In addition, I use the median home value at the tract-year level and use the 2019 CPI to adjust these values to real 2019 USD. Housing variables through the ACS are reported for a large majority of tracts.[10]

2.4 Empirical Strategy

In order to estimate the impact of local land use change on flood insurance claims, I exploit several sources of variation. First, I consider exogenous variation in tract-level precipitation in the cross-section and time series.

Changes in land use vary over time and space, but may be endogenous to claims through the NFIP. For instance, if local amenities change, this may also affect the selection into a neighborhood, changing the composition of individuals that live in a particular area. This may also affect equilibrium home prices and the choice for insurance uptake. In order to mitigate this potential OVB, I saturate the model with tract-by-year housing

[9]This is reported for a subset of these periods: land coverage data are available for 2001, 2003, 2006, 2008, 2011, 2013, and 2016. Impervious surface coverage is available for 2001, 2006, 2011, and 2016.

[10]164 of 5265 tracts have median home values reported in the ACS. 1790 tracts do not have home values reported in the 2000 decennial census for purposes of anonymity.

characteristics as well as changes in the number and value of premiums associated with flood insurance policies in force. I show that the main estimates of interest are economically robust to these inclusions. Finally, in this baseline specification, I show that trends in flood insurance take-up rates are similar across three large categorizations of land use: positive changes in a variable of interest, negative changes, and no change. I report these trends in Figure B.1. If we observed differential trends in policy take-up based on land change, we might be concerned about differential rates of overall realized claims. Fortunately, this is not the case, and the estimates of interest are economically invariant to including these tract-level data on policies in force.

Next, I consider whether effects of land change are symmetric. That is, are increases in a particular type of coverage economically similar to decreases in that coverage? I provide suggestive evidence that effects of increases or decreases in developed, wetland, or water coverage may not be symmetric in this setting. One caveat to this is that the share of tracts with decreasing shares of impervious development or increases in wetland coverage are small, so these effects are identified off of a very small subset of census tracts.

Finally, I present a specification that considers land-use change from neighboring census tracts. I use variation in neighboring tract elevation, assuming that this is uncorrelated with omitted variables that may affect flood insurance claims in a dynamic setting. In this specification, I am able to show evidence of larger spillovers in flood claims from changes in land use in neighboring higher-elevation tracts.

2.4.1 Baseline specification

In the baseline specification, the identifying variation is within-tract land changes and exogenous tract-level anomalous rainfall. The empirical specification is as follows:

$$Y_{imt} = \alpha_i + \gamma_t + \mathbf{R}_{imt}\,\beta_1 + \mathbf{L}_{it}\beta_2 + \mathbf{R}_{imt} \times \mathbf{L}_{it}\,\beta_3 + X_{it}\,\beta_4 + \varepsilon_{imt}, \tag{2.1}$$

where Y_{imt} represents monthly flood insurance claims per policy for census tract i in month-year mt. α_i is a census tract-level fixed effect, which absorbs variation across tract-level characteristics that are constant over time (such invariant tract elevation). γ_t is a year fixed effect. Robust standard errors are calculated at the county-by-year level in order to account for correlation within a county in a specific year. β_3 is the main coefficient of interest, and can be interpreted as the reduced form effect of within-tract land-use change on flood damages, conditional on receiving an exogenous rainfall shock.

Equation 2.1 follows a standard DD specification, which includes additive controls for land use L_{it}. Two-way fixed effects in census tract and year allow us to identify changes in flood insurance claims based on within-tract changes in land use. The inclusion of tract fixed effects mean that we can interpret monthly rainfall as random deviations from average rainfall. Under the assumption that this anomalous tract-by-month rainfall is random, we may choose to relax this specification: in an alternative specification, I impose the structural assumption that land use, absent of rainfall, does not affect flood risk—I identically impose that $\beta_2 = 0$ in Equation 2.1.[11] The results of this specification are reported in Figure 2.1, with detailed results reported in Table B.1.

2.4.2 Asymmetry of land change

Second, I consider the possibility of asymmetric effects of changes in land use. Specifically, what is the effect of a positive change in a particular type of coverage relative to a negative change? To do this, I construct two new variables for each component of $\mathbf{L}$. Following Allison (2019), for each element L_{it} in $\mathbf{L}_{it}$, first I define:

$$z_{it}^{+} = L_{it} - L_{it-1} \text{ if } (L_{it} - L_{it-1}) > 0, \qquad \text{else } L_{it}^{+} = 0,$$

$$z_{it}^{-} = -(L_{it} - L_{it-1}) \text{ if } (L_{it} - L_{it-1}) < 0, \qquad \text{else } L_{it}^{-} = 0.$$

[11]While this seems reasonable, land use could be correlated with non-rainfall related floods. For example, river basin flooding caused by upstream rain or snow melt may be correlated with historical development, in which case this structural assumption would be overly restrictive.

This is a measure of year-over-year changes in land use. Since we are specifically interested in *changes* in the baseline specification, and tract fixed effects absorb mean levels of cover at the tract level, both z_{it}^{+} and z_{it}^{-} are set to zero for the first period. Then, I define:

$$L_{it}^{+} = \sum_{s=1}^{t} z_{it}^{+}, \text{ and}$$
$$L_{it}^{-} = \sum_{s=1}^{t} z_{it}^{-}.$$

Here, L_{it}^{+} is interpreted as all positive changes in a particular land cover up to period t, and L_{it}^{-} is interpreted as all negative changes in a particular land cover up to period t. $\mathbf{L}_{it}^{+}$ and $\mathbf{L}_{it}^{-}$ are defined as the vectors of cumulative positive changes and cumulative negative changes, respectively. This allows me to estimate a model similar to Equation 2.1 that allows for the estimation of asymmetric effects of positive and negative changes in land cover:

$$\begin{aligned} Y_{imt} = \alpha_i + \gamma_t + \mathbf{R}_{imt}\,\beta_1 + \mathbf{L}_{it}^{+}\beta_2^{+} + \mathbf{L}_{it}^{-}\,\beta_2^{-} + \mathbf{R}_{imt} \times \mathbf{L}_{it}^{+}\,\beta_3^{+} + \mathbf{R}_{imt} \times \mathbf{L}_{it}^{-}\,\beta_3^{-} \\ + X_{it}\,\beta_4 + \varepsilon_{imt}. \end{aligned} \tag{2.2}$$

In the case that effects of positive changes are symmetric to negative changes, or $\beta_3^{+} = -\beta_3^{-}$, then the model specified by Equation 2.2 is econometrically equivalent to the model specified by Equation 2.1 (Allison, 2019). The results of this specification are reported in Figure 2.2, with detailed results reported in Table B.2.

2.4.3 Spillovers

From the baseline model, I introduce information on land use for neighboring census tracts in order to explore the existence of spillovers in flood risk. Spatial correlation in land-use patterns and factors affecting flood insurance premiums would imply that simply including information on neighboring land use patterns in the baseline specification may lead to bias in estimates of the importance of neighboring land-use change. In order to address this, I introduce variation in census-tract level elevation to solve two problems. First, under the *ex ante* assumption that spillovers exist, we might expect neighboring land change to matter more for higher elevation than for lower elevation neighbors, conditional on the assumption that rainfall runoff largely runs from higher- to lower-mean elevations. I separate changes in neighboring land use into uphill tracts (neighboring tracts that are on average higher in elevation) and downhill tracts (neighboring tracts that are on average lower in elevation).

Second, variables omitted that are correlated with both changes in land use and flood insurance claims may introduce bias in the estimates of interest. In the baseline specification, I include policies in force, home values and numeracy, and demographic information to minimize OVB. For example, changes in developed or wetland cover may change home values through local amenity changes, which may co-vary with factors that change flood insurance adoption rates. If elevation is not systematically correlated with these factors that introduce OVB in the empirical specification, then this specification with differences in spillovers across tract-level elevation will minimize these potential OVB biases.

For each census tract i, I identify the set of all neighboring tracts whose centroid falls within some radius r of the tract centroid (in the baseline specification, $r = 2.5$km). Then, I separate this set of neighboring census tracts into the set of uphill (U) and downhill (D) tracts, where the elevation for each tract is defined as the mean elevation. I construct the land use characteristics for these uphill and downhill neighboring tracts by taking the size-weighted vector of land characteristics, $\mathbf{L}_{it}^{U}$ and $\mathbf{L}_{it}^{D}$ respectively.

In order to test this spillover hypothesis, the specification is as follows:

$$\begin{aligned} Y_{imt} = \alpha_i + \gamma_t + \mathbf{R}_{imt}\,\beta_1 + \mathbf{L}_{it}\beta_2 + \mathbf{R}_{imt} \times \mathbf{L}_{it}\,\beta_3 + \mathbf{L}^U_{it}\,\beta_4 + \mathbf{L}^D_{it}\,\beta_5 \\ + \mathbf{R}_{imt} \times \mathbf{L}^U_{it}\,\beta_6 + \mathbf{R}_{imt} \times \mathbf{L}^D_{it}\,\beta_7 + X_{it}\,\beta_8 + \varepsilon_{imt}. \end{aligned} \tag{2.3}$$

Here, β_6 is interpreted as the effect of mean changes in land use for uphill neighbors on own flood damages controlling for own tract changes, conditional on receiving an exogenous rainfall shock. β_7 is interpreted similarly but for land-use changes in downhill neighboring tracts. The results of this specification are reported in Figure 2.3, with detailed results reported in Table B.3.

2.5 Results and Discussion

In this section, I discuss the results of each of the main specifications in Section 2.4. In each of these specifications, the estimated coefficients from land use interacted with anomalous rainfall have a narrow interpretation: each should be interpreted as the reduced-form effect of all contributions of land change on changes in flood insurance claims. Particularly, this should be interpreted as the *mean effect* of land change on flood claims. That is, these specifications take the mean impact of land changes in census tracts that, on one end, may combine with mitigation that is unobservable through aggregate land change, and, at the other end, tracts that drastically change the fundamentals of flood risk with observable land change. Additionally, I estimate each of the land-coverage effects additively. While I capture the overall effects of observed changes in land use, this cannot account for interaction effects between different categories of land cover.

2.5.1 Main empirical results

Figure 2.1 shows the results of the baseline specification. Subfigure (a) shows that observed increases in impervious surface coverage increase the numeracy of flood claims for large rain events. Specifically, when a census tract experiences positive changes in impervious development, an extra day with greater than 2 inches of rain increases the expected number of flood insurance claims through the NFIP. Similarly, subfigure (b) shows that an increase in wetlands coverage, when interacted with an extra day of 2–3 inches of rainfall, is associated with a decrease in flood insurance claims. On the far right tail, with an increase in wetlands cover over time, an extra day with 3+ inches of rain appears also to decrease the propensity of flood insurance claims. However, this is estimated imprecisely, and may be evidence of saturation that diminishes mitigation effects. Subfigure (c) demonstrates similar trends in water coverage.

Figure 2.2 shows the results of the baseline specification, allowing for asymmetric effects of land change on flood risk. In line with the trends from the baseline specification, we can observe that, conditional on a large rainfall shock, positive (negative) changes in impervious developed coverage are associated with higher rates of flood insurance claims, and positive (negative) changes in wetlands and water coverage are associated with lower rates of flood insurance claims. Importantly, subfigure (b) implies that there are mitigation benefits to both protecting and restoring existing wetlands.

Finally, Figure 2.3 shows the results of the spillover specification. In each subfigure, we can observe that both the magnitude and significance of higher-elevation neighboring census tracts are larger than the corresponding estimates for lower-elevation neighboring census tracts. I take this as strong evidence that spillovers from land change exist.[12]

2.5.2 Implications for flood insurance premiums

In light of the evidence that changes in land use change the fundamentals of flood risk, it is important to note that a vast majority of flood insurance premiums ($>$ 70

[12]Note that any marginal statistical significance and sign for downhill neighbors does not contradict the main hypothesis that spillovers come largely from uphill census tracts. It is feasible (and likely, for many neighbor pairs) for mean elevation to have one ordinal relationship, while particular locations across census tracts to have the opposite ordinal relationship.

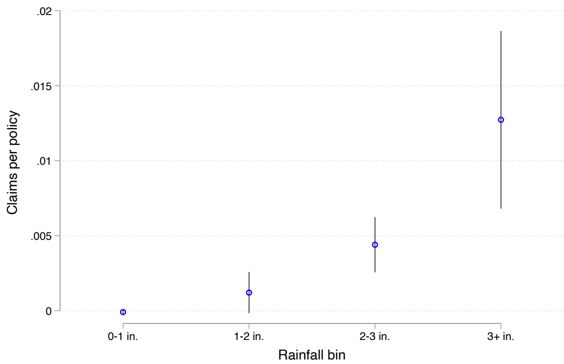

(a) Δ Development

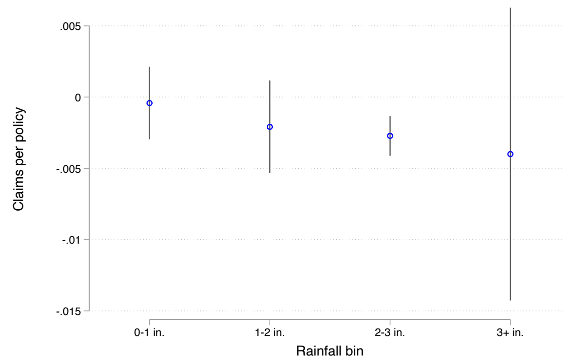

(b) Δ Wetlands

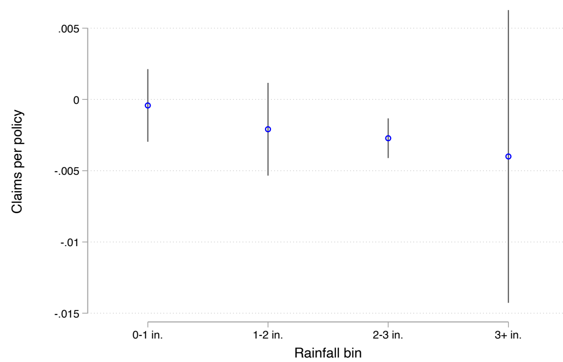

(c) Δ Water

Figure 2.1. Baseline results

Note: this shows the results of the main coefficients of interest from the preferred baseline model presented in Equation 2.1. "Positive" refers to positive changes in the land cover variable of interest, and "negative" refers to negative changes, respectively. This includes all Texas census tracts-by-month from 2010, 2011, 2012, 2015, and 2016.

percent) were static during the course of the baseline study, 2010–2016. That is, under changing fundamentals, premiums did not change to reflect either increases or decreases in flood risk fundamentals. This is especially important given the evidence of spillovers in observed flood claims.

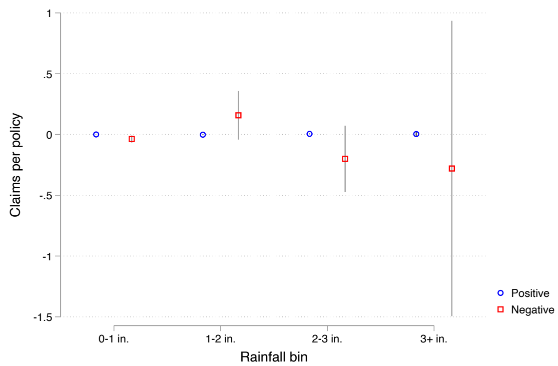

(a) Δ Development

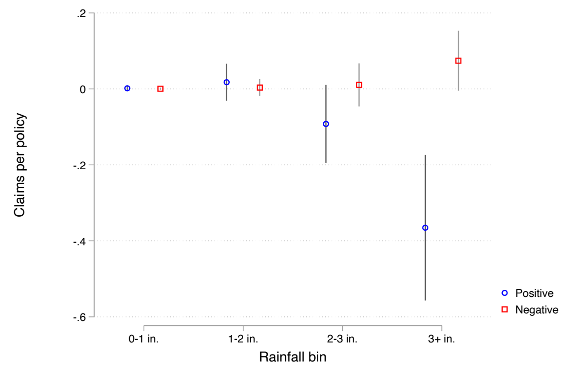

(b) Δ Wetlands

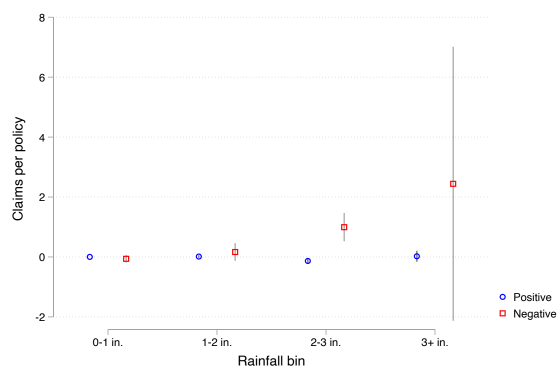

(c) Δ Water

Figure 2.2. Asymmetry results

Note: this shows the results of the main coefficients of interest from the preferred asymmetry model presented in Equation 2.2. This includes all Texas census tracts-by-month from 2010, 2011, 2012, 2015, and 2016.

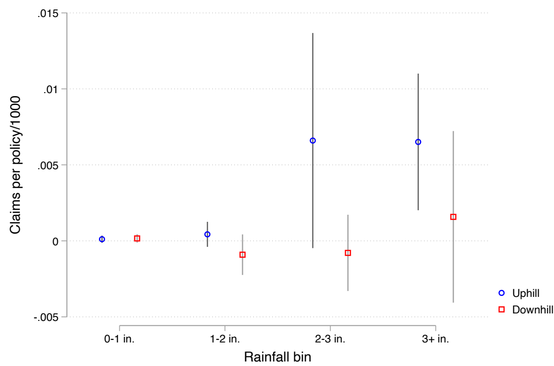

(a) Δ Development

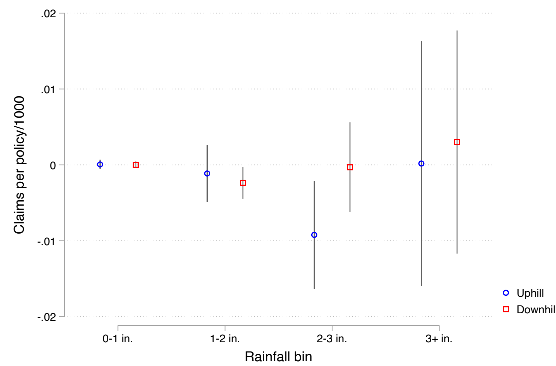

(b) Δ Wetlands

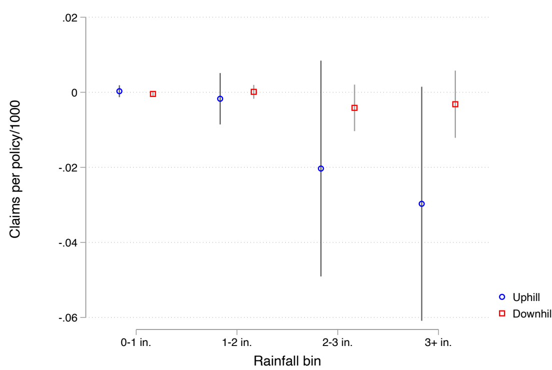

(c) Δ Water

Figure 2.3. Spillovers in flood risk

Note: this shows the results of the main coefficients of interest from the preferred spillover model presented in Equation 2.3. This includes all Texas census tracts-by-month from 2010, 2011, 2012, 2015, and 2016.

2.6 Robustness

In this section, I discuss the robustness of the main results to different methodological choices, and report the results in Section B.2.

First, I report the results from the baseline specification using the spell specification for the rainfall interaction. That is, $\mathbf{R}_{imt}$ is the number of consecutive days within a tract-month with recorded precipitation of two or more inches of rain. The results of this specification are statistically and economically similar to the baseline specification, and are reported in Table B.4. Table B.5 shows the results of the baseline specification, but uses total value of claims paid out at the census-tract-month level. The qualitative results are similar to the baseline specification.

In Table B.6, I report the results of the baseline specification when restricting to non-coastal counties in the state of Texas. Both of these specifications suggest similar impacts of developed, water, and wetland coverage. Importantly, restoration of inland wetland coverage seems to mitigate flood insurance claims in the event of a rainfall shock, and decreases in inland wetland coverage are associated with increases in flood insurance claims.

In Table B.7, I report the results of the spillover specification from Equation 2.3 when changing the neighbor radius from 2.5 km to 5 km. The results are qualitatively and economically similar to the preferred spillover specification. In Table B.8, I increase the size of this radius to 15 km. As expected, estimates are attenuated in magnitude and significance. This tracks with the assumption that as more distant census tracts are included in the spillover specification, we should expect less relationship with "neighboring" land use and own-flood risk.

2.7 Conclusion

I show that aggregate changes in land use can explain significant outlays in claims and public expenditures on national flood insurance, even in inland, non-coastal counties in Texas. This elucidates the importance of aggregation of small land-use changes in urban areas in often under-looked settings. These effects of land change on flood risk fundamentals not only affect very local measures of impacts, but exhibit spillovers on neighboring geographies.

Land management is important for local jurisdictions, but lack of coordination implies large externalities to neighboring geographies. The results of this empirical analysis imply the returns to coordination in land management and present avenues for future research about specific mechanisms and optimal planning.

2.8 Acknowledgements

Chapter 2 is currently being prepared for submission for publication. The dissertation author is the sole author on this chapter.

Chapter 3

Risk Preference Adaptation to Climate Change

3.1 Introduction

Anthropogenic climate change is one of the most serious threats to the collective well-being of humanity in the 21^{st} century. The extent of economic damage it is likely to inflict depends crucially on individuals' ability to adapt to a rapidly changing physical environment. Studies of climate adaptation have documented specific behavioral adjustments in response to climate change, such as the adoption of air conditioning (Barreca, Deschênes and Guldi, 2015; Barreca et al., 2016; Howden, 2021), shifts in time allocation to labor (Graff Zivin and Neidell, 2014), migration (Hauer et al., 2020), and changes to agricultural decision-making (Burke and Emerick, 2016; Kala, 2017). In this paper we study a novel margin of climate adaptation that is domain-general and psychological. Our core hypothesis, which we support theoretically and empirically, is that individual risk preferences may change in response to long-run experiences of climate change, and that such changes can be welfare-improving, and therefore adaptive.

We begin by building a model of risk preference adaptation to climate change (Section 3.2). In our model, an expected utility maximizer is faced each period with a choice from a fixed menu of objective income lotteries known as the *foreground* risk. The agent makes this choice in the presence of an exogenous, unavoidable, and statistically-independent *background* income risk. We assume that the agent's direct utility function, which is defined over the sum of both risks, is risk-vulnerable (Gollier and Pratt, 1996). This means that the two sources of risk are substitutes for the agent—the more risk she believes exists in the environment, the less risk she is willing to take in her individual choices.

We enrich this static background risk framework with a dynamic model of high-dimensional learning over the background risk. We assume the agent is Bayesian, and that she perceives the background risk to be a stationary Gaussian random variable. To capture the deep structural uncertainty inherent to climate change (Weitzman, 2009), we further assume that *both* the mean and the variance of the background risk are unknown to the agent, and are therefore objects of learning. Our agent observes realizations of the background risk over time and updates her beliefs about its moments. As her beliefs about the background risk evolve her risk preferences, captured by the curvature of the indirect utility function over the foreground risk, adapt in turn.

Our model delivers sharp predictions about the effects of new realizations of the background risk on the agent's foreground risk preferences given her existing body of experiences. Our main result is that the effects of the two moments are additive, with the agent's risk aversion decreasing in her posterior mean and increasing in her posterior variance. Intuitively, this means that, unlike in models where risk aversion is monotonic in the shock (Dillenberger and Rozen, 2015), in our model shocks differ in their effects not only by whether they are positive or negative, but also by whether they are large or small.

We further show that our model is equivalent to one where an agent is learning about the mean of a fat-tailed background risk. This extends our results to a class of random processes that have recently gained renewed prominence in the literature on aggregate shocks (Acemoglu, Ozdaglar and Tahbaz-Salehi, 2017; Pomatto, Strack and Tamuz, 2020). Finally, we prove that under constant relative risk aversion (CRRA) utility and complete information, the coefficient governing the relative effects of the mean and variance of background risk on the agent's absolute risk aversion is exactly the agent's prior coefficient of absolute prudence. This provides a link between the agent's higher-order, static risk preferences and her lower-order, dynamic risk preferences.

We next turn our attention to testing the predictions of our model empirically. To do so, we use data from two large longitudinal surveys, one from Indonesia and one from Mexico, each containing two elicited measures of risk aversion for the same individuals years apart. These elicited measures are estimated from choices over objective lotteries, which helps us to sidestep identification challenges that exist in our setting with other kinds of choice data, most notably the potential confounding of estimated preference changes with foreground belief changes. In our main analysis we regress within-person changes in measured risk aversion on changes in the mean and variance of heat and precipitation in subjects' state of birth, from birth to measurement. Our empirical approach allows us to exploit the significant variation that exists within each country in climatic conditions, while providing us with a degree of external validity for the results, given the significant differences between the two countries in most physical, cultural, and socioeconomic dimensions.

In line with the model's predictions, we find (Subsection 3.4.1) that in both countries increases in the experienced lifetime mean of both heat and precipitation induce significant decreases in measured risk aversion. We also find that increases in the experienced variance of heat in Indonesia and the variance of precipitation in Mexico lead to significant increases in measured risk aversion. The estimated magnitudes of the variance effects are approximately 1.6 (Indonesia) and 0.7 (Mexico) times the magnitude of the mean effects, indicating that experienced climatic variance is first-order in its effects on risk aversion in both settings.

