University of Connecticut

Department of Economics Working Paper Series

Risk-Taking Adaptation to Macroeconomic Experiences: Theory and Evidence from Developing Countries

by

Remy Levin
University of Connecticut

Daniela Vidart
University of Connecticut

Working Paper 2021-09
May 2021

365 Fairfield Way, Unit 1063
Storrs, CT 06269-1063
Phone: (860) 486-3022
Fax: (860) 486-4463
http://www.econ.uconn.edu/

Risk-Taking Adaptation to Macroeconomic Experiences: Theory and Evidence from Developing Countries*

Remy Levin

University of Connecticut

Daniela Vidart

University of Connecticut

October 28, 2020

Abstract

How do lifetime experiences of macroeconomic risk shape attitudes towards risk? We study this question theoretically and empirically for individuals in developing countries. We build a Bayesian model of choice in which agents' risk attitude adapts to their evolving beliefs about background risk. Our model predicts that risk aversion will increase monotonically in the variance of the background risk, and will decrease convexly in the mean. We test the model by linking longitudinal surveys from Indonesia and Mexico, containing elicited measures of risk aversion for the same subjects years apart, with state-level real GDP growth time series capturing their lifetime macroeconomic experiences. In both countries measured risk aversion significantly increases in experienced growth volatility and significantly decreases in experienced mean growth. The effect of volatility is 0.9-4.3 times the effect of the mean, indicating that experiences of volatility are first-order drivers of risk attitudes.

Keywords: Risk attitudes, experience effects, macroeconomic volatility, development
JEL Codes: D14, D81, D83, E32, G11, O11, O12

*Corresponding author: Remy Levin, remy.levin@uconn.edu. Vidart: daniela.vidart@uconn.edu. We are deeply grateful for the generous and insightful feedback of numerous colleagues, including Prashant Bharadwaj, Pol Campos-Mercade, Mitch Downey, James Fowler, Simone Galperti, Alex Imas, Sam Krumholz, David Lagakos, Munseob Lee, Mark Machina, Tommaso Porzio, Valerie Ramey, Natalia Ramondo, Wayne Sandholtz, Joel Sobel, Isabel Trevino, and Johannes Wieland, as well as participants at several conferences and numerous UCSD workshops. All mistakes are our own.

"Sometimes, the mean is not the message."

– Zeynep Tufekci

I Introduction

Developing countries are characterized by low per-capita incomes. They also exhibit more macroeconomic volatility, on average, than developed countries. Studies have shown that rates of risk-taking behavior, such as the adoption of new agricultural technology (Munshi (2004), Giné and Yang (2009)), and migration (Bryan, Chowdhury and Mobarak (2014)), are often surprisingly low in developing countries. The standard explanations for the co-occurrence of high volatility and low risk-taking focus on incomplete risk-sharing institutions, credit constraints, and low buffer-stocks of savings in these settings (Rosenzweig and Binswanger (1993), Rosenzweig and Wolpin (1993), Karlan et al. (2014), Deaton (1991)). In this paper we build a theory of- and present empirical evidence for a novel mechanism linking these two phenomena: lifetime experiences of macroeconomic volatility directly change individual risk attitudes.

Our focus on experiences of volatility, in particular, is driven by the following intuition. In expected utility theory risk preference can be thought of as the propensity of an agent to substitute between the mean and the variance of a lottery (Pratt (1964), Arrow (1970)). It stands to reason that if lifetime experiences can change risk attitudes, it would be experiences of both these moments that should matter, since together they represent risk for the agent. The central insight of this paper is that lifetime experiences of risk shape risk attitudes. In other words, individuals' propensity to take risks is affected not only by whether conditions in their environment improve or deteriorate, but also by whether they stabilize or destabilize.

That lifetime experiences, and particularly macroeconomic experiences, can have large and persistent effects on risk attitudes is not a new idea in economics. In their seminal paper Malmendier and Nagel (2011) show that cohorts who lived through the great depression exhibited lower rates of financial risk-taking decades after the fact. This basic relationship between experienced recessions and decreased risk-taking has since been replicated in a number of studies in developed countries (Sahm (2012), Dohmen, Lehmann and Pignatti (2016), Ampudia and Ehrmann (2017), Guiso, Sapienza and Zingales (2018), Shigeoka (2019)). A parallel literature centered on the developing world has shown that traumatic events, particularly experiences of violence (Callen et al. (2014), Jakiela and Ozier (2019), Brown et al. (2019)), and of natural disasters (Cameron and Shah (2015), Brown et al. (2018), Hanaoka, Shigeoka and Watanabe (2018)) can also decrease risk-taking for affected individuals.

We make two primary contributions to this emerging literature. First, we build a theoretical framework for understanding a phenomenon that other papers have, to date, studied chiefly through an empirical lens. Our model explains existing findings on the effects of exogenous shocks on risk-taking as arising from a process of Bayesian learning about risk in the agent's environment. It also generates new testable hypotheses about how individual risk attitudes will change when shocks arrive, given the agent's body of lifetime experiences. Our second contribution is to test the model's main hypotheses empirically. We use large-scale, longitudinal surveys from Indonesia and Mexico to link within-person changes in measured risk aversion to state-level real GDP growth time series capturing subjects' lifetime macroeconomic experiences. In line with the model's predictions, we find that measured risk aversion increases in experienced growth volatility and decreases in experienced mean growth. Our empirical findings complement the existing literature by providing the first robust evidence that macroeconomic fluctuations, and not just experiences of trauma, can have persistent effects on the risk attitudes of individuals in the developing world. These findings are particularly important given the high levels of macroeconomic volatility that are common in these settings.

We begin by developing a dynamic model of choice in which an agent's risk attitude adapts to her evolving beliefs about risk in the environment. In our model an agent chooses an objective income lottery each period (the foreground risk) from a fixed, time-invariant menu. She makes this choice while exposed, each period, to an exogenous and uninsurable income shock. The income shock is a realization of a subjective income lottery that we call the background risk. Imagine, for example, a farmer who must choose each planting season between planting a risky cash crop and a relatively safe food crop, while confronting the possibility of droughts, heatwaves, and price fluctuations that could spell riches or ruins.

We assume that our agent is an expected utility maximizer. Standard conditions on the higher moments of her von-Neumann Morgenstern utility function imply that uncertainty in the foreground and background risks, which are statistically independent by construction, are substitutes for the agent (Kimball (1993), Gollier and Pratt (1996), Eeckhoudt, Gollier and Schlesinger (1996)). In our example, this means that the more risk the agent believes exists in the environment, the less likely she will be to plant the risky cash crop.

To make our framework dynamic we model the evolution of the agent's beliefs about the background risk over time. Departing from the existing background risk literature and the full information benchmark, we do not assume that the agent knows the true data generating process of the income shock. Rather, we think of the agent as using a simple heuristic for the distribution of the background risk that still allows her beliefs about the uncertainty contained in it to change over time. Specifically, we model the agent's beliefs distribution about the likelihood as a stationary Gaussian random variable, for which she knows neither the mean nor the variance. The agent observes realizations of the background risk each period and updates her beliefs about its moments as a Bayesian.[1] As her beliefs about the background risk change, her choices over the endogenous foreground risk change in turn. To close the model, we assume that the agent holds a conjugate prior over the background risk

[1] Bayesian learning about a simplified or incorrect data generating process means that our framework falls into the class of "quasi-Bayesian" models of bounded rationality, by the terminology of Rabin (2013).

distribution in the first period. This allows us to solve for the agent's posterior analytically, making the model highly tractable.[2]

Our model makes several sharp predictions about the effects of new realizations of the background risk on foreground risk-taking, given the agent's previous body of risk experiences. Its three main predictions are that foreground risk aversion will (1) increase monotonically in mean-preserving increases in the variance of the background risk, (2) will decrease in variance-preserving increases in the mean, and (3) will be convex in the posterior mean. Thus, the overall effect of new shocks on risk-taking in our model will depend not only on whether they are good or bad relative to the agent's experiences, but also on whether they are stabilizing or destabilizing.[3]

We test the predictions of our model for lifetime experiences of macroeconomic risk using data from two large developing countries, Indonesia and Mexico. Together, these countries form an advantageous setting for our empirical analysis, because they share a common recent history of volatile and internally-heterogeneous economic growth, but differ along many plausibly relevant dimensions such as geography, language, and dominant religion. This means that our data are likely to contain sufficient variation in macroeconomic experiences to ensure the internal validity of our analysis, while at the same time allaying most concerns about the external validity of our results.

The microdata in our analysis come from the Indonesian Family Life Survey and the Mexican Family Life Survey. Both are high-quality, longitudinal surveys (total n = 25,459),

[2]Our model has the notable property that while the conditional mean distribution is Gaussian, the unconditional mean distribution is a student's-t, which is fat tailed. This means that our framework can also be thought of as a tractable model of the dynamics of risk attitudes due to learning over the mean of a fat-tailed distribution. Fat-tailed distributions have generated renewed interest in the last few years across several fields, including macroeconomics (Gabaix et al. (2006), Morris and Yildiz (2016)), finance (Gabaix et al. (2003), Kelly and Jiang (2014)), and the economics of climate change (Weitzman (2007), Weitzman (2009)). Particularly relevant to our work is recent evidence that macroeconomic variables, especially GDP, are likely distributed with fat tails (Acemoglu, Ozdaglar and Tahbaz-Salehi (2017)).

