# The Global Impacts of Climate Change on Risk Preferences

Wesley Howden
*Wesleyan University*

Remy Levin
*University of Connecticut*

Meng Song*
*University of Connecticut*

January 13, 2025

*Please click **here** for the most recent version*

We study the direct impacts that long-run experiences of climate change have on individual risk preferences. Using panel surveys from Indonesia and Mexico (total N = 25,000), we link within-person changes in elicited risk preferences to state-level, lifetime experiences of climate change. In line with the predictions of a Bayesian model of learning over background climate risk, we find that in both settings increases in the experienced means of temperature and precipitation cause significant decreases in measured risk aversion, while increases in the experienced variance of temperature in Indonesia and the variance of precipitation in Mexico lead to significant increases in measured risk aversion. We replicate this analysis globally using a survey with a representative sample from 75 countries (N = 75,000) containing an elicited measure of risk preference which we link to country-level, lifetime climate experiences. We find significant results for both the means and variances of both climate variables that are consistent with our panel analyses. Across all settings, experiences of climate variance have first-order effects, with coefficient magnitudes of the standard deviation of climate 0.6-2.6 times that of the climate mean. We develop a novel method for estimating the welfare effects of observed risk preference changes using panel data, and find that the climate-induced changes in risk preferences we observe increased welfare in both Indonesia and Mexico by approximately 1%.

**Keywords**: Risk preferences, climate change, experience effects, volatility, welfare
*JEL* Codes: D14, D81, D83, I31, O12, Q54

*Howden: Department of Economics, Wesleyan University. Email: whowden@wesleyan.edu. Levin: Department of Economics, University of Connecticut. Email: remy.levin@uconn.edu. Song: Department of Economics, University of Connecticut. Email: meng.song@uconn.edu. Thanks to Mark Jacobsen, Judd Boomhower, Richard Carson, Josh Graff Zivin, Jeff Shrader, Derek Lemoine, and seminar participants at the University of Arizona AREC, University of Arizona Economics, UC San Diego Environmental Lunch, and Columbia University Workshop in Sustainable Development for discussion and insightful comments. All errors are our own.

# 1 Introduction

Climate change is one of the most important factors affecting the collective well-being of humanity in the 21st century. Numerous studies in the literature have documented its impacts on a variety of economic outcomes. Most of these studies have focused on physical outcomes, such as mortality (**??**), productivity (**??**), income (**??**), and conflict (**???**). A smaller subset of work has examined the impacts of climate change on harder-to-measure cognitive and psychological outcomes, such as human capital (**?**), mental health (**?**), and subjective affect (**?**). In this paper we contribute to this growing literature by studying the direct impacts that long-run experiences of climate change have on individual risk preferences.

Understanding how climate change shapes risk preferences is important for two reasons. First, risk preferences play a central role in a wide variety of economic models of behavior, including (but not limited to) investment decisions, technological adoption, occupational choice, insurance decisions, migration, and health behaviors. As such, understanding the direct impacts of climate change on risk preferences themselves is important for understanding and predicting the behavioral response to climate change across many domains. Second, risk preferences are an important determinant of the optimal policy response to climate change. Traditional integrated assessment models (IAMs), from which the social cost of carbon is conventionally estimated, include societal risk preference over macroeconomic uncertainty as a key parameter (**?**), a role that is only reinforced by recent work that incorporates uncertainty about the evolution of climate change into IAMs (**??**). Since societal risk preference is an aggregate of individual risk preferences, our findings suggest the existence of a heretofore unappreciated causal channel from climate change to the social cost of carbon, via the dynamics of risk-taking psychology.

We begin by building a decision-theoretic model of the impacts of long-run climate experiences on financial risk aversion (section 2). To do this, we extend the risk-taking adaptation model of **?** to the domain of climate experiences. In our model, an expected utility maximizer is faced each period with a single risky choice from a menu of objective income lotteries known as the foreground risk. The agent makes this choice in the presence of an exogenous and unavoidable additive background climate risk that their utility function as a consumption good. Standard assumptions about the agent's higher-order risk attitudes imply that the background climate risk is a substitute for the foreground income risk, so that the more risk exists in the environment, the less risk the agent wants to take in their individual choices. The agent perceives the background risks to be a stationary Gaussian random variable, with unknown mean and unknown variance. We assume the agent is a Bayesian who learns from personal experiences. As climate realizes over the agent's life, their beliefs about its moments update, which leads to the agent's preferences over the foreground income risk to adapt in turn. Our model predicts that increases in the mean of the background climate risk will decrease the agent's foreground risk aversion; that increases in its variance will increase foreground risk aversion; and that the two moment effects will be additive. In other words, agents that experience hotter or wetter conditions relative to their body of experiences will become more risk-seeking, while agents that experience destabilizing weather shocks will become more risk averse.

Having established these theoretical predictions, we next turn to testing them empirically. We use data from panel surveys in Indonesia and Mexico (total $N = 25{,}000$), each containing two elicited measures of risk aversion for the same subjects years apart.[1] In our main panel analysis we regress within-person changes in measured risk aversion on changes in the means and standard deviations of temperature and precipitation in subjects' state of birth, from birth to survey measurement. Our empirical approach allows us to exploit the significant cross-sectional and temporal 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 our model's predictions, we find (subsubsection 4.1.1) that in both countries increases in the experienced lifetime means of both temperature and precipitation cause significant decreases in measured risk aversion. We also find that increases in the experienced standard deviation of temperature in Indonesia and in the standard deviation of precipitation in Mexico lead to significant increases in measured risk aversion. Notably, the estimated magnitudes of the standard deviation 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.

We conduct three additional empirical analyses with our panel data, aimed at understanding the mechanisms underlying our main results. First, we progressively add additional controls into our main regressions, to see whether the results change when measures of other background risks are included (subsubsection 4.1.2). These controls include time-varying individual demographics, changes in household economic circumstances, and measures of behaviorally-relevant experiences, such as exposure to violence, natural disasters, and state-

[1] Our panel survey data come from waves 4 (2007) and 5 (2014) of the Indonesian Family Life Survey, and waves 2 (2006) and 3 (2009) of the Mexican Family Life Survey. Measures of elicited risk aversion in our panel data are estimated from a series of choices over high-stakes, hypothetical, objective income lotteries. These measures allow us to sidestep identification challenges that exist in our setting with other kinds of real-world risky choice data, most notably the potential confounding of estimated preference changes with foreground belief changes, and the potential endogeneity of the foreground choice menu with subjects' lifetime experiences. Additional details on our data, methodology, and identification are available in section 3.

level, lifetime real GDP growth experiences, a direct measure of background income risk. Adding these controls has negligible effects on our estimates, with the exception of lifetime growth experiences, whose inclusion significantly increases the magnitudes of our climate risk estimates in both countries. This suggests that our main results are not driven by spurious correlations between climate risk and other background risk–if anything, the existence of correlated economic risks attenuates our main estimates.

Second, we examine the comparative effects of our main results for relevant subgroups of our subject population (subsubsection 4.1.3), by splitting the sample into binary groups across six observable subject characteristics: occupation in agriculture, urbanicity, consumption level, educational attainment, gender, and age. In both countries we find stronger effects for older subjects, those with lower levels of education, and those employed in agriculture. We also document stronger marginal effects of climate experiences for women in Indonesia and for men in Mexico.

Third, we explore whether the changes we observe in measured risk aversion induced by climate experiences correspond with changes in risk-taking behavior in other domains (subsubsection 4.1.4). We construct a variable measuring predicted change in risk aversion using our main specifications, and examine its correlation with changes in smoking, having ever migrated across province or state lines, self-employment status, and, in Indonesia, the planting of cash crops, a measure of risky agricultural investment. We find that subjects that are predicted to become more risk tolerant significantly increase their plantings of cash crops in Indonesia, and significantly increase their smoking in Mexico.

In addition to our within-country panel analyses, we study the impacts of climate change on risk preferences globally (subsection 4.2). We use data from the Global Preference Survey, a module attached to the 2012 wave of the Gallup World Poll, which collected data for a sample of 75,000 subjects from 75 countries spanning a wide range of climate and economic conditions. Our data contain an experimentally-validated measure of individual risk preferences,[2] which we regress on the country-level lifetime means and standard deviations of temperature and precipitation for each subject, as well as a set of individual and country climate controls. We find results that are highly consistent with our panel analyses: globally, individuals who experience higher mean temperature and precipitation over their lifetimes are significantly less risk averse, while those who experience higher standard deviations of temperature and precipitation are significantly more risk averse. These estimates remain significant (and become considerably larger) when country fixed effects are included in the

[2]The GPS risk aversion measure is a combination of a staircase risk measure using choices over hypothetical, high-stakes income lotteries (very similar to the instruments in the IFLS and the MXFLS), and a self-reported risk aversion Likert scale. For more details on the data, see section 3.

regressions. As in our panel analyses, the marginal effects of climate variance on risk aversion are first-order, with a standard deviation to mean coefficient ratio of 1.6-2.6. Notably, in both our within-country and global analyses, estimated results are strongest where the mean and standard deviation of temperature are strongly negatively correlated, and the mean and standard deviation of precipitation are strongly positively correlated. This indicates that the underlying climatic distributions are skewed (positively for temperature, negatively for precipitation), and suggests that the behavioral effects we observe are driven by unexpected lifetime heatwaves and drought experiences.

In the final part of the paper, we turn to the question of whether the risk preference changes we observe are welfare increasing or decreasing (section 5). To make progress on this question, we develop a new method for estimating the welfare impacts of changes in risk preferences that are driven by changes in subjects' environments, using panel data. Our method is premised on taking the measured risk preferences in our data at face value, and on recent advances in welfare economics that provide a way to calculate welfare measures for groups of risk preference heterogeneous agents given known distributions of preferences and consumption (**?**). Using our method, we derive estimates for the welfare effects of risk preference changes due to all causes, and for the effects of risk preference changes driven by climate change specifically. We find that in Indonesia risk preference changes due to all causes increased group-level welfare by 6.1%; that in Mexico, risk preference changes due to all causes decreased group-level welfare by 8.2%; and that across both settings, risk preference changes due to climate change increased welfare by approximately 1%. This last result suggests that the changes in risk preferences we observe represents a form of psychological climate adaptation.

This paper contributes to several areas of the literature. Most directly, this paper contributes to the climate impacts literature cited above. To our knowledge, this is the first paper in the literature to credibly demonstrate that climate change directly impacts individual economic preferences globally. Since every physical climate impact previously studied almost surely includes individual risk preferences as a key mediator of the observed outcome, our findings should be of use to researchers who seek to understand the mechanisms underlying the observed effects, and to predict novel impacts of climate change. Our findings also bolster recent work that seeks to incorporate climate uncertainty into estimates of the social costs of carbon (**?**) by providing micro evidence that this uncertainty is a first-order driver of the behavioral response to climate change.

This paper is also closely related to the literature on the effects of lifetime experiences on economic preferences, particularly the subset of the literature that studies the effects of natural disasters (**?????**). We contribute to this literature by providing evidence on long-run experiences of climate change (rather than one off disasters); by building a theoretical framework which allows the disentangling of mechanisms underlying the effects, in contrast to all previous papers which are purely empirical; and by estimating the welfare consequences of the observed preference changes. Notably, ours is the first paper in the broader experience effects literature to combine multiple sub-national panel analyses with evidence from a large global cross-section using individual measures of economic preferences.

## 2 Model

Our aim in this paper is to explore the direct impacts that long-run experiences of climate change have on individual risk preferences over monetary lotteries. As a starting point for understanding this relationship theoretically we use the risk preference adaptation model of **?**. The fundamental message of this theoretical framework is that optimal changes in risk preferences over endogenous income lotteries arise from a process of experiential Bayesian learning over exogenous background risk. This basic structure is useful for our purposes, because (1) modeling climate change as a unavoidable background risk maps cleanly onto the concept of a climate impact, or an effect of climate change that cannot be mitigated by conscious behavior; and (2) because the learning process in this model accommodates a first-order role for the variance of the background risk, which allows us to capture a central feature of climate change, namely that it is a process of changes in the variance, and not just the mean, of the weather realization distribution.

To extend the the model in **?** into the domain of climate change we make an important modification to it, by modelling the background risk as a lottery over a climate consumption good rather than as an income lottery. This serves two purposes. First, it allows us to accommodate multiple plausible mechanisms through which climate can affect the agent's utility. These include the impact of climate on the agent's income process, the amenity value (i.e. direct consumption value) of climate, and the potential health or physical effects of weather extremes. Second, it opens the door to modelling the climate background risk as non-monotonic or containing a bliss-point. This is a realistic assumption about the behavioral response to weather that is highlighted in the climate impacts literature (**?**, **?**), and which is far more plausible for a consumption good than for income. Although in the following model we assume a monotonic relationship between the climate good and the agent's utility, work is ongoing on understanding the response of risk preferences to a non-monotonic background risk.

**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 must choose an income lottery $\tilde{x}_t$ from a menu of lotteries $\mathcal{X}_t$. $\tilde{x}_t$, which we call the *foreground risk*, represents a risky choice over income that is under the agent's control, such as the choice between purchasing a risky stock or a risk-free bond, or the decision to plant a safe food crop or a risky cash crop. Without loss of generality we assume that $\mathcal{X}_t$ is a menu of objective lotteries.

In each period, the agent is also exposed to an exogenous and unavoidable *background* climate risk. This background risk, denoted by $\tilde{c}_t$, is a lottery over the level of climate consumption good, which can be thought of as a weather realization at time $t$, for instance in the form of temperature or precipitation. Some parameters of the climate background risk are unknown to the agent, who holds beliefs over them which they update each period as they observe realizations of $\tilde{c}_t$. Denote with $B_t(c)$ the cdf of the agent's beliefs distribution about the outcomes of $\tilde{c}_t$ at time $t$. We assume that $\tilde{x}_t$ and $B_t(c)$ are statistically independent for all $t$. We also assume that our agent exhibits hand-to-mouth consumption, so at the end of each period the realizations of $\tilde{x}_t$, $w$, and $\tilde{c}_t$ are consumed.

**Utility and risk.** The agent is a subjective expected utility maximizer. Following in the footsteps of the background risk literature, we think of the agent as possessing two utility functions. The *background utility function* $u$ is the agent's period utility defined over both foreground and background risks. $u$ takes $\tilde{c}_t$, $\tilde{x}_t$, and $w$ as its arguments, additively:

$$u(w, \tilde{x}_t, \tilde{c}_t) = u(w + \tilde{x}_t + \tilde{c}_t)$$

The *foreground utility function* $v_t$ represents the agent's utility over the foreground risk alone, conditional on their expectations about the background risk:

$$v_t(w, \tilde{x}_t) = \mathbb{E}_{\tilde{c}} u(w, \tilde{x}_t, \tilde{c}_t | B_t(c))$$

An agent with background utility function $u$, who is exposed to background risk $\tilde{c}$ (over which they hold beliefs $B_t$), behaves like an agent with foreground utility function $v_t$ who faces no background risk. $v_t$ thus encodes the agent's preferences over the fully endogenous risk $\tilde{x}_t$, i.e. the risk that is the object of their choice and therefore under their control. While $u$ is fixed, the agent's risk preference in $v_t$ may change as their beliefs about the background risk evolve over time.

Our variable of interest is the agent's Arrow-Pratt coefficient of absolute risk aversion over the foreground utility function $v_t$:[3]

$$r_t(w) = -\frac{v_t''(w)}{v_t'(w)}$$

We assume the background utility function $u$ is four-times differentiable over both consumption and income, and make four assumptions on the signs of its derivatives over consumption: monotonicity ($u' > 0$), diminishing marginal utility ($u'' < 0$), and two additional assumptions included in the following definition:

**Definition 2.1.** (*Ross risk vulnerability*) Let $\tilde{c}$ have a bounded domain on the interval $[a, b]$. Then $u$ is Ross risk vulnerable iff $\exists \lambda > 0$ such that for all $c_1, c_2 \in [a, b]$:

$$-\frac{u'''(w + c_1)}{u''(w + c_2)} \geq \lambda \geq -\frac{u''(w + c_1)}{u'(w + c_2)}$$

and

$$-\frac{u''''(w + c_1)}{u'''(w + c_2)} \geq \lambda \geq -\frac{u''(w + c_1)}{u'(w + c_2)}$$

**?** defines a stronger measure of absolute risk aversion than Arrow-Pratt, suitable for the incomplete insurance or background risk setting. Ross risk vulnerability (**?**) is a natural extension of Ross risk aversion into higher-order risk preferences on $u$. The first condition in the definition above corresponds to decreasing absolute risk aversion in the sense of Ross (i.e. $r^R(w)$ decreasing in $w$). The second condition corresponds to decreasing absolute prudence in the sense of Ross (i.e. $-u'''(w + c_1)/u''(w + c_2)$ decreases in $w$). Collectively, the conditions entail that the background risk and any risk in the foreground will be substitutes for the agent. In other words, the more exogenous climate risk exists in the environment, the less risk the agent will take on in their endogenous risky income choices.