In Subsection 3.4.2, we show that these results are robust to controlling for changes in household demographics and economic constraints, suggesting that the observed estimates are not driven by income effects. We also show that our results are robust to the inclusion of other categories of lifetime experiences that have been shown in the literature to affect risk preferences, such as experiences of violence (Callen et al., 2014; Jakiela and Ozier, 2019; Brown et al., 2019), natural disasters (Cameron and Shah, 2015; Brown et al., 2018; Hanaoka, Shigeoka and Watanabe, 2018), and macroeconomic conditions (Malmendier and Nagel, 2011; Levin and Vidart, 2020).

To examine whether the observed effects represent domain-general shifts in risk attitudes, in Subsection 3.4.3 we estimate correlations between predicted changes in measured risk aversion and observed changes in risk-taking behavior in four domains: migration, smoking, self-employment status, and (in Indonesia) the planting of cash crops. Our results provide suggestive evidence for domain-generalizability. We find strong correlations between increases in predicted risk aversion and decreases in migration and smoking behavior in Indonesia; weak correlations between these variables in Mexico; and no discernible statistical relationship between predicted risk aversion and self-employment or the planting of cash crops in either country. We conclude the main empirical analysis in Subsection 3.4.4 by examining the robustness of our main results to a variety of alternative empirical specifications.

In the final part of the paper we explore whether the climate-change-induced risk preference changes we estimate are, in fact, adaptive. As a rule, the climate adaptation literature generally regards behavioral changes which are causally driven by climate change to be adaptations, under the assumption that they represent re-optimization behavior, and are therefore welfare increasing. A stronger test for climate adaptation would, in theory, estimate the effects of behavioral changes on welfare directly. Welfare analyses are, however, particularly difficult when preferences are not fixed, or the model of choice departs from the neoclassical benchmark (Bernheim and Taubinsky, 2018). Even under expected utility with stable preferences, welfare analyses have generally required an assumption of preference homogeneity, which is unworkable in our setting.

However, in a recent paper Eden (2020) shows that a single welfare measure can be generated for a population with heterogeneous risk preferences under expected utility. Conceptually, this is accomplished by presenting each agent with the population distribution of income and calculating their certainty equivalent. The distribution of certainty equivalents is then presented to each agent to yield another certainty equivalent. Eden (2020) proves that this iterative process converges to a single fixed point, the Equally Distributed Equivalent (EDE), which represents the collective value of an income distribution for a population with heterogeneous risk preferences.

Building on this result, we conduct two empirical welfare exercises. First, we calculate the EDE for the second period consumption distribution under both the first period and second period estimated risk preference distribution in both countries.[1] Here, the difference in the EDE can be thought of as a measure of the sum total of adaptation due to risk preference changes. Second, we calculate the EDE for the second period consumption distribution under the empirical risk preference distribution predicted by our main temperature regression in Indonesia and precipitation regression in Mexico, as well as under the risk preference distribution predicted from these regressions when excluding the climate variables. A comparison of the EDE in both these cases yields a measure the welfare effects of climate-change-induced preference changes, relative to the theoretical counterfactual where climate change had never occurred.

The results of these analyses are presented in Section 3.5. For the first exercise, we find that in Indonesia total risk preference changes result in a 6% increase in welfare, while in Mexico total risk preference changes result in an 8% decrease in estimated welfare. We interpret these findings to indicate that risk preference changes attributable to all causes are adaptive in Indonesia and maladaptive in Mexico. For the second exercise, we estimate that climate-change-induced preferences changes account for a 1% increase in welfare in Indonesia, and a 0.8% increase in welfare in Mexico. Therefore, even though overall

[1]In both exercises we structurally estimate individual risk preferences under the assumption of EU with CRRA utility.

preference changes vary in their effects on welfare across settings, preference changes in response to climate represent (group-level) climate adaptation in both countries.

3.2 The Hunter-Gatherer Model of Risk Preference Adaptation

The core hypothesis of this paper is that individuals' risk preferences adapt to physical changes in risk in their environment. How should we model such a process of *preference adaptation*? A key challenge in answering this question is that a model of this kind entails some departure from dominant paradigms in economic theory. In contrast to the neoclassical paradigm, we are interested in what drives changes in economic preferences, an area of inquiry seldom explored since Becker and Stigler (1977) effectively banished it from standard economic analysis. In contrast to the behavioral paradigm, we are interested in modeling a psychological phenomenon that moves agents closer to optimal decision-making, rather than a bias or heuristic that drives them further away.

Our basic insight is that preference adaptation occurs where physical adaptation ends. If one can adapt to the environment by adopting different physical tools, surely there is no need to undergo fundamental psychological change, which can be relatively difficult and costly to implement.[2] It is only when such physical strategies are absent or prohibitively costly that the agent falls back on adapting their own mental machinery. Necessity is the mother of self-invention.

Thought of through this lens, preference adaptation is a process intrinsically driven by technological constraints. It is instructive, therefore, to consider how such a process might unfold in an environment where technological solutions to problems simply do not exist. Such desperate circumstances have happily become increasingly rare since the advent of the industrial revolution, and in their purest form currently afflict only the poorest of

[2]One need only think of the exorbitant rates that good therapists can command to see evidence of this fact.

the global poor. Conditions of this kind, however, were surely a fact of life for many during the preceding Malthusian epoch, and for most before the advent of agriculture. In building our model we therefore ask what the dynamics of risky choice might have looked like for hunter-gatherers. We seek to answer this question not literally, but figuratively. In the same way that we might learn about the functioning of an exchange economy by thinking about life on Robinson Crusoe's island, we posit that we can learn about the long-run dynamics of risk preferences by thinking about life before the Neolithic revolution.

Consider the following thought experiment. Imagine you are a hunter-gatherer, living somewhere in the world more than 13,000 years ago. What would the landscape of risky choice look like over your lifetime? First, abstracting from seasonality, it stands to reason that the menu of choices facing you would be quite constant from year to year. Since finding sustenance would be your first-order concern, we might profitably simplify your choice environment to a periodic decision between high-risk, high-reward hunting and low-risk, low-reward gathering. Second, you are likely to face significant shocks to your income process due to vagaries of nature that are unavoidable and outside of your control. Without access to a robust savings technology it would be difficult for you to shift income across time, and your consumption would be hand-to-mouth. Because the society in which you live would be relatively small and lacking in insurance mechanisms, constraints on risk sharing would be strongly binding. The existence of uninsurable aggregate shocks would exert a strong selection pressure over time, meaning that systematic departures from optimal risk-taking behavior are unlikely to persist. The existence of such shocks would also imply that you would have a strong incentive to pay attention to and learn about tail draws in your background income process, subject to information and cognitive constraints. Finally, outside sources of news would not be readily available, meaning that the primary source of information you would use to learn about your environment would be your own personal experiences of it.

We capture these features by building a dynamic model of risky choice in which an agent's risk preference adapts to her evolving beliefs about risk in the environment. The foundation of our model is a classical background risk framework. An expected utility maximizer faces a choice each period from a fixed menu containing a risky asset and a safe asset—the foreground risk. She makes this choice in the presence of an exogenous, unavoidable, and statistically independent background risk. We assume that the agent's utility is risk vulnerable, a higher-order analog of risk aversion. Given this assumption, the foreground and background risks are substitutes for the agent: the more risk exists in the environment, the less risk she wants to take in her own individual choices.

We combine this framework with a model of high-dimensional learning over the background risk. We assume that our agent is Bayesian, and that she learns like an econometrician, by observing realizations of the background risk over time and using this data to update her beliefs about its data generating process. We further assume that the agent is boundedly rational, in that she does not know the true data generating process, but perceives it to be a lower dimensional approximation, in our case a stationary Gaussian random variable. To capture the first-order role of tail draws in our setting we assume that the agent knows neither the mean nor the variance of background risk. As she observes realizations of the background risk she updates her beliefs about its moments, which in turn affects her choice over the foreground risk.

Our model makes sharp predictions about the way that the agent's risk preference over the foreground risk changes in response to realizations of the background risk, given the agent's body of experiences. Most importantly, our model predicts that realizations that increase the perceived mean will make the agent less risk averse, while those that increase the perceived variance will make the agent more risk averse. The overall effect of a given realization on the agent's risk preference will be the sum of these two moment effects.

3.2.1 Model

The choice environment

Consider an agent born at time 0. In each period, indexed by $t \in \{1, 2, ..., T\}$, the agent receives a fixed wealth endowment w and is exposed to two sources of risk. First, the agent must choose an income lottery $\tilde{x}$ from a menu of lotteries $\mathcal{X}$. We call $\tilde{x}$ the endogenous or *foreground* risk, and denote its cumulative distribution function (CDF) $F_{\tilde{x}}(x)$ and its probability density function (PDF) $f_{\tilde{x}}(x)$. The menu $\mathcal{X}$ is identical in each period, and consists of a safe lottery x^s, and a risky lottery x^r, such that $\mathbb{E}[x^s] < \mathbb{E}[x^r]$ and $Var[x^s] < Var[x^r]$. To fix ideas, we think of the lotteries in $\mathcal{X}$ as objective gambles for which the agent knows the odds, though $\mathcal{X}$ could also, without loss of generality, consist of several insurance or investment options over which the agent has subjective beliefs, so long as those beliefs do not change over time.

In addition to the endogenous lottery $\tilde{x}$ the agent is exposed in each period to an exogenous *background* income risk $\tilde{y}$, which is a random variable with stationary CDF $F_{\tilde{y}}(y)$. Background risk $\tilde{y}$ is statistically independent of all $\tilde{x} \in \mathcal{X}$ in all t, and is unavoidable by the agent. The agent does not know the parameters of $F_{\tilde{y}}(y)$ but rather has beliefs over them, which she updates each period as she experiences a new realization of $\tilde{y}$. Denote with $B_t(y)$ and $b_t(y)$ the CDF and the PDF, respectively, of the agent's beliefs distribution about the outcomes of $\tilde{y}$ at time t.

Timing

The timing of events in the model is shown in Figure 3.1. The agent enters period t with income endowment w and prior beliefs b_{t-1} about the background risk $\tilde{y}$. She then chooses $\tilde{x}$ before $\tilde{y}$ is realized, given her beliefs. We assume that the agent does not have access to a savings technology, so once $\tilde{x}$ and $\tilde{y}$ realize the agent consumes their period endowment and the combined realization $w + x + y$. At the end of the period the agent updates her prior b_{t-1} to posterior b_t, which forms their prior in the next period.

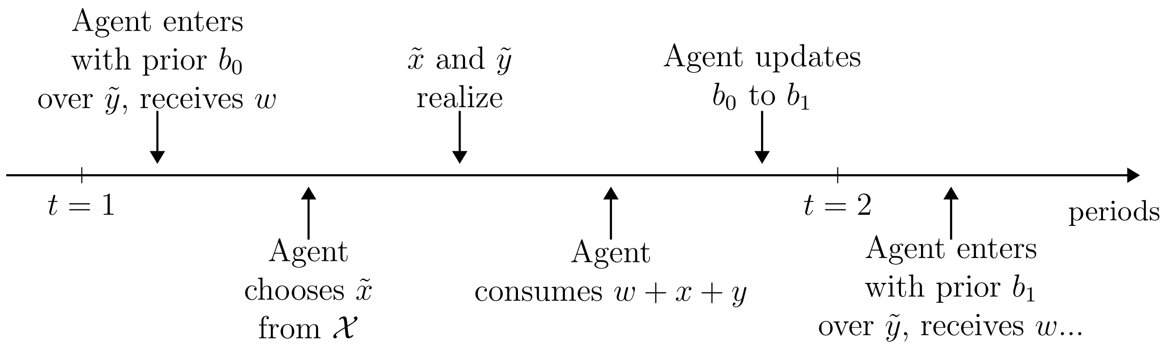

Figure 3.1. Timing of events in the model

Utility and risk

We assume that the agent is a subjective expected utility maximizer and has a four-times-differentiable utility function u, for which $u' > 0$ and $u'' < 0$. u, the agent's *direct* utility function, has as its argument the sum of the wealth endowment, the foreground risk, and the background risk. The agent's objective is therefore to maximize:

$$\begin{aligned}
\mathbb{E}u(w + \tilde{x} + \tilde{y}) &= \int \int u(w + x + y) f_{\tilde{x}}(x) b_t(y) dx dy \\
&= \int \Big[\int u(w + x + y) b_t(y) dy \Big] f_{\tilde{x}}(x) dx \\
&= \mathbb{E}u(w + \tilde{x} | \tilde{y}) \\
&= \mathbb{E}u(w + \tilde{x} | B_t(y)),
\end{aligned}$$

where the second equality follows from the law of iterated expectations. To simplify notation we will use $\tilde{y}_t = \tilde{y}|B_t(y)$ to refer to the background risk that the agent believes they face at time t.

Our measure of risk preference is the Arrow-Pratt coefficient of absolute risk aversion over the foreground risk $r_t(w)$ (Arrow, 1965; Pratt, 1978), here written to depend on the agent's beliefs about $\tilde{y}$, which vary over time:

$$r_t(w) = r(w|B_t(y)) \equiv -\frac{\mathbb{E}_t u''(w|B_t(y))}{\mathbb{E}_t u'(w|B_t(y))}.$$

The coefficient $r_t(w)$ has a well-known behavioral interpretation as the agent's risk premium, or local price for trading off the mean and variance of a risky prospect. Given a choice between a safe and a risky investment option, as in the choice of $\tilde{x}$, an agent with higher $r_t(w)$ will invest a lower amount in (or be less probable to choose, in the discrete case) the risky option.

It is also useful to define two higher-order analogues of $r_t(w)$, the coefficient of absolute prudence $p_t(w) = -\mathbb{E}_t u'''(w|B_t(y))/\mathbb{E}_t u''(w|B_t(y))$ and the coefficient of absolute temperance $q_t(w) = -\mathbb{E}_t u''''(w|B_t(y))/\mathbb{E}_t u'''(w|B_t(y))$. These allow us to define conditions on the third and fourth moments of u that are collectively termed *risk vulnerability*.[3]

Definition 3.2.1. (*Risk-vulnerable utility*) An expected utility maximizer with $u' > 0$ and $u'' < 0$ is risk-vulnerable at time t if $p_t(w) \geq r_t(w)$ and $q_t(w) \geq r_t(w)$.

Risk vulnerability is the feature of the utility function that ensures that background and foreground risks are substitutes for the agent. Intuitively it corresponds to higher-order concavity in the agent's utility function. Note that all HARA utility functions exhibit risk vulnerability.[4] We assume below that the agent is risk vulnerable at all t.

Learning

The agent in our model is a Bayesian who uses personally observed realizations of the background risk to update their belief distribution $B_t(y)$. We make two structural

[3]Risk vulnerability was first defined by Gollier and Pratt (1996) as the condition that any unfair background risk ($\mathbb{E}\tilde{y} \leq 0$) makes risk-averse agents behave in a more risk-averse way. Gollier and Pratt then derive the conditions on u we describe as consequences of their definition. Our definition of risk vulnerability differs from theirs in taking these conditions on u as a starting point. This is important because strong monotonicity of risk aversion in $\mathbb{E}\tilde{y}$ does not hold if $Var\tilde{y}$ can change as well, which is the setting with which we are concerned here.

[4]Hyperbolic Absolute Risk Aversion utility functions are defined as the class of functions for which the reciprocal of the coefficient of absolute risk aversion is linear in wealth. Many utility functions used in applications, including the linear, exponential, power, and logarithmic fall into this class (Merton, 1971).

assumptions about the agent's updating process. First, we assume that the agent believes that the realizations, or signals, are drawn from a stationary Gaussian random variable with unknown mean and unknown variance. Second, we assume that the agent's prior over the mean and variance takes the form of a normal-inverse-chi-squared distribution. We call this learning process Mean-Variance learning, and define it formally in the following definition:

Definition 3.2.2. (*Mean-Variance Learning*) We say that a Bayesian agent is a mean-variance learner if:

1. The agent's perceived likelihood function over a random variable, here the background risk, is a stationary Gaussian random variable:

$$\tilde{y} \sim \mathcal{N}(M, \Sigma^2) \ \ \forall t,$$

   where M and Σ^2 are both scalars that are unknown to the agent.

2. The agent's prior over the mean and variance $p(M, \Sigma^2)$ is a $NI\chi^{-2}$ distribution, that is,

$$\begin{aligned} p(M, \Sigma^2) &= NI\chi^{-2}(\mu_0, \kappa_0, \sigma_0^2, \nu_0) \\ &= \mathcal{N}(M|\mu_0, \Sigma^2/\kappa_0) \times \chi^{-2}(\Sigma^2|\nu_0, \sigma_0^2), \end{aligned}$$

   where μ_0 and σ_0^2 are the agent's point priors over the mean and variance of $\tilde{y}$, and $\kappa_0 > 0$ and $\nu_0 > 2$ are parameters capturing the agent's confidence or precision over the prior mean and variance, respectively.

Given the above prior, it is straightforward to show that the agent's expected values for M and Σ^2 at time 0 are:

$$\mathbb{E}_0[M] = \mu_0 \tag{3.1}$$

$$\mathbb{E}_0[\Sigma^2] = \frac{\nu_0}{\nu_0 - 2}\sigma_0^2. \tag{3.2}$$

The $NI\chi^{-2}$ distribution is the unique conjugate prior of the Gaussian with unknown mean and unknown variance likelihood. This means that the Bayesian agent's posterior distribution upon receiving signals will also be in the $NI\chi^{-2}$ family, with updated parameters. Consequently, the agent's posterior mean and variance have closed form expressions. Let $\mathcal{D}_t = \{y_1, ..., y_t\}$ be a set of t iid draws from $\tilde{y}$. Then these posteriors will be:[5]

$$\mathbb{E}_t[M|\mathcal{D}_t] = \mu_t = \mu_0 + \frac{t}{\kappa_0 + t}(\bar{y}_t - \mu_0) \tag{3.3}$$

$$\mathbb{E}_t[\Sigma^2|\mathcal{D}_t] = \frac{\nu_t}{\nu_t - 2}\sigma_t^2 = \frac{1}{\nu_0 + t - 2}\left[\nu_0\sigma_0^2 + \sum_{i=1}^{t}(y_i - \bar{y}_t)^2 + \frac{t\kappa_0}{\kappa_0 + t}(\bar{y}_t - \mu_0)^2\right], \tag{3.4}$$

where $\bar{y}_t = 1/t\sum_{i=1}^t y_i$ is the sample mean of $\mathcal{D}_t$. It will also be useful to refer to the sample variance of $\mathcal{D}_t$, $s_t^2 = 1/t\sum_{i=1}^t (y_i - \bar{y}_t)^2$.

We will denote the total change in the agent's beliefs about the mean at time t, relative to their prior, as $\Delta_t M = \mathbb{E}_t[M|\mathcal{D}_t] - \mathbb{E}_0[M]$, and about the variance $\Delta_t\Sigma^2 = \mathbb{E}_t[\Sigma^2|\mathcal{D}_t] - \mathbb{E}_0[\Sigma^2]$. These will be distinct quantities in our model from the comparisons that the agent makes between the mean of the data and their prior mean, which we label $\delta_t^m = \bar{y}_t - \mu_0$, and the difference between the sample variance and their prior variance, which we label $\delta_t^v = s_t^2 - \frac{\nu_0}{\nu_0 - 2}\sigma_0^2$.

[5] Degroot (1970) [pg.169] proves this for the parameterization of the normal in terms of mean and precision. Here we use the alternative parameterization for the normal in terms of the mean and variance. This form of the posterior variance follows trivially from replacing the Gamma prior marginal distribution of the precision in Degroot (1970) with an inverse chi squared prior marginal distribution for the variance (Murphy, 2007).

3.2.2 Theoretical results

Proposition 1. (Effect of changes in background risk on absolute risk aversion) *Suppose the agent observes an arbitrary dataset $\mathcal{D}_t$ of draws from the background risk. Then the change in their absolute risk aversion at time t is:*

$$\Delta r_t(w)\Big|_{\mathcal{D}_t} \approx -\frac{tr_0(w)\big(p_0(w) - r_0(w)\big)}{\kappa_0 + t}\delta_t^m + \frac{t\kappa_0 r_0(w)p_0(w)\big(q_0(w) - r_0(w)\big)}{2(\nu_0 + t - 2)(\kappa_0 + t)}(\delta_t^m)^2 + \frac{tr_0(w)p_0(w)\big(q_0(w) - r_0(w)\big)}{2(\nu_0 + t - 2)}\delta_t^v$$

Corollary 3.2.1. *Proposition 1 holds if the background risk is distributed as a $t_{\nu_t}(M|\mu_t, \sigma_t^2/\kappa_t)$, and the agent's objective is to learn only about the mean M.*

Corollary 3.2.2. *Suppose the agent has a CRRA utility function. Then under the complete information case*

$$\lim_{t\to\infty}\big(\Delta r_t(w)\big)\Big|_{\mathcal{D}_t} = -\frac{\eta}{w^2}\delta_t^m + \frac{\eta(\eta+1)}{w^3}\delta_t^v \tag{3.5}$$

3.3 Data and Methodology

We perform our empirical analyses using data from Indonesia and Mexico. These two countries have two advantageous settings for our purposes: first, both countries share a recent history of rapid and volatile economic change. Since both are low- to middle-income, they exhibit significant missing markets in insurance, credit, and risk-sharing. This means that the average individual in both countries is likely to have experienced substantial and unavoidable changes in background risk over their lifetime, which in turn means that we are more likely to detect effects in line with our theoretical predictions in these settings.

Although they have important similarities, the second advantage afforded by studying these two countries is their differences. Indonesia and Mexico offer a distinct contrast along many plausibly important dimensions, including geography, level of development, language, culture, religion, institutions, and other aspects of their history.[6] This aids in establishing both the internal validity and external validity of our results. If we detect common effects in both countries we can be more confident that they are not driven by idiosyncratic characteristics of either setting, and more comfortable in predicting that they will generalize to other settings.

For the Indonesian analysis our source of micro data is the Indonesian Family Life Survey (IFLS) (Strauss et al., 2009; Strauss, Witoelar and Bondan, 2016). The IFLS is a longitudinal study administered by the RAND corporation in 13 provinces in Indonesia in five waves, starting in 1993. For the Mexican analysis our source of micro data is the Mexican Family Life Survey (MxFLS), a longitudinal study administered in 16 states in three waves starting in 2002. The MxFLS was piloted by the RAND corporation, and is now managed by the Iberoamerican University (UIA) and the Center for Economic Research and Teaching (CIDE). Both surveys exhibit high recontact rates (>90%), and contain a wealth of economic and demographic covariates, allowing for a near-complete accounting of the balance sheet for subjects, including household consumption, income, assets, savings and borrowing. Both also contain residence and migration histories, allowing us to link place-based variables like climate experiences, local inflation, and GDP growth to subjects. Crucially for our purposes, the two most recent waves of both the IFLS and the MxFLS (IFLS4: 2007–2008; IFLS5: 2014; MxFLS-2: 2005–2006; and MxFLS-23: 2009–2012) include modules for measuring subject financial risk aversion using hypothetical, high-stakes monetary gambles. We use measures from these modules to construct our

[6] To make a few of these differences concrete: (1) Indonesia straddles the world's largest archipelago, spread out in equatorial waters in south-east Asia, while Mexico comprises a solid landmass in the North American continent; (2) Mexico is about 55% richer in per-capita GDP (PPP) terms than Indonesia as of 2018 ($20,602 vs. $13,230); (3) Indonesia is the world's largest Muslim country in the world, while Mexico is overwhelmingly Christian, primarily Roman-Catholic.

primary dependent variables, which we describe in detail in Subsection 3.3.1.

For the climate change variables in Indonesia we use two reanalyzed gridded temperature and precipitation datasets (Schneider et al., 2011; Willmott and Matsuura, 2001) in addition to the universe of available ground station temperature data, reported by the National Climate Data Center (NOAA-CDO, 2020). In Mexico we use the gridded weather data for North America compiled by Livneh et al. (2015), which contains information on temperature and precipitation. In Subsection 3.3.2 we describe the construction of our climate change experience variables using these data.

The sample for our main analysis is subjects who completed the risk aversion module in both waves of each survey. Focusing on subjects who appear in both waves of each survey allows us to estimate a model with individual fixed effects, which eliminates substantial amounts of noise due to idiosyncratic variation. This results in a primary sample of 16,267 subjects for Indonesia and 8,126 subjects for Mexico, each appearing twice in our data. In some analyses we do not include individual fixed effects, which allows us to expand the sample to all subjects who responded to the risk module in either wave of each survey, for a total of 51,876 subject-year observations in Indonesia and 20,851 subject-year observations in Mexico. Summary statistics for the complete survey samples and the primary samples are available in Section C.6. The geographic distributions of our samples in Indonesia and Mexico are available in Section C.5.

3.3.1 Risk aversion measures

Both surveys include modules for measuring financial risk aversion, from which our main dependent variables are constructed. These modules employ "staircase" instruments, similar to those used in Falk et al. (2018). Staircase instruments have been shown to generate high-quality measures of risk aversion with low subject response burden, which makes them ideal for field applications. In a staircase risk aversion instrument subjects are given a series of hypothetical high-stakes choices between a safe lottery (often a sure amount of money) and a riskier lottery (which generally has a higher mean and a higher variance than the safe option). Lotteries are commonly in the form of fair coin flips. Based on the subject's choice in the first question they are sorted into one of two other questions with different amounts of money for the lotteries. If the subject previously chose the safe (risky) option, risk in the coin flip is reduced (increased) in their subsequent question. This process can then be repeated as many times as necessary to yield as fine a measure of risk aversion as desired. The result is an ordinal binned measure of absolute risk aversion for each subject. Our process for constructing the risk aversion measures from the IFLS and MxFLS data is displayed in Section C.3.