[3]This means that our model can be used to think about the relative effects of large and small shocks on risk-taking, a fact that distinguishes our theoretical results from those of Dillenberger and Rozen (2015), where risk attitudes are affected only by whether experienced shocks are above or below the prior mean.

each containing two elicited measures of risk aversion for the same subjects several years apart (IFLS: 2007 and 2014; MxFLS: 2006 and 2009). Individual risk aversion is measured in these surveys with a staircase instrument similar to the experimentally-validated measure of Falk et al. (2018), composed of a series of hypothetical, high-stakes choices between a sure amount of income and a fair coin toss over a higher and lower amount. The main dependent variables in our analysis are within-subject changes in elicited risk aversion.[4]

To capture lifetime macroeconomic experiences we assign to each individual the time series of real GDP growth from birth to measurement in their state of birth.[5] We then calculate the mean and standard deviation of these time series for each subject from birth to measurement. The main independent variables in our analysis are changes in the lifetime standard deviation, mean, and mean squared of growth between the waves of the surveys. Our preferred specification includes individual and time fixed effects, as well as a variable capturing sub-national inflation to control for changes in the nominal values of the lotteries comprising our risk aversion measures.

The empirical findings in our main analysis (subsubsection C.1) provide strong support for the first two predictions of our model. In both countries increases in the lifetime variance of growth are significantly associated with increases in measured risk aversion [*Indonesia:* $p = 1.78 \times 10^{-15}$; *Mexico:* $p = 1.61 \times 10^{-7}$], while increases in the lifetime mean of growth are

[4]Employing elicited risk aversion measures in our main analysis provides two identification benefits relative to using real-world risky choice data in this setting. First, in real-world data changes in the menu of risky choices may be spuriously correlated with macroeconomic experiences. For example, in the context of portfolio choice, both growth dynamics and available investment vehicles (and therefore individual investment rates) can be driven by government policy-making. In our elicited measures the menu of lotteries offered subjects is independent by construction from their macroeconomic experiences, eliminating this potential source of omitted variable bias. Second, with real-world choice data subject beliefs about the odds and payoffs of the lotteries involved are usually unobserved. This raises the possibility that estimated changes in risk-taking occur due to changes in beliefs about the foreground risk. Since the odds and payoffs of the lotteries in our elicited measures are known to the subjects, we can foreclose this interpretation of our results, and more cleanly identify the background risk mechanism in our model.

[5]We use macroeconomic 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 conditions in state of residence.

associated with significant decreases in measured risk aversion [*Indonesia:* $p = 1.49 \times 10^{-5}$; *Mexico:* $p = 1.70 \times 10^{-7}$]. In both countries the magnitude of the effects of experienced variance is comparable to- or larger than that of the experienced mean (0.9 times as large in absolute terms in Mexico, 4.3 times as large in Indonesia), indicating that experienced variance is a first-order driver of risk attitudes. These results are robust to a variety of alternative specifications.

We find mixed evidence for the model's third prediction. In our Mexican sample the effect of experienced changes in the squared mean of growth on measured risk aversion is positive (in line with the model) and marginally significant [$p = 0.07$], while in the Indonesian sample it is negative, though again marginally significant [$p = 0.09$].

An important question given our main results is whether the observed changes in measured risk aversion of subjects are driven by changes in their personal economic circumstances, or by shifts in their underlying risk attitudes. We examine this question in our secondary analysis (subsubsection C.2), where we add additional controls to the regressions. Controlling for changes in subject income, assets, and savings between the waves of each survey does not meaningfully affect the significance or magnitude of our main results. This indicates that our findings are unlikely to be driven by income effects. We also control for changes in other covariates that could plausibly drive risk attitudes, such as subject demographics (marital status, household size, educational attainment), experiences of violence, and experiences of natural disasters. As with the economic measures, including these controls in the analysis does not meaningfully affect our main results.

Finally, in subsubsection C.3 we examine whether changes in measured risk aversion correlate with changes in subjects' risk-taking behavior in other domains. We construct a variable capturing the predicted change in measured risk aversion according to our main regression model, and examine its association with four kinds of risk-taking behavior in our data: smoking, having ever migrated across state lines, self-employment status, and (in

Indonesia) investment behavior, measured by whether subjects report planting cash crops. Predicted increases in measured risk aversion correlate with decreases in risk-taking behavior in almost all of these domains across both countries. Associations are particularly strong with changes in smoking, migration, and self-employment behavior in Indonesia. These results imply that lifetime experiences of macroeconomic risk affect not just financial risk aversion, as estimated by choices over hypothetical lotteries, but the risk attitudes of subjects more generally. They also suggest that a domain-general propensity to take risks might in fact exist, a hotly contested proposition in the psychology literature on risk-taking (Weber, Blais and Betz (2002), Frey et al. (2017)).

The remainder of the paper proceeds as follows. In section II we describe the structure of our model and its theoretical predictions. In section III we describe our empirical analysis, including details on our data and methodology (subsection A), the identification strategy (subsection B), results (subsection C), and robustness (subsection D)). We conclude in section IV.

II A Model of Adaptive Risk-Taking

We develop a dynamic model of risky choice in which an agent's endogenous risk-taking adapts to evolving beliefs about unknown background risk in their environment. We link changes in the agent's history of experienced shocks to changes in their risky decision-making in two steps. First, we model how changes in the agent's beliefs about the environment affect their decision-making using a static background risk framework. The key feature of this model of utility is that background risk in the environment and foreground risk in the agent's choice problem are substitutes. Second, we model how experienced shocks affect the agent's beliefs about risk using a Bayesian updating process over structural uncertainty. The central aspect of this model of learning is that the agent directly forms beliefs about risk in the environment, meaning that they learn about both the mean and the variance of the background risk from its realizations over their lifetime.

A Model

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

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

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

Figure 1: Timing of events in the model

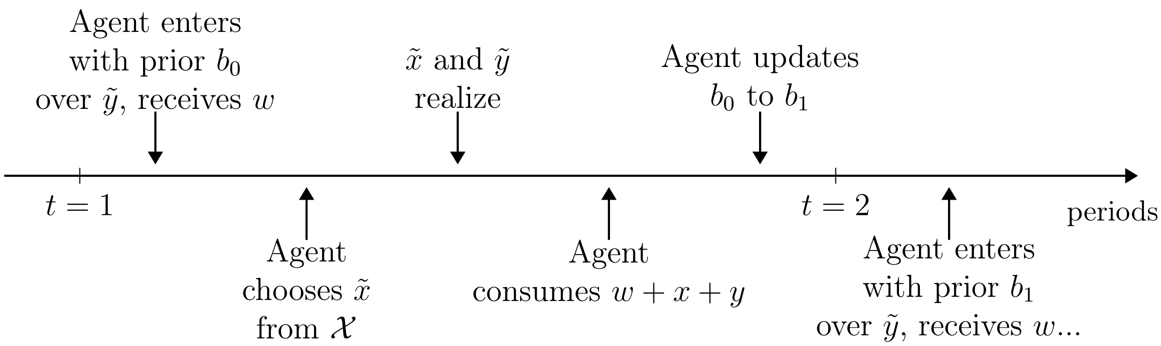

their prior b_{t-1} to posterior b_t, which forms their prior in the next period.

Utility and risk. We assume that the agent is a subjective expected utility maximizer and has a four-times-differentiable utility function u for which $u' > 0$ and $u'' < 0$. u is defined over the wealth endowment, the foreground risk, and the background risk:

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

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

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

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

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

It is also useful to define two higher-order analogues of $r_t(w)$, the coefficient of absolute prudence $p_t(w) = -u'''(w|B_t(y))/u''(w|B_t(y))$ and the coefficient of absolute temperance $q_t(w) = -u''''(w|B_t(y))/u'''(w|B_t(y))$. These allow us to identify conditions on the third and fourth moments of u that are collectively termed *risk vulnerability*.[6]

Definition II.1. (*Risk-vulnerable utility*) An expected utility maximizer with $u' > 0$ and $u'' < 0$ is risk-vulnerable at time t if $p_t(w) \geq r_t(w)$ and $q_t(w) \geq r_t(w)$.

Risk vulnerability is the feature of the utility function that ensures that background and foreground risks are substitutes for the agent. Intuitively it corresponds to higher-order concavity in the agent's utility function. Note that all HARA utility functions exhibit risk vulnerability.[7] We will assume below that the agent is risk vulnerable at all t.

Learning. The agent in our model is a Bayesian who uses personally observed realizations of the background risk to update their belief distribution $B_t(y)$. We make two structural assumptions about the agent's updating process. First, we assume that the agent believes that the realizations, or signals, are drawn from a stationary Gaussian random variable with unknown mean and unknown variance. Second, we assume that the agent's prior over the

[6]Gollier and Pratt (1996).