**Learning.** We combine the above background risk framework with a model of experiential learning over the background climate risk. We assume the agent is a Bayesian who uses personally observed iid realizations of the climate risk to update their belief distribution $B_t(c)$. The agent perceives the background risk to be a stationary Gaussian random variable with unknown mean and unknown variance. For analytical tractability, we further assume that the agent's prior over the moments of the background risk takes the form of the conjugate prior of the Gaussian with unknown mean and unknown variance likelihood, which is a

[3]To make the notation less cumbersome going forward, we will write $v_t$ and $u$ without the $\tilde{x}_t$ argument, with the understanding that they nevertheless determine the choice of $\tilde{x}_t$.

normal-inverse-chi-squared distribution. We call this learning process normal mean-variance learning and describe it formally in the following definition:

**Definition 2.2.** (*Normal mean-variance learning*) We say that a Bayesian agent is a normal mean-variance learner if

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

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

where $M$ and $\Sigma^2$ are time-invariant 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 $\mu_0$ and $\sigma_0^2$ are the agent's point priors over the mean and variance of $\tilde{c}_t$, and $\kappa_0 > 0$ and $\nu_0 > 2$ are parameters capturing the agent's confidence or precision over the prior mean and variance, respectively.

Bayesian learning over the unknown mean of a stationary Gaussian with *known* variance is the benchmark learning model in the economics literature, and has been used in a variety of empirical applications. Normal mean-variance learning departs from this benchmark by allowing the agent to learn directly about the variance of the climate risk as well. This learning process has the appealing feature that it allows the agent's beliefs about the variance of climate to increase or decrease over time, unlike the mean-only model in which the variances of beliefs monotonically decrease to zero as new information arrives. Intuitively, this means that realizations of the background risks can meaningfully change the agent's beliefs about how stable the climate is, in addition to how favorable it is. This captures an important feature of climate change: the fact that it is both a change in level and in the stability of the climate system.

**Agent data.** To simplify the analysis, without loss of generality we will restrict our attention going forward to the three periods $t = 0, 1, 2$. For stating our main result it is useful to define some notation regarding the data available to the agent in different periods. Denote by $m$ the number of iid realizations of $\tilde{c}_t$ the agent observes in period 0. These make up the dataset $\mathcal{D}_1$ the agent uses to form their beliefs $B_1$. Denote by $n$ the number of realizations of $\tilde{c}_t$ the agent observes in period 1. These, together with the previous $m$ observations, compose the dataset $\mathcal{D}_2$ the agent uses to form their beliefs $B_2$.

For dataset $\mathcal{D}_t$, containing $k$ observations, the sample mean is defined as $\bar{c}_t = 1/k \sum_{i=1}^{k} c_i$ and the sample variance as $s_t^2 = 1/k \sum_{i=1}^{k}(c_i - \bar{c}_t)^2$. Thus $\bar{c}_1 = 1/m \sum_{i=1}^{m} c_i$ and $\bar{c}_2 = 1/(m+n) \sum_{i=1}^{m+n} c_i$, while $s_1^2 = 1/m \sum_{i=1}^{m}(c_i - \bar{c}_1)^2$ and $s_2^2 = 1/(m+n) \sum_{i=1}^{m+n}(c_i - \bar{c}_2)^2$.

**Timing.** The agent enters each period with an income endowment and beliefs about the climate background risk. The agent then chooses the foreground risk from the menu available to them before the background risk realizes. Once the foreground risk and the background risk realize, the agent updates their beliefs about the background risk. We assume that she updates about the two moments of the background risk sequentially. This simplifying assumption implies that the change in the agent's risk preferences over the foreground risk will be additive in the changes in her beliefs over the two moments of each of the background risk. This assumption entails minimal loss of generality, because no choices in the model are made in the interim period between the updating of beliefs about the mean and the variance. At the end of the period the agent consumes whatever realizations they receive that period (in line with our hand-to-mouth assumption above), and enters the next period with the appropriate prior and income endowment.

## 2.1 Result

**?** prove the following result for the univariate model with a single source of background income risk, which we here extend to a background consumption risk:

**Proposition 1.** *Let $A, B$ be positive constants. Assume $m$ is large. Then, $\forall w$:*

$$r_2(w) - r_1(w) \approx -A(\bar{c}_2 - \bar{c}_1) + B(s_2^2 - s_1^2) \tag{1}$$

where $\bar{c}_t$ is the period $t$ sample mean of the set of background climate realization observed by the agent, and $s_t^2$ is their period $t$ sample standard deviation. Intuitively, this result states that a novel background climate shock affects the agent because it changes their beliefs about how favorable the exogenous environment is relative to their past experiences, as well as how stable it is, with the two channels having an additive effect. In other words, agents that experience exogenous climate realizations (such as temperature or precipitation) that increase the mean of their lifetime climate experience set will become more risk-seeking, while agents that experience climate realizations that increase the variance of their lifetime climate experience set will become more risk averse. The over effect of any climate experience will be the additive effect of the two moment effects, relative to the agent's preexisting body of experiences.

# 3 Data and methodology

We conduct two primary analyses to estimate the effects that lifetime experiences of climate change have on individual risk preferences. First, we regress within-person changes in measured risk aversion, obtained from panel survey data in Indonesia and Mexico, on changes in the state-level mean and standard deviation of subjects' lifetime monthly temperature and precipitation. Second, we regress measured risk aversion, from a cross-sectional survey containing elicited measures of individual risk aversion for a representative sample of subjects in 76 countries, on country-level measures of the mean and standard deviation of monthly temperature and precipitation over those subjects' lives.

## 3.1 Panel survey data in Indonesia and Mexico

For the Indonesian analysis our source of micro data is the Indonesian Family Life Survey (IFLS) (**??**). The IFLS is a longitudinal study administered by the RAND corporation in 13 states in Indonesia in five waves, starting in 1993.[4] 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 the IFLS and the MXFLS 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 income, assets, savings and borrowing. Both surveys also contain residence and migration histories, allowing us to link place-based variables like sub-national GDP growth to subjects, as well as measures of risk aversion, which we discuss in the next section. We restrict attention to subjects who appear in both of the latest waves of the two surveys, and for whom responses to the risk aversion measure are recorded in both. This results in an initial sample (which we call the "full sample") of 17,183 subjects in Indonesia and 12,152 subjects in Mexico, each observed twice.

[4] Our data in Indonesia is at the province (*provinsi*) level, an administrative unit roughly equivalent in size to a US or Mexican state. To simplify our exposition we refer to Indonesian provinces as Indonesian states throughout the paper.

Table 1: GPS survey sample by country climate and income per capita

| | Hot (45) | Cold (31) |
|---|---|---|
| Wet (17) | (15) Bangladesh, Brazil, Cambodia, Cameroon, Colombia, Costa Rica, Guatemala, Haiti, Indonesia, Nicaragua, Philippines, Sri Lanka, Suriname, Thailand, Vietnam | (2) Japan, South Korea |
| Intermediate (31) | (14) Argentina, Australia, Bolivia, Ghana, India, Kenya, Malawi, Mexico, Nigeria, Portugal, Rwanda, Tanzania, Uganda, Venezuela | (17) Austria, Bosnia Herzegovina, Canada, China, Croatia, Czech Republic, Estonia, France, Georgia, Germany, Italy, Lithuania, Netherlands, Sweden, Switzerland, UK, USA |
| Dry (28) | (16) Algeria, Botswana, Egypt, Greece, Iran, Iraq, Israel, Jordan, Morocco, Pakistan, Peru, Saudi Arabia, South Africa, Spain, UAE, Zimbabwe | (12) Afghanistan, Chile, Finland, Hungary, Kazakhstan, Moldova, Turkey, Poland, Romania, Russia, Serbia, Ukraine |

**Notes**: Temperature split is based on the mean country temperature for 2012 being above/below the NOAA reported world 2012 mean temperature of 14.6°c. Precipitation split is based on the country's mean monthly precipitation tercile for 2008-2012 based on the UDel data. Colors indicate country's tercile of GDP per capita in 2012 based on data from the Penn World Tables, with low-income countries in red, middle-income countries in purple, and high income countries in blue.

## 3.2 Global survey data

Our global data on risk preferences comes from the Global Preference Survey (GPS) (**??**), a module attached to the 2012 wave of the Gallup World Poll (GWP). The 2012 GWP collected data for approximately 1,000 subjects from each of 76 countries, with a sample representative of the resident population aged 15 and older in each country. The countries included in the 2012 GWP sample collectively represent about 90% of global population and global GDP. Conveniently for our purposes, this sample of countries encompasses a wide array of climatic and economic development conditions (Table 1).

The GPS contains experimentally-validated measures of six economic preferences (risk aversion, patience, trust, altruism, positive reciprocity, and negative reciprocity). The public release of the GPS contains information on subject demographics (age, gender, language of survey, a math skills measure, country, sub-national region, and interview month), as well as individual sampling weights provided by Gallup that adjust for probability of selection into the survey and for gender, age, and education or socioeconomic status where available. Further details on the risk preference measures in the GPS, which we use to construct our dependent variable, are in subsection 3.3.

## 3.3 Risk aversion measures

The two most recent waves of the IFLS (IFLS4 (2007 - 2008) and IFLS5 (2014)) and the MXFLS (MXFLS2 (2005-2006) and MXFLS3 (2009-2012)) include modules for measuring financial risk aversion, which we use to construct the main dependent variables in our analysis. In both surveys, the risk aversion measurement modules employ staircase instruments, which have been shown to generate high-quality measures of risk preference with low subject response burden (**?**), making them ideal for field applications. In a staircase risk aversion instrument subjects are presented with 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. Risky lotteries are commonly in the form of fair coin flips, with known probabilities. 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 risk aversion for each subject.

In IFLS4 and IFLS5 subjects answered between two and three questions each, which resulted in a measure with five bins. Each question offered the same sure amount of money (800,000 Indonesian Rupiah), while the amounts of the risky lottery varied between questions. The same module structure and payment amounts were used in both waves of the survey. We code the resulting measure by numbering the bins, with higher numbers indicating a higher degree of risk aversion (Figure 4). The dependent variable in our Indonesian analysis is within-person changes in the measured risk aversion bin between IFLS4 and IFLS5. We present histograms of the measured risk aversion buckets in IFLS4 and IFLS5 in Figure 8.

One complicating factor with the IFLS risk aversion module is that the first question in both waves offered subjects a choice between a sure payment amount and a coin flip over the same amount and a higher payment. A significant fraction of subjects in the sample are gamble averse–that is, they chose the certain option, which is first-order stochastically dominated, even after being prompted to reconsider a second time.[5] Gamble aversion in our setting may represent an extreme form of risk aversion, or it may be driven by other factors, with the most likely culprits being subject misunderstanding, inattentiveness, or religiosity (due to the prohibition on gambling in Islam, for instance). To deal with this issue, and

[5]The gamble aversion (GA) questions were deployed as part of the risk aversion measurement module in both IFLS4 and IFLS5, as well as in MXFLS3 (see Appendix A). Of the full sample of 17,183 subjects in Indonesia, 42.1% were GA in IFLS4, 31% in IFLS5, 14.8% were GA in both waves, and 58.2% were GA in at least one wave. Of the full sample of 12,152 subjects in Mexico, 14.1% were GA in MXFLS3.

separate the two groups of gamble averse subjects, we use the procedure outlined in **?** to exclude gamble averse subjects from the sample who are risk-seeking in a second instrument in the IFLS. This reduces the number of subjects from 16,942 to 15,044 or about 11% of the sample.

In MXFLS2 subjects answered between two and five questions each, which resulted in a measure with seven bins. Most questions offered subjects a choice between a sure amount of money (1,000 Mexican pesos) and a riskier coin flip, with the amounts of the riskier coin flips generally changing between questions. Some questions in the MXFLS2 instrument offered subjects a choice between two coin flips, with the riskier lottery having a higher mean and a higher variance. We code the resulting measure by numbering the bins, with higher numbers indicating a higher degree of risk aversion (Figure 5). To maintain consistency with MXFLS3, we combine subjects in the three highest risk-seeking categories (bins 5, 6, and 7) into a single group, yielding a measure with five bins.[6]

In MXFLS3 subjects answered between two and six questions each, resulting again in a measure with seven bins. All questions offered a choice between a sure amount of money (2,500 Mexican pesos in most of them), while the amounts of the risky lottery varied between questions. A "gamble averse" option was also offered in this instrument. We code the resulting measure by numbering the bins, with higher numbers indicating a higher degree of risk aversion, and the top two bins indicating gamble aversion (Figure 6). We drop subjects in the two gamble averse bins from the analysis, resulting in a measure with five bins. Histograms of the measured risk aversion buckets in MXFLS2 and MXFLS3 can be found in Figure 8. The dependent variable in our Mexican analysis is within-person changes in the measured risk aversion bin between MXFLS2 and MXFLS3.

The risk aversion measure in the GPS is a combination of two measures, one quantitative (a staircase risk aversion instrument) and one qualitative (self-reported risk aversion). The GPS uses a staircase with five steps, resulting in a measure with 32 bins (Figure 7). Subjects in the GPS also answered a question indicating their self-assessed propensity to take risks, captured on an 11-point Likert scale, another common field instrument for measuring risk aversion. Scores for the two measures were then averaged, using weights derived from an experimental validation procedure,[7] to yield an overall risk aversion index for each subject. We normalize the GPS risk aversion index so that higher values indicate more risk aversion

[6]If we did not combine these bins in MXFLS2, then the most risk-seeking subjects in bins 6 and 7 would mechanically appear to become more risk averse in the analysis, because the maximum value that their risk aversion can take in MXFLS3 is 5.

[7]The experimentally-derived weights for the risk aversion instrument placed roughly even weight on the two measures. The resulting index was further validated by examining its correlation with measures of risk-taking behavior in the GWP (**?**).

and so that its values are all positive. This results in a continuous measure of risk aversion that runs from 0 to 4, which we use as the dependent variable in our global analysis.

### 3.4 Climate change experience variables

To construct our main independent variables for analyses, we begin by constructing state-month time series for both temperature and precipitation in Indonesia and Mexico, and country-month time series for these variables for the global analysis. 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 (**??**).[8] 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 state-month series of temperature and precipitation, we take the average of all pixels that fall within a state boundary (**?**).[9]

In Mexico, we use data from the gridded weather product for the Continental US and North America created by **?**. 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 a state-by-day panel by averaging over the pixel-by-day observations that fall within a specific state boundary. With these state-by-day observations, we construct a state-by-month level time series by averaging the state-day values of temperature and precipitation, following a similar procedure to **?**.

For the global analysis, we use monthly-level air surface temperature (°C) and cumulative precipitation (cm/month) gridded at a 0.5° resolution available from the University of Delaware Terrestrial Air Temperature and Terrestrial Precipitation (UDEL) from 1914–2012 (**?**).[11] We

[8] 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.

[9] In an alternative specification, we use the universe of available ground station data, reported by the National Oceanic and Atmospheric Administration (**?**). In this station-based analysis, we construct monthly temperature series for 1976 onwards. Following concerns about entry and exit of weather stations (**?**), 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.[10] 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 state-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 state-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 **??**.

[11] UDel_AirT_Precip data provided by the NOAA/OAR/ESRL PSL, Boulder, Colorado, USA, from their Web site at https://psl.noaa.gov/data/gridded/data.UDel_AirT_Precip.html.

combine these with a year 2000 global population raster in order to weight the climate grid by population location within each country boundary at a 2.5 minute resolution (**?**).[12] This generates, for each country in our sample, a monthly population-weighted observation of temperature (°C) and cumulative precipitation (cm/month).[13]

Having constructed state-month and country-month time series for climatic variables, we match them to subjects in our data by their state or country and year of birth. Subjects in our IFLS and MXFLS samples are matched with a time series for their state of birth starting in January of the year following their birth year.[14][15] Subjects in our GPS sample are matched with time series for their country of residence starting in the January following their birth 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, a subject in our IFLS sample born in East Java in 1981, for instance, will be assigned the statistics for the East Java temperature and precipitation time series from January 1982 to December 2007 (the year of IFLS4) and from January 1982 to 2014 (the year of IFLS5).[16] Similarly, a subject in our GPS sample born in Afghanistan in 1985 will be assigned the statistics for the Afghanistan time series from January 1986 to December 2012. Let $c_{is}$ be the climatic variable assigned to person $i$ in year $s$ (with $c \in \{\mathbf{T}\text{emperature}, \mathbf{P}\text{recipitation}\}$). Then for month of measurement $t$ (with $t$ corresponding to the number of months from January of the subject's birth year to December of the year of measurement in the respective wave of the IFLS, MXFLS, and GPS) these statistics are:

[12] Since the climate and population grids do not perfectly overlap, we resample the population raster over the UDEL weather raster.