In IFLS4 and IFLS5 subjects answered between two and three questions each, which resulted in measure with five bins. Each question offered the same fixed safe amount of money, while the amounts of the risky lottery varied between questions. The same module with the exact same amounts per question was used in both waves of the survey. We code the resulting measure with higher numbers (1–5) indicating more risk aversion. One complicating factor with the IFLS risk aversion module is that the first question offered subjects a choice between a sure amount and a coin flip over two higher amounts. Between 28% and 40% of the sample chose the dominated, certain option, even after being prompted to reconsider a second time (see Section C.10 for the sample distribution of the risk aversion measure). It is unclear whether these "gamble averse" subjects are extremely risk averse (or certainty seeking), or whether another factor, like subject misunderstanding or aversion to gambling generally is driving these choices. In our main analysis we include these subjects and code them as having the highest rate of risk aversion.

In MxFLS-2 subjects answered between two and five questions each, which resulted in a measure with five bins. Questions offered subjects a choice between a safe coin flip and a riskier coin flip, with the amounts of the riskier coin flips generally changing between questions. We code the resulting measure with higher numbers (1–5) indicating more risk aversion. The staircase instrument was changed for MxFLS-3 to align more closely with the instrument in the IFLS. In MxFLS-3 subjects answered between two and five questions each, resulting in a measure with six bins. Each question offered the same fixed safe amount of money, while the amounts of the risky lottery varied between questions. A "gamble averse" option was offered in this instrument. Since gamble aversion only appears in one wave of the MxFLS we drop subjects who chose this option in MxFLS-3 from our sample. We code the resulting measure (1–5) in the same way as the other measures.

A pervasive concern with all elicited measures of financial risk aversion is the high degree of noise that they exhibit, which often means their predictive power for real-world risky behavior is quite low (Yariv, Gillen and Snowberg, 2019). This raises the possibility that any detected effects on measured risk aversion will be due to noise, and will not translate to real changes in risk-taking behavior by the subjects. We address this in Subsection 3.4.3, where we show suggestive evidence that subjects who became more risk averse by our measures also became less risk-taking in their economic behavior. We can also examine the predictive capacity of our measures in the cross-section. In Section C.7 we present the results of regressing our measures of risk aversion on a host of demographic covariates and economic variables capturing risk-taking behavior in our samples, without including individual fixed effects in the regression. For subjects for whom we have complete data for all covariates, our IFLS risk aversion measure, unlike many in the literature, exhibits significant correlations with risk-taking behavior like self-employment and migration, and demographic measures like age and gender in expected ways, both in primary (panel) sample and in the broader sample. Our measure of risk aversion from the MxFLS is noisier than that in the IFLS, and consequently only exhibits significant correlations with smoking and age.

3.3.2 Climate experience variables

To construct our main independent variables for the analysis, we begin by constructing province/state-month time series for both temperature and precipitation in Indonesia and Mexico. In Indonesia, we use the Global Historical Climatology Network Climate Anomaly Monitoring System (GHCN CAMS) gridded temperature dataset and gridded rainfall data from the Global Precipitation Climatology Centre (Schneider et al., 2011; Willmott and Matsuura, 2001).[7] These are each gridded historical reanalyses and report average monthly temperature in degrees Celsius (C) and precipitation in centimeters (cm) on a 0.5 degree grid. In our main specification, we use the entire catalogue of monthly data from 1901 to 2014. To construct the province-month series of temperature and precipitation, we take the average of all pixels that fall within a province boundary (ESRI, 2018).

In an alternative specification, we use the universe of available ground station data, reported by the National Oceanic and Atmospheric Administration (NOAA-CDO, 2020). In this station-based analysis, we construct monthly temperature series for 1976 onwards. Following concerns about entry and exit of weather stations (Dell, Jones and Olken, 2014), we use 1976 as a cutoff and restrict to 61 stations that do not exit during the extent of our panel, from 1976 to 2014. These stations report daily mean temperatures in degrees C.[8] To reduce the incidence of measurement noise, we winsorize this station-day at the 1–99 level over the universe of station-day observations. We then take the median of these station-day means over all stations in a province-month to produce this station-based series. Some measurement error exists in earlier years due to stations going offline. Reassuringly, for earlier years in the data, less than 1% of province-month observations are missing. Since the data generating process illuminates potential error in the gridded data we use in the main specification, we also consider using a subset of the gridded data where we restrict the series to 1976 to 2014. We demonstrate robustness to alternate specifications of the Indonesia temperature data in Section C.1.

In Mexico, we use data from the gridded weather product for the Continental US

[7]UDel_AirT_Precip and Temp data provided by the NOAA/OAR/ESRL PSL, Boulder, Colorado, USA, from their website: http://climate.geog.udel.edu/~climate/html_pages/download.html.

[8]Following convention, this daily mean is the mean of the daily maximum and minimum readings.

and North America created by Livneh et al. (2015). These data contain temperature (degrees C) and precipitation (cm) at a 6 kilometer-pixel resolution at the daily frequency from 1950 to 2013. We use daily mean temperature and total rainfall, and construct pixel-month time series by averaging daily pixel means within each month (or summing, in the case of precipitation). We match pixels to Mexican states using the GIS layer of Mexican administrative state boundaries from the Department of the Interior (DOI, 2020) (in Subsection 3.4.4 we show robustness to a different state matching procedure using inverse distance weighting of pixels from the state centroid). With these matched pixels we construct state-month level time series by averaging the pixel-month values of temperature and precipitation for each pixel that falls within a state's administrative boundary, following a similar procedure to Auffhammer and Rubin (2018).

Once we obtain province/state-month time series for climatic variables, we match them to subjects in our data by their state and year of birth. Subjects born in a given year are matched with a time series for their province/state of birth starting in January of the next year. Once the time series are assigned we calculate for each individual the mean (A_{it}) and the standard deviation (V_{it}) of their climatic time series from birth to year of measurement in the corresponding survey. Thus, an individual born in East Java in 1981, for instance, will be assigned the statistics for the East Java temperature time series from January 1982 to 2007 (the year of IFLS4) and from 1982 to 2014 (the year of IFLS5). In Mexico, since MxFLS-2 was administered between 2005 and 2007, and MxFLS-3 was administered between 2009 and 2013, subjects are assigned time series that extend from birth to their exact measurement year. Let c_{is} be the climatic variable assigned to person i in year s (with $c \in \{\text{temperature, precipitation}\}$). Then for month of measurement t (with $t = 1$ for January of the subject's birth) these statistics are:

$$A_{it} = \frac{1}{t - b_i} \sum_{s=b_i+1}^{t} c_{is} \tag{3.6}$$

$$V_{it} = \sqrt{\frac{1}{t - b_i - 1} \sum_{s=b_i+1}^{t} (c_{is} - A_{it})^2} \tag{3.7}$$

where

$$b_i = \begin{cases} \text{BirthYear}_i & \text{if BirthYear}_i > B \\ B & \text{if BirthYear}_i \leq B, \end{cases}$$

and

$$B = \begin{cases} 1901 & \text{if Country}_i = \text{Indonesia \& main spec} \\ 1976 & \text{if Country}_i = \text{Indonesia \& restricted spec} \\ 1950 & \text{if Country}_i = \text{Mexico.} \end{cases}$$

Significant variation exists in these experienced climate variables, as can be seen from Figure 3.2 and Figure C.1.

3.3.3 Empirical specification

Our baseline empirical specification is a two-way fixed effects model where we regress the individual risk aversion measure R_{it} on A_{it}, V_{it}, a constant α_{FE}, and individual and time fixed effects:

$$R_{it} = \alpha_{FE} + \alpha_i + \alpha_t + \beta_1 A_{it} + \beta_2 V_{it} + \gamma_1 PriceLevel_p + \gamma_2 X_{it} + \varepsilon_{it}, \tag{3.8}$$

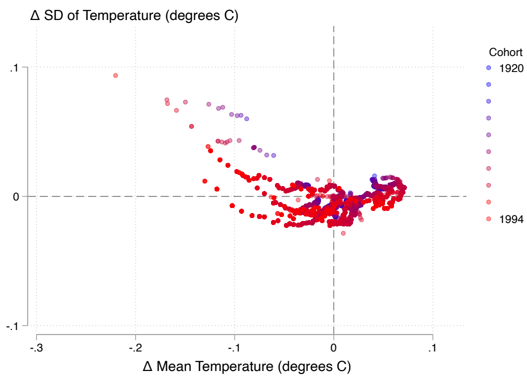

(a) Temperature, Indonesia

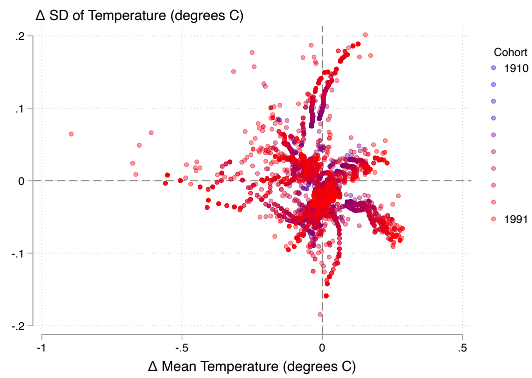

(b) Temperature, Mexico

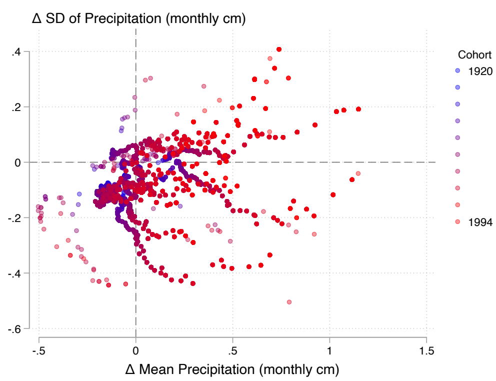

(c) Precipitation, Indonesia

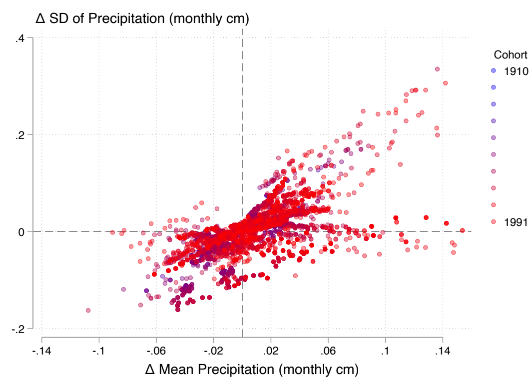

(d) Precipitation, Mexico

Figure 3.2. Moment correlations for birth-province/state cohorts

Note: this figure displays the raw distributions of our main explanatory variables ΔA_{it} and ΔV_{it} graphed against each other in each country for each independent variable. These scatterplots demonstrate that substantial variation exists not only in climate conditions across provinces/states, but also in the dynamics of climate experiences at the individual level. Cohorts are plotted by color, with oldest cohorts in blue and youngest cohorts in red. The plots for Indonesia report the baseline specification.

where A_{it} is either individual experiences of average monthly temperature or total precipitation, and V_{it} is corresponding individual experience of the respective volatility. The individual fixed effect α_i absorbs variation due to time-invariant idiosyncratic heterogeneity, whereas the time fixed effect α_t nets out the effect of aggregate time trends. X_{it} is a vector of extra individual controls (we discuss these specifically in Subsection 3.4.2).

Since we have two periods in our analysis (the first and second waves of each survey), our two-way fixed effects specification is econometrically equivalent to a first-difference specification:

$$\Delta R_{it} = \alpha_{FD} + \beta_1 \Delta A_{it} + \beta_2 \Delta V_{it} + \gamma_1 Inflation_p + \gamma_2 \Delta X_{it} + \varepsilon_{it}. \tag{3.9}$$

For exposition, we present the results below for the first-difference specification.

3.4 Results

This section contains the findings from our three primary empirical analyses. In Subsection 3.4.1, we present the results from regressing within-subject changes in measured risk aversion on subjects' experienced mean temperature change and temperature volatility change in Indonesia and Mexico, as well as the corresponding analysis for precipitation in each country. These regressions, which include no controls aside from subnational inflation, are the most direct tests of the predictions of our model. In Subsection 3.4.2, we demonstrate the robustness of our main findings to the inclusion of controls for changes in subjects' economic constraints and experiences of violence, natural disasters, and macroeconomic growth. In Subsection 3.4.3, we present correlations between changes in several kinds of risky behaviors and predicted change in risk-taking across the distribution of predicted risk preference changes.

3.4.1 Effects of climate experiences on measured risk aversion

Our main empirical findings are presented in Table 3.1. Column 1 displays the result of regressing changes in measured risk aversion on mean changes in experienced lifetime heat and precipitation (separately) in Indonesia. In line with the model's predictions, increases in the mean of each climate variable result in significant decreases in measured risk aversion in Indonesia. In column 2, changes in risk aversion are regressed on changes in the standard deviation of the climate variables, but results are not significant for either variable. Column 3 presents the results of regressing measured risk aversion on both the mean and the standard deviation of the climate variables in Indonesia. The mean effects here remain highly significant, while the variance effect of heat now becomes significant and positive, in line with our theoretical predictions.

Columns 4–6 present the results of the parallel analysis in Mexico. In column 4, changes in measured risk aversion are regressed (separately) on changes in the experienced lifetime mean of heat and precipitation. We find that the effect of mean temperature is significant and negative, while that of mean precipitation is negative but not significant. In column 5, changes in measured risk aversion are regressed on changes in the standard deviation of the climate variables. Here, the effect of the standard deviation of precipitation is significant and positive, while the effect of temperature standard deviation of temperature are not significant. Finally, in column 6, measured risk aversion is regressed on both the mean and the standard deviation of the climate variables. In line with the model's predictions, both mean variables have negative and highly significant effects, while the standard deviation of precipitation has a positive and highly significant effect.

Overall, the results are strongly consistent with the predictions of the model on the direction of the effects of changes in the mean and variance of experienced climate on measured risk aversion. Two additional observations are worth noting at this juncture. First, even though the effects of mean climatic variables are highly consistent in both settings, only the variance of heat in Indonesia and precipitation in Mexico have significant effects in our analysis. Why the effects of climate volatility differ between the two settings is unclear, though it might be attributable to differences in the correlation structure between the moments of climate variable and the dominant income process in each setting.[9] Second, it is notable that in the specifications in which the variance effect is significant, their magnitudes are approximately 1.6 (Indonesia temperature) and 0.7 (Mexico precipitation) times the magnitudes of the respective mean climate effects. This suggests that the effects of experienced climate variance are first order on risk preferences.

Table 3.1. Main results

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)
	Indonesia			**Mexico**		
Δ Mean Temp	-3.75††		-4.23††	-1.16†		-1.19††
	(.49)		(.57)	(.22)		(.22)
Δ Std. Dev. Temp		1.54	6.82**		-0.10	-0.35
		(2.32)	(2.37)		(.48)	(.49)
Δ Mean Precip	-0.25**		-0.21*	-1.14		-3.99**
	(.09)		(.10)	(.93)		(1.15)
Δ Std. Dev. Precip		-0.44	-0.27		1.17*	2.58***
		(.25)	(.28)		(.55)	(.69)
Observations	16267	16267	16267	8126	8126	8126

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014) and MxFLS (2006–2012). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

[9] For instance, heat waves may be especially damaging to agricultural yields in the Indonesian context, and floods may be especially damaging to such yields in Mexico.

3.4.2 Additional controls

Our main results are estimated without the inclusion of any additional controls aside from subnational inflation, though there are well-founded reasons to include additional covariates. Theoretically, changes in subjects' income, wealth, buffer stocks of savings, consumption, or other economic circumstances might be expected to influence their measured risk aversion. Empirically, previous studies have shown that exposure to traumatic experiences like natural disasters and violence can change measured risk aversion. In a similar context, Levin and Vidart (2020) show that macroeconomic experiences significantly change measured risk aversion.

In the main specification, we choose to omit these controls because they are endogenous to risk aversion itself. This means that their inclusion could threaten the causal interpretation of our results. Nevertheless, we interpret the changes we observe in measured risk aversion as representing changes in underlying risk attitudes, or merely as driven by changes in personal economic circumstances. Further, it is useful to directly test whether experiences of climate change are in fact driving the observed changes or whether other kinds of experiences whose incidence may be correlated with climate dynamics are in fact playing a central role.

We provide some evidence on these points in Table 3.2, where we progressively add in additional controls to the specification for the last column in Table 3.1. These include time-varying demographics, like marital status, educational attainment, and household size; changes in income, assets, savings, and consumption; self-reported exposure to violence and natural disasters; and measured GDP growth experiences from Levin and Vidart (2020) (full details on the controls are available in Section C.9). In both countries our results are highly robust to the inclusion of this rich set of covariates. Overall, this is suggestive that the changes we estimate in measured risk aversion are driven by lifetime experiences of climate. The only covariates that substantially attenuate or increase the magnitude of the measured effect are the inclusion of macroeconomic experience variables, suggesting correlation between changes in the climate and growth and perception on underlying risk.

Table 3.2. Additional controls

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
					Indonesia				
Δ Mean Temp	-4.23††	-4.39††	-4.39††	-4.37††	-4.37††	-4.43††	-4.41††	-4.38††	-3.61†
	(.57)	(.56)	(.56)	(.56)	(.56)	(.58)	(.58)	(.59)	(.55)
Δ Std. Dev. Temp	6.82**	6.93**	6.88**	6.90**	6.90**	6.90**	6.86**	6.63*	12.56†
	(2.37)	(2.36)	(2.36)	(2.36)	(2.36)	(2.43)	(2.44)	(2.44)	(2.32)
Δ Mean Precip	-0.21*	-0.20	-0.20	-0.20	-0.20	-0.22*	-0.22*	-0.21*	-0.19
	(.10)	(.10)	(.10)	(.10)	(.10)	(.11)	(.11)	(.11)	(.12)
Δ Std. Dev. Precip	-0.27	-0.27	-0.27	-0.27	-.28	-.24	-0.24	-0.31	-0.51
	(.28)	(.28)	(.28)	(.28)	(.28)	(.28)	(.28)	(.28)	(.27)
Observations	16267	16267	16263	16263	16263	14974	14974	14974	14974
					Mexico				
Δ Mean Temp	-1.19†	-1.18†	-1.18†	-1.18†	-1.17†	-1.17†	-1.17***	-1.15†	-0.95***
	(.22)	(.22)	(.22)	(.22)	(.22)	(.22)	(.22)	(.22)	(.23)
Δ Std. Dev. Temp	-0.35	-0.38	-0.39	-0.39	-0.38	-0.36	-0.36	-0.39	-1.12*
	(.49)	(.49)	(.49)	(.49)	(.49)	(.49)	(.49)	(.49)	(.50)
Δ Mean Precip	-3.99***	-4.10***	-4.14***	-4.13***	-4.11***	-4.10***	-4.10***	-4.46***	-7.95†
	(1.15)	(1.14)	(1.14)	(1.14)	(1.15)	(1.15)	(1.15)	(1.14)	(1.20)
Δ Std. Dev. Precip	2.58***	2.59***	2.59***	2.59***	2.58***	2.59***	2.58***	2.70***	4.63†
	(.69)	(.69)	(.69)	(.69)	(.69)	(.69)	(.69)	(.68)	(.73)
Observations	8126	8126	8126	8126	8126	8126	8126	8126	8126
Inflation	X	X	X	X	X	X	X	X	X
Δ Demographics		X	X	X	X	X	X	X	X
Δ Income			X	X	X	X	X	X	X
Δ Assets				X	X	X	X	X	X
Δ Savings					X	X	X	X	X
Δ Consumption						X	X	X	X
Δ Violence							X	X	X
Δ Natural Disasters								X	X
Δ Growth experiences									X

Note: measured risk aversion reported from 1–5, 5 being highest. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Demographics include marital status, household size, and household size squared. Violence and natural disasters variables from self-reported exposure. Growth experiences include the mean, mean squared, and standard deviation of province/state level real GDP growth in subjects' province/state of birth. Standard errors clustered at the cohort by province of birth level in parenthesis. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

3.4.3 Correlations with risky behavior

Another issue of interpretation of our results is the question of whether changes in measured risk aversion capture changes in actual risk-taking behavior for subjects. We study this question by constructing a variable measuring predicted change in risk aversion ($\widehat{\Delta R_{it}}$) using our preferred specifications (columns 3 and 6 of Table 3.1) and examining its correlation with changes in downstream risky behaviors in our data. We focus on behaviors commonly examined in relation to risk-taking in the literature for which we have data: smoking, having ever migrated across province or state line, self-employment status, and, in Indonesia, whether subjects report that their land is planted with at least one cash crop.[10] We do this exercise for both temperature and precipitation in Indonesia and Mexico.

Results of this analysis for temperature and precipitation in Indonesia are presented in Figure 3.3 and Figure 3.4; and results for temperature and precipitation in Mexico are presented in Figure 3.5 and Figure 3.6. These figures display the average value of each downstream variable for each quartile of the $\widehat{\Delta R_{it}}$ distribution. Here, lighter blue bars represent individuals who are predicted to become measurably less risk averse by our empirical model, while dark blue bars represent subjects who are predicted to become measurably more risk averse. We report 95% confidence intervals for each quartile, and we use the first to fourth interquartile range as an empirical benchmark to run a two-sided t-test to check the statistical significance of the difference between the average values of the outcome. Importantly, we acknowledge that this is not a causal exercise, as we only claim to capture the part of change in measured risk aversion due to changes in our climate variables.

Our results provide some suggestive evidence that increases in measured risk aversion predicted by climate change experiences are correlated with overall decreases in risk-taking behavior in Indonesia, and less but some suggestive evidence that this relationship exists in the Mexican case.

In the case of Indonesia, we observe several cases where increases in risk aversion

[10]Cash crops included in the IFLS include coconut, coffee, cloves, rubber, and other hard-stemmed plants.

are correlated with decreases in risk-taking behavior. Four of the eight first-to-fourth quartile declines are significant at conventional levels: ever migrated across province lines by temperature (1.3 percentage point increase to 0.5 percentage point increase; p=0.001); ever migrated across province lines by precipitation (1.3 pp increase to 0.3 pp increase; p=0.0001); rate of smoking by temperature (2.1 pp increase to 1.0 pp increase; p=0.046); rate of smoking by precipitation (3.8 pp increase to 0.1 pp increase; p=0.0001). However, these relationships are not all monotonic over the intermediate quartiles. Additionally, for self-employment, we observe no statistically or economically significant relationship. And, for the case of cash crops by precipitation, we observe an upward trend and marginally significant first-to-fourth quartile difference (0.7pp increase to 2.0 pp increase; p=0.63). Overall, we take this to be suggestive evidence that climate-predicted increases in measured risk aversion to be correlated with decreases in observed risk-taking behavior in Indonesia.

In the case of Mexico, we observe little evidence of overall statistical change in risky behavior attributable due to changes in experienced temperature or precipitation. For the temperature analysis, we find no significant differences across the distribution of these quartiles of predicted risk increase. However, there exists a positive trend between decreases in risk aversion and increases in migration. For the precipitation analysis, again we find some baseline relationship between predicted rates of decreases in risk aversion and increases in migration, but we find no statistical differences for changes in these internal rates of migration or smoking, and even find a weakly *increasing* relationship for changes in self-employment with predicted changes in risk aversion. One potential explanation for what we observe is that there are other significant contributions to observed changes in risk attitudes during the time of the study. For example, Brown et al. (2019) find significant impacts of violent crime during this period on measures of risk aversion.