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

mean and variance takes the form of a normal-inverse-chi-squared distribution. We call this learning process Bayesian updating over structural uncertainty, and describe it formally in the following definition:

Definition II.2. (*Bayesian updating over structural uncertainty*) We say that a Bayesian agent is updating over structural uncertainty if:

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

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

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

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

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

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

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

$$\mathbb{E}_0[M] = \mu_0 \tag{1}$$

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

The $NI\chi^{-2}$ distribution is the unique conjugate prior of the Gaussian with unknown mean and unknown variance likelihood. This means that the Bayesian agent's posterior distribution upon receiving signals will also be in the $NI\chi^{-2}$ family, with updated parameters. Consequently, the agent's posterior mean and variance have closed form expressions. Let $\mathcal{D}_t = \{y_1, ..., y_t\}$ be a set of t iid draws from $\tilde{y}$. Then these posteriors will be:[8]

$$\mathbb{E}_t[M|\mathcal{D}_t] = \mu_t = \mu_0 + \frac{t}{\kappa_0 + t}(\bar{y}_t - \mu_0) \tag{3}$$

$$\mathbb{E}_t[\Sigma^2|\mathcal{D}_t] = \frac{\nu_t}{\nu_t - 2}\sigma_t^2 = \frac{1}{\nu_0 + t - 2}\left[\nu_0\sigma_0^2 + \sum_{i=1}^{t}(y_i - \bar{y}_t)^2 + \frac{t\kappa_0}{\kappa_0 + t}(\bar{y}_t - \mu_0)^2\right], \tag{4}$$

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

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

B Discussion of the model

Models of background risk, and their use in analyzing choice under multiple sources of risk, have a long tradition in economics, stemming from their introduction in Pratt and Zeckhauser (1987). Foundational work in this area identified conditions on the higher moments of

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

the agent's von-Neumann Morgenstern utility function that guarantee "natural" behavioral responses of risk-taking to independent sources of risk (Kimball (1990), Kimball (1993), Gollier and Pratt (1996), Caballé and Pomansky (1996)). These conditions, often termed prudence and temperance, generally imply that more risk in the environment will lead to lower endogenous risk-taking by the agent, as they do in our model, even if the two sources of risk are statistically independent. Substantial empirical evidence from experimental studies exists for these conditions (Noussair and Trautmann (2014), Deck and Schlesinger (2014), Ebert and Wiesen (2014)).

The theoretical literature on background risk has generally assumed that the odds and payoffs of these risks are known to the agent when they are making decisions. In reality, however, individuals often make high-stakes risky choices in the presence of substantial background risk over which they have limited information and considerable uncertainty. A wage worker deciding whether to start their own business must contend with the possibility of an economic downturn over the medium-term; a renter choosing whether to buy a house (and how much of it to buy) needs to consider the likely path of the housing market over the next few years; a farmer choosing between planting a risky cash crop and a relatively safe food crop has to confront the possibility of droughts, heatwaves, and price fluctuations that could spell riches or ruins. These risks are substantial, largely outside of individuals' control, and often difficult or impossible to forecast even for experts with access to sophisticated models and the best data.

We contribute to the background risk literature by building and studying a model of individual risky choice where the exact parameters of the background risk are unknown to the agent. Rather the agent in our model has priors over these parameters, observes realizations of the background risk over time, and updates as a Bayesian in the face of new information. While "risk" has many definitions, in its most classical formulation it involves a trade-off between the mean and variance of a gamble (Pratt (1964), Arrow (1970)). Since we are interested in the consequences of uncertainty and updating over the background risk, we model the agent as having beliefs over both its mean and variance that evolve over time.

The result is a model in which the history of shocks an agent experiences shapes their risk-taking over time. The theoretical literature on history-dependent risk attitudes is markedly scant, with the notable exception of Dillenberger and Rozen (2015). In that paper the authors construct a model where an agent recursively evaluates compound lotteries to classify experienced realizations as disappointing or elating relative to some threshold. The authors show that assuming choice consistency, the agent in their model exhibits a "reinforcement effect," where risk-taking will always increase after an elation and decrease after a disappointment. Importantly, the agent in Dillenberger and Rozen (2015) does not consider how big the realizations they observe are, but only whether they are good or bad. We contribute to this nascent literature by building the first model, to our knowledge, wherein the magnitude of an experienced realization affects risk attitudes. As we show below, because of the variance beliefs channel in our model, the magnitude of a realization can matter just as much as its sign for the risk-taking of the agent. This means that our model can be used to think about the relative effects of large and small shocks, and not just the relative effects of positive and negative realizations.

This model also contributes to a growing literature on fat-tailed distributions in economics. Fat-tailed distributions have generated renewed interest in the last few years across several fields, including macroeconomics (Gabaix et al. (2006), Morris and Yildiz (2016)), finance (Gabaix et al. (2003), Kelly and Jiang (2014)), and the economics of climate change (Weitzman (2007), Weitzman (2009))[9]. Recent evidence suggests that macroeconomic variables, in particular GDP, are likely distributed with fat tails (Acemoglu, Ozdaglar and Tahbaz-Salehi (2017)). An interesting feature of our model is that while the conditional

[9] We adopt the term "structural uncertainty" from Weitzman's work, who used a related updating framework to model societal uncertainty over the tail-risk of climate change.

marginal distribution of the agent's mean is normal, its unconditional marginal distribution is in fact a student's-t[10], which is fat-tailed. Our model can therefore be interpreted as examining the dynamics of risk-taking under updating over a fat-tailed mean distribution.

C The effect of mean-preserving changes in variance on risk-taking

We first examine the effect of mean-preserving changes in the variance of the background risk on the agent's risk-taking. These are, in a sense, the "purest" changes in risk that the agent experiences. We use *mean-preserving* here to refer to an aspect of the data, or signals, when compared to the agent's prior. Since this usage is slightly different than that commonly employed, we define it below:

Definition II.3. (*Mean-preserving dataset*) Let $\mathcal{D}_t = \{y_1, ..., y_t\}$ be a set of t iid draws from $\tilde{y}$. We say D_t is a mean-preserving dataset for the agent if $\bar{y}_t = \mu_0$.

With this definition in hand we can now state and prove our first result:

Proposition 1. (Monotonicity in mean-preserving variance changes) *Suppose the agent observes a mean-preserving dataset of signals $\mathcal{D}_t$. Then $r_t(w) > (\leq)\ r_0(w)$ iff $\delta_t^v > (\leq)\ 0$.*

Proof. Since $\mathcal{D}_t$ is a mean-preserving dataset, the agent's posterior mean equals their prior mean μ_0, and $\Delta_t M = 0$. The agent's posterior variance, meanwhile, reduces to the term $\frac{1}{\nu_0+t-2}\left[\nu_0\sigma_0^2 + \sum_{i=1}^t (y_i - \bar{y}_t)^2\right]$. Therefore, $\Delta_t \Sigma^2 > (\leq)\ 0$ iff $\delta_t^v > (\leq)\ 0$. This results in a mean-preserving spread (scrunch) in the agent's belief distribution $B_t(y)$, represented by $\tilde{y}_t$.

[10] A $t_{\nu_t}(M|\mu_t, \sigma_t^2/\kappa_t)$ distribution, to be precise.

Consider the effect of $\tilde{y}_t$ on the agent's absolute risk aversion:

$$
\begin{aligned}
\left.\frac{r_t(w)-r_0(w)}{r_0(w)}\right|_{\Delta_t\Sigma^2|\tilde{y}_t} &= [r_0(w)]^{-1}\left[\frac{-\mathbb{E}_t u''(w+\tilde{y}_t)}{\mathbb{E}_t u'(w+\tilde{y}_t)} - r_0(w)\right] \\
&= [r_0(w)]^{-1}\left[\frac{-\mathbb{E}_t[u''(w)+\tilde{y}_t u'''(w)+0.5\tilde{y}_t^2 u''''(w)]}{\mathbb{E}_t[u'(w)+\tilde{y}_t u''(w)+0.5\tilde{y}_t^2 u'''(w)]} - r_0(w)\right] \\
&= \frac{1+\frac{tp_0(w)q_0(w)}{2(\nu_0+t-2)}\delta_t^v}{1+\frac{tr_0(w)p_0(w)}{2(\nu_0+t-2)}\delta_t^v} - 1 \\
&= \frac{\frac{tp_0(w)\left(q_0(w)-r_0(w)\right)}{2(\nu_0+t-2)}\delta_t^v}{1+\frac{tr_0(w)p_0(w)}{2(\nu_0+t-2)}\delta_t^v},
\end{aligned}
$$

where the second equality follows from the normality of $\tilde{y}_t$.[11] For a small change in variance the change in $r_t(w)$ is well approximated by

$$
\left.r_t(w)-r_0(w)\right|_{\Delta_t\Sigma^2|\tilde{y}_t} \approx \frac{tr_0(w)p_0(w)\left(q_0(w)-r_0(w)\right)}{2(\nu_0+t-2)}\delta_t^v. \tag{5}
$$

Since the agent is risk averse $r_0(w) > 0$. By risk vulnerability $p_0(w) \geq r_0(w)$ and $q_0(w) \geq r_0(w)$, so the sign of $r_t(w) - r_0(w)$ is the same as the sign of δ_t^v.[12] □

Proposition 1 has a straightforward interpretation. Holding mean changes constant, our agent compares the sample variance of the realizations of the background risk that they observe to their prior beliefs about its variance. If this sample variance is larger than

[11] If $\tilde{y}_t$ was not normal, we could derive this equality approximately from a second order Taylor approximation of the relevant derivatives of $u(w+\tilde{y}_t)$ about w, provided that all moments of $F_{\tilde{y}}$ of third degree or higher are $o(\tilde{y}^2)$.

[12] Note that the logic in this proof is valid for any second order stochastic change in $B_t(y)$, though the form of risk vulnerability required is somewhat stronger, as shown by Eeckhoudt, Gollier and Schlesinger (1996):

1. there exists a scalar c such that $p_t(w+(x+y)) \geq c \geq r_t(w+(x+y)')$ $\forall w$ and $\forall (x+y), (x+y)'$ in t.
2. there exists a scalar d such that $q_t(w+(x+y)) \geq d \geq r_t(w+(x+y)')$ $\forall w$ and $\forall (x+y), (x+y)'$ in t.