[13] As a point of reference, the population-weighted temperature for Mexico in November 2017 is 17.4°C , with the simple mean being 19.1°C.

[14] We use climate conditions in the state of birth, rather than state of residence, to control for endogenous migration in our baseline specification. Our results are robust to using lifetime climate conditions in state of residence.

[15] Subjects born before the first year for which data are available are assigned time series as if the first year of data availability is their birth year. Our results are robust to dropping these older subjects entirely.

[16] 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.

$$A_{it}^c = \frac{1}{t - b_i} \sum_{s=b_i+1}^{t} c_{is} \tag{2}$$

$$V_{it}^c = \sqrt{\frac{1}{t - b_i - 1} \sum_{s=b_i+1}^{t} (c_{is} - A_{it})^2} \tag{3}$$

where

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

and

$$B = \begin{cases} 1901 & \text{if } IFLS \\ 1950 & \text{if } MXFLS \\ 1914 & \text{if } GPS. \end{cases}$$

For our IFLS and MXFLS analyses we use the changes in these lifetime climate experience variables ($\Delta A_{it}^T$, $\Delta V_{it}^T$, $\Delta A_{it}^P$, and $\Delta V_{it}^P$) between the waves of the respective survey as the primary independent variables in the analyses. For our global analysis we use $A_{it}^T$, $V_{it}^T$, $A_{it}^P$, and $V_{it}^P$ as the main independent variables. We plot these lifetime climate experience variables against each other by cohort in each setting in Figure 1. Two features of the climate experience data are notable. First, a significant amount of variation exists across cohorts in all climate experience moments in all settings. This variation is our main source of identification.

Second, the first two moments of temperature are strongly negatively correlated in Indonesia and globally, while the first two moments of precipitation are strongly positively correlated in Mexico and globally. These correlations indicate that the underlying climate variable distributions in these settings are skewed. In the case of the negative temperature moment correlation, the skew is positive, implying relatively many unexpected heatwave shocks, while for the positive precipitation moment correlations the skew is negative, implying relatively many drought shocks.[17] Notably, the strongest effects in both the panel and

[17]Intuitively, a negative correlation between the mean and the variance of temperature implies that when the mean is low, variance is high, and when the mean is high, variance is low. This implies an underlying distribution where hot temperatures are associated with few low temperature shocks, while low temperatures are associated with relatively many high temperature shocks. To yield these conditions, the underlying distribution must be positively skewed, that is have a long right tail, which implies unexpected heatwaves. A

global analyses are in settings with the skewed climate distribution, suggesting that the main climatic channels for our observed behavioral effects are through unexpected heatwaves and droughts.

### 3.5 Empirical specifications

Our baseline empirical specification for the analyses using the IFLS and MXFLS is a first-difference regression:

$$\Delta R_{it} = \alpha + \beta_1 \Delta A_{it}^T + \beta_2 \Delta V_{it}^T + \beta_3 \Delta A_{it}^P + \beta_4 \Delta V_{it}^P + \gamma Inflation_p + \epsilon_{it}, \tag{4}$$

where we regress within-person changes in measured risk aversion $R_{it}$ on changes in the lifetime state-level experienced mean and standard deviation of temperature and precipitation. We include sub-national inflation occurring between the waves of the respective survey in all our regressions to reduce noise and address potential bias.[18] We regress $R_{it}$ on the temperature and precipitation statistics jointly (rather than climatic variable by climatic variable) to account for spatial correlation patterns between temperature and precipitation that could potentially generate spurious results in regressions that examine these variables separately. Standard errors $\epsilon_{it}$ are clustered at the state-of-birth by birth-year level, which is the level of treatment in our analysis.

For the global analysis using the GPS data we regress individual risk aversion $R_i$ on the mean and standard deviation of lifetime temperature and precipitation:

$$R_i = \alpha + \beta_1 A_i^T + \beta_2 V_i^T + \beta_3 A_i^P + \beta_4 V_i^P + \gamma_1 X_i + \gamma_2 X_c + \varepsilon_i. \tag{5}$$

we include controls for individual ($X_i$) and country ($X_c$) characteristics, as detailed in subsection 4.2. In some specifications we include country fixed effects, which means the results we observe rely only on between-cohort variation. Standard errors in this analysis are clustered at the sub-national region level, as recorded in the GPS.

similar argument applies to the positive moment correlation in the precipitation distribution, which implies a negative skew, or long left tail - that is, relatively many unexpected droughts.

[18]Since $R_{it}$ is measured off of nominal hypothetical lotteries, national-level inflation can introduce noise into the analysis by changing the real value of the prizes offered in the risk elicitation task between waves. Inflation can also introduce bias into our estimates if it varies substantially at the state level, and if it correlates with state-level growth and with risk aversion. To address these concerns we include a measure of sub-national inflation for administrative unit $p$ as a control variable in all baseline specifications. This takes the form of a the change in a consumer price index normalized to 100 during the year of the first wave (IFLS4 and MXFLS2) of the respective survey. In Indonesia $p$ is state, whereas in Mexico due to data constraints $p$ is region.

Figure 1: Climate change experience data

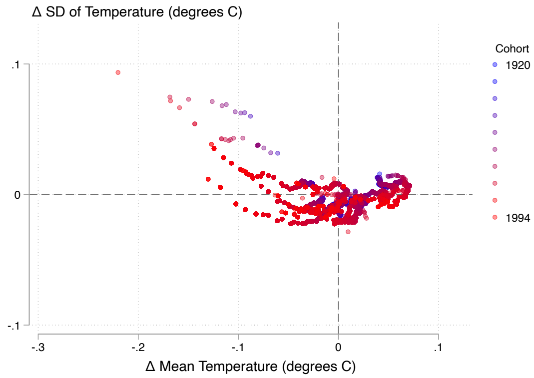


(a) Temperature – **Indonesia**

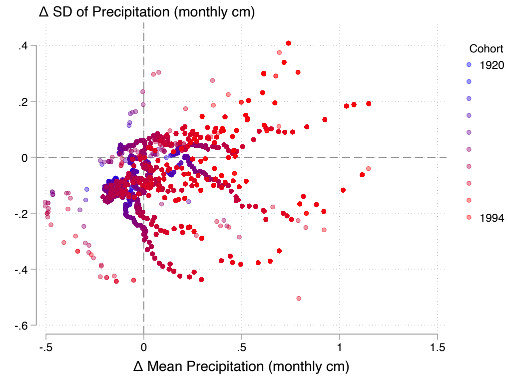


(b) Precipitation – **Indonesia**

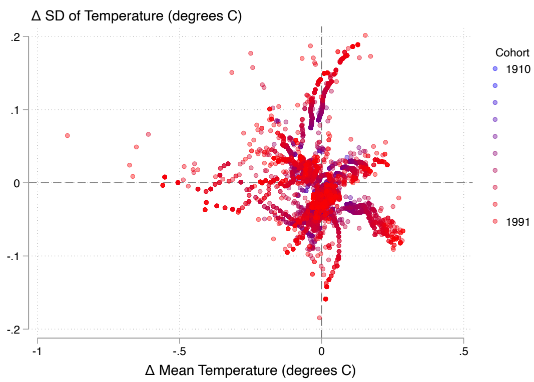


(c) Temperature – **Mexico**

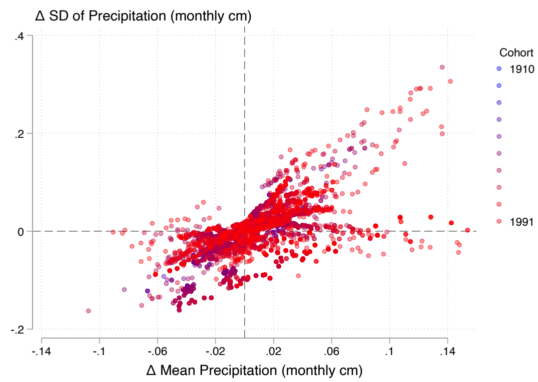


(d) Precipitation – **Mexico**

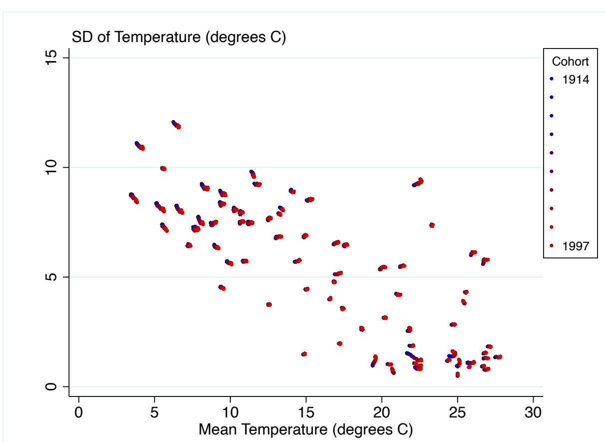


(e) Temperature – **Global**

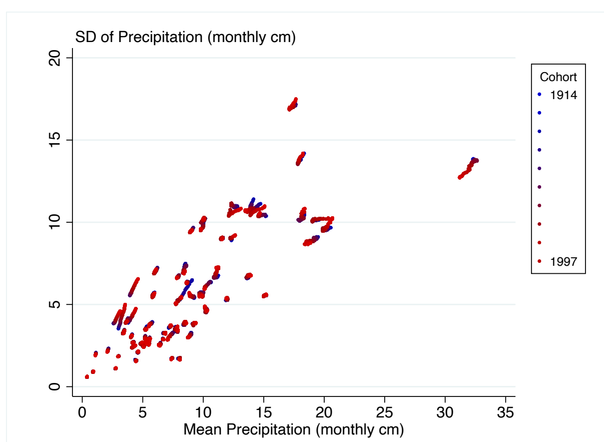


(f) Precipitation – **Global**

These figures present the distributions of our main explanatory variables ($\Delta A_{it}$ and $\Delta V_{it}$ in Indonesia and Mexico, $A_{it}$ and $V_{it}$ globally) plotted against each other. Each point is a cohort, where cohorts are defined by birth year and state in the panels and birth year and country in the global analysis, with older cohorts in blue and younger cohorts in red. Note the significant variation that exists in climate experiences across cohorts in the data, as well as the strong negative correlations between the moments of temperature in Indonesia and globally, and the strong positive correlations between the moments of precipitation in Mexico and globally.

### 3.6 Identification

The aim of our empirical analysis is to cleanly identify the causal marginal effects of long-run changes in the mean and variance of experienced climate on risk aversion over independent income lotteries. Our empirical set-up has two key advantages for achieving this task, one stemming from the nature of our risk aversion measures, and the other from the identifying variation in our climate experience variables.

To see the advantages inherent in our risk aversion measure for the task at hand, consider an alternative approach using a commonly employed measure of risky decision-making, portfolio allocations. Using changes in portfolio allocations to identify the causal effects in question raises two key issues. First, to attribute changes in portfolio allocations to changes in underlying risk preferences, we would have to somehow eliminate the alternative explanation that behavioral changes are occurring due to changes in agent beliefs about the returns of different portfolio compositions. This would be impossible to do without additional, extensive data on agent beliefs. Since our measures of risk aversion are estimated from choices over objective lotteries whose distributions are known by subjects, we can be confident that the changes in choice behavior we observe are not driven by changes in subject beliefs about the distribution of payouts of the lotteries themselves.

The second major identification advantage of using our risk aversion measures concerns the independence assumption between background and foreground risk in our conceptual framework. This assumption codifies the separation between the agent's experiences and choices, and is required to yield the behavioral predictions of the framework. For this assumption to hold in practice, it must be the case that the choice menu from which risk aversion is estimated empirically is independent of subjects' own background income risk experiences. This assumption does not plausibly hold for portfolio allocations: it is quite likely that the menus of assets that subjects can allocate to their portfolios are endogenous to their lifetime experiences of climate shocks (for instance, if luckier subjects, whose income is affected by climate realizations and are wealthier ex post, have access to assets with higher returns and fixed costs of investment). This assumption does hold, however, for our risk aversion measures, as the menu of lotteries from which subjects are choosing is by construction independent of their own lifetime experiences.

The other fundamental identification issue in our setting is whether the correlations we observe between lifetime macroeconomic experiences and risk attitudes represent the causal effects of these experiences on risk attitudes, or are driven by some other variable that correlates both with changes in growth experiences and changes in risk attitudes. We bring two primary tools to bear on solving this issue. First, since our empirical set-up involves estimating within-person changes in experiences and measured risk aversion, we can be confident that our results are not driven by time-invariant unobserved individual heterogeneity. Second, our empirical approach involves exploiting significant sources of variation in climate experiences, as can be seen in Figure 1. To our knowledge, ours is the only the second paper in the empirical experience effects literature (together with **?**) to exploit sub-national variation in experience data from multiple countries combined with repeat observations of the outcome of interest for the same individuals. Our paper is the first in the literature to combine this sub-national panel analysis with evidence from a large global cross-sectional containing individual measures of economic preferences. **?** use the national-level time series of stock market returns and examine stock market participation and elicited risk aversion in a repeated cross-section specification. **?** exploits within-country variation in macroeconomic conditions in Japan, but its empirical analysis does not include within-person changes or cross-country data. **?** and **?** use cross-country data from 13 Eurozone countries in their estimates of the effects of economic shocks on risk taking, but do not include within-person changes or sub-national macroeconomic data in their analyses. A number of papers in the development literature present data containing repeat measures of risk aversion for the same individuals, and exploit fine-grained sub-national variation in violence (**??**), or a natural disaster (**??**), but these studies generally focus on a discrete set of events rather than cumulative lifetime experiences of risk, and none contains evidence from multiple countries. The closest paper in the literature to ours is **?**, who examine the effects of historical rainfall shocks on the dynamics of elicited risk preferences for a sample of 900 farming households in Ethiopia. Relative to their paper, we provide evidence for both temperature and precipitation for a combined sample of 100,000 subjects distributed globally.

# 4 Empirical results

In this section we describe the results from our empirical analyses of the effects of lifetime experiences of climate change on individual risk aversion. We first present results from analyses using panel data (the IFLS and MXFLS) in Indonesia and Mexico (subsection 4.1). After introducing our main findings in subsubsection 4.1.1, we present results from three additional analyses aimed at understanding the mechanisms underlying our findings: regressions where controls for changes in income and other experiences are included (subsubsection 4.1.2); an examination of the heterogeneity of the main findings by relevant subgroups of subjects (subsubsection 4.1.3); and an analysis of the correlation between predicted changes in financial risk aversion and observed changes in other domains of risky behavior (subsubsection 4.1.4). We then present our results on the relationship between climate experiences and risk aversion globally using data from the Global Preference Survey (subsection 4.2).

## 4.1 Panel analyses in Indonesia and Mexico

### 4.1.1 Main results

Our primary empirical findings from the panel data analyses are presented in Table 2. Column 1 displays the results of regressing changes in measured risk aversion on mean changes in experienced lifetime temperature and precipitation in Indonesia. We find that 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 highly significant and positive, in line with our theoretical predictions on the effects of increases in background risk.

Columns 4–6 present the results of the parallel analysis in Mexico. In column 4, changes in measured risk aversion are regressed 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, neither the variance effect of temperature nor the variance effect of precipitation are significant, though are suggestive of a weak relationship between increased variance and risk aversion. Finally, in column 6, measured risk aversion is regressed on both the mean and the standard deviation of the climate variables. Both mean variables have negative and highly significant effects, while the standard deviation of precipitation has a positive and highly significant effect. Increases in the variance of temperature are also positively correlated with increases in measured risk aversion, but are not significant at conventional levels.

We can draw several conclusions from this set of results. First, we find consistent evidence that increases in the lifetime mean of climatic variables cause decreases in financial risk aversion. The associations underlying this conclusion are highly significant in both settings for both temperature and precipitation. Second, we find evidence that increases in the lifetime variance of climatic variables increases financial risk aversion. This finding is in line with the predictions of our conceptual framework, where increases in background climate risk increase measured risk aversion. Notably, variance effects appear to be more context-dependent than mean effects, with temperature risk driving the results in Indonesia, while precipitation risk drives the results in Mexico. Third, where variance effects on risk aversion are detected, they are first-order. In Indonesia, the marginal effect of temperature variance is 1.7 times that of the marginal effect of the temperature mean, while in Mexico the comparable number is 0.6

for precipitation.[19]

Finally, it is notable that the strongest results in our analysis – for temperature in Indonesia and precipitation in Mexico – occur in settings where the underlying climate variable moments are strongly correlated (see Figure 1). As we discuss in subsection 3.4, these correlations indicate skew in the underlying climate distributions. This suggests that the main climatic channels through which the observed changes in risk preferences occur are via unexpected lifetime heatwaves (particularly in Indonesia), and unexpected lifetime droughts (particularly in Mexico).