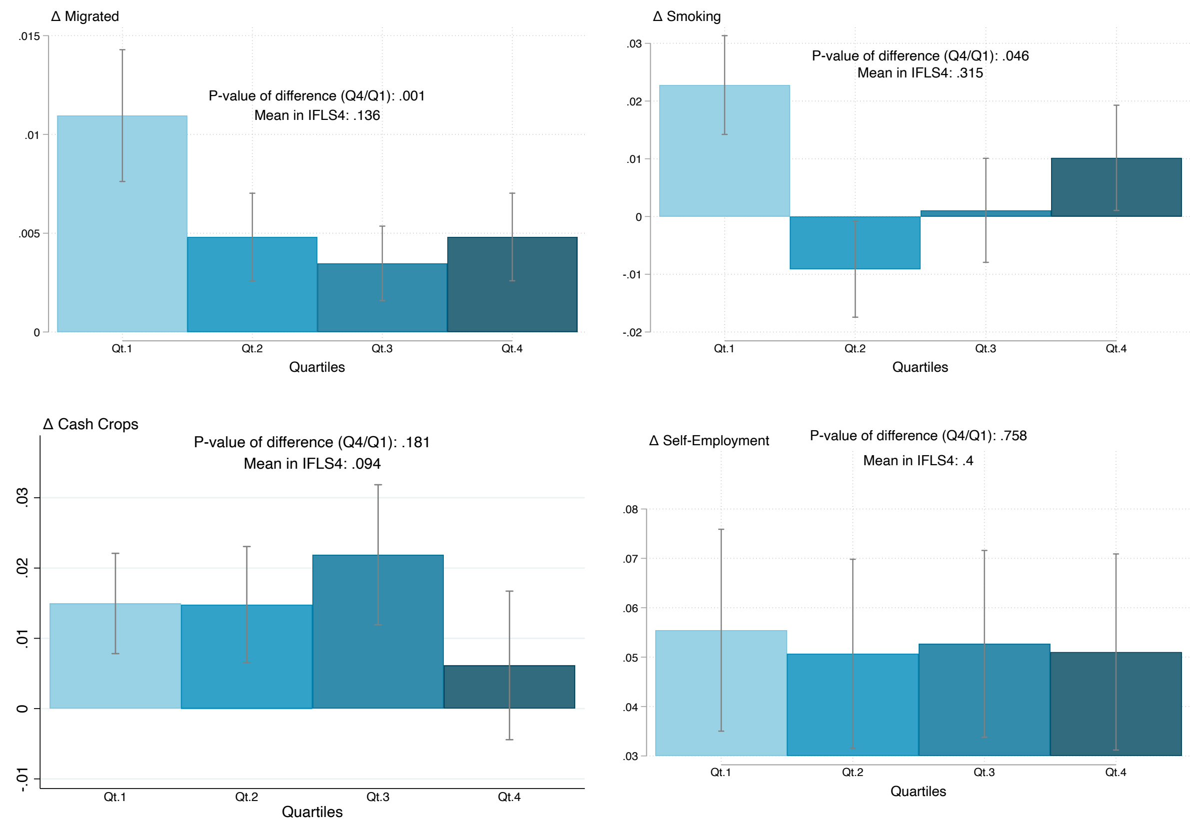

Indonesia—Temperature

Figure 3.3. Correlations of changes in risky behaviors with predicted temperature increase in risk aversion, Indonesia

Note: bars represent quartiles of the predicted change in risk aversion distribution. Light blue is the bottom quartile of this distribution, representing the agents who are predicted to experience a decrease (or smaller increase) in risk aversion.

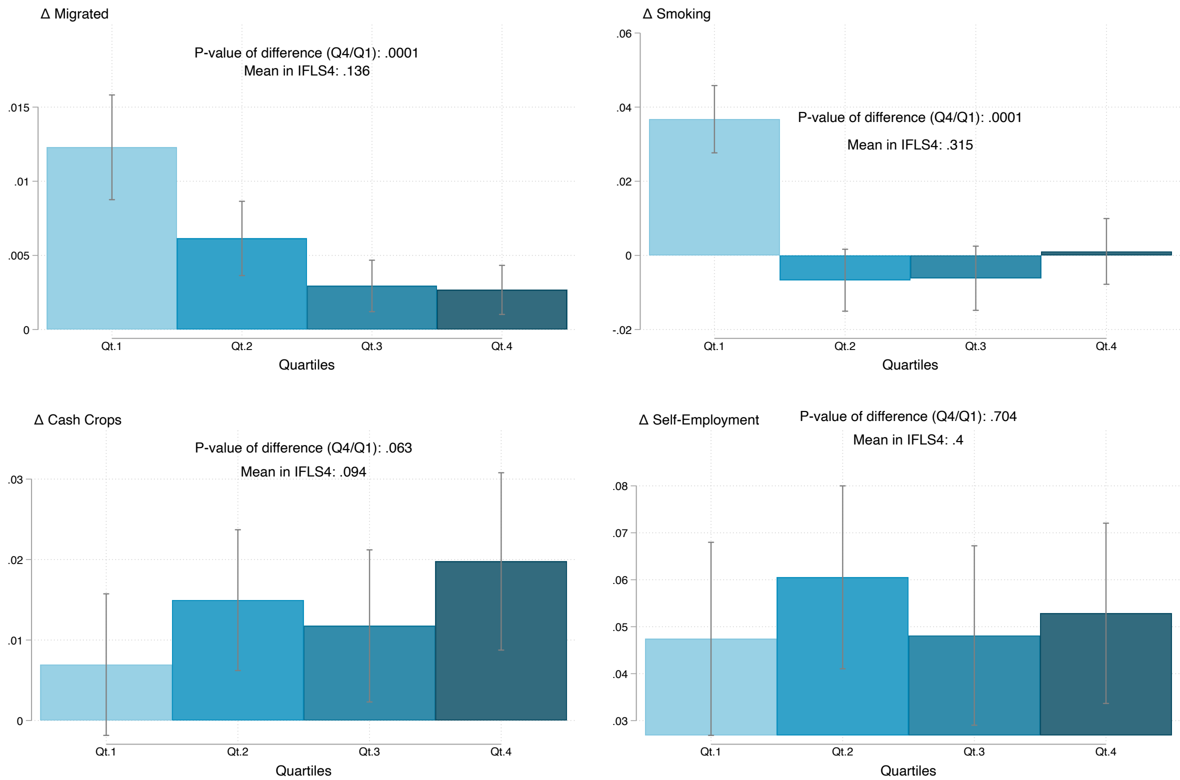

Indonesia—Precipitation

Figure 3.4. Correlations of changes in risky behaviors with predicted precipitation increase in risk aversion, Indonesia

Note: bars represent quartiles of the predicted change in risk aversion distribution. Light blue is the bottom quartile of this distribution, representing the agents who are predicted to experience a decrease (or smaller increase) in risk aversion.

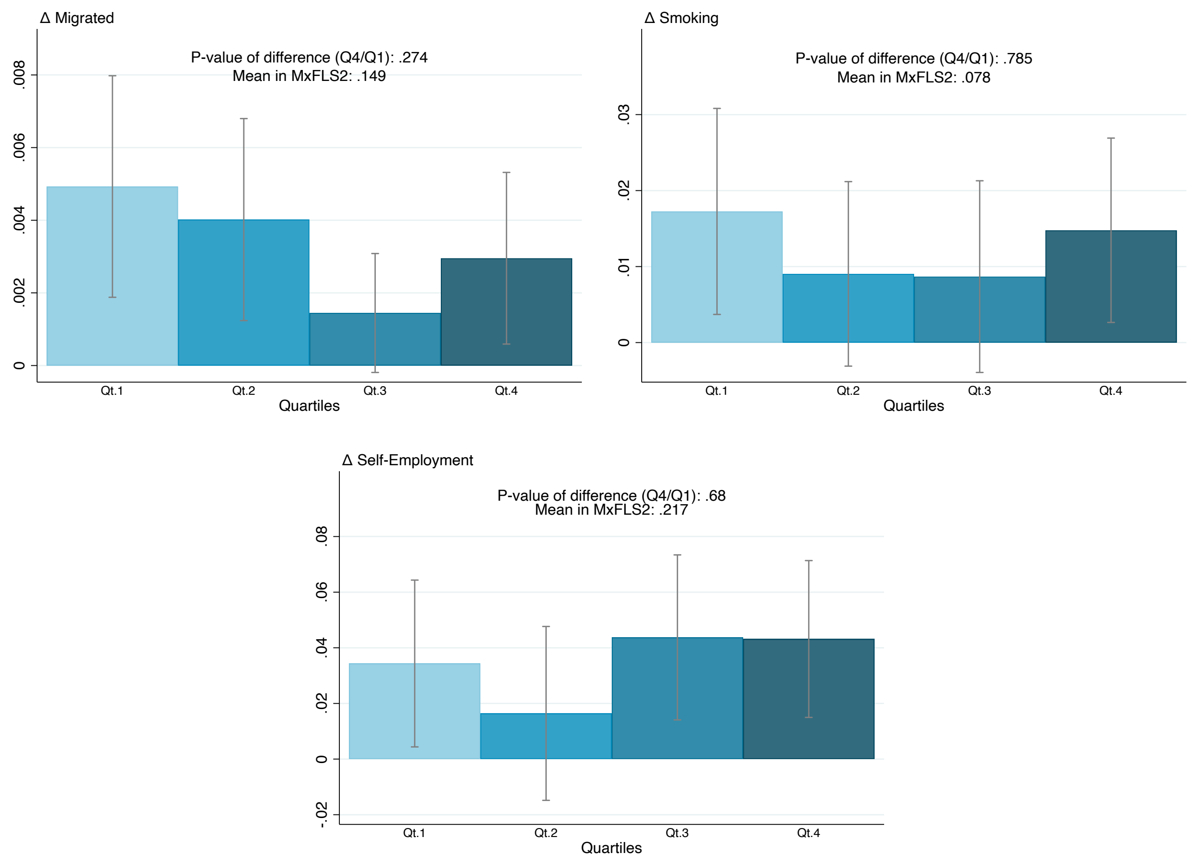

Mexico—Temperature

Figure 3.5. Correlations of changes in risky behaviors with predicted temperature increase in risk aversion, Mexico

Note: bars represent quartiles of the predicted change in risk aversion distribution. Light blue is the bottom quartile of this distribution, representing the agents who are predicted to experience a decrease (or smaller increase) in risk aversion.

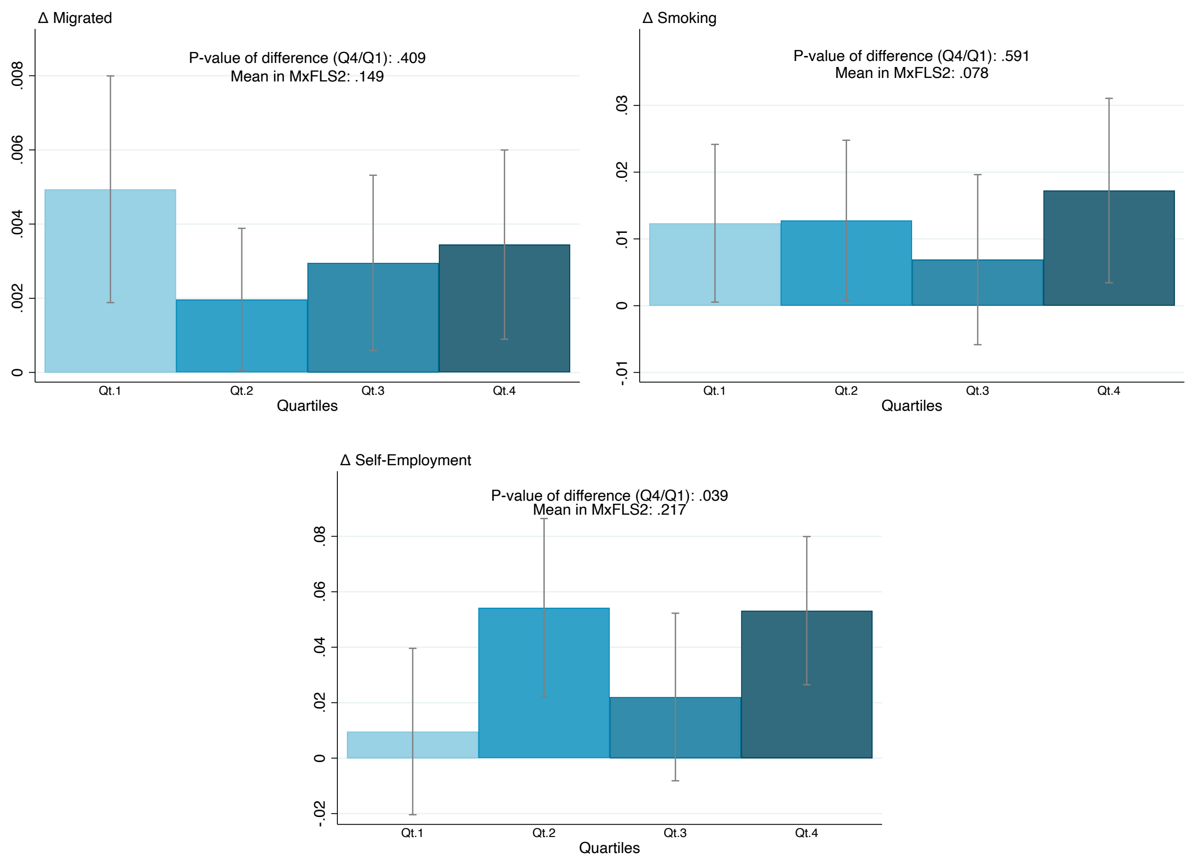

Mexico—Precipitation

Figure 3.6. Correlations of changes in risky behaviors with predicted precipitation increase in risk aversion, Mexico

Note: bars represent quartiles of the predicted change in risk aversion distribution. Light blue is the bottom quartile of this distribution, representing the agents who are predicted to experience a decrease (or smaller increase) in risk aversion.

3.4.4 Robustness

We test the robustness of our main results to varying methodological choices in our analysis in Appendix C.1.

First, in Table C.1 we present the results of running our main analysis with alternate sample compositions. In particular, we limit the analysis to individuals born after 1976 in Indonesia and 1950 in Mexico (for whom we have full lifetime climate histories). The results are qualitatively very similar for each of these samples, though the linear mean term in the Mexico precipitation regression becomes marginally significant when the sample is restricted to those born after 1950.

In Table C.2 and Table C.3, we present the results of our main analysis for alternate specifications of measured risk aversion. For both Indonesia and Mexico, we repeat the analysis with (1) a binarized measure of risk aversion (instead of using the 5 buckets of measured risk aversion, we set buckets 1 and 2 to be 0, and buckets 3, 4 and 5 to be 1); and (2) using an ordered probit specification. The latter specification accounts explicitly for the ordinal nature of our risk aversion measure, though its results should be interpreted with care as the ordered probit with two-way fixed effects estimator is known to be biased. For both specifications results are qualitatively quite similar to the baseline.

In Table C.4 we present the results of our main analysis using data on climate conditions in subjects' province or state of residence at the time of the first survey, rather than their province or state of birth. These data more closely match the intuitive notion of climate change experiences, but, as discussed above, suffer from a potential identification problem due to endogenous migration. Results are again qualitatively similar to the main analysis.

In Table C.5 we present the results for the station-based specification for Indonesia temperature variables. The construction of the Indonesian ground station data is discussed in Subsection 3.3.2. Because of concerns of bias due to endogenous station placement, we use this robustness check to benchmark the reliability of the gridded data we use in the baseline. In Table C.6, we restrict the weather series to start at 1976 (assigning individuals born within a province prior to 1976 the same climate variables). Since the data generating process for the gridded data is based on station data plus interpolation from climate models, we take the similar qualitative results across these specifications to be suggestive that our analysis is not primarily driven by bias or noise from the gridded baseline specification.

In Table C.7 we present the results of our main analysis with standard errors clustered at the province/state of birth level, using the wild bootstrap method of Cameron, Gelbach and Miller (2008). Estimates of our coefficients are mostly not significant at conventional levels under this scheme, though the coefficient of the mean coefficient in Mexico and the volatility coefficient in Indonesia are significant in some specifications.

In Table C.8 we conduct our main analysis using a repeated cross-section specification that drops the individual fixed effects from the regression. We perform this analysis restricting to our primary sample. The magnitude of the estimated coefficients drops considerably in both settings under this specification. In Indonesia the mean temperature coefficients remain the same sign and highly significant. In Mexico, the signs on estimated coefficients are conserved for those that were originally significant.

In Table C.9 and Table C.10, we report the results of the main specification, but instead we use changes in the baseline structural CRRA parameters that we computed in Subsection 3.5.1 as the outcome variable. The sign and significance when using our structurally computed risk parameters our similar to our main analysis. Note that the mapping from risk buckets to CRRA risk parameters is nonlinear and over a different range.

3.5 Adaptation and Counterfactual Exercise

In this section, we establish two main contributions. First, following Eden (2020), we use the concept of an equally-distributed equivalent (EDE) to measure societal-level preferences over ex-post outcomes (here, consumption) when individuals have heterogeneous risk preferences. We apply this to our two-period setting, where we observe individual preferences over a risky lottery change over time, to estimate how welfare changes along this dimension relative to fixing preferences over these gambles. We call this risk preference adaptation.[11]

Then, we decompose this risk preference adaptation into the portion that is explained by changes in the climate distribution. We use our baseline linear model to predict counterfactual risk distributions with individual risk preferences that adapt to changes in background risk with and without climate effects. In Indonesia, we find that overall risk adaptation is welfare improving. Compared to a baseline distribution with fixed preferences, we show that about one sixth of the six percentage point increase in welfare is attributable to this climate risk preference adaptation channel. In the Mexican setting, we find that overall risk adaptation is welfare decreasing, but that the climate risk preference adaptation offsets about one sixteenth of the overall eight percentage point decrease in welfare.

3.5.1 General risk preference adaptation

In Subsection 3.4.1, we showed that individual risk preferences respond to changes in individual-experienced climate. Here, we present the results of a structural exercise to

[11]While we use the adaptation terminology, we acknowledge an important difference from its regular use within economics. When risk preferences are fixed, we often refer to adaptation as the re-optimization effect, which in this neoclassical setting should only lead to increases in welfare. In this scenario, we relax the static risk preference assumption (conditional on the background risk). *A priori*, the theoretical motivation gives us no reason to think that observable changes in risk preferences should be "adaptive" (increasing welfare) or "maladaptive" (decreasing welfare). Under this scenario, even with re-optimization, the net effect on welfare is ambiguous.

determine how this observed change in individual risk preferences is associated with social welfare.

The measure of welfare we use is the EDE of the observed consumption distribution in 2014 Indonesia and 2012 Mexico, respectively. Following Eden (2020), this gives us a measure that can be interpreted similarly to an individual certainty equivalent for a given distribution of consumption outcomes. Like a certainty equivalent, higher risk aversion measures are associated with a lower EDE. Unlike an individual certainty equivalent, the EDE allows us to calculate a single measure of social welfare when individuals have heterogeneous risk preferences, and can be interpreted as a measure of how much society collectively would be willing to pay to avoid uncertainty in the distribution of consumption. We then calculate the EDE for different distributions of risk preferences: the preferences in the first wave of the IFLS and MxFLS; the preferences in the second wave; and finally a counterfactual distribution in the second wave, for which we construct a distribution where we net-off predicted changes attributed to the climate.

To do this, We assume that individuals have isoelastic utility, as defined in Equation 3.10, and allow the CRRA parameter θ_j to be individual-specific.

$$u_j(s) = \begin{cases} \frac{s^{1-\theta_j}-1}{1-\theta_j}, & \theta_j \neq 1 \\ \log(s), & \theta_j = 1 \end{cases} \tag{3.10}$$

In order to create a mapping from the risk bins in the IFLS and the MxFLS to individual structural parameters, we consider a variety of specifications with different structural assumptions. First, we consider how individuals bracket the trade-off between each risky gamble and the sure payoff. In Equation 3.10, let $s = x + p$, where p is the payoff of the lottery. For these lotteries, it is unclear whether subjects are weighing-off only the specific prize of the lottery (narrow bracketing, $x = 0$), or if they integrate that with consumption (broad bracketing, $x =$consumption). Second, we calculate the upper- and lower-bounds on each θ_j assuming that individuals are maximizing their expected utility. We assign the average of these bounds to each θ_j. For individuals that select into the most and least risk averse bins, we cannot observe a finite upper-bound or lower-bound respectively.[12] To bound these, we make either a tight or broad assumption on these bounds.[13] Finally, for the IFLS, we consider different choices of θ_j construction for individuals that exhibit gamble aversion. We either exclude them, treat them identically to the most risk averse (but non-gamble averse), or assume they are more risk averse than this previous group.[14]

We consider both per capita and household real consumption as the relevant distribution of outcomes for calculating the EDE. We report the EDE ratio as the ratio of the EDE to the mean level of consumption, which can be interpreted as the fraction of total consumption that society would be willing to accept for eliminating uncertainty in the consumption distribution. The numerical procedure is described in Appendix C.8.

3.5.2 Climate risk preference adaptation

We use our constructed measure of risk preference adaptation and combine this with our empirical model to decompose the portion of this adaptation that is attributable to individual response to changes in the experienced climate. To do this, we estimate the baseline linear model in Equation 3.9 using individual-specific changes in the structural CRRA risk parameter θ as the outcome variable. Results of this specification are economically and statistically consistent with our main results and are reported in Table C.9 and Table C.10.

Using the distribution of initial θ's, we consider two counterfactual distributions in the later survey wave. First, we use predicted individual-specific changes in θ to construct

[12] That is, for the least risk averse bin, we cannot rule out that individuals are infinitely risk loving, so that $\theta = -\infty$, or that individuals in the most risk averse bin are infinitely risk averse, so that $\theta = \infty$.

[13] We calculate the mean interval between the upper- and lower-bounds for individuals that select into the interior bins (2 and 3 in the IFLS, and 2, 3, and 4 in the MxFLS). Under the tight assumption, we create a lower (upper) bound for the least (most) risk averse of one interval from the known upper (lower) bound. Under the broad assumption, we extend this to two full intervals.

[14] Using the same intervals used to calculate the bounds for the least and most risk averse bins, we add one extra interval length to the upper bound relative to the most risk averse, but non-gamble averse individuals.

a distribution that reflects full risk preference adaptation. Second, we use the predicted changes net of climate effects. This second counterfactual distribution allows for changes in our observed measure of risk that are not attributable to changes in individual-experienced climate.

With the initial distribution of θ's and the two counterfactual distributions, we calculate three measures of the EDE ratio over the observed distribution of consumption in the later survey wave. In the main Indonesia specification, the fixed preference EDE ratio is 0.367. The EDE ratio under full adaptation is 0.428, or an overall increase in welfare over the consumption distribution. The EDE ratio under adaptation net of climate effects is 0.418. That is, when we allow individuals to adapt to changes in experienced climate, we find an overall increase in our measure of welfare. The magnitude and interpretation of this exercise are similar across a majority of our specifications. The results of alternative specifications are reported in Figure 3.7 and Figure 3.8.

In the main Mexico specification, the fixed preference EDE ratio is 0.379. The EDE ratio under full adaptation is 0.297, which we interpret as an overall decrease in welfare over the consumption distribution. The EDE ratio under adaptation net of climate effects is 0.291. Again, we interpret this as evidence of increases in overall welfare through the climate adaptation channel. While the calculation of the EDE ratio under full adaptation varies across alternative specifications, our decomposition of the climate adaptation channel is consistent across specifications. The results of the alternative specifications for this exercise are reported in Figure 3.9 and Figure 3.10.

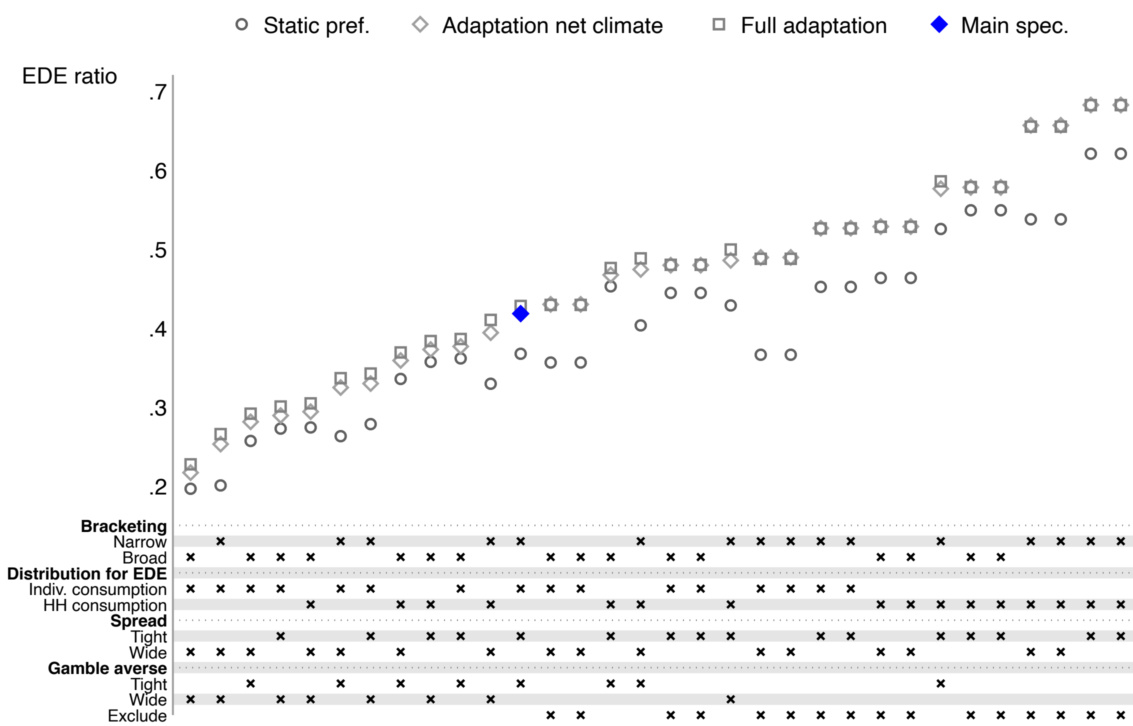

Figure 3.7. Climate and risk adaptation, Indonesia

Note: EDE ratio is defined as the counterfactual EDE divided by the mean level of consumption. Static preferences refer to the distribution of preferences implied by the 2007 IFLS4.