These conditions were first articulated by Ross (1981) to accommodate an ordering of individuals by risk aversion when all objects of choice are risky lotteries.

their prior variance they reduce their endogenous risk-taking, while if it is smaller they increase their risk-taking. In other words, the agent's changing beliefs about pure risk in the environment monotonically drive their own risk-taking, even though this background risk is statistically independent of the returns on the foreground risk that is their object of choice.

D The effects of variance-preserving mean changes on risk-taking

We next turn our attention to the effects of small mean changes on risk-taking, by which we refer to the effects of deviations of the sample mean of background risk signals $\tilde{y}_t$ from the agent's prior μ_0. From equation 3 we can see that δ_t^m shifts the agent's posterior mean. A shift in mean beliefs functions as a deterministic shift in wealth for the agent, so that

$$r_t(w) - r_0(w)\Big|_{\Delta_t M|\delta_t^m} \approx (\Delta_t M|\delta_t^m) r_0'(w) = -\frac{t r_0(w)\big(p_0(w) - r_0(w)\big)}{\kappa_0 + t}\delta_t^m. \tag{6}$$

Given decreasing absolute risk aversion (the first condition in risk vulnerability) we can sign this effect consistently: $\delta_t^m > 0$ implies that $r_t(w) - r_0(w) < 0$. However, as can be seen from equation 4, δ_t^m also enters the agent's posterior variance as a quadratic term. Assuming that the sample variance is at the posterior-neutral point ($\delta_t^v = 0$, or $s_t^2 = \mathbb{E}_0[\Sigma^2]$), the effect of the mean via the variance on risk-taking is

$$r_t(w) - r_0(w)\Big|_{\Delta_t \Sigma^2|\delta_t^m} \approx \frac{t\kappa_0 r_0(w) p_0(w)\big(q_0(w) - r_0(w)\big)}{2(\nu_0 + t - 2)(\kappa_0 + t)}(\delta_t^m)^2 \tag{7}$$

Given risk vulnerability we can also sign this effect: if $(\delta_t^m)^2$ increases $r_t(w) - r_0(w)$ increases. Combining these two effects yields

$$r_t(w) - r_0(w)\Big|_{\delta} \approx -\frac{tr_0(w)\big(p_0(w) - r_0(w)\big)}{\kappa_0 + t}\delta_t^m + \frac{t\kappa_0 r_0(w)p_0(w)\big(q_0(w) - r_0(w)\big)}{2(\nu_0 + t - 2)(\kappa_0 + t)}(\delta_t^m)^2 \quad (8)$$

This leads us to our next result:

Proposition 2. (Local convexity in mean changes) *Suppose the agent observes a dataset of signals $\mathcal{D}_t$ for which $\delta_t^v = 0$ and δ_t^m is small. Then $r_t(w) - r_0(w)$ is a decreasing and convex function of δ_t^m.*

Proof. The proposition follows directly from the quadratic functional form of $r_t(w) - r_0(w)$ in δ_t^m in equation 8. □

Proposition 2 is a significant departure from the mean-only model. Even if the agent observes no change in variance, and even for small deviations of the data mean from the prior mean, $r_t(w) - r_0(w)$ is convex (and decreasing), rather than linear in δ_t^m. This means that the marginal effect on risk-taking in absolute terms for negative shocks will be larger for the agent than for equal-sized positive shocks. This asymmetry of positive and negative shocks is driven not by assumptions on the utility function, as in, for instance, the case of loss aversion in prospect theory, but rather by the interdependence of the agent's posterior mean and variance in our learning model.

E The combined effects of mean and variance

We are now ready to examine the combined effects of mean and variance changes on the agent's risk-taking. The overarching theme of the results in this section is that unlike in the mean-only model, in our model risk-taking is non-monotonic in the mean displacement from the prior.

Assume that $\tilde{y}_t$ is a background risk involving a small change in the agent's mean plus a small mean preserving-spread in the agent's variance relative to the agent's prior. Let $\delta_t^v = (s_t^2 - \frac{\nu_0}{\nu_0 - 2}\sigma_0^2)$. Then the change in the agent's risk-taking due to his posterior variance is:

$$r_t(w) - r_0(w)\Big|_{\Delta_t \Sigma^2 | \tilde{y}_t} \approx \frac{t r_0(w) p_0(w)\big(q_0(w) - r_0(w)\big)}{2(\nu_0 + t - 2)} \Big(\delta_t^v + \frac{\kappa_0}{t + \kappa_0}(\delta_t^m)^2\Big), \tag{9}$$

which we obtain from (equation 5) by including the appropriate terms for $\Delta_t \Sigma^2$, given non-zero δ_t^v and δ_t^m. This results in a total change in risk-taking for the agent, given both mean and variance, of

$$\begin{aligned} r_t(w) - r_0(w)\Big|_{\tilde{y}_t} \approx -\frac{t r_0(w)\big(p_0(w) - r_0(w)\big)}{\kappa_0 + t}\delta_t^m &+ \frac{t\kappa_0 r_0(w) p_0(w)\big(q_0(w) - r_0(w)\big)}{2(\nu_0 + t - 2)(\kappa_0 + t)}(\delta_t^m)^2 \\ &+ \frac{t r_0(w) p_0(w)\big(q_0(w) - r_0(w)\big)}{2(\nu_0 + t - 2)}\delta_t^v \end{aligned}$$

We will next assume that the agent has a CRRA utility function ($u(w) = (w^{1-\eta} - 1)/(1 - \eta)$). This will allow us to maximize intuition in the following results. Note, however, that our results would be qualitatively similar for any utility function, so long as it exhibits risk vulnerability.

Given CRRA utility we have that $r_0(w) = \eta/w$, $p_0(w) = (\eta + 1)/w$, and $q_0(w) = (2 + \eta)/w$. The above equation simplifies to:

$$r_t(w) - r_0(w)\Big|_{\tilde{y}_t} \approx -\frac{t\eta}{(\kappa_0 + t)w^2}\delta_t^m + \frac{t\kappa_0\eta(\eta + 1)}{(\nu_0 + t - 2)(\kappa_0 + t)w^3}(\delta_t^m)^2 + \frac{t\eta(\eta + 1)}{(\nu_0 + t - 2)w^3}\delta_t^v \tag{10}$$

The next two propositions relate to this function:

Proposition 3. (Variance dominating threshold) *There exists a threshold value $\delta_t^{v*} > 0$ such that for small $\tilde{y}_t$, if $\delta_t^v > \delta_t^{v*}$, then $r_t(w) - r_0(w) > 0$ for all δ_t^m. If the agent's utility function is CRRA this threshold is:*

$$\delta_t^{v*} = \Big(\frac{\nu_0 + t - 2}{2\kappa_0}\Big)^2 \Big(\frac{\kappa_0 + t}{\kappa_0}\Big)^{-1} \Big(\frac{\eta + 1}{w}\Big)^{-2} \tag{11}$$

Proof. As can be see from equation 10, $r_t(w) - r_0(w)$ is a quadratic function in δ_t^m, where the coefficient of the δ_t^m term is negative, the coefficient of the $(\delta_t^m)^2$ term is positive, and the coefficient of the δ_t^v term is positive (these signs are the same for any risk-vulnerable utility function). Since δ_t^v is unbounded above, δ_t^{v*} must exist. The form of δ_t^{v*} in 11 follows directly from solving for the discriminant in equation 10. □

Proposition 3 states that for an agent who is updating beliefs about both the mean and variance of the background risk, there exist high enough realizations of the variance so as to entirely overwhelm the effects of the mean, and make the agent more risk averse regardless of its realization. This is in stark contrast to a mean only model of updating, which typically implies strong monotonicity on the part of the mean, with increases leading to lower risk aversion and decreases leading to higher risk aversion (Dillenberger and Rozen (2015)). This is a direct consequence of the local convexity of the mean established in the previous proposition. Without the countervailing effect of $(\delta_t^m)^2$ no such threshold on δ_t^v would exist. Thus, although this result is written as a threshold on the variance, it is the consequence more strictly of the interdependence via the mean of the posterior mean and variance in our model.

The form of the threshold δ_t^{v*} in 11 provides instructive intuition for how the agent's utility and beliefs interact in our model. The threshold is composed of three terms. The first is, to a scale factor, the ratio of the agent's posterior confidence about the variance and their prior confidence about the mean. The threshold is increasing in this term. The second is the ratio of the agent's posterior confidence about the mean and their prior confidence about the mean. The third is the agent's coefficient of absolute prudence. The threshold is decreasing in the latter two terms. Intuitively, higher values of δ_t^{v*} correspond to lower sensitivity of the agent's risk-taking to the variance relative to the mean. Thus, 11 tells us that in our model the agent is relatively more sensitive to the variance when they are less confident about its value ex post relative to the value of the mean ex ante; and relatively more sensitive to the variance when they are more confident about the value of the mean ex post relative to the value of the mean ex ante. In other words, the agent is more sensitive to the moment for which uncertainty declines relatively less in their posterior. The agent is also relatively more sensitive to the variance the more prudent they are.

Proposition 3 identifies an upper bound on δ_t^v beyond which the agent becomes more risk averse for every value of δ_t^m. In the next proposition we examine how the combined effects of the mean and variance when the variance is below this critical value.