To test the robustness of our main panel findings repeat the main analysis with variety of alternative specifications: binarizing the risk aversion measures, using an ordered probit estimator, repeated cross-section specification, clustering the main analysis at the state level, and weighting the regressions using provided survey weights for the IFLS and MXFLS. Results for these alternative specifications are available in Appendix C. Our results are qualitatively similar for all alternative specifications, though in some, particularly where power is lacking (such as the repeated cross-section or state-level clustering), some coefficients are no longer significant.

### 4.1.2 Additional controls

We next turn to analyses aimed at understanding the mechanisms underlying our primary findings. Our conceptual framework suggests that experiences of climate change may correlate with observed changes in risk aversion over independent monetary lotteries through three possible channels: (1) background climate risk may correlate with background income risk (or other behaviorally-relevant background risk), rendering the observed associations between experienced climate change and risk aversion non-causal; (2) background climate risk may function as an independent income risk itself; and (3) background climate risk may directly shape foreground financial risk aversion, through one of the non-monetary channels discussed in section 2. Our analysis in this section is aimed at testing, to the best of our ability, whether channel (1) is in operation in our data, as it presents the most direct threat to a causal interpretation of our results.

To provide evidence for this question, we progressively add additional controls to the specifications in columns (3) and (6) in Table 2. These controls, which are meant to capture time-varying exposure to other sources of background risk, fall into three categories: (1) time-varying individual demographics (marital status, household size, educational attainment);

[19] Bolstering this conclusion is the fact that the signs and significance of the coefficients of both moment effects meaningfully change once the other moment is included in the regression. This is particularly evident in the cases of temperature variance in Indonesia, as well as precipitation in Mexico, where both the mean and variance effects are mostly obscured until the variance terms are added.

Table 2: Main results – panel analyses

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | **Indonesia** | | | **Mexico** | | |
| Δ Mean Temp | -3.98$^{\dagger\dagger}$ | | -4.86$^{\dagger\dagger}$ | -1.22$^{***}$ | | -1.17$^{***}$ |
| | (0.48) | | (0.58) | (0.25) | | (0.25) |
| Δ Std. Dev. Temp | | 2.07 | 8.25$^{***}$ | | 0.87 | 0.91 |
| | | (2.14) | (2.33) | | (0.51) | (0.54) |
| Δ Mean Precip | -0.40$^{***}$ | | -0.28$^{**}$ | -1.14 | | -3.61$^{**}$ |
| | (0.09) | | (0.10) | (0.95) | | (1.12) |
| Δ Std. Dev. Precip | | -0.15 | -0.40 | | 1.04 | 2.18$^{**}$ |
| | | (0.23) | (0.27) | | (0.58) | (0.70) |
| Observations | 15044 | 15044 | 15044 | 10218 | 10218 | 10218 |

Δ *Risk Aversion*: within-subject changes in measured risk aversion in IFLS (2007–2014) and MXFLS (2006–2012). Binned measure (1–5), 5 most risk averse. Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey measurement. State (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort-by-state-of-birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^\dagger$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

(2) changes in household balance sheets between the waves of the respective survey (income, assets, savings, and consumption); and (3) changes in measures of behaviorally-relevant experiences: exposure to violence (**??**), natural disasters (**??**), and state-level real GDP growth experiences, a direct measure of background income risk (**?**).[20]

Results for this analysis in Indonesia are presented in Table 3, and for Mexico in Table 4. We see consistent patterns of results in both countries. Controlling for changes in personal demographics, economic circumstances, and experiences of violence and natural disasters have negligible effects on our estimates. Controlling for changes in lifetime GDP growth experiences, our most direct measure of background income risk, have large effects on the estimates, but generally reinforce the main findings. In Indonesia, including growth experiences in the regression decreases the magnitude of the mean temperature coefficient by 22% in absolute terms, while increasing the magnitude of the temperature variance coefficient by 59% (and decreasing its p-value by 9 orders of magnitude). Adding these controls also attenuates to zero the (previously small but significant) coefficient of mean precipitation in this setting. In

[20]We do not include these controls in our main analysis because these variables, especially those in categories (1) and (2), are potentially endogenous to subjects' risk aversion, and therefore may introduce bias into the analysis.

Mexico, including growth experiences in the regression increases the magnitude of the mean precipitation coefficient by 78% in absolute terms (and decreases its p-value by 4 orders of magnitude), while increasing the magnitude of the precipitation variance coefficient by 82% (and decreases its p-value by 4 orders of magnitude). Adding these controls does somewhat attenuate the mean temperature coefficient in Mexico, but it remains negative and highly significant.

Overall, the results of this analysis suggest that, with the possible exception of the mean temperature effects in Indonesia, the results we observe in the main analysis are not driven by other observable sources of background risk. In fact, the preponderance of the evidence points to the effects of climate risk being obscured in our baseline empirical analysis by their (likely negative) correlation with independent background income risk–once empirical measures of lifetime experiences of background income risk are added to the analysis, our results for climate risk generally become stronger.

Table 3: Additional controls – **Indonesia**

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---|---|---|---|---|---|---|---|---|
| Δ Mean Temp | $-4.86^{\dagger\dagger}$ | $-4.92^{\dagger\dagger}$ | $-4.91^{\dagger\dagger}$ | $-4.91^{\dagger\dagger}$ | $-4.90^{\dagger\dagger}$ | $-4.89^{\dagger\dagger}$ | $-4.84^{\dagger\dagger}$ | $-4.89^{\dagger\dagger}$ | $-3.83^{\dagger}$ |
| | (0.58) | (0.58) | (0.58) | (0.58) | (0.58) | (0.58) | (0.58) | (0.58) | (0.54) |
| Δ Std. Dev. Temp | $8.25^{***}$ | $8.49^{***}$ | $8.43^{***}$ | $8.42^{***}$ | $8.43^{***}$ | $8.40^{***}$ | $8.30^{***}$ | $9.06^{***}$ | $14.40^{\dagger}$ |
| | (2.33) | (2.33) | (2.33) | (2.33) | (2.34) | (2.34) | (2.33) | (2.47) | (2.56) |
| Δ Mean Precip | $-0.28^{**}$ | $-0.26^{*}$ | $-0.26^{*}$ | $-0.26^{*}$ | $-0.26^{*}$ | $-0.26^{*}$ | $-0.25^{*}$ | $-0.25^{*}$ | -0.00 |
| | (0.10) | (0.11) | (0.11) | (0.11) | (0.11) | (0.11) | (0.11) | (0.11) | (0.12) |
| Δ Std. Dev. Precip | -0.40 | -0.42 | -0.41 | -0.41 | -0.41 | -0.41 | -0.40 | -0.45 | -0.39 |
| | (0.27) | (0.27) | (0.27) | (0.27) | (0.27) | (0.27) | (0.27) | (0.27) | (0.25) |
| Observations | 15044 | 15043 | 15040 | 15040 | 15040 | 15040 | 15040 | 15040 | 15040 |
| 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 |

Δ *Risk Aversion*: within-subject changes in measured risk aversion in IFLS (2007–2014). Binned measure (1–5), 5 most risk averse. Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey measurement. Inflation: state-level. Demographics: marital status, hh size, hh size$^2$, educational attainment. See Table 19 for details on additional controls. Standard errors clustered at the cohort-by-state-of-birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

Table 4: Additional controls – **Mexico**

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---|---|---|---|---|---|---|---|---|
| Δ Mean Temp | -1.17*** | -1.16*** | -1.16*** | -1.17*** | -1.16*** | -1.16*** | -1.16*** | -1.14*** | -0.83*** |
| | (0.25) | (0.25) | (0.25) | (0.25) | (0.25) | (0.25) | (0.25) | (0.25) | (0.25) |
| Δ Std. Dev. Temp | 0.91 | 0.87 | 0.87 | 0.87 | 0.88 | 0.90 | 0.90 | 0.89 | 0.23 |
| | (0.54) | (0.54) | (0.54) | (0.54) | (0.54) | (0.54) | (0.54) | (0.54) | (0.56) |
| Δ Mean Precip | -3.61** | -3.64** | -3.67** | -3.67** | -3.65** | -3.66** | -3.68** | -4.06*** | -7.23† |
| | (1.12) | (1.12) | (1.12) | (1.12) | (1.12) | (1.12) | (1.12) | (1.11) | (1.19) |
| Δ Std. Dev. Precip | 2.18** | 2.17** | 2.17** | 2.16** | 2.16** | 2.17** | 2.17** | 2.30*** | 4.20† |
| | (0.70) | (0.70) | (0.70) | (0.70) | (0.70) | (0.70) | (0.70) | (0.70) | (0.76) |
| Observations | 10218 | 10217 | 10217 | 10217 | 10217 | 10217 | 10217 | 10217 | 10217 |
| 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 |

Δ *Risk Aversion*: within-subject changes in measured risk aversion in MXFLS (2006–2012). Binned measure (1–5), 5 most risk averse. Temperature (°C) and precipitation (cm) moments calculated from monthly measurements in subjects' state of birth, from birth to year of survey measurement. Demographics: marital status, hh size, hh size$^2$, educational attainment. See Table 19 for details on additional controls. Standard errors clustered at the cohort-by-state-of-birth level in parentheses. * p < .05, ** p < .005, *** p < .0005, † $p < 5 \times 10^{-7}$, †† $p < 5 \times 10^{-13}$.

### 4.1.3 Subgroup heterogeneity

Our second analysis aimed at understanding the mechanisms underlying the main results is examining the comparative effects for subgroups of our subject population. We split the sample into binary groups across six observable subject characteristics that have either been highlighted in the literature as mediators of experience effects, or that can serve as proxies for causal channels in our model: occupation in agriculture, urbanicity, consumption level, educational attainment, gender, and age. Results of these analyses are presented in Figure 2, and in the form of tables in Appendix D.

Starting with the split by agricultural occupation, we see a consistent pattern of stronger effects of climate experiences for farmers in both countries, across both climate variables. Particularly noteworthy is the significance and magnitude of the temperature standard deviation coefficient for farmers in Mexico, which is three times larger than the corresponding (non-significant) coefficient for the overall Mexican sample. These results suggest that the effects of climate experiences on agricultural income likely an important channel for their effects on financial risk aversion (though alternative explanations, like farmers paying more attention to the climate than non-farmers, are also possible). However, these results also indicate that the agricultural income channel is unlikely to be the only one in operation, as non-farmers also exhibit significant effects in line with our main results. This conclusion is also reinforced by the results of the urban/rural split, where no consistent patterns in the strength of the results across urbanicity are evident.

Splitting our sample by consumption, our best measure of economic circumstances, yields inconclusive results. We see no differences across groups in the Indonesian context, and mostly stronger results for lower consumption subjects in Mexico, though this pattern is reversed in the case of the mean temperature coefficient. Stronger results for poorer subjects might follow from a variety of reasons, including a lack of resources to aid in climate adaptation for these subjects, more exposure to income shocks, or differences in the sources and processing of information. This last interpretation is reinforced by our results from the education split, where we see stronger effects for subjects with lower levels (below high school) of educational attainment. This distinction is particularly strong for the effects of temperature in Indonesia.

Turning to the gender split, we see heterogeneity across the two contexts: the effects of climate experiences are generally stronger for women in Indonesia and for men in Mexico. Although the correlation between static risk aversion and gender is well-established in the literature (?), the experience effects literature on risk preference dynamics, particularly in the wake of natural disasters, has found mixed results, with stronger effects for women due to hurricane Katrina in the US (?), but stronger effects for men of the East Japan earthquake in Japan (?). To our knowledge, our set of findings is the first to demonstrate heterogeneity in preference dynamics across gender due to the same type of experiences. More research is necessary to understand the drivers of this heterogeneity.

Finally, we split the sample by age during the second survey wave, with the old defined as subjects older than 40. We find that across both settings, results are generally stronger for older subjects, though they are still significant across most coefficients for the young as well. This is a striking pattern of results, given the Bayesian underpinning of our conceptual framework. Since in the empirical analysis we are estimating the effects of the set of experiences between the waves of the surveys (with a set time interval) to subject experiences before the first wave of the survey, in general the later set of experiences forms a larger subset of total experiences for the young in the sample. If all subjects are Bayesian and the object of learning is static, we might reasonably expect that the marginal effects of new experiences will be stronger for the young, because they contain more information for them. However, this is not the case if the object of learning is time-varying, since older subjects will then be surprised more by new information than younger subjects. We therefore interpret stronger results for the old in our sample as suggestive evidence that the underlying changes in experienced climate are sufficiently strong as to overpower the general weight of Bayesian learning towards stronger effects for the young.

Overall, our subgroup heterogeneity analysis reveals some consistent patterns and some striking heterogeneity. The effects of climate experiences appear to be stronger for farmers, the less educated, and older subjects across both settings. Results are heterogeneous by setting for gender and consumption, and indeterminate for urbanicity.

#### 4.1.4 Changes in risk-taking behavior

In this section we explore whether the changes we observe in measured financial risk aversion induced by climate experiences correlate with changes in risk-taking behavior in other domains. Our objective with this analysis is to better understand the behavioral mechanisms driving changes in risky behavior in a variety of behavioral domains, and how these changes are similar or different from those in the domain of financial risk aversion. While an active debate exists in the literature on the domain-specificity of economic risk preferences when considered statically (**?**), to our knowledge these results are only the second presented in the literature (after **?**) on domain-specificity of the dynamics of risk-taking.

To conduct this analysis, we construct a variable measuring predicted change in risk aversion ($\widehat{\Delta R_{it}}$) using our preferred specification (columns (3) and (6) of Table 2), and examine its correlation with changes in risk-taking behaviors for our subjects. We focus on behaviors commonly examined in relation to risk-taking in the literature, and for which we have data: smoking, having ever migrated across province or state lines, self-employment status, and, in Indonesia, whether subjects report that their land is planted in at least one cash crop, a measure of risky agricultural investment.[21]

Results for this analysis are presented in Figure 3, which displays the average value of changes in each risky behavior for different sub-groups of the $\widehat{\Delta R_{it}}$ distribution. Here, the green bars represent subjects who are predicted to become less risk averse, while the purple bars represent subjects who are predicted to become more risk averse. Darker colored bars within each of these groups represent subjects with a larger predicted change in measured risk aversion: for instance, dark green bars represent subjects whose decrease in predicted risk aversion is larger (in absolute terms) than the median, while light green bars represent subjects whose change in predicted risk aversion is smaller (in absolute terms) than the median among individuals who have a predicted decrease in measured risk aversion.

We draw four conclusions from these results. First, choice of cash crops responds strongly and consistently to climate experiences in the same way as financial risk aversion in our Indonesian data. Subjects who become more risk tolerant significantly increase their plantings of cash crops, and subjects who become more risk averse significantly reduce (or do not

[21] Cash crops asked about in the IFLS include coconut, coffee, cloves, rubber, and other hard stem plants.

change) their planting behavior. Differences in behavioral changes between these two groups are highly statistically significant. These results are also consistent with our findings in the previous section that the main effects are stronger for farmers in our sample.

Second, we find evidence that rates of smoking increased with climate induced risk tolerance in Mexico. In addition, we find that within these individuals, those with a larger predicted increase risk tolerance have a larger increase in smoking. However, this relationship does not appear to hold in Indonesia, where rates of smoking are flat (or increasing) in increases in risk aversion. More research on the relationship between the dynamics of smoking and risk aversion is clearly warranted.

Third, we find no evidence that rates of migration increase due to climate-induced increases in risk tolerance in either country. This finding is in stark contrast to the findings in **?**, which found such a gradient in both countries (with non-linearities in Indonesia) for changes in risk aversion due to macroeconomic experiences. The most likely explanation for this disparity is that migration is more likely to be induced in these countries over the period of study by direct changes in economic circumstances, rather than climate change itself.

Finally, we find across the board null effects or increases in self-employment across all bins in both settings, with no significant differences in behavioral change in this domain by predicted changes in financial risk tolerance. This finding bolsters a set of recent findings in the literature that distinguish between measures of self-employment and entrepreneurship in developing countries, and argue that the former do not capture the latter in, given high rates of subsistence self-employment in these settings (**?**).

## 4.2 Results from global analysis using the GPS

In this section we present results from our analysis of the effects of climate experiences on risk aversion in the Global Preference Survey (GPS). The GPS collected experimentally-validated economic preference measures for a sample of approximately 1,000 subjects in 75 countries in 2012. Using data from the public release of the GPS, we regress measured risk preferences on statistics capturing lifetime climate experiences for each subject, calculated based on their country of residence and birth year. Our GPS analysis closely mirrors our empirical approach in the IFLS and MXFLS, but with a single measurement of risk aversion instead of two, and with climate experiences at the country-birth-year level instead of the state-birth-year level.

Table 5 presents the results of this analysis. In column (1) we present the results of regressing measured risk aversion on the mean and the standard deviation of lifetime temperature and precipitation for the sample. This specification is identical to that in columns (3) and (6) of Table 2, except in levels instead of differences. We find results consistent with our panel analyses for two coefficients – a negative and significant effect of mean lifetime temperature, and a positive and significant effect of the lifetime precipitation variance.