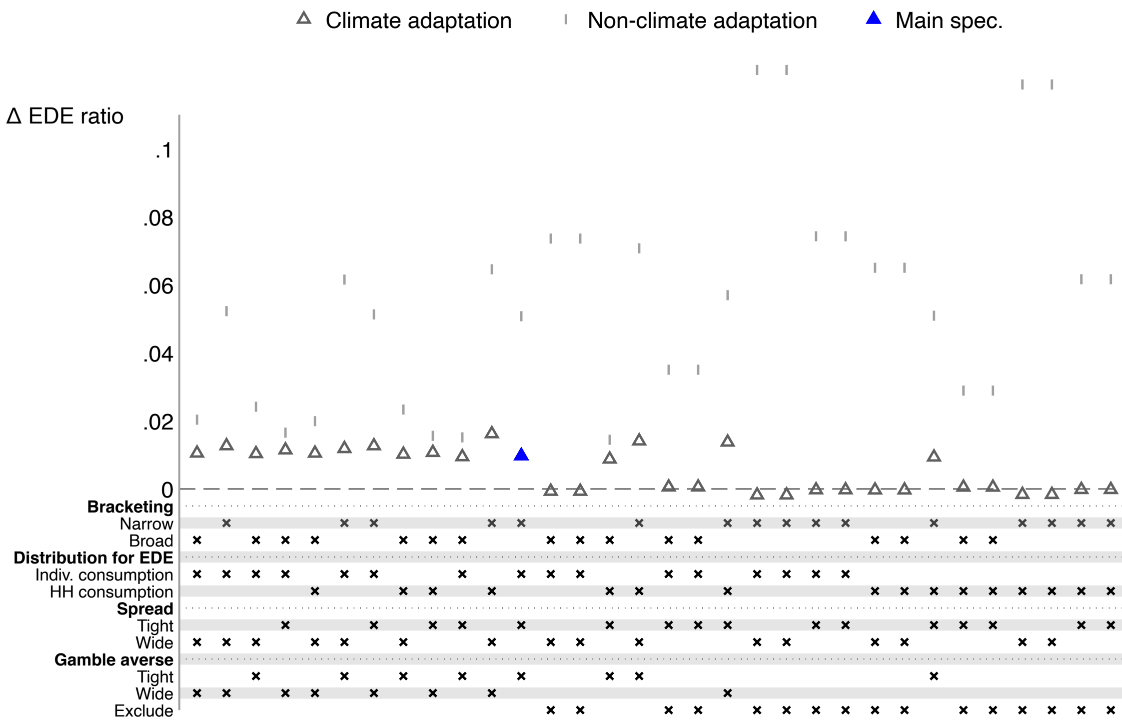

Figure 3.8. Climate and non-climate adaptation, Indonesia

Note: EDE ratio is defined as the counterfactual EDE divided by the mean level of consumption. Static preferences refer to the distribution of preferences implied by the 2007 IFLS4.

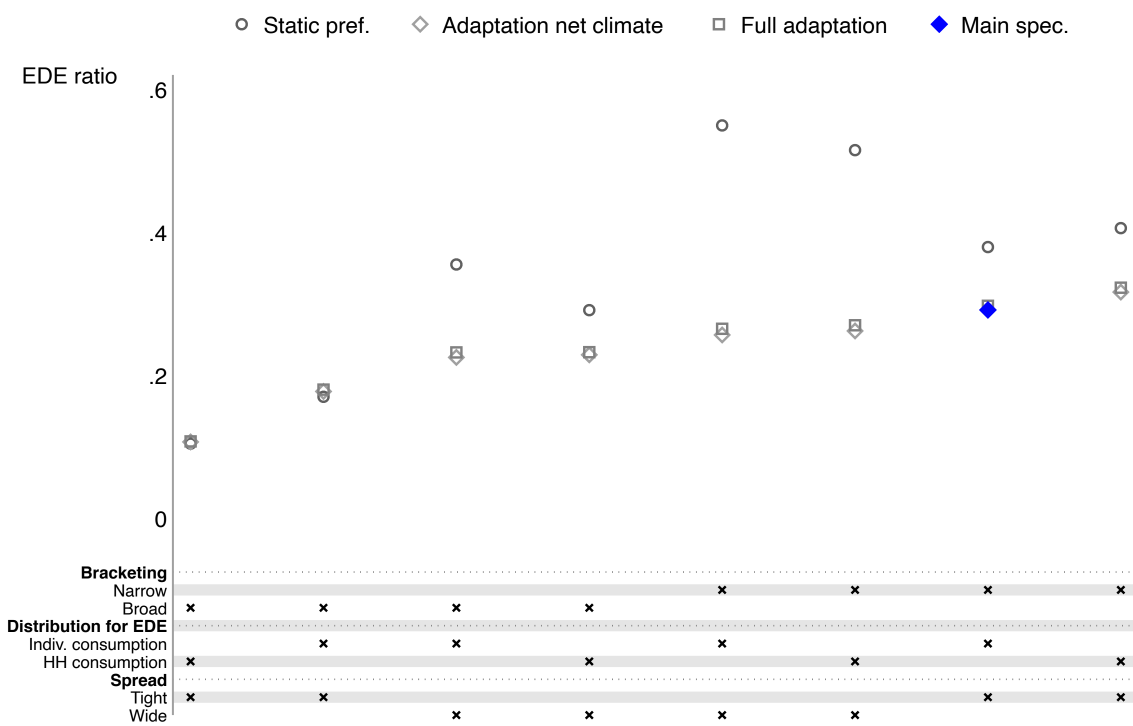

Figure 3.9. Climate and risk adaptation, Mexico

Note: EDE ratio is defined as the counterfactual EDE divided by the mean level of consumption. Static preferences refer to the distribution of preferences implied by the 2005 MxFLS-3-2.

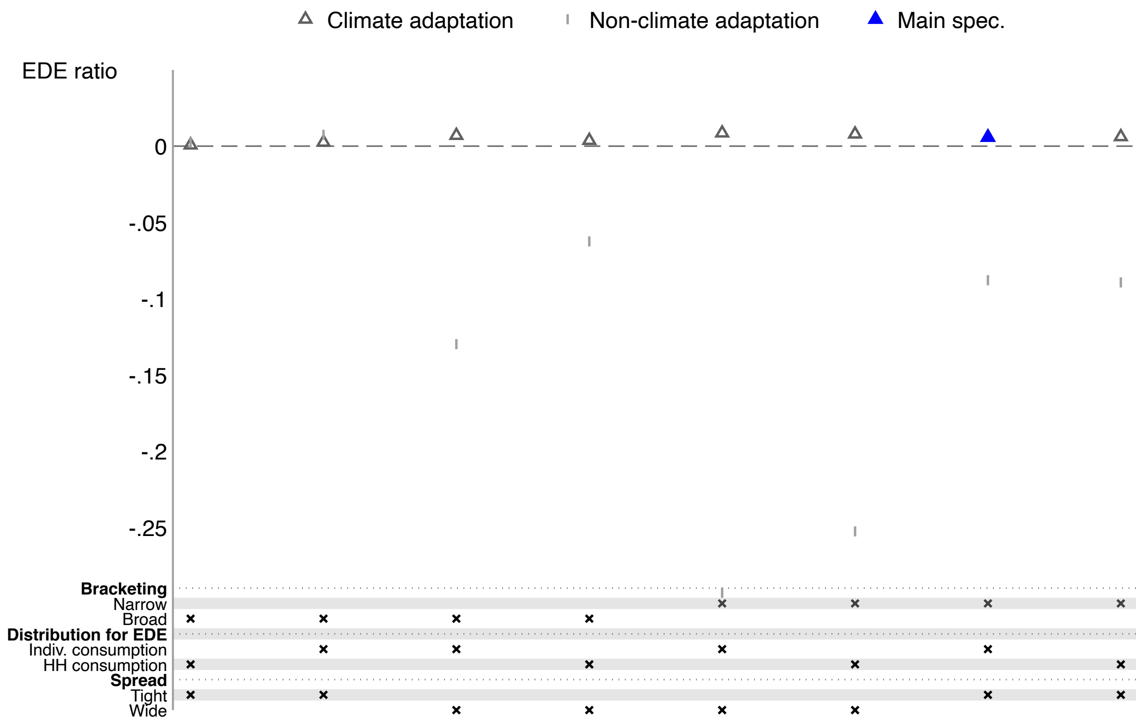

Figure 3.10. Climate and non-climate adaptation, Mexico

Note: EDE ratio is defined as the counterfactual EDE divided by the mean level of consumption. Static preferences refer to the distribution of preferences implied by the 2005 MxFLS-2.

3.6 Conclusion

Our analysis provides significant evidence that lifetime experiences of climate change shape risk-taking for individuals in Indonesia and Mexico. Using micro data containing elicited risk aversion for the same subjects years apart, linked to subnational climate statistics capturing subjects' climate experiences, we find strong support for the hypotheses of our model about the adaptation of risk-taking to changes in the mean and variance of background risk. These findings are robust to the inclusion of a rich set of controls for changes in economic circumstances and other categories of experiences. Changes in measured risk aversion correlate with substantial changes risk-taking in the domains of migration, health, and investment behavior.

3.7 Acknowledgements

Chapter 3 is co-authored with Remy Levin and is currently being prepared for submission for publication. The dissertation author is a principle researcher on this chapter.

Bibliography

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2017. "Microeconomic origins of macroeconomic tail risks." *American Economic Review*, 107(1): 54–108.

Albouy, David, Walter Graf, Ryan Kellogg, and Hendrik Wolff. 2016. "Climate amenities, climate change, and American quality of life." *Journal of the Association of Environmental and Resource Economists*, 3(1): 205–246.

Allison, Paul D. 2019. "Asymmetric fixed-effects models for panel data." *Socius*, 5: 2378023119826441.

Arrow, Kenneth J. 1965. *Aspects of the theory of risk-bearing.* Yrjö Jahnssonin Säätiö.

Asadoorian, Malcolm O, Richard S Eckaus, and C Adam Schlosser. 2008. "Modeling climate feedbacks to electricity demand: The case of China." *Energy Economics*, 30(4): 1577–1602.

Auffhammer, Maximilian. 2014. "Cooling China: The weather dependence of air conditioner adoption." *Frontiers of Economics in China*, 9(1): 70–84.

Auffhammer, Maximilian, and Anin Aroonruengsawat. 2011. "Simulating the impacts of climate change, prices and population on Californias residential electricity consumption." *Climatic change*, 109(1): 191–210.

Auffhammer, Maximilian, and Edward Rubin. 2018. "Natural gas price elasticities and optimal cost recovery under consumer heterogeneity: evidence from 300 million natural gas bills."

Auffhammer, Maximilian, and Erin T Mansur. 2014. "Measuring climatic impacts on energy consumption: A review of the empirical literature." *Energy Economics*, 46: 522–530.

Bakkensen, Laura A, and Lint Barrage. 2017. "Flood risk belief heterogeneity and coastal home price dynamics: Going under water?" National Bureau of Economic

Research.

Bakkensen, Laura A, Xiaozhou Ding, and Lala Ma. 2019. "Flood risk and salience: New evidence from the Sunshine State." *Southern Economic Journal*, 85(4): 1132–1158.

Barrage, Lint, and Jacob Furst. 2019. "Housing investment, sea level rise, and climate change beliefs." *Economics letters*, 177: 105–108.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S Shapiro. 2016. "Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the twentieth century." *Journal of Political Economy*, 124(1): 105–159.

Barreca, Alan, Olivier Deschênes, and M Guldi. 2015. "Maybe Next Month? Temperature Shocks, Climate Change, and Dynamic Adjustments in Birth Rates."

Baylis, Patrick, and Judson Boomhower. 2021. "Building codes and community resilience to natural disasters."

Becker, Gary S., and George J. Stigler. 1977. "De Gustibus Non Est Disputandum." *American Economic Review*, 67(2): 76–90.

Bernheim, B Douglas, and Dmitry Taubinsky. 2018. "Behavioral Public Economics." In *Handbook of Behavioral Economics*. Vol. 1, 381–516. Elsevier B.V.

Biddle, Jeff. 2008. "Explaining the spread of residential air conditioning, 1955–1980." *Explorations in Economic History*, 45(4): 402–423.

Brown, Phillip, Adam Daigneault, Emilia Tjernström, and Wenbo Zou. 2018. "Natural Disasters, Social Protection, and Risk Perceptions." *World Development*, 104: 310–325.

Brown, Ryan, Verónica Montalva, Duncan Thomas, and Andrea Velásquez. 2019. "Impact of violent crime on risk aversion: Evidence from the Mexican drug war." *Review of Economics and Statistics*, 101(5): 892–904.

Brueckner, Jan K. 1997. "Infrastructure financing and urban development:: The economics of impact fees." *Journal of Public Economics*, 66(3): 383–407.

Burke, Marshall, and Kyle Emerick. 2016. "Adaptation to climate change: Evidence from US agriculture." *American Economic Journal: Economic Policy*, 8(3): 106–140.

Busse, Meghan R, Devin G Pope, Jaren C Pope, and Jorge Silva-Risso. 2015. "The psychological effect of weather on car purchases." *The Quarterly Journal of Eco-*

nomics, 130(1): 371–414.

Callen, Michael, Mohammad Isaqzadeh, James D Long, and Charles Sprenger. 2014. "Violence and Risk Preference : Experimental Evidence from Afghanistan." *American Economic Review*, 104(1): 123–148.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller. 2008. "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics*, 90(3): 414–427.

Cameron, Lisa, and Manisha Shah. 2015. "Risk-taking behaviour in the wake of natural disasters." *Journal of Human Resources*, 50(2): 484–515.

Cappiella, Karen, William P Stack, Lisa Fraley-McNeal, Cecilia Lane, and Gerard McMahon. 2012. "Strategies for managing the effects of urban development on streams."

Cleetus, Rachel. 2013. "Overwhelming risk: Rethinking flood insurance in a world of rising seas." *August. Union of Concerned Scientists. Online at http://www. ucsusa. org/assets/documents/global_warming/Overwhelming-Risk-Full-Report. pdf. Accessed December*, 23: 2013.

Degroot, Morris H. 1970. *Optimal Statistical Decisions.* McGraw-Hill.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken. 2014. "What do we learn from the weather? The new climate-economy literature." *Journal of Economic Literature*, 52(3): 740–98.

Deschênes, Olivier, and Michael Greenstone. 2011. "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the US." *American Economic Journal: Applied Economics*, 3(4): 152–85.

Dewitz, J. 2019. "National Land Cover Database (NLCD) 2016 Products: U.S. Geological Survey data release." *https://doi.org/10.5066/P96HHBIE*.

Dillenberger, David, and Kareen Rozen. 2015. "History-dependent risk attitude." *Journal of Economic Theory*, 157: 445–477.

DOI. 2020. "Digital map of the state (political) boundaries of Mexico." *https://catalog.data.gov/dataset/digital-map-of-the-state-political-boundaries-of-mexico*.

Donat, Markus G, Andrew L Lowry, Lisa V Alexander, Paul A O?Gorman, and Nicola Maher. 2016. "More extreme precipitation in the world?s dry and wet regions." *Nature Climate Change*, 6(5): 508–513.

Eden, Maya. 2020. "Welfare Analysis with Heterogeneous Risk Preferences." *Journal of Political Economy*, 128(12): 4574–4613.

ESRI. 2018. "Indonesia Provinsi Boundaries 2018." *https://esri.maps.arcgis.com/home/item.html?id=512f074c60674df39ff1766a89684908*.

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, and David Huffman. 2018. "Global Evidence on Economic Preferences." *Quarterly Journal of Economics*, 133(4): 1645–1692.

Federal Emergency Management Agency. 2002. "National Flood Insurance Program: Program Description."

Frimpong, Eugene, Daniel R Petrolia, Ardian Harri, and John H Cartwright. 2020. "Flood insurance and claims: The impact of the Community Rating System." *Applied Economic Perspectives and Policy*, 42(2): 245–262.

Gallagher, Justin. 2014. "Learning about an infrequent event: evidence from flood insurance take-up in the United States." *American Economic Journal: Applied Economics*, 206–233.

Garnache, Cloe, and Todd Guilfoos. 2019. "Does the Salience of Risk Affect Large, Risky Asset Purchases?" *Risky Asset Purchases*.

Gibson, Matthew, and Jamie T Mullins. 2020. "Climate risk and beliefs in new york floodplains." *Journal of the Association of Environmental and Resource Economists*, 7(6): 1069–1111.

Gollier, Christian, and John W. Pratt. 1996. "Risk Vulnerability and the Tempering Effect of Background Risk." *Econometrica*, 64(5): 1109–1123.

Graff Zivin, Joshua, and Matthew Neidell. 2014. "Temperature and the Allocation of Time: Implications for Climate Change." *Journal of Labor Economics*, 32(1): 1–26.

Hanaoka, Chie, Hitishi Shigeoka, and Yasutora Watanabe. 2018. "Do Risk Preferences Change? Evidence from the Great Eastern Japan Earthquake." *American Economic Journal: Applied Economics*, 10(2)(2): 298–330.

Hauer, Mathew E., Elizabeth Fussell, Valerie Mueller, Maxine Burkett, Maia Call, Kali Abel, Robert McLeman, and David Wrathall. 2020. "Sea-level rise and human migration." *Nature Reviews Earth & Environment*, 1(1): 28–39.

Howden, Wesley. 2021. "Adaptation to Weather Shocks and Household Beliefs on Climate: Evidence from California."

Howe, Peter D, Matto Mildenberger, Jennifer R Marlon, and Anthony Leiserowitz. 2015. "Geographic variation in opinions on climate change at state and local scales in the USA." *Nature climate change*, 5(6): 596–603.

Jakiela, Pamela, and Owen Ozier. 2019. "The Impact of Violence on Individual Risk Preferences: Evidence from a natural experiment." *Review of Economics and Statistics*, 101(3).

Kala, Namrata. 2017. "Learning, Adaptation, and Climate Uncertainty: Evidence from Indian Agriculture." 1–52.

Keenan, Jesse M, and Jacob T Bradt. 2020. "Underwaterwriting: from theory to empiricism in regional mortgage markets in the US." *Climatic Change*, 162(4): 2043–2067.

Knowlton, Kim, Miriam Rotkin-Ellman, Galatea King, Helene G Margolis, Daniel Smith, Gina Solomon, Roger Trent, and Paul English. 2009. "The 2006 California heat wave: impacts on hospitalizations and emergency department visits." *Environmental health perspectives*, 117(1): 61.

Konrad, Christopher Peter. 2003. "Effects of urban development on floods."

Kriesel, Warren, and Craig Landry. 2004. "Participation in the National Flood Insurance Program: An empirical analysis for coastal properties." *Journal of Risk and Insurance*, 71(3): 405–420.

Kunreuther, Howard. 1996. "Mitigating disaster losses through insurance." *Journal of risk and Uncertainty*, 12(2): 171–187.

Kunreuther, Howard, and Erwann O Michel-Kerjan. 2009. "Managing catastrophes through insurance: Challenges and opportunities for reducing future risks." *Philadelphia, Pa.: Wharton Risk Management and Decision Processes Center, Working Paper*, 11–30.

Levin, Remy, and Daniela Vidart. 2020. "Risk-Taking Adaptation to Macroeconomic Experiences: Theory and Evidence from Developing Countries."

Livneh, Ben, Theodore J Bohn, David W Pierce, Francisco Munoz-arriola, Bart Nijssen, Russell Vose, Daniel R Cayan, and Levi Brekke. 2015. "A spatially comprehensive, hydrometeorological data set for Mexico, the U .S., and Southern Canada 1950 ? 2013." *Scientific Data*, 150042: 1–12.

Malmendier, Ulrike, and Stefan Nagel. 2011. "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" *The Quarterly Journal of Economics*, 126(1): 373–416.

McCright, Aaron M, and Riley E Dunlap. 2011. "Cool dudes: The denial of climate change among conservative white males in the United States." *Global environmental change*, 21(4): 1163–1172.

McNeil, Michael A, and Virginie E Letschert. 2010. "Modeling diffusion of electrical appliances in the residential sector." *Energy and Buildings*, 42(6): 783–790.

Merton, Robert C. 1971. "Optimum Consumption and Portfolio Rules in a Continuous-Time Model." *Journal of Economic Theory*, 3: 373–413.

Murphy, K. P. 2007. "Conjugate Bayesian Analysis of the Gaussian Distribution."

NOAA-CDO. 2020. "National Climate Data Center." *https://www.ncdc.noaa.gov/cdo-web/datatools/selectlocation.*

Pomatto, Luciano, Phillip Strack, and Omer Tamuz. 2020. "Stochastic Dominance under Independent Noise." *Journal of Political Economy*, 128(5): 1877–1900.

Pratt, John W. 1978. "Risk aversion in the small and in the large." In *Uncertainty in economics.* 59–79. Elsevier.

PRISM Climate Group, Oregon State University. 2021. *http://prism.oregonstate.edu.*

Schlenker, Wolfram, and Charles A Taylor. 2019. "Market expectations about climate change." National Bureau of Economic Research.

Schlenker, Wolfram, and Michael J Roberts. 2009. "Nonlinear temperature effects indicate severe damages to US crop yields under climate change." *Proceedings of the National Academy of sciences*, 106(37): 15594–15598.

Schneider, Udo, Andreas Becker, Peter Finger, Anja Meyer-Christoffer, Bruno Rudolf, and Markus Ziese. 2011. "GPCC full data reanalysis version 6.0 at 0.5: Monthly land-surface precipitation from rain-gauges built on GTS-based and historic data." *GPCC Data Rep., doi*, 10.

Shukla, PR, J Skea, E Calvo Buendia, V Masson-Delmotte, HO Pörtner, DC Roberts, P Zhai, Raphael Slade, Sarah Connors, Renée Van Diemen, et al. 2019. "IPCC, 2019: Climate Change and Land: an IPCC special report on climate change, desertification, land degradation, sustainable land management, food security, and greenhouse gas fluxes in terrestrial ecosystems."

Sleeter, Benjamin M, Terry L Sohl, Thomas R Loveland, Roger F Auch, William Acevedo, Mark A Drummond, Kristi L Sayler, and Stephen V

Stehman. 2013. "Land-cover change in the conterminous United States from 1973 to 2000." *Global Environmental Change*, 23(4): 733–748.

Strauss, John, Firman Witoelar, and Sikoki Bondan. 2016. "The Fifth Wave of the Indonesia Family Life Survey: Overview and Field Report."

Strauss, John, Kathleen Beegle, Bondan Sikoki, Agus Dwiyanto, Yulia Herawati, and Firman Witoelar. 2009. "The Fourth Wave of the Indonesia Family Life Survey: Overview and Field Report." *RAND Working Paper*, WR-675/1-N(April 2009): i–82.

Sun, Fanglin, and Richard T Carson. 2020. "Coastal wetlands reduce property damage during tropical cyclones." *Proceedings of the National Academy of Sciences*, 117(11): 5719–5725.

TX. 2020. *Senate Bill 8, 86th Texas Legislature.*

USGS EROS. 1999. "National Elevation Dataset."

Wagner, Katherine. 2020. "Adaptation and Adverse Selection in Markets for Natural Disaster Insurance." *Available at SSRN 3467329.*

Weitzman, Martin L. 2009. "On modeling and interpreting the economics of catastrophic climate change." *Review of Economics and Statistics*, 91(1): 1–19.

Willmott, Cort J, and Kenji Matsuura. 2001. "Terrestrial air temperature and precipitation: Monthly and annual time series (1950–1999) Version 1.02." *Center for Climatic Research, University of Delaware, Newark.*

Wolfram, Catherine, Orie Shelef, and Paul Gertler. 2012. "How will energy demand develop in the developing world?" *The Journal of Economic Perspectives*, 26(1): 119–137.

Woodruff, Jonathan D, Jennifer L Irish, and Suzana J Camargo. 2013. "Coastal flooding by tropical cyclones and sea-level rise." *Nature*, 504(7478): 44–52.

Yariv, Leeat, Ben Gillen, and Erik Snowberg. 2019. "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study." *Journal of Political Economy*, 127(4): 1826–1863.

Chapter 1 Appendix

A.1 Alternative Specifications

Table A.1. Central and standalone unit specification

	(1) All units b/se	(2) All units b/se	(3) Electricity b/se	(4) Electricity winter b/se
Anomaly × 2009	0.0001*		0.0404**	0.0152
	(0.00006)		(0.02017)	(0.01256)
Q1 interaction		-0.0000		
		(0.00017)		
Q2 interaction		0.0000		
		(0.00006)		
Q3 interaction		0.0001*		
		(0.00004)		
Q4 interaction		0.0001		
		(0.00007)		
Controls	X	X	X	X
Utility FE	X	X	X	X
City FE	X	X	X	X
N	38581	38581	35734	33503

Note: standard errors clustered at the ZIP-code level. In column 3, electricity corresponds to household electricity demand for the billing cycle covering mostly July in kWh. In column 4, this corresponds to electric billing cycle covering most of February. "Most coverage" is necessary because of the staggered billing cycles across households. *p< 0.1, **p< 0.05, ***p< 0.01.

Table A.2. Clustering at the city level

	(1) Central air b/se	(2) Central air b/se	(3) Electricity b/se	(4) Electricity winter b/se
Anomaly × 2009	0.0001***		0.0602**	0.0214
	(0.00007)		(0.03059)	(0.01999)
Q1 interaction		-0.0000		
		(0.00036)		
Q2 interaction		0.0001		
		(0.00001)		
Q3 interaction		0.0002**		
		(0.00009)		
Q4 interaction		-0.0000		
		(0.00019)		
Controls	X	X	X	X
UtilityFE	X	X	X	X
CityFE	X	X	X	X
N	38581	38581	35734	33503

Note: standard errors clustered at the municipal level. In column 3, electricity corresponds to household electricity demand for the billing cycle covering mostly July in kWh. In column 4, this corresponds to electric billing cycle covering most of February. "Most coverage" is necessary because of the staggered billing cycles across households. *p< 0.1, **p< 0.05, ***p< 0.01.