Proposition 4. (Combined effects of mean and variance changes) *Suppose an agent has a CRRA utility function. Suppose further that $\delta_t^v < \delta_t^{v*}$. Then there exists a threshold mean value*

$$\delta_t^{m*} = \frac{\nu_0 + t - 2}{2\kappa_0}\left(\frac{\eta+1}{w}\right)^{-1} - \sqrt{\left(\frac{\nu_0 + t - 2}{2\kappa_0}\right)^2\left(\frac{\eta+1}{w}\right)^{-2} - \frac{\kappa_0 + t}{\kappa_0}\delta_t^v} \tag{12}$$

and we have the following cases:

1. *$\delta_t^v = 0$: then $\delta_t^{m*} = 0$, and $r_t(w) - r_0(w) > 0$ iff $\delta_t^m < 0$.*

2. *$0 < \delta_t^v < \delta_t^{v*}$: then $\delta_t^{m*} > 0$, and $r_t(w) - r_0(w) > 0$ iff $\delta_t^m < \delta_t^{m*}$.*

3. *$-\frac{\nu_0}{\nu_0 - 2}\sigma_0^2 \leq \delta_t^v < 0$: then $\delta_t^{m*} < 0$, and $r_t(w) - r_0(w) < 0$ iff $\delta_t^m > 0$*

Proof. The proposition follows directly from solving for the smaller root of δ_t^m in equation 10. $\square$

Proposition 4 illustrates that the simple monotonic relationship between changes in mean and the agent's risk-taking in the mean-only model is an edge case in the mean-variance model, and breaks down even for moderate changes in variance. In particular, when the agent experiences an increase in variance relative to their prior, their risk-taking may decrease even under small increases in the mean. Conversely, agents who experience decreases in their variance relative to the prior may exhibit increases in risk-taking even under negative mean changes.

Note that here, as in the previous proposition, the threshold value that determines in which direction the agent's risk-taking changes under the countervailing effects of the experienced mean and variance depends both on the agent's information and on their utility. Under both cases 2 and 3 the absolute value of δ_t^{m*} captures the size of interval over which the variance effect dominates the mean. Aside from increasing in the absolute value of δ_t^v, in both cases it also increases in the ratio of the agent's posterior mean confidence to their prior mean confidence, and decreases in the ratio of variance posterior confidence to mean prior confidence and in the agent's absolute prudence.

Unlike in the previous proposition, however, the current result does not hinge on the quadratic form of δ_t^m in the $r_t(w) - r_0(w)$ function. Even under linearity in the mean a threshold δ_t^{m*} with similar properties would exist. This fact will prove important when we consider the effects of time in the next section.

F Limiting behavior

Up to now we have assumed that the time period t is fixed for the agent. In the next proposition we examine how the change in risk-taking function is affected by t. To do this we could take a derivative in t, but the intuition of the result is clearer if we examine the limiting behavior of the function instead, as we do in the next proposition:

Proposition 5. (Limit in time) *Suppose an agent has a CRRA utility function. Then*

$$\lim_{t\to\infty} \big(r_t(w) - r_0(w)\big)\Big|_{\tilde{y}_t} = -\frac{\eta}{w^2}\delta_t^m + \frac{\eta(\eta+1)}{w^3}\delta_t^v \tag{13}$$

Proof. The proposition follows from applying L'Hôpital's rule once to equation 10 before taking its limit in t. □

Given a fixed frequency of realizations we can think of t as being the age of the agent. Proposition 5 can then be interpreted as speaking to the role of mean and variance changes to affect the agent's risk-taking when they are old. Two points can be seen from the form of the limit. First, the effect of the quadratic term in δ_t^m declines faster in time than that of the linear terms on the mean and variance. This is in line with the intuition in the previous section that the agent is affected more by quantities over which they have higher posterior uncertainty. Intuitively, the agent learns about the quadratic term from both the posterior mean and the posterior variance. Functionally, this would mean that we would expect younger individuals to exhibit a more asymmetric marginal response to the mean than older individuals.

The second point is that at the limit our learning model converges to a fully deterministic functional form, with the agent's utility function alone driving their response to new realizations of the background risk.[13] This is due to the stationarity of the underlying data generating process. Behaviorally, this corresponds to declining learning and increasing stability of risk-taking in the face of background risk as individuals age.

[13] In fact, this is exactly the functional form of the marginal effects of mean and variance in Gollier and Pratt (1996), as can be seen from converting the CRRA representations of the coefficients of absolute risk aversion, prudence, and temperance back to their original forms.

G Discussion of the results

Our model paints a comprehensive picture of how risk aversion adapts to learning about unknown background risk. We reach several conclusions. Risk aversion monotonically increases in mean-preserving spreads in the data. Risk aversion is convex in mean changes, even for small shocks, and declining across the prior mean boundary. This means that the marginal effect of small negative shocks on risk aversion is larger than the marginal effect of small positive shocks. This asymmetry in response declines in the agent's age. The relative sensitivity of the agent to the two moments depends both on their information and their utility, with the moment containing more posterior uncertainty having a proportionally larger effect on risk aversion. When the effects of both the mean and the variance are taken into account, the simple monotonic link posited in prior work (Dillenberger and Rozen (2015)) between mean changes and changes in risk aversion breaks down. Increases in risk aversion are sometimes associated with an increasing mean and decreases in risk aversion are sometimes associated with an decreasing mean, depending on the value of the variance and the agent's relative sensitivity to the two moments.

In order to reach these conclusions we have made several simplifying assumptions. First, we have completely abstracted from any savings, insurance, or hedging decision that the agent might pursue, through our assumption of the non-existence of a savings technology and hand-to-mouth consumption. This assumption considerably simplifies our analysis, since it reduces both the choices that the agent makes and their beliefs function to essentially a static problem. This assumption means, however, that the results of our current analysis are likely to be more applicable to settings with missing insurance markets and significant barriers to savings, as is the case in much of the developing world today and was the case in developed countries in the not-too-distant past. How savings behavior and insurance choices interact with learning over background risk is important for extending our results to a broader class of settings and is a key direction for future work in this area.

A second set of important simplifying assumptions we make are on the agent's learning process. In our model learning is entirely passive, meaning that the agent makes no decisions that affect the information they receive. Understanding how adaptive risk-taking interacts with information search is important if we are to extend our analysis to examine risky decisions that also involve information gathering, like migration. Furthermore, we assume in the model that the agent's only source of information is personally-observed realizations of the background risk. This abstracts from processes of social learning. By relaxing this assumption we should be able to study important margins in the long-run dynamics of risk-taking, particularly intergenerational transmission of risk attitudes and the consequent effect of shocks on the cultural evolution of risk-taking.

III Empirical Evidence on Risk-Taking Adaptation to Macroeconomic Experiences

Our model of adaptive risk-taking makes sharp predictions about the effects of changes in background risk on individual risk-taking over time. In this section we test these predictions for macroeconomic risk using two large panel survey data sets from two large and diverse developing countries, Indonesia and Mexico. Both surveys contain repeated measures of absolute risk aversion for the same individuals several years apart. We link within-person changes in these measures to changes in real GDP growth statistics in individuals' state of birth over their lifetime.

Let R_{it} be our empirical measure of individual risk aversion, with higher values indicating lower propensity to take risks. We term the experienced mean of real GDP growth A_{it}, and the experienced volatility V_{it}. Our model makes three primary testable predictions for the correlation of growth experiences with individual measured risk aversion: (1) R_{it} *is increasing* *in* V_{it}; (2) R_{it} *is decreasing in* A_{it}; (3) R_{it} *is increasing in* A_{it}^2. The focus of this section is on testing these three predictions empirically, examining their behavioral consequences, and establishing whether they represent causal relationships.

A Data and Methodology

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

The second reason is their differences. Indonesia and Mexico offer a distinct contrast along many plausibly important dimensions, including geography, level of development, language, culture, religion, institutions, and other aspects of their history.[14] This aids in establishing the both the internal validity and external validity of our results. If we detect common effects in both countries we can be more confident that they are not driven by idiosyncratic characteristics of either setting, and more comfortable in predicting that they will generalize to other settings.

For the Indonesian analysis our source of micro data is the Indonesian Family Life Survey (IFLS) (Strauss et al. (2009), Strauss, Witoelar and Bondan (2016)). The IFLS is a longitudinal study administered by the RAND corporation in 13 provinces in Indonesia in five

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

waves, starting in 1993. For the Mexican analysis our source of micro data is the Mexican Family Life Survey (MXFLS), a longitudinal study administered in 16 states in three waves starting in 2002. The MXFLS was piloted by the RAND corporation, and is now managed by the Iberoamerican University (UIA) and the Center for Economic Research and Teaching (CIDE). Both surveys exhibit high recontact rates (>90%), and contain a wealth of economic and demographic covariates, allowing for a near-complete accounting of the balance sheet for subjects, including household income, assets, savings and borrowing. Both also contain residence and migration histories, allowing us to link place-based variables like GDP growth to subjects. Crucially for our purposes, the two most recent waves of both the IFLS and the MXFLS (IFLS4 (2007 - 2008), IFLS5 (2014), MXFLS2 (2005-2006), and MXFLS3 (2009-2012)) include modules for measuring subject financial risk aversion using hypothetical, high-stakes monetary gambles. We use measures from these modules to construct our primary dependent variables, which we describe in detail in subsubsection A.1.