Since we do not observe within-person changes in measured risk aversion in the GPS, we rely on including individual and country-level controls in the regressions to increase the precision of our estimates. We first add controls for demographic variables that are available in the GPS: subject gender, age, math scores, interview language, and the month the interview was conducted. Results for this specification are presented in column (2). Though the coefficient of mean temperature attenuates, and is no longer significant at conventional levels, including these controls results in positive and significant coefficients for the variance of both temperature and precipitation, again in line with our panel results.

We next add controls for the other economic preference measures available in the GPS: patience (i.e. discounting), trust, altruism, positive reciprocity, and negative reciprocity.[22] Previous work in the literature has found that economic preferences are often correlated between individuals (**?**), particularly risk aversion and discounting, which suggests we can increase the precision of our estimates through their inclusion. Results for this specification are presented in column (3), where we see that the estimates of three of the four coefficients are significant and in line with our previous analyses.

Having exhausted all individual-level controls available in the public release of the GPS, we turn to country-level controls. Our results from Indonesia and Mexico suggest that heterogeneity in the response to climate experiences can be moderated by local climate, with responses to temperature shocks stronger in Indonesia, and to precipitation shocks stronger in Mexico. This is an especially relevant concern in our global analysis, which spans a much larger variety of climates than our country panels (Table 1). To account for this potential source of noise in our regressions, we include variables capturing the country's position in the global temperature and precipitation distribution. Results for the specification are presented in column (4). We find that, with these controls included, all four coefficients in the regressions are significant, with signs in the expected directions.

One important concern with empirical analyses using data from many countries is that observed effects might be driven by unobserved differences between countries, rather than by the explanatory variables in the analysis. To address this concern, in our final specification we include country fixed effects instead of country climate controls. Doing so shuts down all between-country variation, which means that the only variation in the regressions is from differences between cohorts within countries. Results of this analysis are presented in column (5). Three of the four coefficients are still highly significant and have the same sign as in the previous specifications. The fourth, on the effects of mean precipitation, is still

[22] Details on all measures of economic preferences in the GPS can be found in **?**.

negative, but not significant. Notably, the magnitudes of the three significant coefficients in this specification are an order of magnitude larger than in the specification without country fixed effects. This suggests that unobservable between-country differences are attenuating the main results, rather than driving them. Put another way, this last specification suggests that the variation driving our results is primarily that due to the time series, that is the variation that is driven by climate change within each location.

Overall, our results in the global analysis paint a highly coherent picture on the effects of experiences of climate change on risk preferences. Across cohorts, countries, and climate variables, higher levels of the lifetime mean of the climate variable are correlated with higher risk tolerance, and higher variables of the lifetime variance are correlated with lower risk tolerance. These results are highly consistent with our panel results from Indonesia and Mexico. Further, since the global temperature distribution exhibits a positive skew, and the global precipitation distribution exhibits a negative skew, the most likely environmental channel for these effects is through unexpected heatwaves and droughts.

Table 5: Main results – **Global** Analysis

| Dep Var: Risk Aversion | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Mean Temp | -0.017*** | -0.006 | -0.009** | -0.012* | -0.238* |
| | (0.004) | (0.003) | (0.003) | (0.005) | (0.103) |
| Std. Dev. Temp | 0.013 | 0.017* | 0.022** | 0.032*** | 0.569*** |
| | (0.008) | (0.008) | (0.007) | (0.007) | (0.136) |
| Mean Precip | 0.005 | 0.004 | 0.007 | -0.014** | -0.022 |
| | (0.006) | (0.006) | (0.006) | (0.005) | (0.043) |
| Std. Dev. Precip | 0.040*** | 0.031*** | 0.028** | 0.023** | 0.132* |
| | (0.010) | (0.009) | (0.009) | (0.008) | (0.062) |
| $R^2$ | 0.016 | 0.090 | 0.151 | 0.156 | 0.209 |
| Demographics | | X | X | X | X |
| Preferences | | | X | X | X |
| Country Climate | | | | X | |
| Country FEs | | | | | X |
| Observations | 78451 | 76999 | 74673 | 74673 | 74673 |

*Risk aversion:* continuous measure 0-4, higher values indicating higher risk aversion, from the Global Preference Survey (GPS). *Temperature* and *precipitation* variables are moments of country-level, population-weighted lifetime (birth year to 2012) monthly time series for each subject. *Demographics*: subject gender, age, math score, interview language, interview month. *Preferences*: patience, trust, altruism, positive reciprocity, negative reciprocity. *Country Climate*: above/below mean global temperature 2012, tercile of country precipitation distribution. Standard errors in parentheses clustered at the sub-national region level. All regressions weighted by individual sampling weights for the GPS. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure 2: Heterogeneity of main effects by subgroup

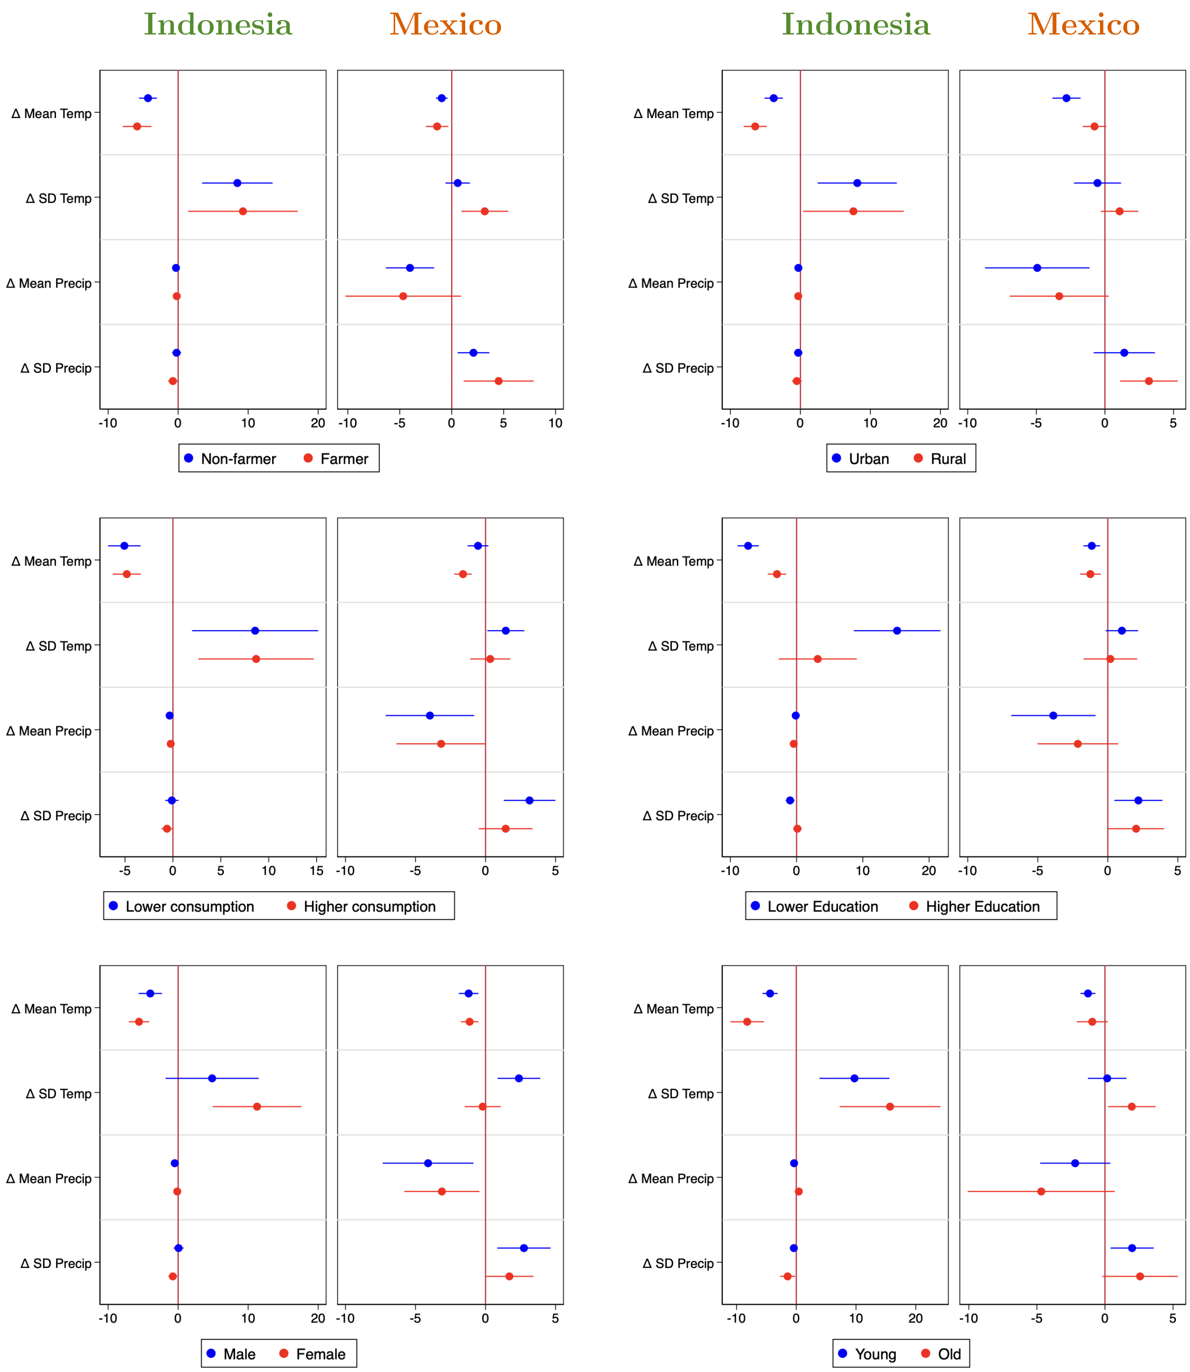


Regression results from columns (3) and (6) in Table 2, for subgroups defined by binary splits of listed covariates. *Farmers* are subjects who report spending a nonzero amount of labor on farm work over the course of the previous year. *Low consumption* subjects are in the bottom 50% of the consumption distribution. *Low education* are subjects with less than a high school degree. *Old* subjects are older than 40. All splits are defined on observable characteristics reported in the second wave of the respective survey. Bars represent 95% confidence intervals. Corresponding regression tables are available in Appendix D.

Figure 3: Correlations of changes in risky behaviors with predicted change in risk aversion

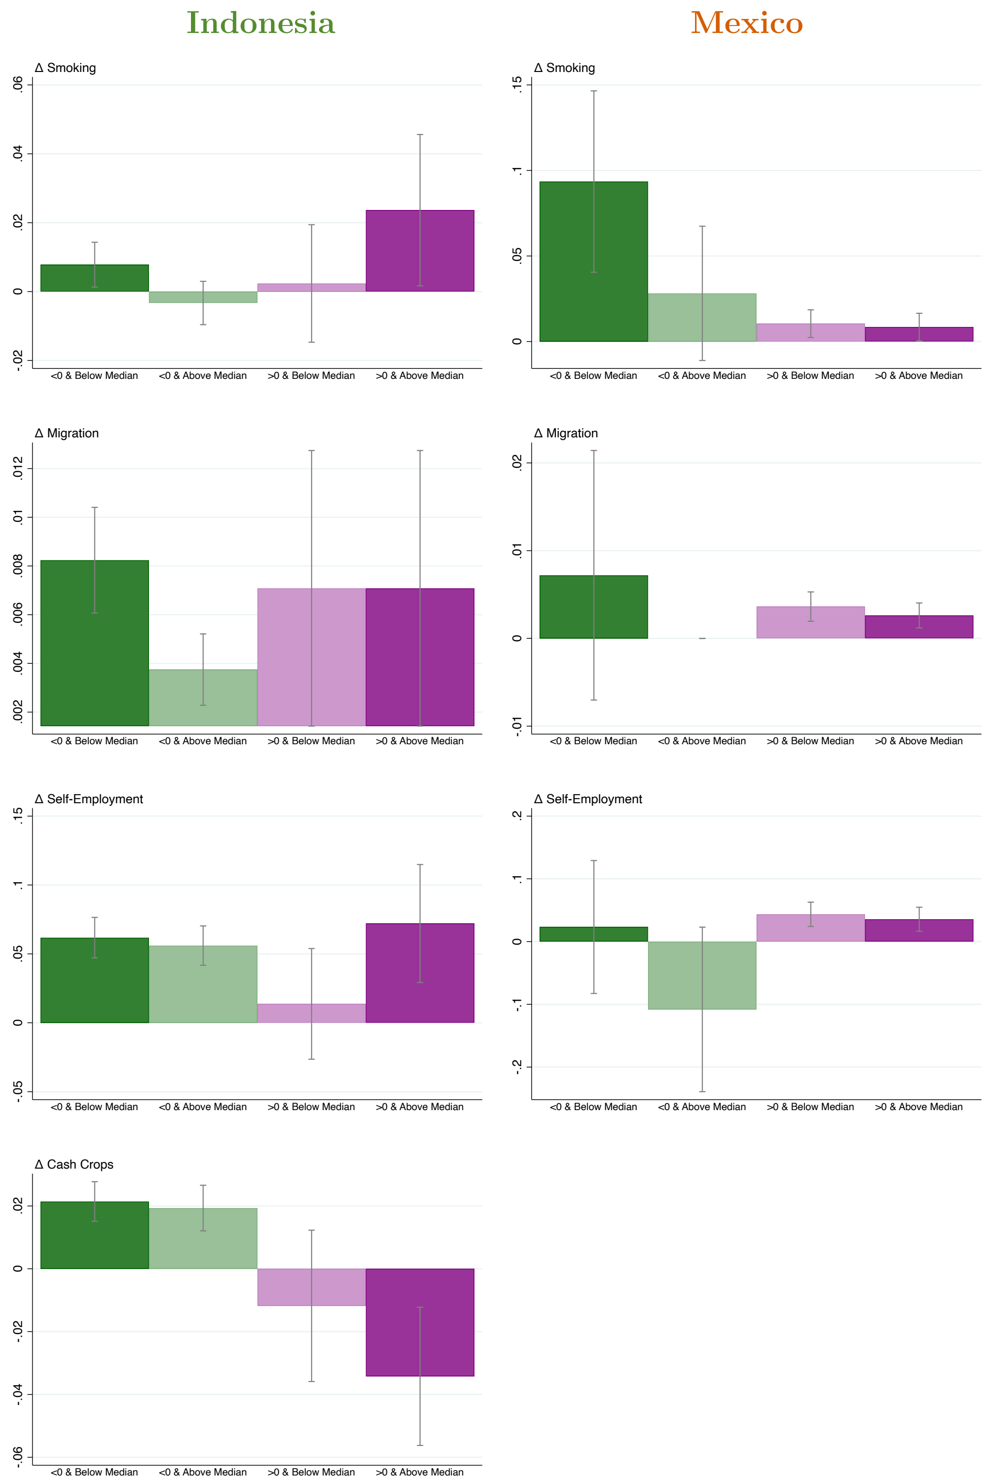


Bars represent results of mean group changes in the indicated behavior for different bins of the predicted change in risk aversion distribution. Green bars are subjects who are predicted to become less risk averse, purple bars are subjects who are predicted to become more risk averse. Darker colored bars within each of these groups represent subjects whose change in predicted risk aversion is larger (in absolute terms) than the median among individuals in the group. 95% confidence intervals are indicated for each group.

Table 6: Results from Main Global Analysis - staircase measure

| | (1) riskaver_strcse | (2) riskaver_strcse | (3) riskaver_strcse | (4) riskaver_strcse | (5) riskaver_strcse |
|---|---|---|---|---|---|
| Mean Temp | -0.073 | 0.008 | -0.036 | -0.050 | -1.684 |
| | (0.041) | (0.044) | (0.044) | (0.058) | (1.077) |
| SD Temp | -0.129 | -0.070 | -0.096 | -0.040 | 1.120 |
| | (0.091) | (0.085) | (0.087) | (0.093) | (1.731) |
| Mean Precip | 0.283*** | 0.257*** | 0.255*** | 0.104 | 0.370 |
| | (0.067) | (0.065) | (0.063) | (0.062) | (0.439) |
| SD Precip | 0.002 | -0.039 | -0.036 | -0.070 | 0.559 |
| | (0.110) | (0.111) | (0.111) | (0.111) | (0.610) |
| $R^2$ | 0.020 | 0.038 | 0.060 | 0.062 | 0.129 |
| Demographics | | X | X | X | X |
| Preferences | | | X | X | X |
| Country Climate | | | | X | |
| Country FEs | | | | | X |
| Cluster | Region | Region | Region | Region | Region |
| Observations | 70550 | 69504 | 68020 | 68020 | 68020 |

**Note**: *Risk aversion* from the Global Preference Survey. *Temperature* and *precipitation* variables are moments of country-level, population-weighted lifetime (birth year to 2012) monthly time series for each subject, from UDel data. *Demographics*: subject gender, age, math score, interview language, interview month. *Preferences*: patience, trust, altruism, positive reciprocity, negative reciprocity from Falk et. al. 2018. *Country Climate*: above/below mean global temperature 2012, tercile of country precipitation distribution. All regressions weighted by sampling weights for the GPS. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors in parentheses.