Table A.3. Clustering at the county level

	(1) Central air b/se	(2) Central air b/se	(3) Electricity b/se	(4) Electricity winter b/se
Anomaly × 2009	0.0001**		0.0602*	0.0214
	(0.00002)		(0.0356)	(0.02867)
Q1 interaction		-0.0000		
		(0.00022)		
Q2 interaction		0.0001		
		(0.00011)		
Q3 interaction		0.0002*		
		(0.0001)		
Q4 interaction		-0.0000		
		(0.00021)		
Controls	X	X	X	X
UtilityFE	X	X	X	X
CityFE	X	X	X	X
N	38581	38581	35734	33503

Note: standard errors clustered at the county level. In column 3, electricity corresponds to household electricity demand for the billing cycle covering mostly July in kWh. In column 4, this corresponds to electric billing cycle covering most of February. "Most coverage" is necessary because of the staggered billing cycles across households. *p< 0.1, **p< 0.05, ***p< 0.01.

Table A.4. Homeowner restricted specification

	(1) Central air b/se	(2) Central air b/se	(3) Electricity b/se	(4) Electricity winter b/se
Anomaly × 2009	0.0001*		0.0612***	0.0214*
	(0.00006)		(0.00952)	(0.01281)
Q1 interaction		-0.0001		
		(0.00026)		
Q2 interaction		0.0002		
		(0.00014)		
Q3 interaction		0.0002*		
		(0.00011)		
Q4 interaction		0.0001		
		(0.00007)		
Controls	X	X	X	X
Utility FE	X	X	X	X
City FE	X	X	X	X
N	30118	30118	28444	26826

Note: standard errors clustered at the ZIP-code level. In column 3, electricity corresponds to household electricity demand for the billing cycle covering mostly July in kWh. In column 4, this corresponds to electric billing cycle covering most of February. "Most coverage" is necessary because of the staggered billing cycles across households. *p< 0.1, **p< 0.05, ***p< 0.01.

Table A.5. Varying Fixed Effects

	(1) Central air b/se	(2) Central air b/se	(3) Central air b/se
Anomaly × 2009	5.2e-5	0.0001***	6.5e-5*
		(0.0003)	(3.4e-5)
Controls	X	X	X
UtilityFE	X	X	X
CityFE		X	
CountyFE			X
N	38581	38581	38581

Note: standard errors clustered at the ZIP-code level. *p< 0.1, **p< 0.05, ***p< 0.01.

A.2 Supplementary Figures

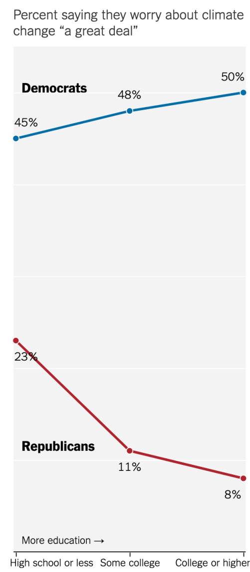

Figure A.1. Age by partisan climate change attitudes

Note: retrieved from:
https://news.gallup.com/poll/182159/college-educated-republicans-skeptical-global-warming.aspx on August 23, 2019.

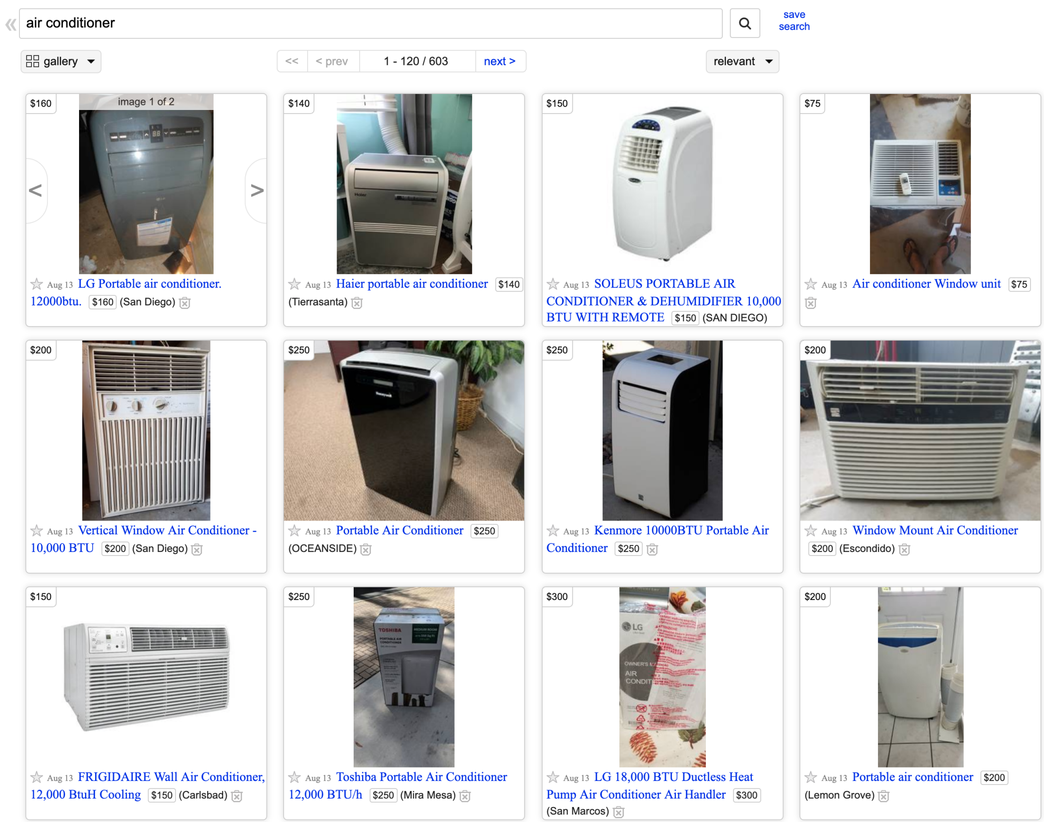

Figure A.2. Secondary market for portable air conditioners

Note: an example of the secondary market for portable air conditioners on an online resale website in San Diego. This displays a screenshot of a popular American classifieds website accessed on August 13, 2019.

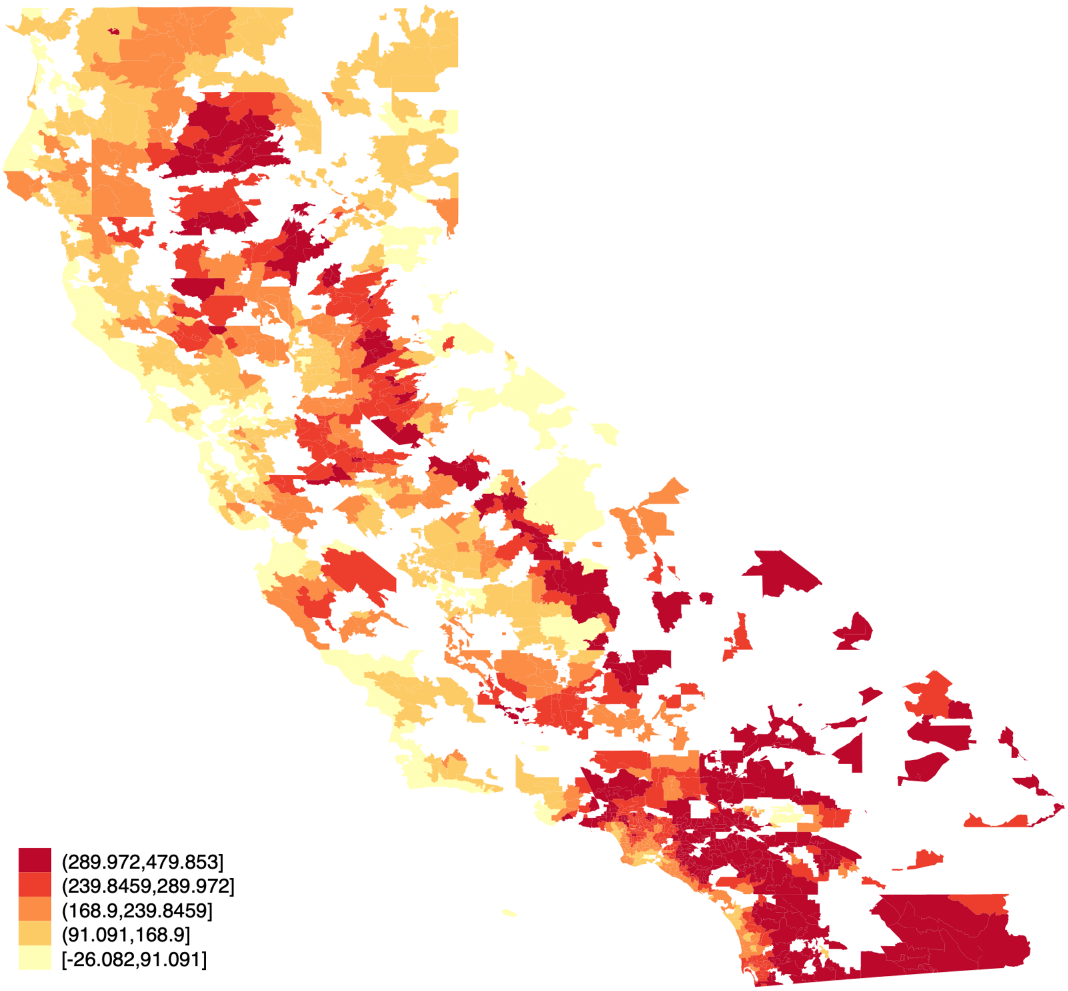

Figure A.3. 2006 CDD anomaly plot

Note: this shows the the CDD anomaly for each ZIP code in 2006. Reported ZIP codes are from SDG&E, PG&E, SCE, and LADWP.

Chapter 2 Appendix

B.1 Detailed Results from Main Specifications

Table B.1. Baseline specification

Dep. Var: Claims/policy	(1)	(2)
Land×rainfall		
Impermeable development		
0–1 inch	-.00009	-.00008
	(.00009)	(.00009)
1–2 inch	.00121*	.00121*
	(.0007)	(.0007)
2–3 inch	.00440***	.00442***
	(.00094)	(.00089)
3+ inch	.01273***	.01262***
	(.00302)	(.00415)
Wetlands		
0–1 inch	-.00042	-.00073
	(.00130)	(.00215)
1–2 inch	-.00209	-.00245
	(.00166)	(.00179)
2–3 inch	-.00272***	-.00272***
	(.00071)	(.00069)
3+ inch	-.00400	-.00401
	(.00524)	(.00526)
Water		
0–1 inch	-.00017	-.00016
	(.00013)	(.00014)
1–2 inch	.00079	.00079
	(.00097)	(.00095)
2–3 inch	-.00258*	-.00288**
	(.00145)	(.00145)
3+ inch	-.00933	-.00846
	(.00595)	(.00694)
Observations	313608	313608
Rain FE	X	X
Year FE	X	X
Land characteristics	X	

Note: standard errors clustered at the county-by-year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. 0 inches of rainfall is the omitted rainfall bin. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall bin.

Table B.2. Asymmetric effects of land change

Dep. Var: Claims/policy	(1)		(2)	
Land×rainfall				
Impermeable development	+ changes	- changes	+ changes	- changes
0–1 inch	.00012	-.03727***	.00012	-.03184**
	(.00034)	(.01469)	(.00040)	(.01500)
1–2 inch	-.00171	.15754	-.00182	.14879
	(.00134)	(.10177)	(.00152)	(.10674)
2–3 inch	.00497**	-.19949	.00497**	-.17853
	(.00214)	(.13879)	(.00213)	(.13802)
3+ inch	.00354	-.27987	.00346	-.28579
	(.01054)	(.61967)	(.01286)	(.59893)
Wetlands				
0–1 inch	.00154	.00044	.00153	.00044
	(.00297)	(.00163)	(.00296)	(.00163)
1–2 inch	.01756	.00356	.01748	.00360
	(.02481)	(.01135)	(.02579)	(.01151)
2–3 inch	-.09221*	.01047	-0.09213*	.01185
	(.05234)	(.02890)	(.05158)	(.02900)
3+ inch	-.36543***	.07414*	-.35988***	.07513*
	(.09779)	(.04017)	(.09794)	(.03942)
Water				
0–1 inch	.00236	-.06130	.00236	-.06131
	(.00230)	(.06347)	(.00230)	(.06348)
1–2 inch	.01367	.16527	.01400	.17859
	(.02101)	(.15113)	(.02099)	(.15254)
2–3 inch	-.13358	.99401***	-.13358	.98048***
	(.04446)	(.24078)	(.04582)	(.24975)
3+ inch	.02262	2.44088	.02262	2.18233
	(.09327)	(2.3322)	(.08959)	(2.4598)
Observations	313608		313608	
Rain FE	X		X	
Year FE	X		X	
Land characteristics	X			

Note: standard errors clustered at the county-by-year level. * p < 0.1, ** p< 0.05, *** p< 0.01. 0 inches of rainfall is the omitted rainfall bin. Positive changes are cumulative positive changes in a census tract from 2010 forward, and negative changes are cumulative negative changes from 2010 forward. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall bin.

Table B.3. Spillover results

Dep. Var: Claims/policy	(1)		(2)	
Land×rainfall				
Impermeable development	Uphill	Downhill	Uphill	Downhill
0–1 inch	.00011	.00016	.00012	.00016
	(.00012)	(.00014)	(.00013)	(.00014)
1–2 inch	.00043	-.00091	.00042	-.00092
	(.00042)	(.00068)	(.00042)	(.00069)
2–3 inch	.00661*	-.00079	.00599*	-.00081
	(.00361)	(.00128)	(.00345)	(.00139)
3+ inch	.00651***	.00158	.00614**	.00159
	(.002293)	(.00288)	(.00314)	(.00260)
Wetlands				
0–1 inch	.00006	.00001	.00006	.00001
	(.00032)	(.00023)	(.00031)	(.00023)
1–2 inch	-.00113	-.00236**	-.00121	-.00237**
	(.00193)	(.00107)	(.00192)	(.00111)
2–3 inch	-.00922**	-.00031	-.00913**	-.00029
	(.00363)	(.00302)	(.00374)	(.00319)
3+ inch	.00018	.00301	.00019	.00301
	(.00822)	(.00750)	(.00813)	(.00691)
Water				
0–1 inch	.00029	-.00044	.00030	.00001
	(.00081)	(.00034)	(.00083)	(.00031)
1–2 inch	-.00172	.00012	-.00169	.00009
	(.00349)	(.00094)	(.00350)	(.00089)
2–3 inch	-.02032	-.00415	-.02357	-.00390
	(.01467)	(.00316)	(.00398)	(.00318)
3+ inch	-.02971*	-.00318	-.03010*	-.00320
	(.01591)	(.00457)	(.01698)	(.00419)
Observations	130884		130884	
Rain FE	X		X	
Year FE	X		X	
Land characteristics	X			

Note: standard errors clustered at the county-by-year level. * p < 0.1, ** p< 0.05, *** p< 0.01. 0 inches of rainfall is the omitted rainfall bin. Uphill refers to neighboring census tracts within a 2.5 km radius that have a higher mean elevation than the tract of interest. Downhill refers to neighboring census tracts in the same radius that have lower mean elevation than the tract of interest. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall bin.

B.2 Alternative Specifications

Table B.4. Main results, spell specification

Dependent Variable: Claims/policy	(1)
Land×rainfall	
Impermeable development	
1 day	.00003
	(.00005)
2 days+	.00215**
	(.00111)
Wetlands	
1 day	-.00071
	(.00068)
2 days+	-.00314**
	(.00143)
Water	
1 day	-.00054
	(.00043)
2 days+	-.00215*
	(.00132)
Observations	313608
Rain FE	X
Year FE	X
Land characteristics	X

Note: standard errors clustered at the county-by-year level. * p < 0.1, ** p< 0.05, *** p< 0.01. 0 inches of rainfall is the omitted rainfall bin. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall spell bin.

Table B.5. Main results, value of claims

Dependent Variable: Value of claims/1000 of 2019 USD	(1)
Land×rainfall	
Impermeable development	
0–1 inch	.10
	(.08)
1–2 inch	.11**
	(.05)
2–3 inch	.17**
	(.08)
3+ inch	.19***
	(.05)
Wetlands	
0–1 inch	-.01
	(.08)
1–2 inch	-.11
	(.07)
2–3 inch	-.27**
	(.10)
3+ inch	-.19
	(.13)
Water	
0–1 inch	-.03
	(.06)
1–2 inch	.05
	(.06)
2–3 inch	-.21**
	(.11)
3+ inch	-.08
	(.08)
Observations	313608
Rain FE	X
Year FE	X
Land characteristics	X
Policies	X

Note: standard errors clustered at the county-by-year level. * p < 0.1, ** p< 0.05, *** p< 0.01. 0 inches of rainfall is the omitted rainfall bin. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall bin. Claims are reported at the monthly level.

Table B.6. Main results, non-coastal

Dependent Variable: Claims/policy	(1)
Land×rainfall	
Impermeable development	
0–1 inch	.00001
	(.00003)
1–2 inch	.00143**
	(.00070)
2–3 inch	.00415**
	(.00202)
3+ inch	.01106***
	(.00510)
Wetlands	
0–1 inch	.00014
	(.00022)
1–2 inch	-.00028
	(.00019)
2–3 inch	-.00215***
	(.00071)
3+ inch	-.00414
	(.00498)
Water	
0–1 inch	.00001
	(.00015)
1–2 inch	.00062
	(.00088)
2–3 inch	-.00310**
	(.00139)
3+ inch	-.00837
	(.00713)
Observations	241860
Rain FE	X
Year FE	X
Land characteristics	X
Policies	X

Note: standard errors clustered at the county-by-year level. * p < 0.1, ** p< 0.05, *** p< 0.01. 0 inches of rainfall is the omitted rainfall bin. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall bin.

Table B.7. Spillovers with 5 km radius

Dependent Variable: Claims/policy	(1)	
Land×rainfall		
Impermeable development	Uphill	Downhill
0–1 inch	.00009	.00013
	(.00013)	(.00018)
1–2 inch	.00031	-.00010
	(.00039)	(.00072)
2–3 inch	.00515*	-.00057
	(.00311)	(.00134)
3+ inch	.00573**	.00231
	(.00291)	(.00234)
Wetlands		
0–1 inch	.00006	.00002
	(.00036)	(.00031)
1–2 inch	-.00099	-.00198*
	(.00134)	(.00118)
2–3 inch	-.00777*	-.00333
	(.00469)	(.00484)
3+ inch	.00019	.00017
	(.00921)	(.00580)
Water		
0–1 inch	.00018	-.00032
	(.00089)	(.00044)
1–2 inch	-.00098	.00010
	(.00345)	(.00089)
2–3 inch	-.01978	-.001231
	(.01849)	(.00312)
3+ inch	-.03010*	-.00266
	(.01603)	(.00548)
Observations	231060	
Rain FE	X	
Year FE	X	
Land characteristics	X	
Policies	X	

Note: standard errors clustered at the county-by-year level. * p < 0.1, ** p< 0.05, *** p< 0.01. 0 inches of rainfall is the omitted rainfall bin. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall bin.

Table B.8. Spillovers with 15 km radius

Dependent Variable: Claims/policy	(1)	
Land×rainfall		
Impermeable development	Uphill	Downhill
0–1 inch	.00006	.00003
	(.00045)	(.00034)
1–2 inch	.00021	-.00005
	(.00389)	(.00693)
2–3 inch	.00384	.00012
	(.00413)	(.00185)
3+ inch	.00414*	.00238
	(.00222)	(.00308)
Wetlands		
0–1 inch	.00000	.00001
	(.001235)	(.00683)
1–2 inch	-.00023	-.00099
	(.00264)	(.00193)
2–3 inch	-.00458	-.00224
	(.00316)	(.00475)
3+ inch	.00012	.00018
	(.01143)	(.00876)
Water		
0–1 inch	.00019	-.00012
	(.00092)	(.00076)
1–2 inch	-.00042	-.00009
	(.00462)	(.00098)
2–3 inch	-.01482	-.00095
	(.02154)	(.00852)
3+ inch	-.02875*	-.00194
	(.01578)	(.01296)
Observations	282060	
Rain FE	X	
Year FE	X	
Land characteristics	X	
Policies	X	

Note: standard errors clustered at the county-by-year level. * p < 0.1, ** p< 0.05, *** p< 0.01. 0 inches of rainfall is the omitted rainfall bin. Land×rainfall refers to each category of land cover (impermeable development, wetlands, water) interacted with each rainfall bin.

B.3 Supplementary Figures

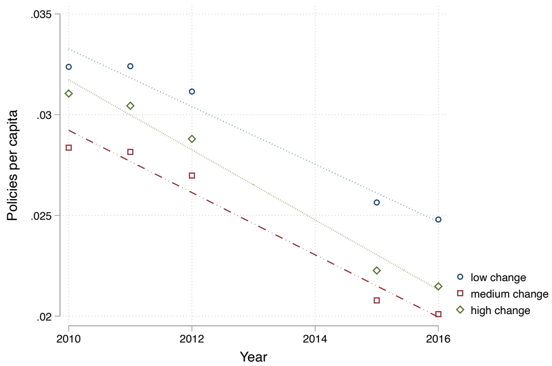

(a) Δ Development

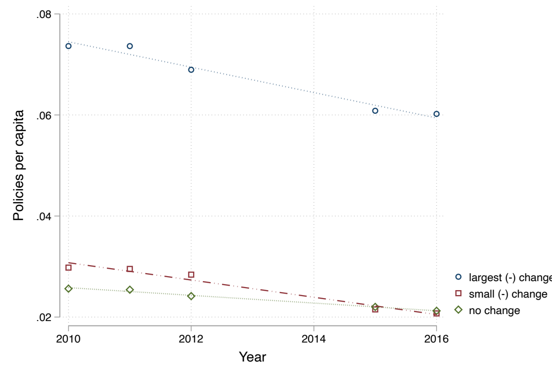

(b) Δ Wetlands

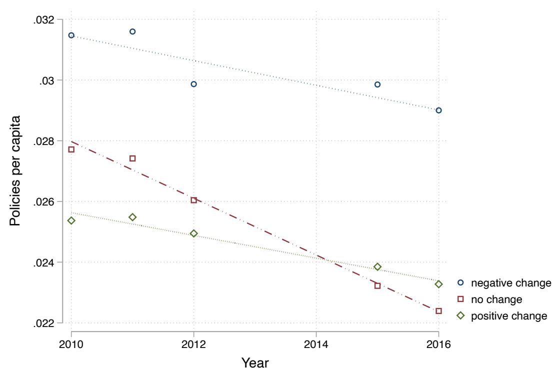

(c) Δ Water

Figure B.1. Trends in policies in force by land change

Note: evolution of policies by splitting Census tracts into positive, negative, and no change. The downward trend is consistent with drop-off in salience of flood risk as homeowners let policies lapse when not exposed to a storm event (e.g. hurricane).

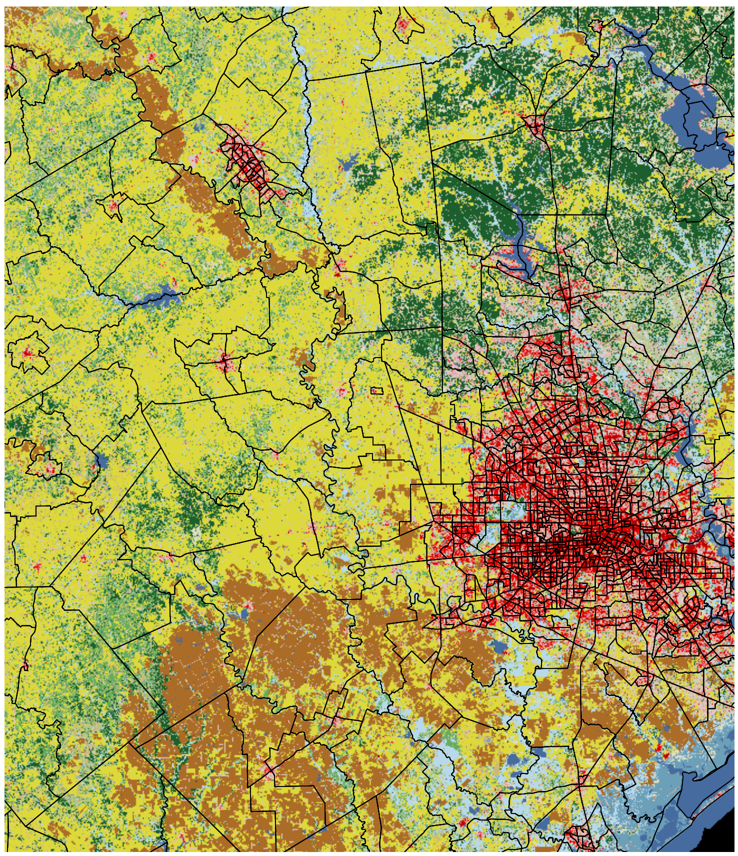

Figure B.2. Houston land use 2011

Note: this shows the grid of land use pixels for Houston county in 2011. Land use pixels are reported by the National Land Cover Dataset. Black lines show census tract boundaries.