We use sub-national measures of real GDP growth to construct measures capturing subject lifetime macroeconomic experiences, which are the primary independent variables in our analysis. Our data on GDP growth at the province level in Indonesia (equivalent to the state level in the United States) comes from the Indonesian Bureau of statistics (BPS) via the World Bank's INDO-DAPOER database, and from the BPS's statistical yearbooks for the years 2012-2015. These data exist at the province level starting in 1977. For Mexico, we obtain state level growth data from German-Soto (2005), who construct the GDP series using historical data from the National Institute on Statistics and Geography (INEGI). These data are available starting in 1941. We describe how we construct macroeconomic experience variables and assign them to subjects in detail in subsubsection A.2.

The sample for our main analysis is subjects who completed the risk aversion module in both waves of each survey. Focusing on subjects who appear in both waves of each survey allows us to estimate a model with individual fixed effects, which eliminates substantial amounts of noise due to idiosyncratic variation. This results in a primary sample of 17,302 subjects for Indonesia and 8,187 subjects for Mexico, each appearing twice in our data. In some analyses we do not include individual fixed effects, which allows us to expand the sample to all subjects who responded to the risk module in either wave of each survey, for a total of 55,820 subject-year observations in Indonesia and 25,005 subject-year observations in Mexico. Summary statistics for the complete survey samples and the primary samples are available in Appendix A. The geographic distributions of our samples in Indonesia and Mexico are available in Appendix B.

A.1 Risk Aversion Measures

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

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

In MXFLS2 subjects answered between two and five questions each, which resulted in a measure with five bins. Questions offered subjects a choice between a safe coin flip and a riskier coin flip, with the amounts of the riskier coin flips generally changing between questions. We code the resulting measure with higher numbers (1-5) indicating more risk aversion. The staircase instrument was changed for MXFLS3 to align more closely with the instrument in the IFLS. In MXFLS3 subjects answered between two and five questions each, resulting in a measure with six bins. Each question offered the same fixed safe amount of money, while the amounts of the risky lottery varied between questions. A "gamble averse" option was offered in this instrument. Since gamble aversion only appears in one wave of the MXFLS we drop subjects who chose this option in MXFLS3 from our sample. We code the resulting measure (1-5) in the same way as the other measures.

A pervasive concern with all elicited measures of financial risk aversion is the high degree of noise that they exhibit, which often means their predictive power for real-world risky behavior quite low (Yariv, Gillen and Snowberg (2019)). This raises the possibility that any detected effects on measured risk aversion will be due to noise, and won't translate to real changes in risk-taking behavior by the subjects. We address this concern head on in subsubsection C.3, where we show that subjects who became more risk averse by our measures also became less risk-taking in their economic behavior. We can also examine the predictive capacity of our measures in the cross-section. In Appendix E we present the results of regressing our measures of risk aversion on a host of demographic covariates and economic variables capturing risk-taking behavior in our samples, without including individual fixed effects in the regression. For subjects for whom we have complete data for all covariates our IFLS risk aversion measure, unlike many in the literature, exhibits significant correlations with risk-taking behavior like self employment and migration, and demographic measures like age and gender in expected ways, both in primary (panel) sample and in the broader sample. Our measure of risk aversion from the MXFLS is noisier than that in the IFLS, and consequently only exhibits significant correlations with smoking and age.

A.2 Macroeconomic Experience Variables

As a broad-based measure of macroeconomic conditions we use real GDP growth at the lowest administrative unit for which all pertinent data is available, the Indonesian province and Mexican state (both roughly equivalent to US states). Our data contain time series from 25 Indonesian provinces spread out over 9 major islands, and all 32 Mexican states, corresponding to the birth provinces reported by subjects in the IFLS and MXFLS.

Two complications with the growth data are worth noting. First, the boundaries of administrative units in Indonesia have not remained constant over our period of measurement. Following the 1998 collapse of the Suharto regime, Indonesia underwent a rapid (and still ongoing) process of decentralization. As a consequence, many administrative units at all levels of the state split, with the number of provinces increasing from 27 to 34 from 1993 to 2015. Since our analysis requires a consistent definition of administrative units, we mapped back all province-level variables (including GDP, population, and inflation) to provinces as they were defined in 1993. This is possible to do because in all cases a larger province split into multiple provinces, and in no cases did they recombine into novel provinces. Thus, every province in 2015 has exactly one corresponding ancestor province in 1993. To avoid confusion we refer to Indonesian provinces throughout by their names in 1993. Second, measurement error for sub-national macroeconomic variables in both countries is likely to be substantial. To reduce the effects of noise due to measurement error on our results we winsorize the province- and state-level GDP measures at the 5-95 level.

To construct our macroeconomic experience variables, we first assign to each individual the province/state real GDP growth time series in their birth province. Subjects who are born after 1976 in Indonesia and 1940 in Mexico, the first years for which subnational GDP data is available in each country, respectively, are assigned time series starting in their year of birth. Subjects born before these years are assigned time series starting in these initial years for their province of birth.

Once the time series are assigned we calculate for each individual the mean (A_{it}) and the standard deviation (V_{it}) of their time series from birth to year of measurement in the corresponding survey. Thus an individual born in East Java in 1981, for instance, will be assigned the statistics for the East Java time series from 1981 to 2007 (the year of IFLS4) and from 1981 to 2014 (the year of IFLS5). In Mexico, since MXFLS2 was administered between 2005 and 2007, and MXFLS3 was administered between 2009 and 2013, subjects are assigned time series that extend from birth to their exact measurement year. Let g_{is} be the growth rate assigned to person i in year s. Then for year of measurement t these statistics are:

$$A_{it} = \frac{1}{t - b_i} \sum_{s=b_i+1}^{t} g_{is} \tag{14}$$

$$V_{it} = \sqrt{\frac{1}{t - b_i - 1} \sum_{s=b_i+1}^{t} (g_{is} - A_{it})^2} \tag{15}$$

where

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

and

$$B = \begin{cases} 1976 & \text{if Country}_i = \text{Indonesia} \\ 1940 & \text{if Country}_i = \text{Mexico.} \end{cases}$$

To estimate the effects of experienced growth an alternative approach we might have employed, rather than calculating the above statistics, is to regress our measure of risk-taking on growth in each year for each individual. We elect to use the above method for three reasons. First, the year-by-year analysis would result in an unbalanced panel structure for our data, with all the attendant difficulties. Second, we would likely be under-powered to estimate the large number of parameters in such an analysis (Malmendier and Nagel (2011)). Third, the statistics we calculate correspond closely to relevant quantities in our model, allowing for direct tests of the theoretical predictions in our data.

A.3 Empirical Specification

Our baseline empirical specification is a two-way fixed effects model where we regress the individual risk aversion measure R_{it} on A_{it}, A_{it}^2, and V_{it}, as well as a constant α_{FE} and individual and time fixed effects:

$$R_{it} = \alpha_{FE} + \alpha_i + \alpha_t + \beta_1 A_{it} + \beta_2 A_{it}^2 + \beta_3 V_{it} + \gamma_1 PriceLevel_p + \gamma_2 X_{it} + \epsilon_{it}. \tag{16}$$

The individual fixed effect α_i absorbs variation due to time-invariant idiosyncratic heterogeneity, whereas the time fixed effect α_t nets out the effect of aggregate time trends.

A salient concern about regressing R_{it} on the macroeconomic experience variables is that inflation could bias the estimates of the effects, for two reasons. First, since R_{it} is measured off of nominal hypothetical lotteries, the real value of the prizes offered in these lotteries can change between waves. This means that inflation, even at the national level, can introduce noise into the repeated measurement of risk aversion using these survey instruments. The second concern is that inflation can introduce bias into the analysis if it varies significantly at the subnational level, as we in fact observe in the data. This is because subnational inflation might correlate with province-level growth and growth volatility, as well as R_{it}, creating a potential omitted variable bias problem. To address these concerns we include a measure of the price level subnational for administrative unit p as a control variable in all baseline specifications. This takes the form of 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 the province level, whereas in Mexico, due to data constraints, p is region.

Our baseline specifications, which are meant to yield as close to a clean estimate of the causal effect of macroeconomic experiences on risk-taking, include no additional controls. This is due to concerns about the endogeneity of variables like income or assets with risk-taking, which could affect identification. In subsubsection C.2 we demonstrate the robustness of our results to including additional time-varying controls, which are represented by X_{it} above. The errors ϵ_{it} are clustered at the province of birth by birth year level in our baseline specification, which is the level of treatment in our analysis.

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

$$\Delta R_{it} = \alpha_{FD} + \beta_1 \Delta A_{it} + \beta_2 \Delta A_{it}^2 + \beta_3 \Delta V_{it} + \gamma_1 Inflation_p + \gamma_2 \Delta X_{it} + \epsilon_{it}. \tag{17}$$

For expositional reasons we present the results below for the first-difference specification.

B Identification

The goal of our empirical analysis is to yield clean causal estimates of the effects of lifetime experiences of growth dynamics on individuals' attitudes towards risk. This aim is complicated by several econometric issues. We discuss these issues and our solutions to them in this section.

The first issue arises from the nature of the data used to estimate individuals' risk-taking in economics. By definition, subjects' propensity to take risks is a mental process, which we have not yet figured out how to observe directly. Rather, we observe choices that an individual makes and infer something about the underlying attitudes that are driving them. In the case of risk-taking a vast array of different instruments are used in the literature, ranging from real-world behaviors like portfolio, employment, or crop choices, to experimental instruments using monetary lotteries, to surveys of self-reported attitudes.