Table 7: Results from Main Global Analysis - self-reported measure

| | (1) riskaver_selfrep | (2) riskaver_selfrep | (3) riskaver_selfrep | (4) riskaver_selfrep | (5) riskaver_selfrep |
|---|---|---|---|---|---|
| Mean Temp | -0.041*** | -0.019 | -0.018 | -0.031* | -0.577 |
| | (0.011) | (0.011) | (0.010) | (0.014) | (0.300) |
| SD Temp | 0.084** | 0.088** | 0.120*** | 0.148*** | 2.076*** |
| | (0.027) | (0.027) | (0.025) | (0.025) | (0.408) |
| Mean Precip | -0.043* | -0.043** | -0.027 | -0.083*** | -0.172 |
| | (0.017) | (0.016) | (0.015) | (0.015) | (0.155) |
| SD Precip | 0.158*** | 0.131*** | 0.115*** | 0.102*** | 0.442* |
| | (0.027) | (0.025) | (0.022) | (0.021) | (0.199) |
| $R^2$ | 0.022 | 0.094 | 0.154 | 0.159 | 0.202 |
| Demographics | | X | X | X | X |
| Preferences | | | X | X | X |
| Country Climate | | | | X | |
| Country FEs | | | | | X |
| Cluster | Region | Region | Region | Region | Region |
| Observations | 70550 | 69504 | 68020 | 68020 | 68020 |

**Note**: *Risk aversion* from the Global Preference Survey. *Temperature* and *precipitation* variables are moments of country-level, population-weighted lifetime (birth year to 2012) monthly time series for each subject, from UDel data. *Demographics*: subject gender, age, math score, interview language, interview month. *Preferences*: patience, trust, altruism, positive reciprocity, negative reciprocity from Falk et. al. 2018. *Country Climate*: above/below mean global temperature 2012, tercile of country precipitation distribution. All regressions weighted by sampling weights for the GPS. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Standard errors in parentheses.

# 5 Welfare analysis

A key question left open by our analysis up to this point is whether the climate impacts on risk aversion that we detect are beneficial or detrimental to subjects' welfare. For many outcomes examined in the climate impacts literature, like mortality or income, this question is second-order, because common sense and economic theory make it easy to sign the welfare impacts of changes in the outcome. Risk aversion is different. There is no straightforward intuitive sense in which higher risk aversion is always better or worse. And economic theory, having mostly ignored the very possibility of preference changes since **?**, is mostly silent on how to think about welfare in situations where preferences are dynamic.

To make progress on this question, we develop a new method for estimating the welfare impacts of changes in risk preferences that are driven by changes in subjects' environments. Our method is premised on two key ideas. The first is deceptively simple: we take the measured risk preferences in our data at face value, as a measure of subjects' true preferences. This means that we attribute observed changes in risk aversion in our data to changes in subjects' "real" underlying preferences, rather than changes in their beliefs or constraints. Theoretically, this is equivalent to the statement that the agent's foreground utility function in our model, rather than the background utility function, captures the agent's real preferences.

The second key idea underlying our method is that, given data on the distribution of risk preferences and the distribution of consumption of a group of heterogeneous subjects, we can calculate a population-level measure of welfare for the group. This is the primary finding of **?**, who proves that such a measure always exists if agents in the group are assumed to be expected utility maximizers. This measure, known as an Equally Distributed Equivalent (EDE), corresponds to the counterfactual equal level of consumption that all agents behind the veil of ignorance would be willing to accept, given the unequal consumption distribution which exists in reality.[23] As such, the EDE can be interpreted as a measure of how much society collectively would be willing to pay to avoid uncertainty in the distribution of consumption. To yield our welfare measure, we calculate the ratio between a given EDE and the contemporaneous mean level of group consumption.

We proceed in two stages. First, we structurally estimate the risk preferences of subjects in our IFLS and MXFLS samples in both waves, under the assumption that they are expected

[23] An EDE is derived in the following way. Take a group of expected utility maximizers with heterogeneous risk preferences, and elicit for each of them the certainty equivalent of the lottery corresponding to the the consumption distribution of the group. This will result in a distribution of certainty equivalents. Next elicit for each subject the certainty equivalent of the distribution of certainty equivalents, resulting in a second-order distribution of certainty equivalents. Repeat this process iteratively. **?** proves that this iterative process always converges to fixed point under the assumption of expected utility. Intuitively this follows from the close connection between risk aversion and preferences for inequality in expected utility theory.

utility maximizers with CRRA utility.[24] Using these estimated preferences, we then calculate the period 2 EDE ratio using both the period 1 preferences and the period 2 preferences. The difference between these two EDE ratios captures the total welfare change from risk preference changes due to all causes. In our preferred specification for Indonesia, the EDE ratio with the period 1 preferences is 0.367, while the EDE ratio with the period 2 preferences is 0.428. This means that risk preference changes due to all causes increased welfare for our Indonesian sample by 6.1%. In our preferred specification for Mexico, the EDE ratio with the period 1 preferences is 0.379, while the EDE ratio with the period 2 preferences is 0.297. Therefore, we estimate that for our Mexican sample risk preference changes due to all causes decreased welfare by 8.2%.

Next, we examine the welfare effects of risk preference changes that are due only to climate change. We use the results from our main regressions[25] to construct the counterfactual risk preference distribution in period 2 had no climate change occurred (relative to agents' lifetime body of experiences) between the periods. We then calculate the EDE ratio using the period 2 consumption distribution under the counterfactual preferences, and compare it to the period 2 EDE ratio under the realized preference distribution. In Indonesia, the EDE ratio difference between the two corresponds to a 1% increase in welfare, while the EDE ratio difference in Mexico corresponds to a 0.8% increase in welfare. This suggests that the changes in risk preferences that we observe in our panel analyses may be a form of psychological climate adaptation, where changes in agents' risk-taking psychology enables them to better match the dynamics of their environments.

# 6 Conclusions

In this paper we provide theoretical and empirical evidence that long-run experiences of climate change shape individual risk preferences. Our theoretical framework suggests that increases in the mean of experienced climate variables will decrease individual risk aversion, while increases in the variance of experienced climate will increase risk aversion. Our empirical results largely confirm these predictions. Across all settings, increases in the means of temperature and precipitation lead to decreases in measured risk aversion. In Indonesia and globally increases in the variance of temperature lead to increases in risk aversion. And in Mexico and globally increases in the variance of precipitation lead to increases in risk

[24]In order to conduct this estimation from the binned data in the panel surveys we must make a series of methodological choices on bracketing of the lotteries, the treatment of the edge bins, and the treatment of gamble averse. We detail our method and explore sensitivity to these methodological choices in Appendix F.

[25]More precisely, for this exercise we use the effects of mean and variance of temperature in Indonesia, and mean and variance of precipitation in Mexico.

aversion. Notably, in the settings where climate variance affects preferences, its effects are first-order: the ratio of the standard deviation coefficient to the mean coefficient is 0.6 for precipitation in Mexico, 1.7 for temperature in Indonesia, and 1.6 - 2.6 globally. We estimate that climate-induced changes to risk preferences increased collective welfare in Indonesia and Mexico by approximately 1%. This suggests that our results are driven by psychological climate adaptation.

# Appendix

## A Construction of Risk aversion measures

Figure 4: Construction of risk aversion measure in IFLS2 and IFLS3

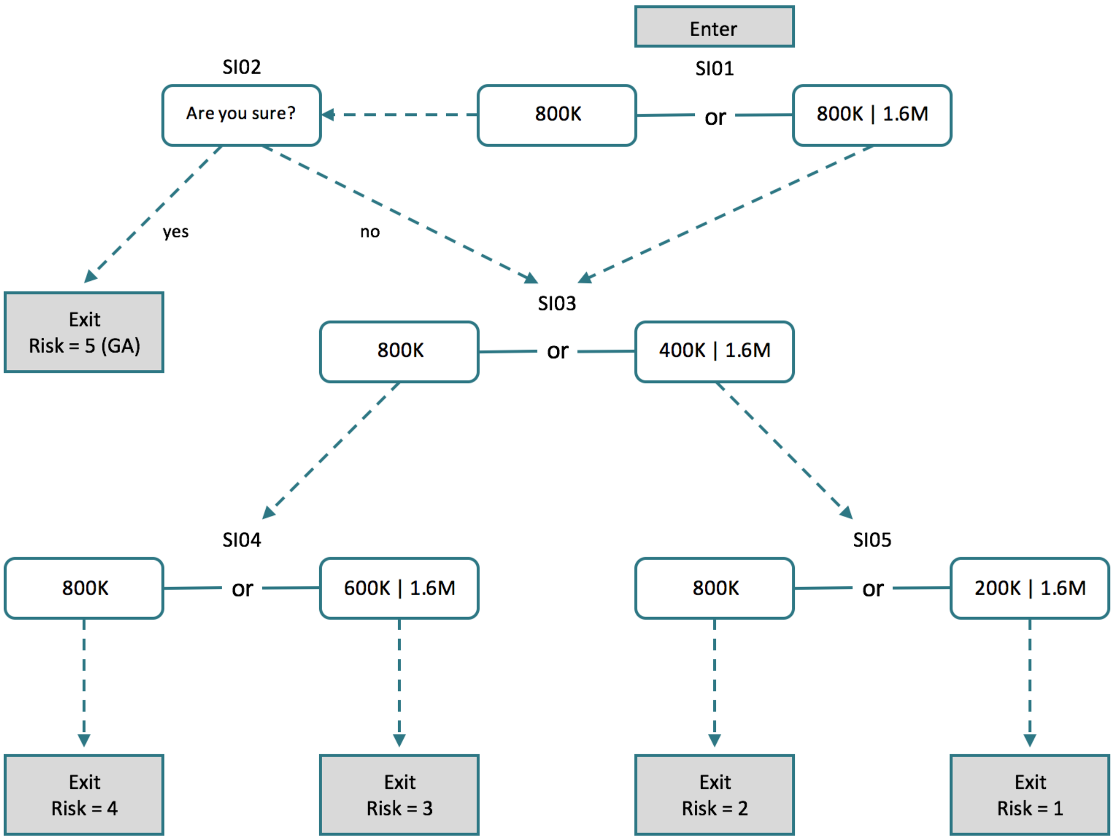


**Notes:** Hypothetical lottery values are in Indonesian Rupiah. All lotteries offer a 50% probability of winning each of the two prizes. Higher values for “Risk” indicate a higher level of measured risk aversion (Risk = 1 is the most risk-seeking choice).

Figure 5: Construction of risk aversion measure in MXFLS2

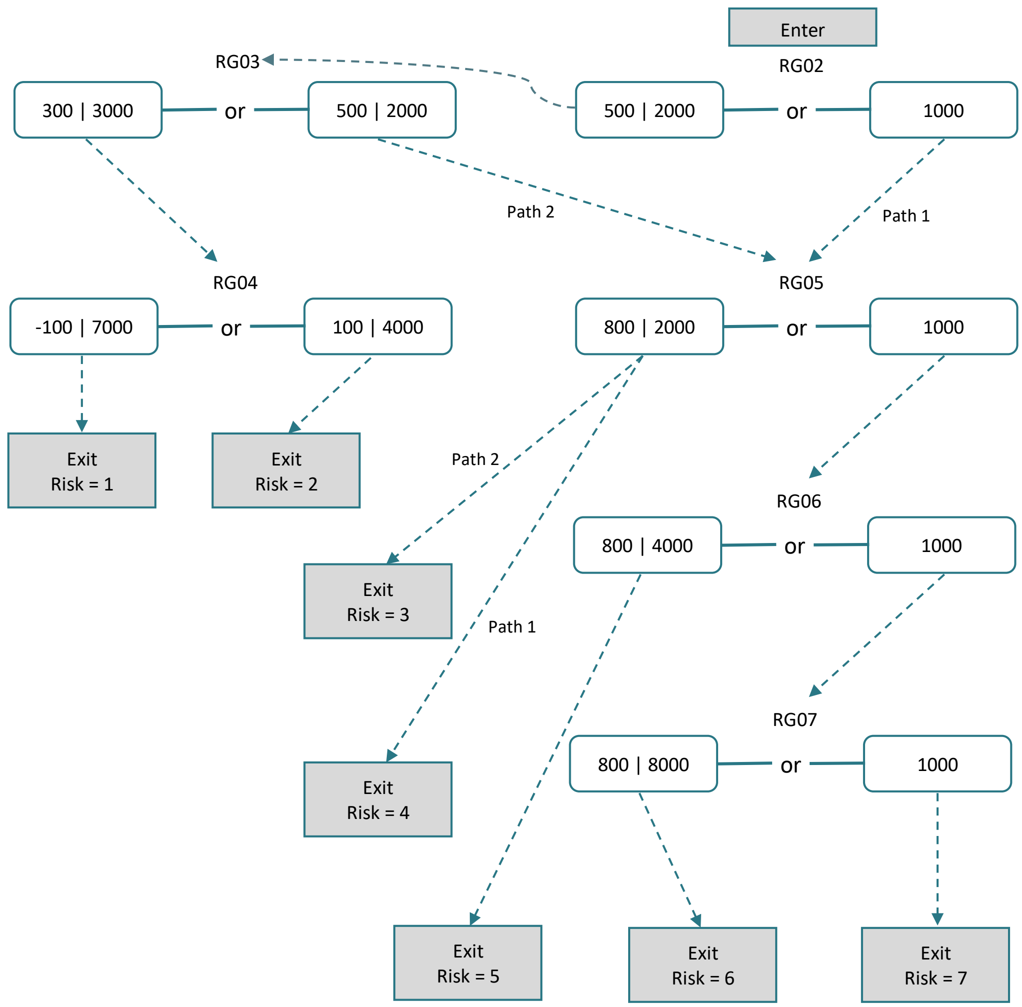


**Notes:** Hypothetical lottery values are in Mexican Pesos. All lotteries offer a 50% probability of winning each of the two prizes. The safe option in each stage is depicted on the right side of the stage. Higher values for "Risk" indicate a higher level of measured risk aversion (Risk = 1 is the most risk-seeking choice).

Figure 6: Construction of risk aversion measure in MXFLS3

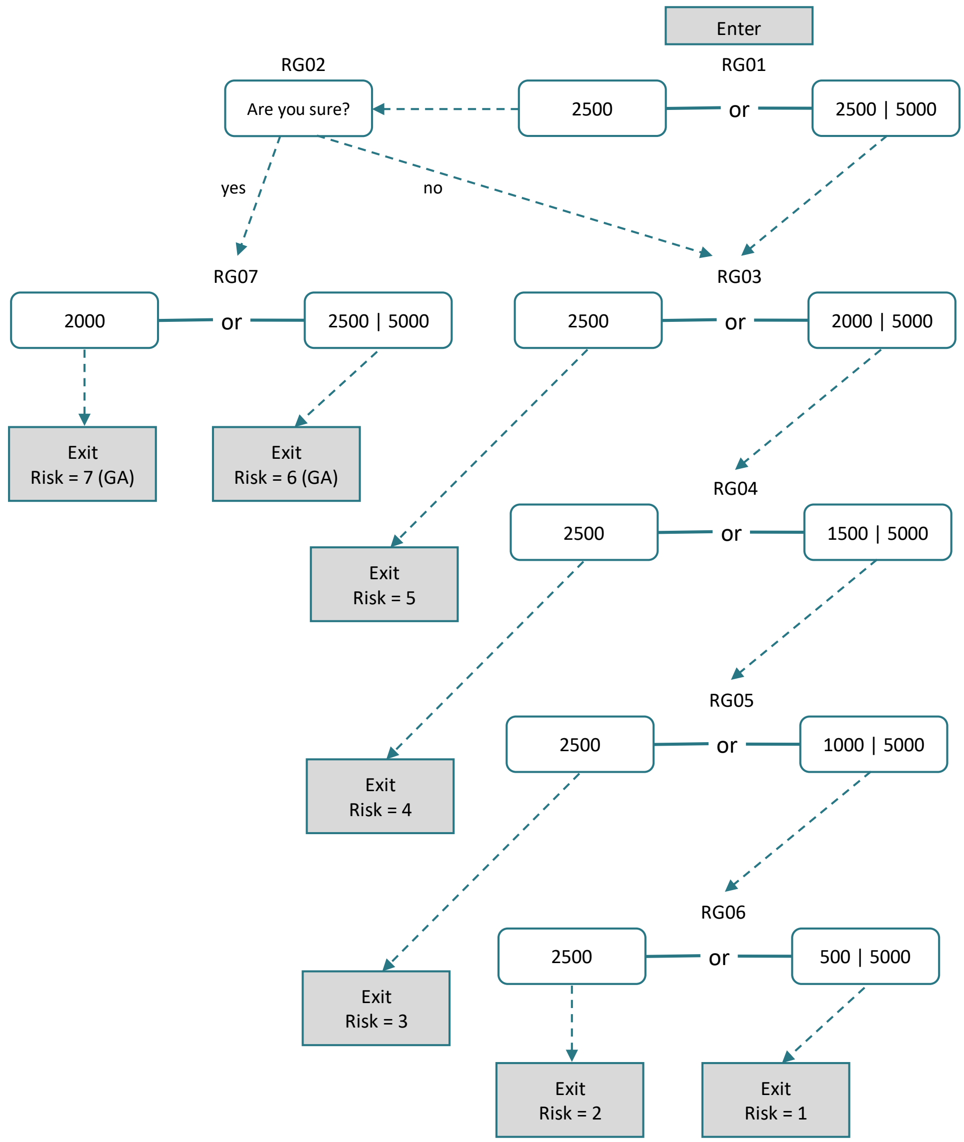


**Notes:** Hypothetical lottery values are in Mexican Pesos. All lotteries offer a 50% probability of winning each of the two prizes. The safe option in each stage is depicted on the left side of the stage. Higher values for "Risk" indicate a higher level of measured risk aversion (Risk = 1 is the most risk-seeking choice).