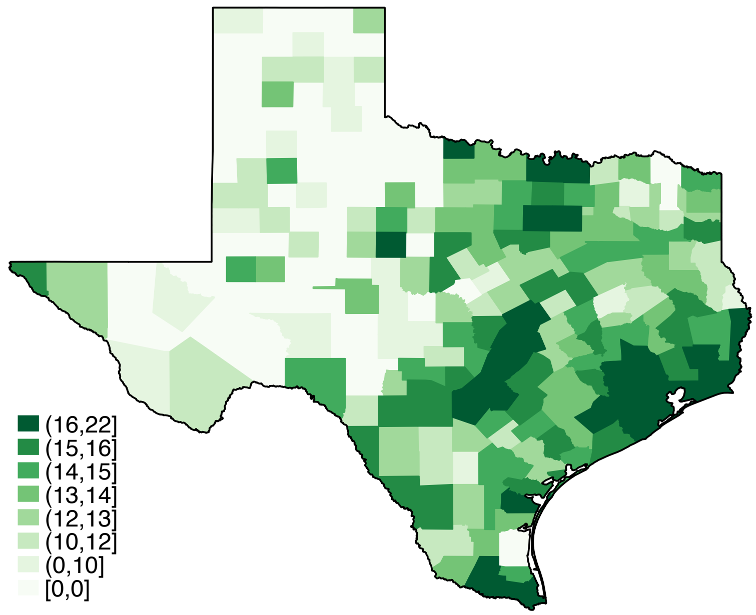

Figure B.3. County level payouts for Texas, 2010–2016

Note: this shows the natural log of county level flood insurance payouts in the state of Texas from 2010–2016.

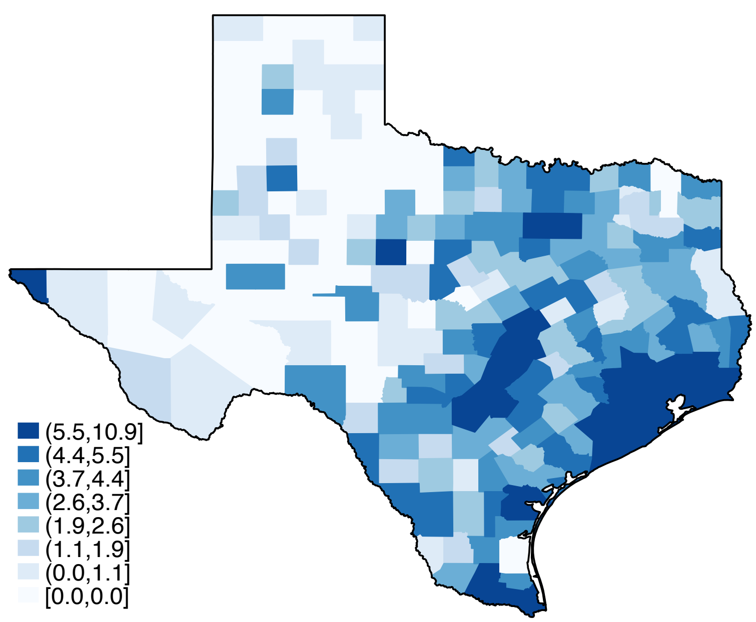

Figure B.4. County level claims for Texas, 2010–2016

Note: this shows the natural log of county level flood insurance claims in the state of Texas from 2010–2016.

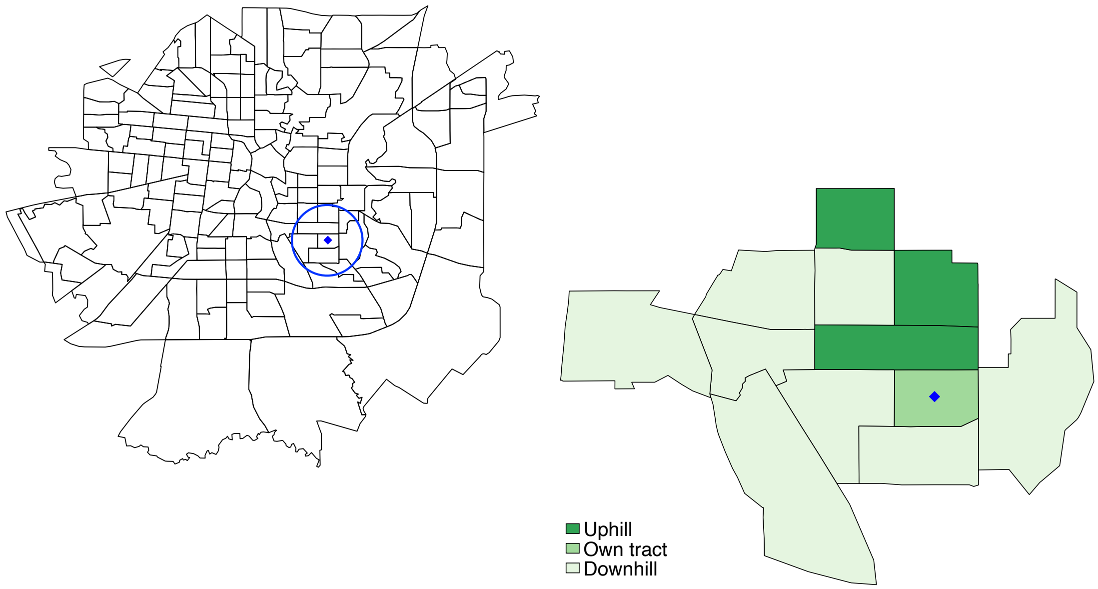

Figure B.5. Construction of tract neighbors

Note: example of neighbor construction. This shows a census tract in Bexar County, Texas. Neighbors are defined as uphill if they fall within a certain radius of the target census tract and have a mean elevation than the target tract and defined as downhill if they have a mean elevation lower than the target tract.

CRS PREMIUM DISCOUNTS

Class	Discount	Class	Discount
1	45%	6	20%
2	40%	7	15%
3	35%	8	10%
4	30%	9	5%
5	25%	10	---

SFHA (Zones A, AE, A1-A30, V, V1-V30, AO, and AH): Discount varies depending on class.

SFHA (Zones A99, AR, AR/A, AR/AE, AR/A1-A30, AR/AH, and AR/AO): 10% discount for Classes 1-6; 5% discount for Classes 7-9.*

Non-SFHA (Zones B, C, X, D): 10% discount for Classes 1-6; 5% discount for Classes 7-9.

*In determining CRS Premium Discounts, all AR and A99 zones are treated as non-SFHAs.

Figure B.6. CRS schedule

Note: this shows the schedule of NFIP discounts based on CRS participation.

B.4 Note on NFIP Data Retrieval

Information for policies in force also exist at the ZIP-code level in an unredacted form, and were retrieved through FEMA FOIA Case Number 2020-FEFO-00734. These data have not been used because of a reporting issue with the dataset.

Chapter 3 Appendix

C.1 Alternative Baseline Specifications

Table C.1. Restricting the sample by birthyear

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)
	Indonesia			**Mexico**		
Δ Mean Temp	-3.74†		-3.90††	-1.22††		-1.25††
	(.63)		(.66)	(.23)		(.23)
Δ Std. Dev. Temp		7.02*	8.59*		-0.31	-0.56
		(3.29)	(3.24)		(.52)	(.54)
Δ Mean Precip	-0.34*		-0.28*	-0.57		-3.59**
	(.13)		(.14)	(.99)		(1.18)
Δ Std. Dev. Precip		-0.47	-0.27		1.66**	2.89**
		(.13)	(.32)		(.57)	(.71)
Observations	6987	6987	6987	6723	6723	6723

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014) and MxFLS (2006–2012). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.2. Binarized measure of risk aversion

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)
	Indonesia			**Mexico**		
Δ Mean Temp	-1.09††		1.14†	-0.29***		-0.30***
	(.15)		(.79)	(.07)		(.07)
Δ Std. Dev. Temp		-0.64	0.79		-0.02	-0.08
		(.70)	(.69)		(.15)	(.15)
Δ Mean Precip	-0.04		-0.02	-0.50		-1.34***
	(.03)		(.03)	(.28)		(.35)
Δ Std. Dev. Precip		-0.14	-0.12		0.28	0.76***
		(.07)	(.08)		(.17)	(.21)
Observations	16267	16267	16267	8126	8126	8126

Note: *risk aversion*: 1 or 0, 1 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014) and MxFLS (2006–2012). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.3. Ordered probit with two-way fixed effects

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)
	Indonesia			**Mexico**		
Δ Mean Temp	-1.81††		-2.06††	-0.54***		-0.55***
	(.24)		(.28)	(.12)		(.12)
Δ Std. Dev. Temp		0.94	3.49†		-0.19	-0.31
		(1.15)	(1.19)		(.24)	(.25)
Δ Mean Precip	-0.13**		-0.11*	-0.52		-1.99***
	(.04)		(.05)	(.46)		(.57)
Δ Std. Dev. Precip		-0.22	-0.13		0.62*	1.32***
		(.13)	(.14)		(.27)	(.34)
Observations	16267	16267	16267	8126	8126	8126

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014) and MxFLS (2006–2012). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.4. Results for province/state of residence

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)
	Indonesia			**Mexico**		
Δ Mean Temp	-4.85*		-5.30*	-0.12		-0.12
	(1.69)		(2.15)	(.07)		(.08)
Δ Std. Dev. Temp		-2.55	5.21		0.09	-0.02
		(7.99)	(7.99)		(.17)	(.19)
Δ Mean Precip	-0.20		-0.15	-0.23		-3.45***
	(.20)		(.23)	(.21)		(.77)
Δ Std. Dev. Precip		-0.40	-0.29		0.16	2.59***
		(.99)	(1.06)		(.18)	(.60)
Observations	14223	14223	14223	8268	8268	8268

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014) and MxFLS (2006–2012). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.5. Alternative weather station specification

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)
		Indonesia	
Δ Mean Temp	-1.22**		-1.79***
	(.40)		(.40)
Δ Std. Dev. Temp		7.48†	8.30†
		(1.16)	(1.28)
Observations	16267	16267	16267

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014). Temperature (°C) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) inflation included in all regressions. Standard errors clustered at the cohort by province of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.6. Restricted weather series specification

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)
		Indonesia	
Δ Mean Temp	-3.58††		-4.22††
	(.45)		(.53)
Δ Std. Dev. Temp		3.16	8.98***
		(2.16)	(2.27)
Δ Mean Precip	-0.26*		-0.21*
	(.09)		(.10)
Δ Std. Dev. Precip		-0.42	-0.27
		(.23)	(.24)
Observations	16267	16267	16267

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) inflation included in all regressions. Standard errors clustered at the cohort by province of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.7. Province/state-level clustering

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)
		Indonesia			**Mexico**	
Δ Mean Temp	-2.75*		-4.23*	-1.16*		-1.19*
	(1.42)		(1.82)	(.43)		(.41)
Δ Std. Dev. Temp		1.55	6.82		-0.10	-0.35
		(7.23)	(6.98)		(1.46)	(1.57)
Δ Mean Precip	-0.25		-0.21	-1.14		-3.99
	(.17)		(.22)	(1.46)		(2.20)
Δ Std. Dev. Precip		-0.44	-0.27		1.17	2.58
		(.99)	(1.07)		(1.58)	(1.96)
Observations	16269	16269	16269	8126	8126	8126

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014) and MxFLS (2006–2012). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.8. Repeated cross-section specification

Dep. Var: Δ Risk Aversion	(1)	(2)	(3)	(4)	(5)	(6)
	Indonesia			**Mexico**		
Δ Mean Temp	-2.59††		-2.41†	-0.12**		-0.11*
	(.32)		(.30)	(.04)		(.04)
Δ Std. Dev. Temp		4.45***	3.27*		0.19	0.10
		(1.27)	(1.18)		(.13)	(.13)
Δ Mean Precip	0.08		0.12*	-0.01		-0.57*
	(.04)		(.05)	(.11)		(.22)
Δ Std. Dev. Precip		-0.06	-0.16		0.07	0.37*
		(.10)	(.11)		(.06)	(.13)
Observations	51876	51876	51876	20851	20851	20851

Note: measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014) and MxFLS (2006–2012). Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.9. Structural specification, Indonesia

Dep. Var: Δ θ	(1)	(2)	(3)
		Indonesia	
Δ Mean Temp	-2.73*		-2.35*
	(.80)		(.82)
Δ Std. Dev. Temp		9.06*	6.53*
		(3.26)	(3.28)
Observations	16267	16267	16267

Note: *risk aversion* is the structurally calculated θ, where θ is the CRRA risk parameter. Dependent variables: within-subject changes in measured risk aversion in IFLS (2007–2014). Temperature (°C) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) inflation included in all regressions. This shows the main specification, where we use narrow bracketing over per capita consumption with tight bound assumptions. Gamble averse individuals are included. Standard errors clustered at the cohort by province of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

Table C.10. Structural specification, Mexico

Dep. Var: Δ θ	(1)	(2)	(3)
		Mexico	
Δ Mean Precip	-1.63		-5.55**
	(1.27)		(1.51)
Δ Std. Dev. Precip		1.58*	3.54**
		(.75)	(.91)
Observations	8126	8126	8126

Note: *risk aversion* is the structurally calculated θ, where θ is the CRRA risk parameter. Dependent variables: within-subject changes in measured risk aversion in MXLS (2006–2012). Temperature (°C) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey. Independent variables: changes in lifetime moments between waves of respective survey. Province (Indonesia) inflation included in all regressions. This shows the main specification, where we use narrow bracketing over per capita consumption with tight bound assumptions. Standard errors clustered at the cohort by province of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

C.2 Raw Climate Data

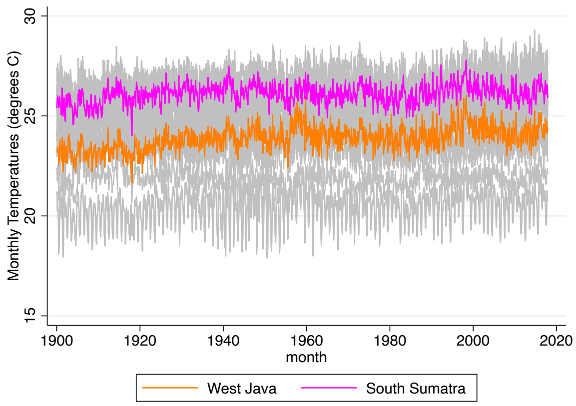

(a) Temperature, Indonesia

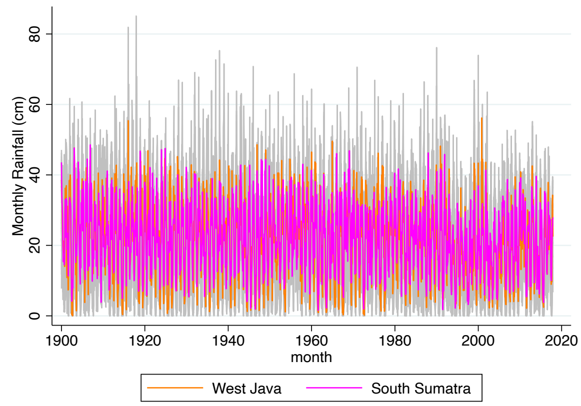

(b) Precipitation, Indonesia

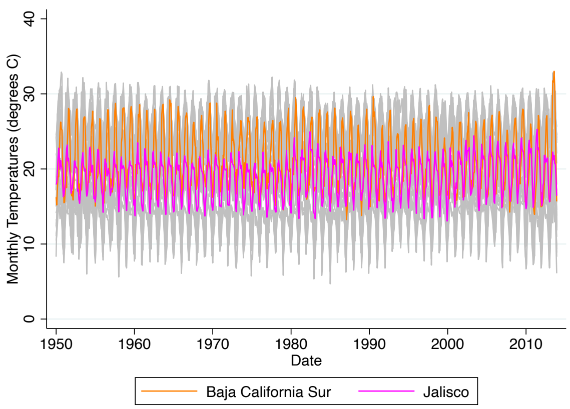

(c) Temperature, Mexico

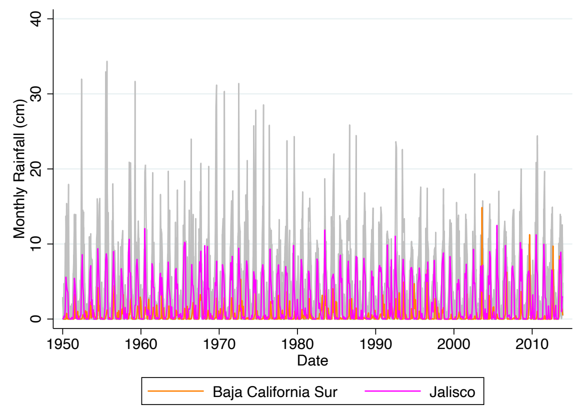

(d) Precipitation, Mexico

Figure C.1. Province/state level time series of climate variables

Note: this figure displays the monthly temperature and precipitation time series for all 25 Indonesian provinces (1993 definitions) and 32 Mexican states in our data. As can be seen these time series exhibit substantial variation both in the cross section and over time.

C.3 Construction of Risk Aversion Measures

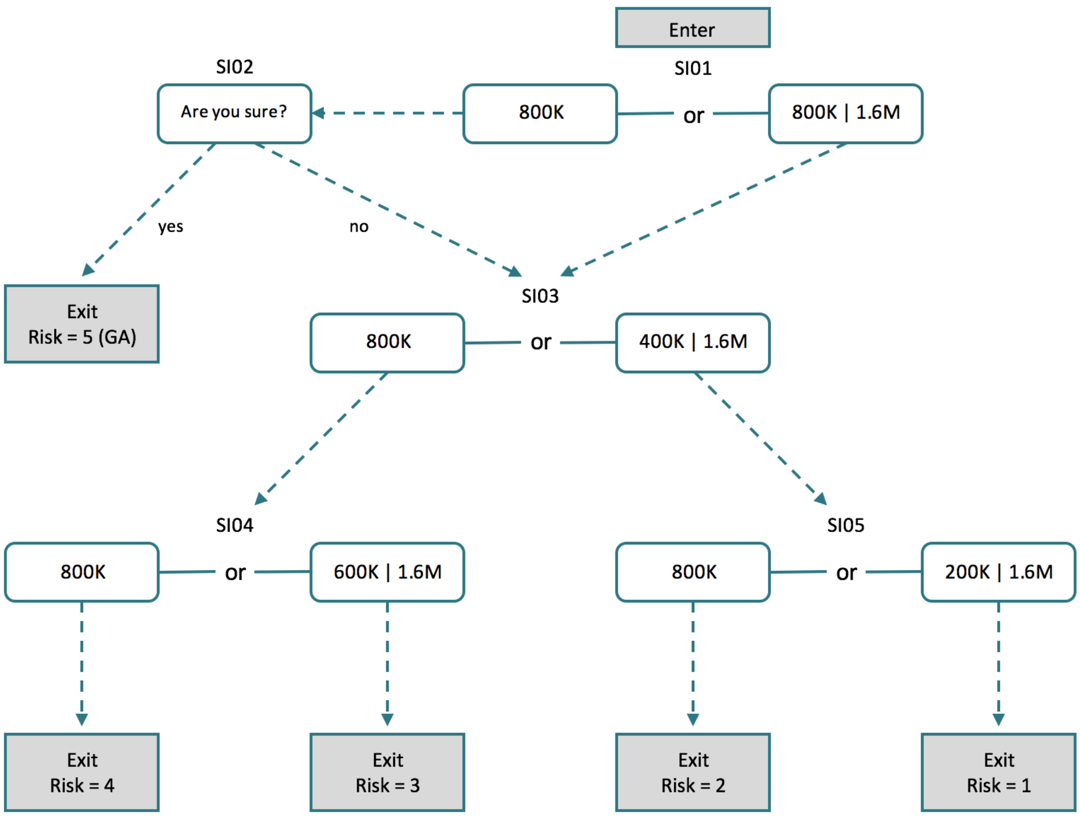

Figure C.2. Construction of risk aversion measure in IFLS2 and IFLS3

Note: higher numbers for "Risk" indicate a higher rate of measured risk aversion. Values are in Indonesian Rupiah.

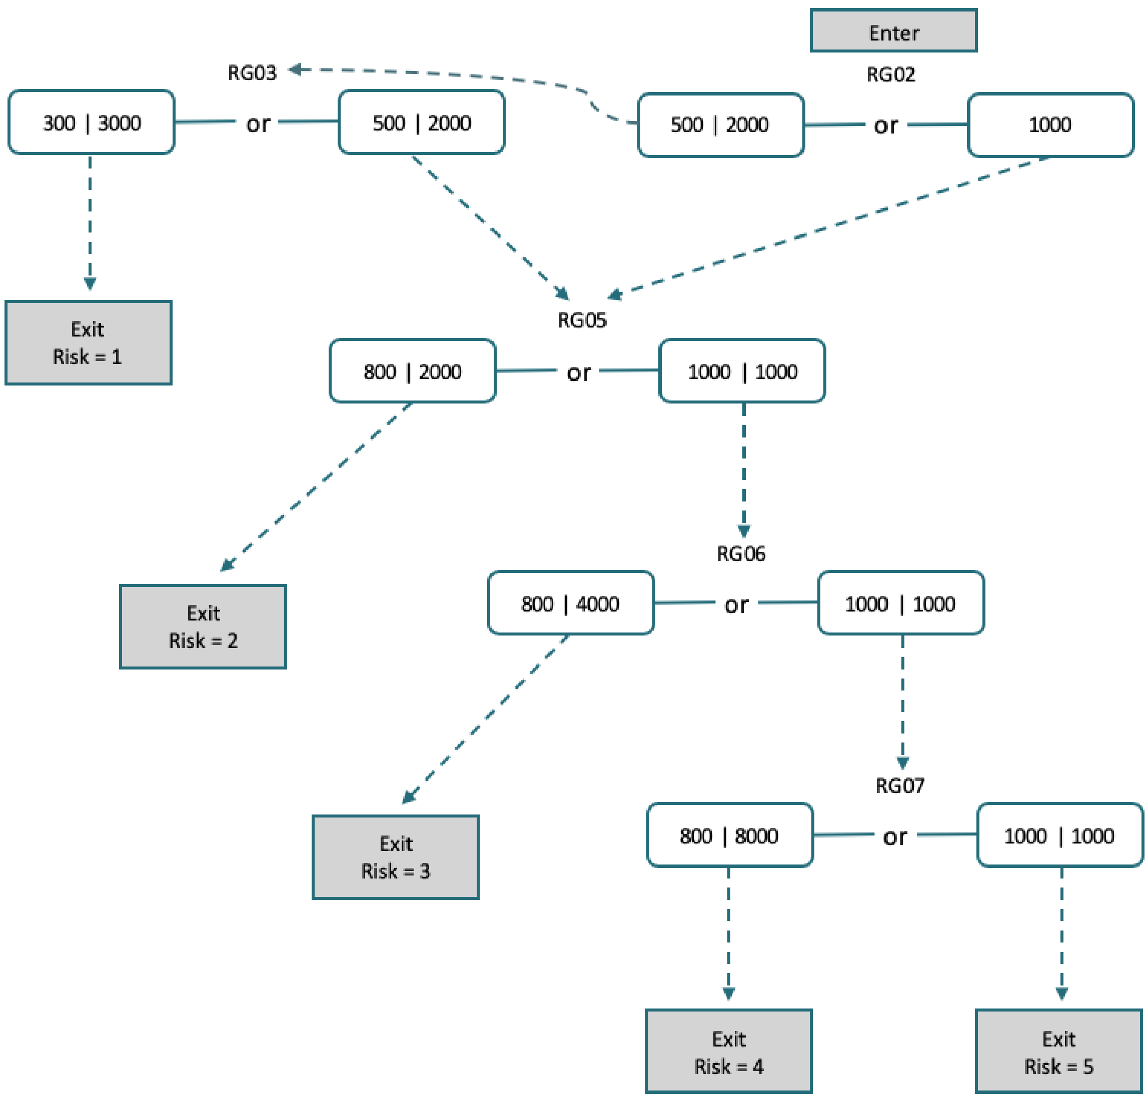

Figure C.3. Construction of risk aversion measure in MxFLS-2

Note: higher numbers for "Risk" indicate a higher rate of measured risk aversion. Values are in Mexican Pesos.