In order to be able to identify changes in attitudes towards risk from data on changes in risky choice one confound we must account for is that observed changes in choices might reflect changes in the properties of the good over which the choices are being made, rather than the attitudes themselves. A second, related concern is that since we are linking changes in choices to changes in the agent's environment, changes in the choices might reflect changes in the relationship between the object of choice and the environment, rather than the effect of the environment on the underlying attitudes.

As an illustration of these points consider the potential use of portfolio choice as a measure of risk-taking in our setting. Imagine that we observe a subject investing 100% of their wealth in stocks in period one, and 100% of their wealth in (safer) bonds in period two. We know that macroeconomic volatility increased from period one to period two. Can we say that this increase in volatility made the agent more risk averse? Perhaps. But something else that could be happening is that the characteristics of the stocks in the agent's portfolio have independently changed over time–perhaps the companies issuing them came out with negative earnings statements between period one and two. Another possibility is that the relationship between macroeconomic volatility and stock returns has changed between the two periods, as it is known to do over the business cycle. In either one of these cases we might think that the agent's attitudes towards risk have changed, when in reality the properties of the choice we use to measure risk-taking have changed instead.

Our empirical measure of risk-taking is uniquely suited to overcoming these issues. The hypothetical lotteries offered the subjects have known odds and prizes in both surveys, and are fixed between the waves of the survey in the IFLS. This means that changes in subject choice are unlikely to be driven by changes in the characteristics of the lotteries, especially once we control for inflation. Furthermore, the prizes in the lotteries are statistically independent and exogenous of the growth experiences of the subjects, meaning that observed changes in choice are unlikely to be driven by changes in the relationship of the object of choice to the environment.

A second concern for identification is the possibility of reverse causality, where changes in subjects' propensity to take risks changes the macroeconomic environment to which they are exposed, instead of the other way around. The primary channel through which this might occur in our setting is endogenous migration, as migration could lead to exposure to different macroeconomic conditions, and is thought to be a risky choice itself (Bryan, Chowdhury and Mobarak (2014)). For instance, it might be the case that individuals whose propensity to take risks increases differentially migrate to provinces experiencing increases in growth. To address this concern in our baseline analysis we use macroeconomic conditions in an individual's province or state of birth over their lifetime, rather than their province or state of residence. By doing so we are, in essence, instrumenting for presence in a given location with their location of birth, which is exogenous under the assumption that location of birth does not correlate with later in life risk-taking dynamics or growth dynamics.

A third possible threat to identification is macro-level omitted variable bias, or the existence of an unobserved aggregate variable that could be correlated with growth dynamics and driving changes in individual risk attitudes. One might imagine, for instance, that instances of natural disasters or the outbreak of violence could increase risk aversion (Cameron and Shah (2015), Callen et al. (2014)) and be correlated with growth volatility. We employ a three-pronged approach to deal with this issue. First, our inclusion of a time fixed effect in the regression ensures that aggregate trends at the national level do not contaminate our estimates. Second, we employ subnational variation in macroeconomic conditions within two countries over time. This means that any potential omitted macro variable driving the results would have to correlate with growth dynamics over both space and time in order to generate our estimated effects. While this is possible, we believe that it is unlikely given the substantial variation that exists in our data (see Figure 2). This identifying variation comes from three sources: heterogeneity in growth experiences between cohorts, within cohorts over time, and within cohort in the cross-section. Third, we control directly for some of the most likely suspects, particularly exposure to violence and natural disasters, in subsubsection C.2

and see no meaningful change in the results.

A fourth and final concern is micro-level omitted variable bias, or the possibility that an unobserved individual-level covariate is erroneously driving the correlation we observe between growth dynamics and measured risk aversion. Since our outcome variable is within-person changes in measured risk aversion, we can rest easy that static individual-level covariates, like subjects' income level, are not likely to be driving our results. Nevertheless it is possible that changes in individual-level variables could be driving the result. The most likely culprits, according to theory, are household-level economic variables like changes in income and assets. In subsubsection C.2 we control for these directly and find that our results are robust to their inclusion.

To our knowledge, ours is the first paper in the empirical experience effects literature to include data from multiple countries, subnational variation in each, and repeat observations of the outcome of interest for the same individuals. Malmendier and Nagel (2011) 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. Shigeoka (2019) exploits within-country variation in macroeconomic conditions in Japan, but its empirical analysis does not include individual fixed effects or cross-country data. Malmendier and Shen (2019) and Ampudia and Ehrmann (2017) use cross-country data from 13 Eurozone countries in their estimates of the effects of economic shocks on risk taking, but do not include individual fixed effects or subnational 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 subnational variation in the occurrence of violence (Jakiela and Ozier (2019), Brown et al. (2019)) or a natural disaster (Cameron and Shah (2015), Hanaoka, Shigeoka and Watanabe (2018)), 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.

Figure 2: Variation in Macroeconomic Data

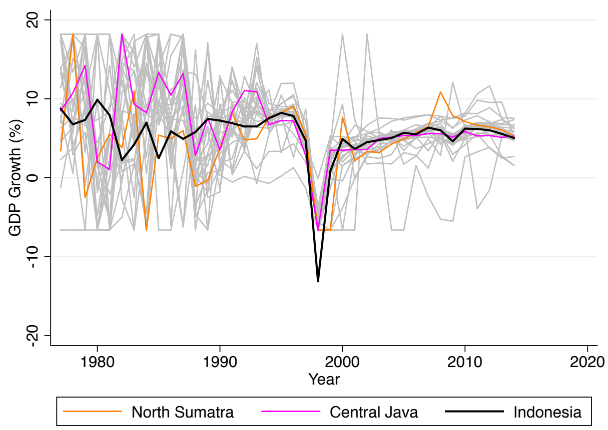

(a) Indonesian provinces, 1977-2015

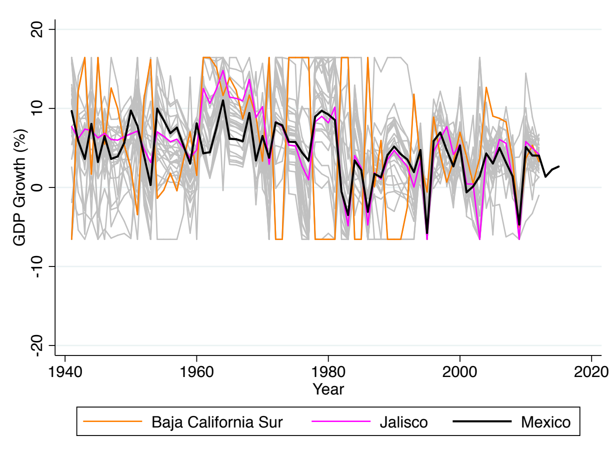

(b) Mexican states, 1940-2015

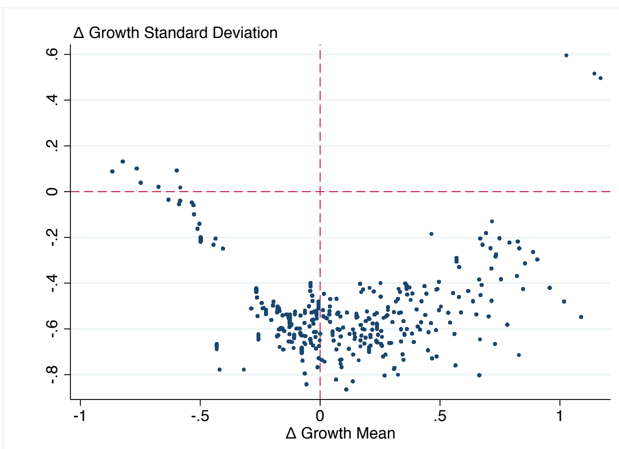

(c) Indonesia birth-province cohorts

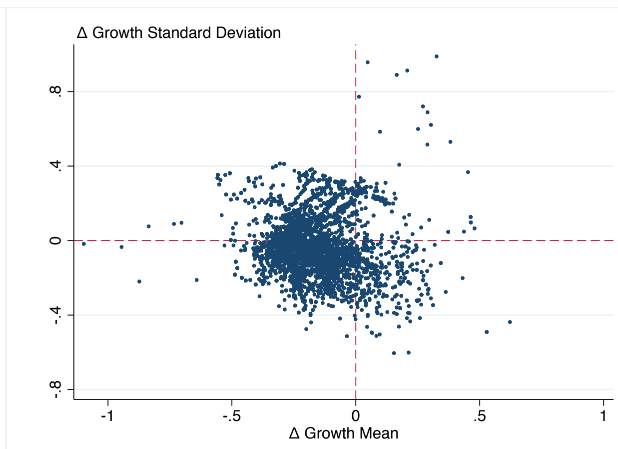

(d) Mexican birth-state cohorts

Notes: This figure displays our macroeconomic data in two ways. The top two panels display the real GDP growth time series for all 25 Indonesian provinces (1993 definitions) and 32 Mexican states in our data (winsorized at the 5-95 level to reduce measurement error), as well as the national GDP time series. As can be seen these time series exhibit substantial variation both in the cross section and over time. The bottom two panels display the raw distributions of our main explanatory variables ΔA_{it} and ΔV_{it} graphed against each other in each country. These scatterplots demonstrate that substantial variation exists not only in macroeconomic conditions across provinces/states, but also in the dynamics of macroeconomic experiences at the individual level.