Figure 7: Construction of risk aversion staircase measure in GPS

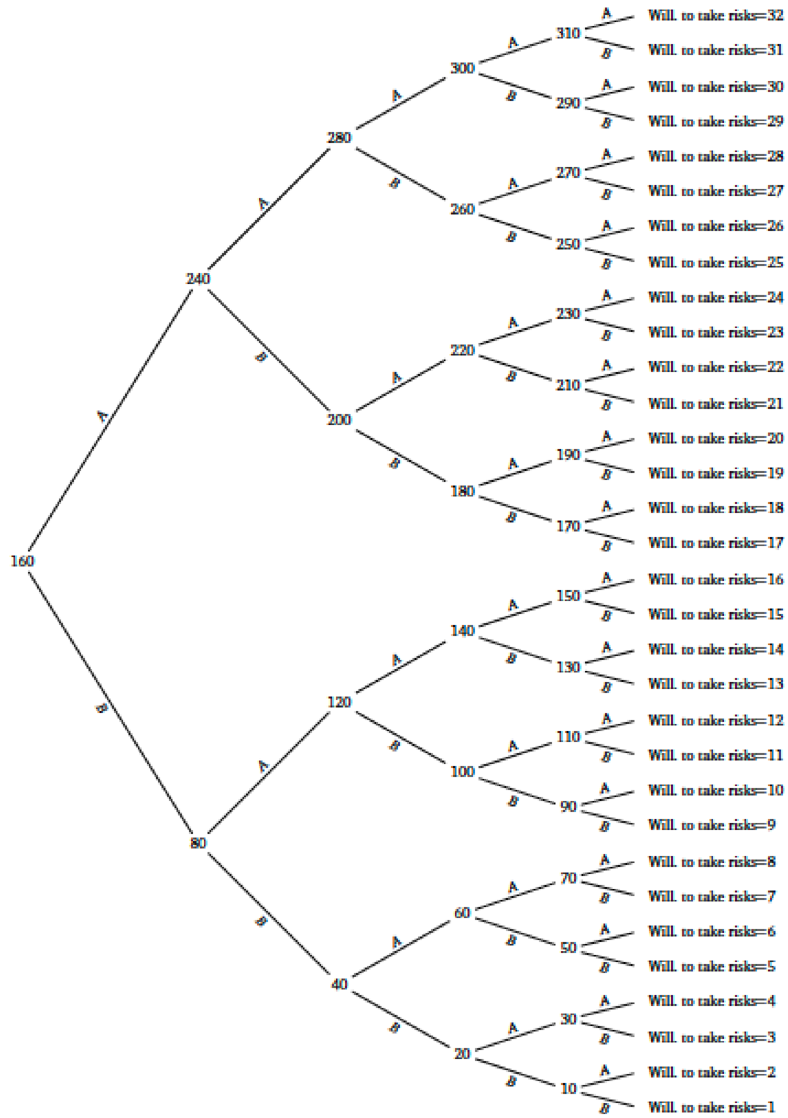


**Notes:** Tree for the staircase risk task (numbers = sure payment, A = choice of sure payment, B = choice of lottery). The staircase procedure worked as follows. First, each respondent was asked whether they would prefer to receive 160 euros for sure or whether they preferred a 50:50 chance of receiving 300 euros or nothing. In case the respondent opted for the safe choice (“B”), the safe amount of money being offered in the second question decreased to 80 euros. If, on the other hand, the respondent opted for the gamble (“A”), the safe amount was increased to 240 euros. Working further through the tree. Image and notes from **?**. follows the same logic..

# B Sample distribution for risk aversion measures

Figure 8: Histograms of Measured Risk Aversion buckets across IFLS4 and IFLS5, and MxFLS4 and MxFLS5

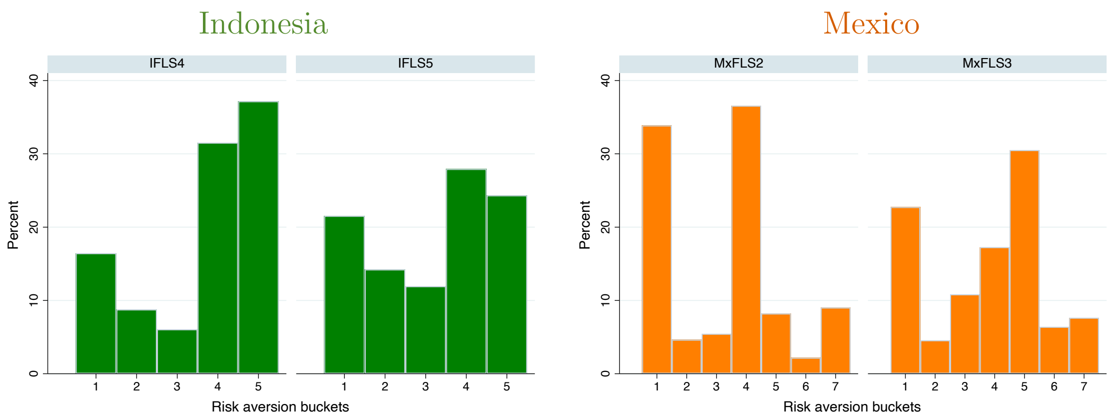


*Measured Risk Aversion*: 1-5 (Indonesia), 1-7 (Mexico), 5 or 7 highest measured risk aversion. Distributions for the primary sample, comprising individuals who: (1) have risk information on both waves of the survey; (2) are born after 1960 in Indonesia and 1925 in Mexico; and (3) either do not report being "gamble averse" in our main measure, or report being "gamble-averse" and also report the highest level of risk aversion in a secondary measure (Indonesia), or do not report being gamble averse in the second wave of the survey (Mexico).

# C Alternative Baseline Specifications

Table 8: Binarized

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | Indonesia | | | Mexico | | |
| Δ Mean Temp | -1.15$^{††}$ | | -1.32$^{††}$ | -0.37$^{***}$ | | -0.34$^{***}$ |
| | (0.15) | | (0.18) | (0.08) | | (0.08) |
| Δ Std. Dev. Temp | | -0.66 | 1.17 | | 0.59$^{***}$ | 0.59$^{***}$ |
| | | (0.64) | (0.69) | | (0.15) | (0.16) |
| Δ Mean Precip | -0.09$^{***}$ | | -0.05 | -0.24 | | -0.92$^{**}$ |
| | (0.03) | | (0.03) | (0.29) | | (0.32) |
| Δ Std. Dev. Precip | | -0.07 | -0.16$^{*}$ | | 0.29 | 0.58$^{**}$ |
| | | (0.06) | (0.08) | | (0.17) | (0.20) |
| Observations | 15044 | 15044 | 15044 | 10218 | 10218 | 10218 |

Measured risk aversion reported from 0 to 1, where 1–3 and 4–5 are mapped to 0 and 1 respectively. 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 9: Ordered probit

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | **Indonesia** | | | **Mexico** | | |
| Δ Risk Aversion | | | | | | |
| Δ Mean Temp | -1.94$^{\dagger\dagger}$ | | -2.40$^{\dagger\dagger}$ | -0.53$^{***}$ | | -0.51$^{***}$ |
| | (0.24) | | (0.29) | (0.11) | | (0.11) |
| Δ Std. Dev. Temp | | 1.40 | 4.43$^{***}$ | | 0.27 | 0.30 |
| | | (1.08) | (1.19) | | (0.23) | (0.24) |
| Δ Mean Precip | -0.21$^{***}$ | | -0.14$^{**}$ | -0.59 | | -1.73$^{***}$ |
| | (0.04) | | (0.05) | (0.42) | | (0.50) |
| Δ Std. Dev. Precip | | -0.08 | -0.20 | | 0.46 | 1.02$^{**}$ |
| | | (0.12) | (0.14) | | (0.26) | (0.32) |
| Observations | 15044 | 15044 | 15044 | 10218 | 10218 | 10218 |

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. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

Table 10: Clustering at state (province) level

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | **Indonesia** | | | **Mexico** | | |
| Δ Mean Temp | -3.98* | | -4.86* | -1.22* | | -1.17* |
| | (1.46) | | (1.84) | (0.50) | | (0.49) |
| Δ Std. Dev. Temp | | 2.07 | 8.25 | | 0.87 | 0.91 |
| | | (6.24) | (5.25) | | (1.49) | (1.67) |
| Δ Mean Precip | -0.40* | | -0.28 | -1.14 | | -3.61 |
| | (0.19) | | (0.25) | (1.65) | | (2.30) |
| Δ Std. Dev. Precip | | -0.15 | -0.40 | | 1.04 | 2.18 |
| | | (0.84) | (1.02) | | (1.34) | (1.48) |
| Observations | 15044 | 15044 | 15044 | 10218 | 10218 | 10218 |

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 province/state of birth level in parentheses. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

Table 11: Cross-section

| Dep. Var: Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | **Indonesia** | | | **Mexico** | | |
| Mean Temp | -1.00*** | | -1.11$^{\dagger}$ | -0.11* | | -0.07 |
| | (0.20) | | (0.21) | (0.05) | | (0.05) |
| Mean Precip | -0.02 | | -0.00 | 0.13 | | -0.20 |
| | (0.02) | | (0.03) | (0.11) | | (0.26) |
| Std. Dev. Temp | | 0.33 | -0.36 | | 0.29* | 0.22 |
| | | (0.39) | (0.42) | | (0.14) | (0.15) |
| Std. Dev. Precip | | -0.00 | -0.06 | | 0.11 | 0.23 |
| | | (0.03) | (0.04) | | (0.06) | (0.14) |
| Observations | 48263 | 48263 | 48263 | 24400 | 24400 | 24400 |

Measured risk aversion reported from 1–5, 5 being highest measured risk aversion. Dependent variables: 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. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

Table 12: Main results (with survey weights)

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| | Indonesia | | | Mexico | | |
| Δ Mean Temp | -3.69$^{\dagger\dagger}$ | | -4.40$^{\dagger\dagger}$ | -1.71$^{***}$ | | -1.70$^{\dagger}$ |
| | (0.47) | | (0.54) | (0.34) | | (0.33) |
| Δ Std. Dev. Temp | | 0.10 | 6.25$^{**}$ | | 1.89$^{**}$ | 1.71$^{**}$ |
| | | (2.04) | (2.22) | | (0.60) | (0.63) |
| Δ Mean Precip | -0.40$^{***}$ | | -0.30$^{**}$ | 0.90 | | -2.14 |
| | (0.09) | | (0.10) | (1.27) | | (1.45) |
| Δ Std. Dev. Precip | | -0.16 | -0.32 | | 1.92$^{**}$ | 2.59$^{**}$ |
| | | (0.22) | (0.24) | | (0.73) | (0.86) |
| Observations | 15044 | 15044 | 15044 | 9895 | 9895 | 9895 |

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. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

# D Subgroup heterogeneity tables

Table 13: Agricultural occupation split

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Indonesia | | Mexico | |
| Δ Mean Temp | -4.30$^{\dagger}$ | -5.85$^{\dagger}$ | -0.96$^{***}$ | -1.41$^{*}$ |
| | (0.65) | (1.04) | (0.29) | (0.55) |
| Δ Std. Dev. Temp | 8.46$^{**}$ | 9.26$^{*}$ | 0.58 | 3.18$^{**}$ |
| | (2.56) | (3.98) | (0.60) | (1.15) |
| Δ Mean Precip | -0.30$^{*}$ | -0.20 | -4.02$^{***}$ | -4.67 |
| | (0.13) | (0.15) | (1.18) | (2.83) |
| Δ Std. Dev. Precip | -0.22 | -0.74$^{*}$ | 2.10$^{**}$ | 4.52$^{**}$ |
| | (0.32) | (0.36) | (0.78) | (1.72) |
| Observations | 9945 | 5097 | 8219 | 1976 |
| Non-farmer | X | | X | |
| Farmer | | X | | X |

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. $^{*}$ $p < .05$, $^{**}$ $p < .005$, $^{***}$ $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

Table 14: Urban/rural split

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | **Indonesia** | | **Mexico** | |
| Δ Mean Temp | -3.81$^{\dagger}$ | -6.45$^{\dagger\dagger}$ | -2.80$^{\dagger}$ | -0.76 |
| | (0.67) | (0.84) | (0.52) | (0.44) |
| Δ Std. Dev. Temp | 8.13$^{**}$ | 7.58$^{*}$ | -0.54 | 1.07 |
| | (2.88) | (3.66) | (0.87) | (0.70) |
| Δ Mean Precip | -0.29$^{*}$ | -0.31$^{*}$ | -4.93$^{*}$ | -3.33 |
| | (0.13) | (0.14) | (1.94) | (1.84) |
| Δ Std. Dev. Precip | -0.31 | -0.50 | 1.41 | 3.21$^{**}$ |
| | (0.31) | (0.37) | (1.14) | (1.07) |
| Observations | 8841 | 6203 | 3346 | 4814 |
| Urban | X | | X | |
| Rural | | X | | X |

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. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

Table 15: Consumption levels split

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | **Indonesia** | | **Mexico** | |
| Δ Mean Temp | -5.07$^{\dagger}$ | -4.81$^{\dagger}$ | -0.53 | -1.61$^{***}$ |
| | (0.86) | (0.75) | (0.38) | (0.32) |
| Δ Std. Dev. Temp | 8.58$^{*}$ | 8.69$^{**}$ | 1.45$^{*}$ | 0.34 |
| | (3.35) | (3.07) | (0.67) | (0.73) |
| Δ Mean Precip | -0.34$^{*}$ | -0.23 | -3.97$^{*}$ | -3.17 |
| | (0.15) | (0.13) | (1.61) | (1.63) |
| Δ Std. Dev. Precip | -0.11 | -0.62$^{*}$ | 3.15$^{***}$ | 1.44 |
| | (0.35) | (0.30) | (0.94) | (0.98) |
| Observations | 7522 | 7522 | 5109 | 5109 |
| Lower 50% | X | | X | |
| Upper 50% | | X | | X |

Sample cut based on consumption levels in second wave of the panel. Lower 50% (higher 50%) refers to individuals in the lower half (upper half) of the household consumption distribution. 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, $^{\dagger}$ p $< 5 \times 10^{-7}$, $^{\dagger\dagger}$ p $< 5 \times 10^{-13}$.