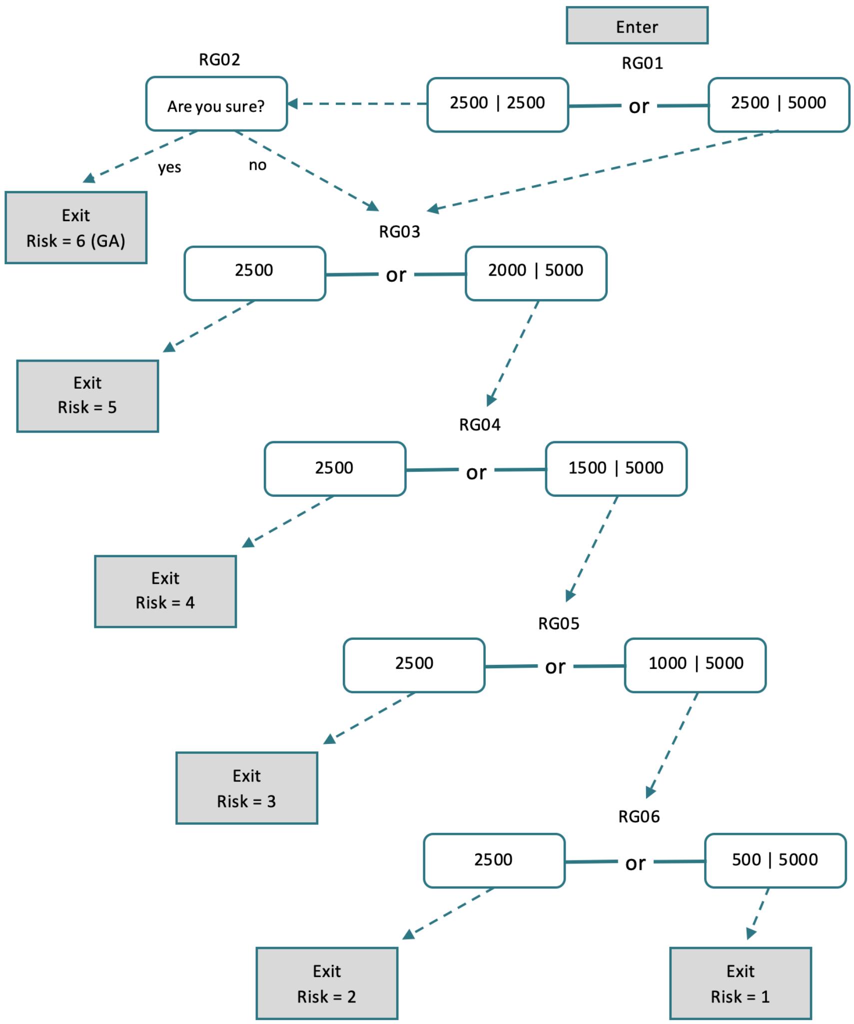

Figure C.4. Construction of risk aversion measure in MxFLS3

Note: higher numbers for "Risk" indicate a higher rate of measured risk aversion. Values are in Mexican Pesos.

C.4 Distribution of Structural Risk Parameters

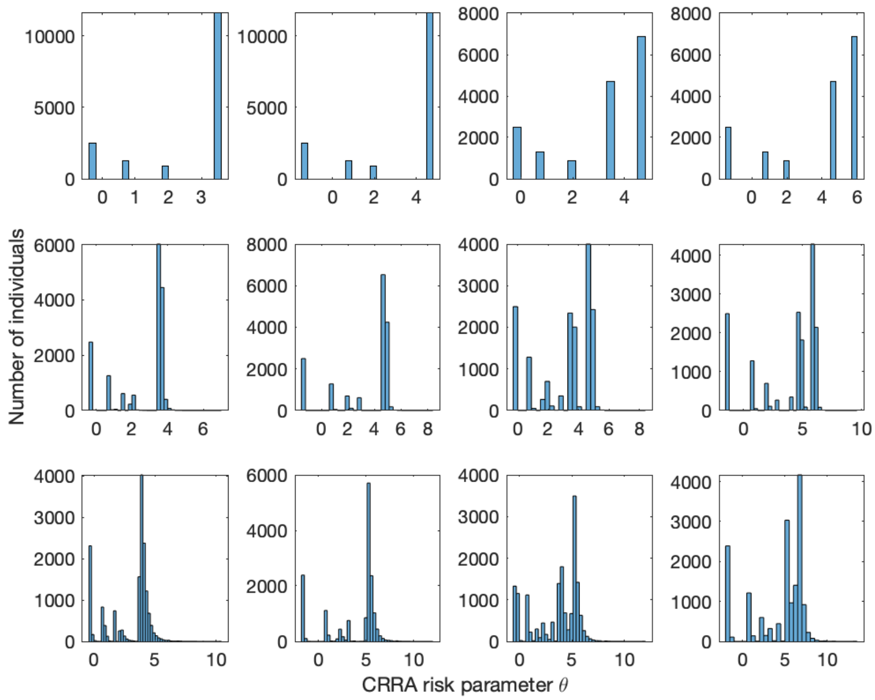

Figure C.5. Structural risk parameters, IFLS4

Note: distribution of the structurally-recovered risk parameters for the welfare exercise. Each row represents narrow, broad-bracketing over per capita consumption, and broad-bracketing over household consumption respectively. Each column corresponds to 1 of 4 bounding assumptions.

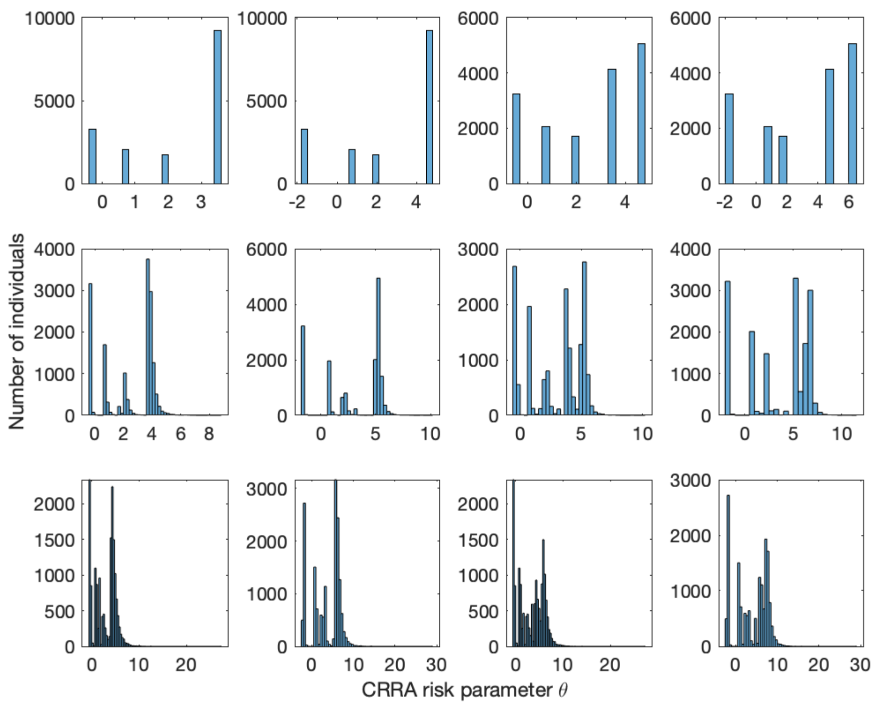

Figure C.6. Structural risk parameters, IFLS4

Note: distribution of the structurally-recovered risk parameters for the welfare exercise. Each row represents narrow, broad-bracketing over per capita consumption, and broad-bracketing over household consumption respectively. Each column corresponds to 1 of 4 bounding assumptions.

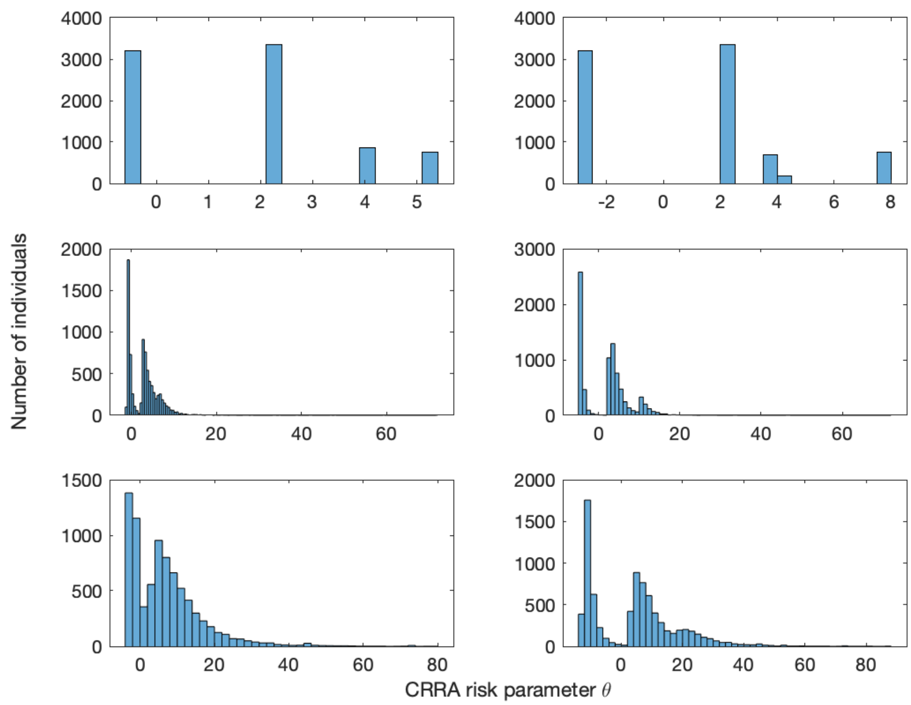

Figure C.7. Structural risk parameters, MxFLS-2

Note: distribution of the structurally-recovered risk parameters for the welfare exercise. Each row represents narrow, broad-bracketing over per capita consumption, and broad-bracketing over household consumption respectively. Each column corresponds to 1 of 2 bounding assumptions.

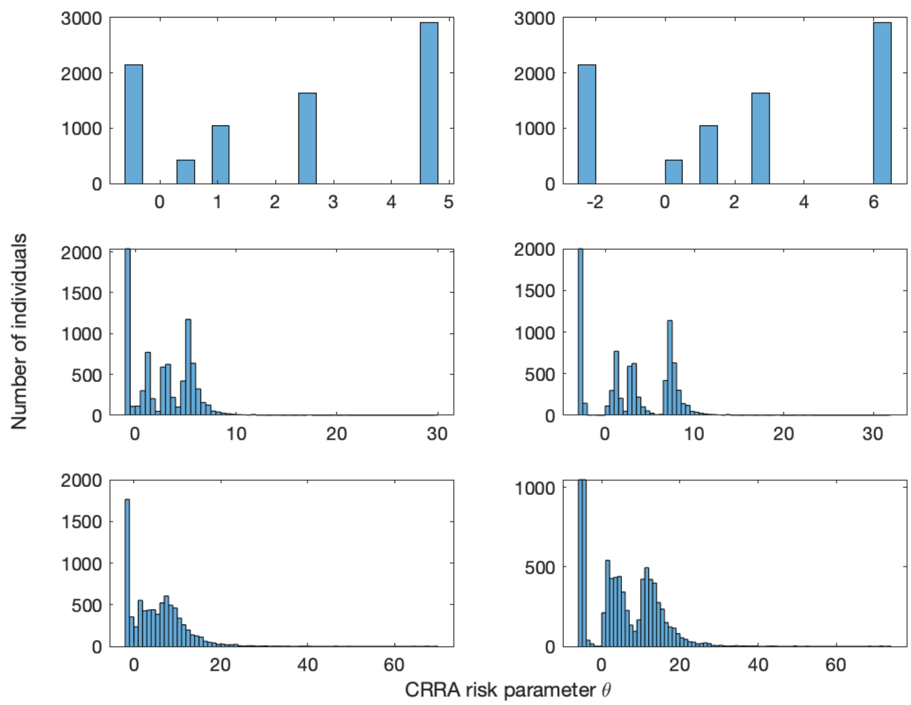

Figure C.8. Structural risk parameters, MxFLS-3

Note: distribution of the structurally-recovered risk parameters for the welfare exercise. Each row represents narrow, broad-bracketing over per capita consumption, and broad-bracketing over household consumption respectively. Each column corresponds to 1 of 2 bounding assumptions.

C.5 Geographic Distribution of Survey Samples in the Data

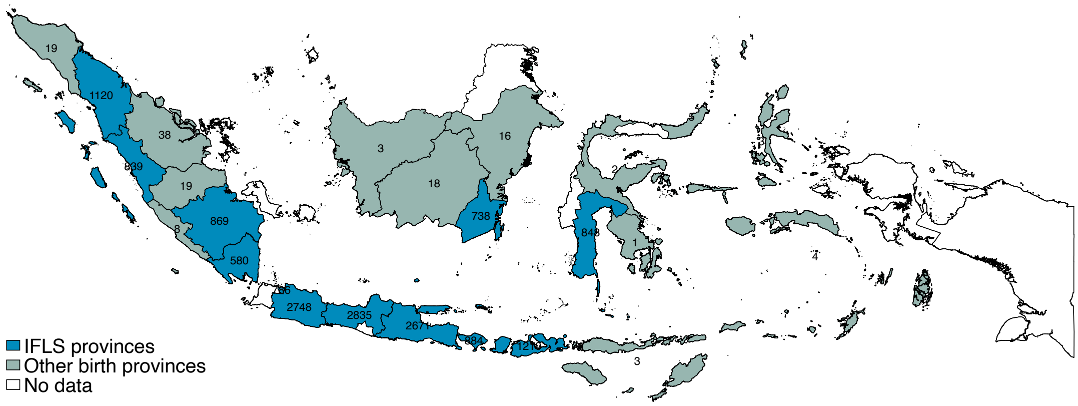

Figure C.9. Distribution of the primary sample in Indonesia by province of birth

Note: provinces in blue are ones in which the IFLS has been deployed. Provinces in green are non-IFLS provinces in which some subjects in our primary sample were born.

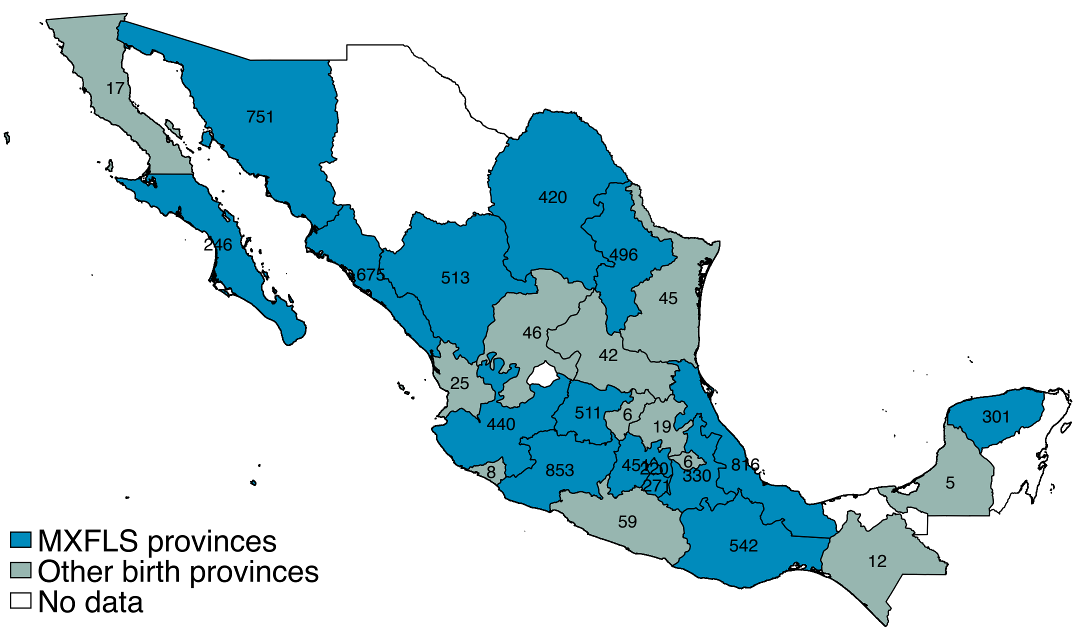

Figure C.10. Distribution of the primary sample in Mexico by province of birth

Note: states in blue are ones in which the MxFLS has been deployed. States in green are non-MxFLS states in which some subjects in our primary sample were born.

C.6 Summary Statistics

The summary statistics for both the IFLS and the MxFLS are reported in Table C.11.

Table C.11. Summary statistics for the sample mean

	Indonesia		**Mexico**	
Sample:	Primary sample	Full sample	Primary sample	Full sample
Measured Risk aversion	3.55	3.52	2.67	2.41
Woman	0.55	0.53	0.58	0.59
Age	40.35	37.32	42.44	42.02
Married	0.89	0.81	0.67	0.65
Household Size	5.23	5.21	5.62	5.65
Comp. Elementary	0.42	0.35	0.51	0.50
Comp. Junior High	0.19	0.20	0.25	0.25
Comp. High School	0.27	0.32	0.13	0.13
Above High School	0.12	0.14	0.11	0.11
Self-employed	0.42	0.39	0.23	0.22
Currently smoke	0.32	0.32	0.08	0.08
Ever migrated	0.13	0.17	0.15	0.15
Income/month	10.60	8.868	4,922	4,816
Consumption/month	3.02	2.45	3,746	3,636
Savings	8.10	8.76	10,248	9,522
Borrowing	3.47	2.68	13,058	12,620
Observations	32534	55820	16252	25005

C.7 Correlates of Risk Aversion Measures in the Cross Section

Table C.12. Correlates of risk preference measures

	Indonesia		**Mexico**	
Dep. Var:	Measured Risk Aversion		Measured Risk Aversion	
Sample:	X-Sec	Panel	X-Sec	Panel
Self-employed	-0.11***	-0.10***	0.01	0.03
	(0.018)	(0.021)	(0.03)	(0.04)
Migrated	-0.10***	-0.08**	0.02	0.03
	(0.023)	(0.033)	(0.034)	(0.039)
Income	1.52e-06***	1.75e-06***	0.07	0.08*
	(3.39e-07)	(3.86e-07)	(0.05)	(0.04)
Consumption	-0.015***	-0.018***	-0.16	-0.05
	(0.004)	(0.005)	(0.24)	(0.26)
Total assets	-3.47e-05	-3.39e-05	0.009	0.01
	(3.04e-05)	(3.82e-05)	(0.01)	(0.01)
Borrowing	-0.001**	-0.001**	-0.07	-0.1
	(0.0004)	(0.0005)	(0.21)	(0.23)
Savings	-0.0003	-0.0002	0.06	0.18
	(0.0002)	(0.0003)	(0.32)	(0.33)
Smoker	0.09***	0.07*	-0.17***	-0.15***
	(0.0 30)	(0.038)	(0.05)	(0.06)
Cigs/day	-0.06***	-0.04**	0.001	0.001
	(0.02)	(0.02)	(0.0007)	(0.0008)
Woman	0.28***	0.26***	0.04	0.02
	(0.023)	(0.028)	(0.03)	(0.03)
Age	-0.015***	-0.014***	-0.012**	-0.012**
	(0.004)	(0.005)	(0.005)	(0.006)
Age^2	0.002***	0.002***	0.0001**	0.0001**
	(4.25e-05)	(5.64e-05)	(5.55e-05)	(6.22e-05)
Observations	35848	23995	11740	9335
R-squared	0.052	0.055	0.18	0.168

Note: coefficients from regressions of dependent variables on all covariates. Monthly income and consumption. Income, consumption, assets, borrowing, and savings at household level. Standard errors clustered at the cohort by province of birth in parenthesis. Observations are at the individual by year level. Controls: Time FE, Province FE, HH size, marital status education dummies, and religiosity dummies (religiosity dummies only for Indonesia). Monetary variables in millions of rupiah and pesos. *** p<0.01, ** p<0.05, * p<0.1. "X-SEC" refers to subjects appearing in at least one wave; "Panel" refers to those who appear in both. Note that the sample size for this analysis is smaller than in the baseline results, due to missing data in variables of interest for some subjects.

C.8 Numerically Calculating the Equally-Distributed Equivalent

For an individual $j \in J$, let $s_j \in S$ be an individual realization of the outcome variable of interest. Given an individual CRRA parameter θ_j, we want to find the certainty equivalent for individual j over the distribution of S in society. Let individuals be indexed by i. Then individual $j \in J$ has a certainty equivalent given by the following:

$$CE_j = \left(\frac{1}{N} \sum_{i \in I} s_i^{1-\theta_j} \right)^{\frac{1}{1-\theta_j}}. \tag{11}$$

Repeat this calculation for each individual in J. Using the resulting distribution of certainty equivalents, replace the distribution of S with these certainty equivalents. We iterate on the process until converging to an equally-distributed equivalent. The numerical procedure is as follows:

1. For the initial distribution of outcomes and specifications, define $S_{N \times N_S}$, where N is the number of individuals and N_S is the number of specifications. CE' is undefined.

2. Define $S \equiv CE'$ if CE' has been defined.

3. Define $CE_{N \times N_S}$ as the certainty-equivalent matrix over S, where individual values of θ_j are calculated as defined in Equation 11.

4. Define the new matrix $S' \equiv CE$.

5. Calculate the new certainty equivalent CE' over the distribution of outcomes in S'.

6. While the difference between the max and min cells of CE' for a particular specification exceeds a given tolerance, iterate over steps 2–5.

7. Each specification converges to a uniform column of certainty equivalents (Eden, 2020). We interpret this as the equally-distributed equivalent for the distribution of

per capita or household income.

This procedure gives us a single statistic for each specification. The interpretation of each statistic is an aggregate measure of how much society is willing to trade *ex ante* uncertain consumption for guaranteed consumption.

C.9 Details of Additional Controls

Table C.13. Description of controls included in Table 3.2

Category	Variables Included
Demographics (**Indonesia** and **Mexico**)	Married Household Size Household Size Squared
Income, Assets, Savings, Consumption (**Indonesia** and **Mexico**)	Total household income Total value of household assets Net Households Savings (Savings-Borrowing) Total household consumption (yearly)
Violence (**Indonesia**)	Perceived safety level of village Perceived safety of walking in village alone at night Civil strife in HH region, last 5 years Civil strife severe enough to cause death, major injury, direct financial loss, or relocation of any member of HH
Violence (**Mexico**)	Perceived safety level of village Feels safe at home Fear of assault during the day Fear of assault at night No. of times robbed, assaulted, kidnapped Family/friend robbed, assaulted, kidnapped in last 12 month
Natural Disasters (**Indonesia**)	Natural disaster in HH region, last 5 years Natural disaster severe enough to cause death, major injury, direct financial loss, or relocation of any member of HH
Natural Disasters (**Mexico**)	Household/business lost due to natural disaster
Growth Experiences (**Indonesia** and **Mexico**)	Province/state-level GDP growth

C.10 Sample Distribution for Risk Aversion Measures

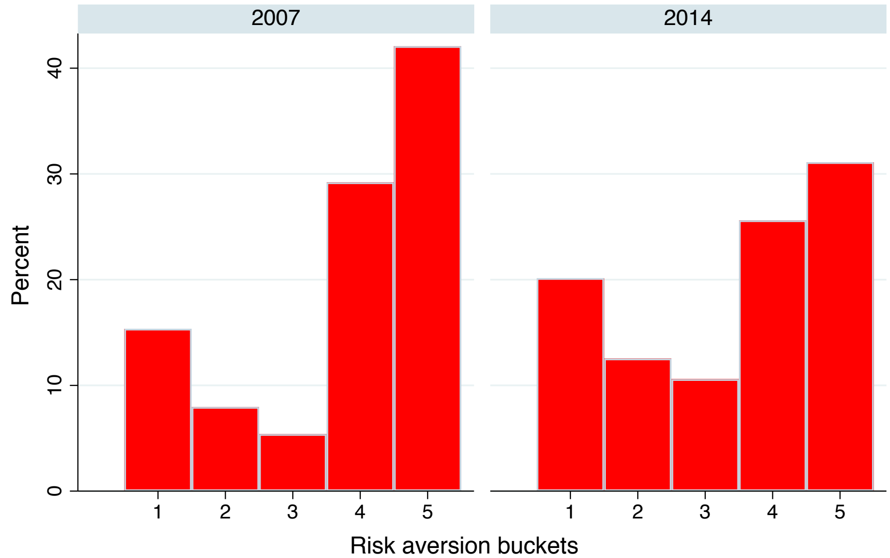

Figure C.11. Histogram of measured risk aversion buckets in IFLS4 and IFLS5

Note: measured risk aversion reported from 1–5, 5 being the highest measured risk aversion. Distributions for individuals in main regressions: present in both 2007 and 2014 surveys.

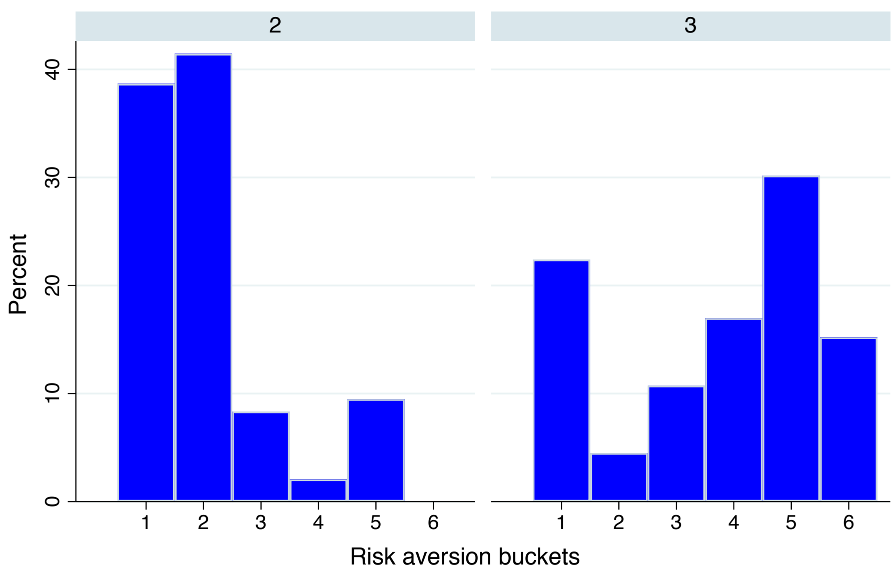

Figure C.12. Histogram of measured risk aversion buckets in MxFLS-2 and MxFLS-3

Note: measured risk aversion reported from 1–6, 6 being highest measured risk aversion. Distributions for individuals in main regressions: present in both 2005 and 2009 surveys. Individuals in bucket 6 in Mexico are not included due to inconsistencies in survey design across waves.