Table 16: Educational attainment split

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Indonesia | | Mexico | |
| Δ Mean Temp | -7.32$^{\dagger\dagger}$ | -2.96$^{***}$ | -1.14$^{***}$ | -1.24$^{***}$ |
| | (0.82) | (0.71) | (0.30) | (0.38) |
| Δ Std. Dev. Temp | 15.18$^{***}$ | 3.19 | 1.01 | 0.18 |
| | (3.33) | (3.00) | (0.59) | (0.98) |
| Δ Mean Precip | -0.13 | -0.43$^{**}$ | -3.88$^{*}$ | -2.14 |
| | (0.13) | (0.13) | (1.54) | (1.47) |
| Δ Std. Dev. Precip | -0.98$^{**}$ | 0.12 | 2.18$^{*}$ | 2.02$^{*}$ |
| | (0.34) | (0.32) | (0.87) | (1.02) |
| Observations | 8184 | 6860 | 7531 | 2687 |
| Low | X | | X | |
| High | | X | | X |

Low: defined as individuals with less than high school education. High: high school education or greater. 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 17: Gender split

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | **Indonesia** | | **Mexico** | |
| Δ Mean Temp | -3.98*** | -5.59†† | -1.20*** | -1.14*** |
| | (0.85) | (0.74) | (0.36) | (0.32) |
| Δ Std. Dev. Temp | 4.86 | 11.28*** | 2.39** | -0.20 |
| | (3.38) | (3.23) | (0.78) | (0.66) |
| Δ Mean Precip | -0.48** | -0.12 | -4.10* | -3.11* |
| | (0.15) | (0.13) | (1.65) | (1.37) |
| Δ Std. Dev. Precip | 0.06 | -0.75* | 2.74** | 1.70 |
| | (0.36) | (0.32) | (0.97) | (0.88) |
| Observations | 6681 | 8362 | 4182 | 6019 |
| Male | X | | X | |
| Female | | X | | X |

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 18: Age split

| Dep. Var: Δ Risk Aversion | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | **Indonesia** | | **Mexico** | |
| Δ Mean Temp | -4.38$^{\dagger}$ | -8.22$^{\dagger}$ | -1.25$^{***}$ | -0.93 |
| | (0.65) | (1.43) | (0.28) | (0.58) |
| Δ Std. Dev. Temp | 9.74$^{**}$ | 15.69$^{***}$ | 0.17 | 1.97$^{*}$ |
| | (2.97) | (4.29) | (0.72) | (0.89) |
| Δ Mean Precip | -0.36$^{**}$ | 0.43 | -2.19 | -4.68 |
| | (0.13) | (0.27) | (1.31) | (2.75) |
| Δ Std. Dev. Precip | -0.41 | -1.44$^{*}$ | 2.00$^{*}$ | 2.58 |
| | (0.32) | (0.64) | (0.80) | (1.41) |
| Observations | 6922 | 8122 | 4601 | 5617 |
| Young | X | | X | |
| Old | | X | | X |

Young: defined as individuals up to 40. Old: defined as older than 40. 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. * $p < .05$, ** $p < .005$, *** $p < .0005$, $^{\dagger}$ $p < 5 \times 10^{-7}$, $^{\dagger\dagger}$ $p < 5 \times 10^{-13}$.

# E Details of additional controls

Table 19: Description of controls included in regressions in subsubsection 4.1.2

| Category | Variables Included |
|---|---|
| Demographics (**Indonesia** and **Mexico**) | Married<br>Household Size<br>Household Size Squared<br>Educational Attainment |
| Income, Assets and Net Yearly Savings (**Indonesia** and **Mexico**) | Total household income<br>Total value of household assets<br>Net Households Savings (Savings-Borrowing) |
| Violence (**Indonesia**) | Perceived safety level of village<br>Perceived safety of walking in village alone at night<br>Occurrence of civil strife in household's region of residence in last 5 years<br>Civil strife severe enough to cause death, major injury,<br>direct financial loss, or relocation of any member of HH |
| Violence (**Mexico**) | Feels safe at home<br>Fear of assault during the day<br>Fear of assault at night<br>No. of times robbed, assaulted, kidnapped<br>No. of family/friends robbed, assaulted, kidnapped in last 12 months<br>Homicide Rate in Last 12 Months in Municipality<br>of Residence in MxFLS2 (built by **?**) |
| Natural Disasters (**Indonesia**) | Occurrence of natural disaster in household's region of residence in last 5 years<br>Natural disaster severe enough to cause death, major injury,<br>direct financial loss, or relocation of any member of HH |
| Natural Disasters (**Mexico**) | Household/business lost due to natural disaster |
| Macroeconomic Experiences (**Indonesia** and **Mexico**) | State-level real GDP growth mean and standard deviation<br>in subject's state of birth, from birth year to survey measurement<br>year (built by **?**) |

Income, assets and net yearly savings are at the household level, and inflation-adjusted to local currency in the first wave of the survey (millions of rupiah of 2007 in Indonesia and pesos of 2005 in Mexico).

# F Details and robustness of structural estimation for welfare analysis

To calculate the EDEs we assume that individuals have isoelastic utility, as defined in Equation 6, and allow the CRRA parameter $\theta_j$ to be individual-specific.

$$u_j(s) = \begin{cases} \frac{s^{1-\theta_j}-1}{1-\theta_j}, & \theta_j \neq 1 \\ \log(s), & \theta_j = 1 \end{cases} \tag{6}$$

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 6, 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 $\theta_j$ assuming that individuals are maximizing their expected utility. We assign the average of these bounds to each $\theta_j$. For individuals that select into the most and least risk averse bins, we cannot observe a finite upper-bound or lower-bound respectively.[26] To bound these, we make either a tight or broad assumption on these bounds.[27] Finally, for the IFLS, we consider different choices of $\theta_j$ 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.[28]

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.

Once the risk aversion parameters are assigned to individuals, we calculate the EDE as follows. 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 $\theta_j$, we want to find the certainty equivalent for individual $j$ over the distribution of $S$ in society. Let individuals be indexed

[26] 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 $\theta = \infty$.

[27] 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.

[28] 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.

by $i$. Then individual $j \in J$ has a certainty equivalent given by the following:

$$CE_j = \left( \frac{1}{N} \sum_{i \in I} s_i^{1-\theta_j} \right)^{\frac{1}{1-\theta_j}}. \tag{7}$$

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 $\theta_j$ are calculated as defined in Equation 7.

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 (**?**). We interpret this as the equally-distributed equivalent for the distribution of per capita or household income.

This procedure gives us a single EDE for each specification. We present results for these statistics across the universe of specifications we consider in the following figures. Although the estimates of the EDEs (and their differences) vary by the methodological choices we make, qualitatively they are quite similar to the results in our preferred specification.

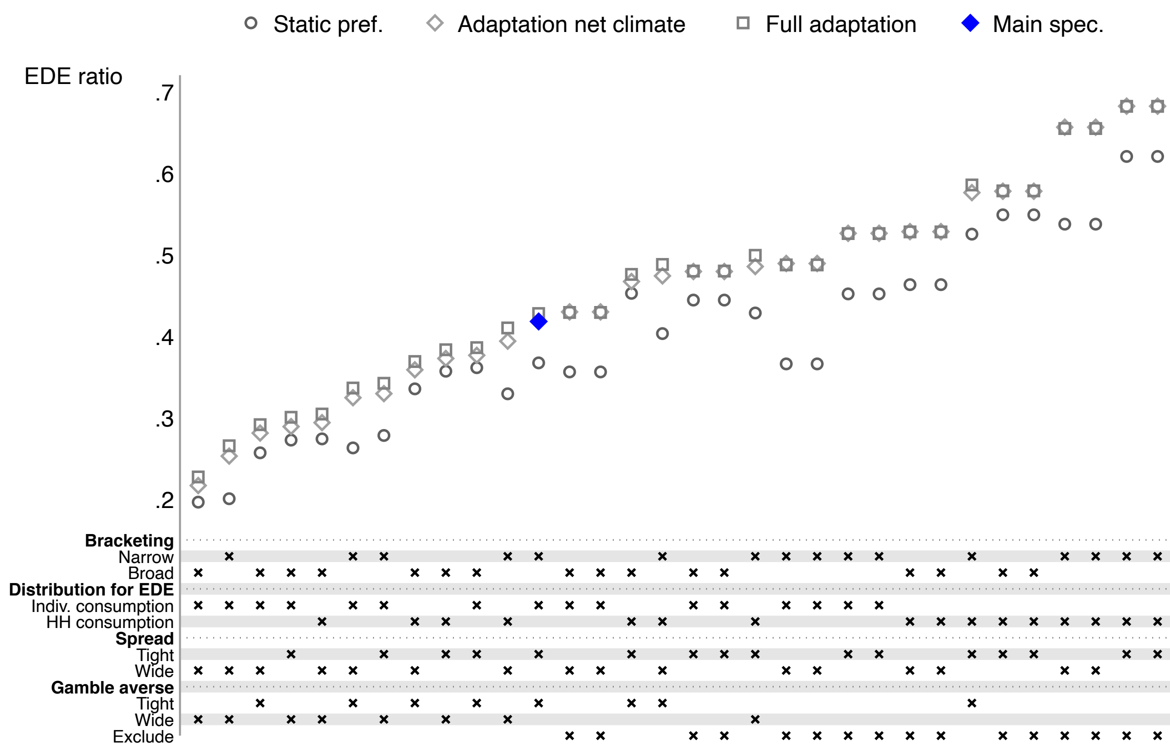


Figure 9: Climate and risk adaptation, Indonesia

**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.

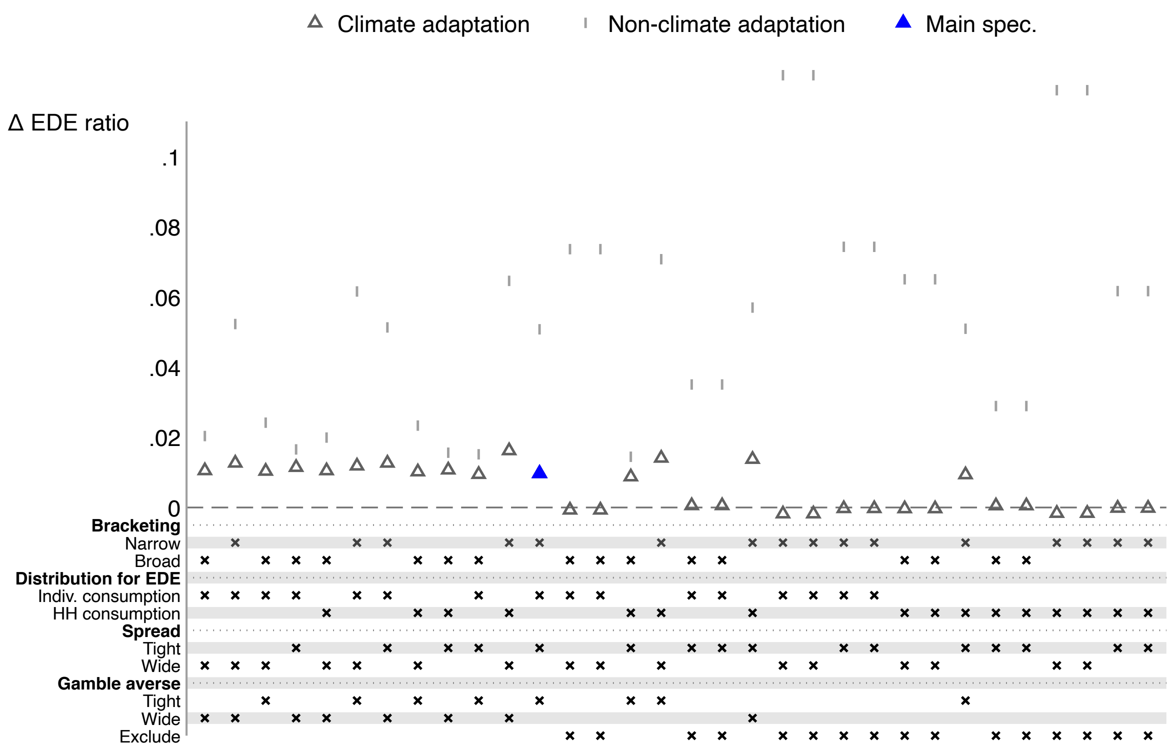


Figure 10: Climate and non-climate adaptation, Indonesia

**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.

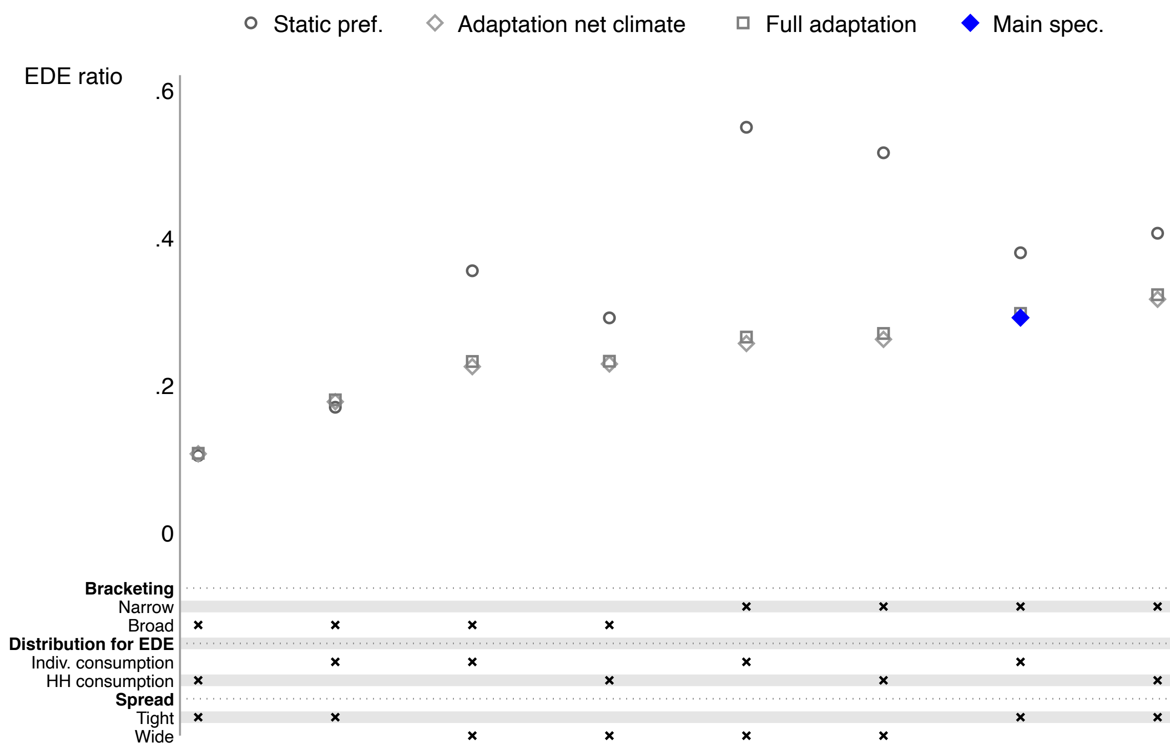


Figure 11: Climate and risk adaptation, Mexico

**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.

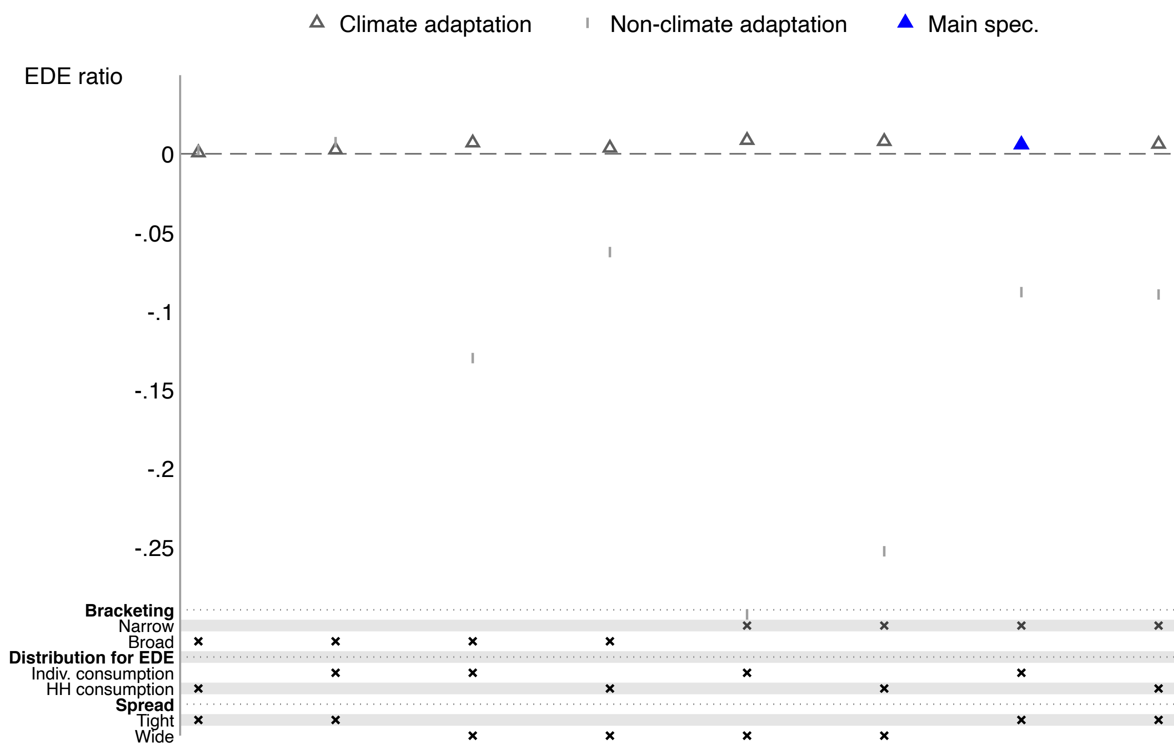


Figure 12: Climate and non-climate adaptation, Mexico

**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.