C Results

This section contains the findings from our three primary empirical analyses. In subsubsection C.1 we present the results from regressing within-subject changes in measured risk aversion on subjects' experienced mean real GDP growth (linear and squared) and growth volatility. These regressions, which include no controls aside from subnational inflation, are the most direct tests of the three model-generated hypotheses discussed at the beginning of this chapter. In subsubsection C.2 we demonstrate the robustness of these findings to the inclusion of controls for changes in subjects' economic constraints and experiences of violence and natural disasters. In subsubsection C.3 we display the correlation between changes in several kinds of risky behaviors and predicted change in risk-taking based on the model we estimated in subsubsection C.1.

C.1 Effects of macroeconomic experiences on measured risk aversion

Our main empirical findings are presented in Table 1. Column 1 displays the result of regressing changes in measured risk aversion on mean changes in experienced growth in subjects' province or state of birth. In line with hypothesis 2, the estimated effect of the mean in both countries is negative and highly significant.[15] In column 2 we show the results of regressing changes in measured risk aversion on changes in the standard deviation of growth. In line with hypothesis 1 the estimated effect of changes in volatility is positive and highly significant in both settings. These findings hold when we regress changes in measured risk aversion on both mean and volatility, as can be seen in column 3.

In column 4 we display the results of adding the change in squared mean term to the previous specification. As before, the coefficient of the volatility term remains highly significant

[15]This is the specification most closely analogous in our main analysis to Malmendier and Nagel (2011), who examine the differential effects of mean changes in macroeconomic conditions (like stock market returns) on stock market participation and elicited risk aversion. Our first result here is broadly consistent with their findings.

and positive in both settings. In the Mexican data the coefficient of the linear mean term remains negative, highly significant, and its magnitude increases by approximately 75%, while the coefficient of the mean squared term is positive and marginally significant, in line with hypothesis 3. In the Indonesian data in this specification the coefficient of the linear mean term becomes statistically insignificant, while the coefficient of the mean squared is actually marginally significant and negative.

Overall our results provide strong evidence for hypotheses 1 and 2, but only weak or mixed evidence for hypothesis 3. While the coefficient of the mean squared term is of the correct sign in Mexican data, it is only marginally significant, and the corresponding coefficient in the Indonesian data is actually of the opposite sign (though also marginally significant). It is possible that these results (or lack thereof) are driven by the stronger dependence of the mean squared on the subjects' age, as the theory suggests. Notably, however, the magnitude of the linear mean term in both settings changes considerably once the mean squared term is introduced. This suggests that the square of the mean might play an important role in the overall effect of the mean on measured risk aversion, though ultimately our data are too noisy to reliably estimate this effect consistently.

It is also notable that the magnitude of the volatility term is .9 (Mexico) to 4.3 (Indonesia) times as large as the linear mean term in specification 3, where their magnitudes are most directly comparable. This tells us that the marginal effect of changes in variance are not second-order relative to that of the mean, but are rather as important if not significantly more so for driving changes in subjects' measured risk aversion.

C.2 Additional controls

Our main results are estimated without the inclusion of any additional controls aside from subnational inflation, though there are well-founded reasons to include additional covariates. Theoretically, changes in subjects' income, wealth, buffer stocks of savings, or other economic

Table 1: Main results

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.35***		-0.30***	0.42
	(0.08)		(0.07)	(0.46)
Δ Growth Mean^2				-0.07*
				(0.04)
Δ Growth Volatility		1.36***	1.30***	1.21***
		(0.17)	(0.16)	(0.18)
Observations	17302	17302	17302	17302
Mexico				
Δ Growth Mean	-1.02***		-0.97***	-1.69***
	(0.20)		(0.19)	(0.44)
Δ Growth Mean^2				0.10*
				(0.05)
Δ Growth Volatility		0.91***	0.87***	0.86***
		(0.17)	(0.17)	(0.17)
Observations	8187	8187	8187	8187

Measured Risk Aversion: 1-5, 5 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

circumstances might be expected to influence their measured risk aversion. Empirically, previous studies have shown that exposure to traumatic experiences like natural disasters and violence can change measured risk aversion.

We do not include these controls in our main analysis because they are endogenous to risk aversion itself. This means that their inclusion could threaten the causal interpretation of our results. Nevertheless, we would like to know whether we can interpret the changes we observe in measured risk aversion as representing changes in underlying risk attitudes, or merely as driven by changes in personal economic circumstances. Further, it would be useful to directly test whether macroeconomic experiences are in fact driving the observed changes or whether other kinds of experiences whose incidence may be correlated with growth dynamics are in fact playing a central role.

We provide some evidence on these points in Table 2, where we progressively add in additional controls to the specification for the last column in Table 1. These include time-varying demographics, like marital status, educational attainment, and household size; changes in income, assets, and savings; and self-reported exposure to violence and natural disasters (full details on the controls are available in Appendix F). In both countries our results are highly robust to the inclusion of this rich set of covariates, suggesting that the changes we estimate in measured risk aversion are driven by experienced growth dynamics and represent changes in underlying attitudes towards risk.

C.3 Changes in Risky Behavior

Another issue of interpretation of our results is the question of whether changes in measured risk aversion capture changes in actual risk-taking behavior for subjects. We study this question by constructing a variable measuring predicted change in risk aversion ($\widehat{\Delta R_{it}}$) using our preferred specification (column 4 of Table 1), and examining its correlation with changes in downstream risky behaviors in our data. We focus on behaviors commonly examined in

Table 2: Additional controls

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Indonesia							
Δ Growth Mean	0.42	0.33	0.32	0.33	0.34	0.34	0.42
	(0.43)	(0.44)	(0.45)	(0.44)	(0.45)	(0.44)	(0.44)
Δ Growth Mean^2	-0.07*	-0.06	-0.06	-0.06	-0.06	-0.06	-0.07*
	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)	(0.04)
Δ Growth Volatility	1.21***	1.24***	1.24***	1.24***	1.24***	1.23***	1.22***
	(0.18)	(0.19)	(0.19)	(0.19)	(0.19)	(0.18)	(0.18)
Observations	17,302	16,086	16,082	16,082	16,082	16,082	16,082
Mexico							
Δ Growth Mean	-1.69***	-1.78***	-1.77***	-1.75***	-1.74***	-1.79***	-1.71***
	(0.44)	(0.45)	(0.45)	(0.45)	(0.45)	(0.46)	(0.45)
Δ Growth Mean^2	0.10*	0.11*	0.11*	0.10*	0.10*	0.10*	0.09*
	(0.05)	(0.05)	(0.05)	(0.05)	(0.05)	(0.06)	(0.05)
Δ Growth Std. Dev.	0.86***	0.85***	0.86***	0.86***	0.86***	0.82***	0.80***
	(0.17)	(0.17)	(0.17)	(0.17)	(0.17)	(0.17)	(0.17)
Observations	8,187	8,046	8,046	8,046	8,046	7,996	7,996
Inflation	X	X	X	X	X	X	X
Δ Demographics		X	X	X	X	X	X
Δ Income			X	X	X	X	X
Δ Assets				X	X	X	X
Δ Savings					X	X	X
Δ Violence						X	X
Δ Natural Disasters							X

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Province (Indonesia) or regional (Mexico) inflation included in all regressions. Demographics include marital status, household size, household size squared, and educational attainment. Violence and natural disasters variables from self-reported exposure. Standard errors clustered at the cohort by province of birth level in parenthesis. *** p<0.01, ** p<0.05, * p<0.1.

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.[16] This last measure captures risky investment behavior in our data.

Results for this analysis are presented in Figure 3, which displays the average value of each downstream variable for each quartile of the $\widehat{\Delta R_{it}}$ distribution. Here light blue bars represent subjects who are predicted to become measurably less risk averse by our model, while dark blue bars represent subjects who are predicted to become measurably more risk averse. 95% confidence intervals are indicated for each quartile. We use the first to fourth interquartile range as an empirical benchmark and run a two-sided t-test to check if the difference between the average values of the outcome in these quartiles is statistically significant.

Our results provide strong evidence that increases in measured risk aversion predicted by experienced growth dynamics are associated with decreases in risk-taking behavior in Indonesia, and moderate and somewhat conflicting evidence of the same in Mexico. Increases in predicted risk aversion are associated with decreases in risk-taking behavior for six out of the seven variables examined. Three of these declines are significant at conventional levels, all in Indonesia: Smoking (4.3 percentage point increase to 1.2 percentage point emphdecrease ($p = .0001$)), Ever Migrated across Province Lines (+1.3pp $\rightarrow$ +0.2pp, ($p = .0001$)), and Self Employed (+7.3pp $\rightarrow$ +4.7pp, ($p = .053$)). Three decreases are not significant at conventional levels: planting of cash crops in Indonesia (+2.1pp $\rightarrow$ +1.5pp, ($p = .3$)), smoking in Mexico (+2pp $\rightarrow$ +0.9pp, ($p = .236$)), and Ever Migrated across State Lines in Mexico (+0.5pp $\rightarrow$ +0.3pp, ($p = .274$)). One outcome, self-employment in Mexico, exhibits a marginally significant increase (+0.9pp $\rightarrow$ +4.6pp, ($p = .087$)).

The decreases in risk-taking behavior predicted by increases in measured risk aversion are large and economically significant. In Indonesia, the first to fourth interquartile range

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

of $\widehat{\Delta R_{it}}$ for smoking represents a 17.5 percent decline relative to the IFLS4 baseline, for migration an 8.1% decline, for self employment a 6.5% decline, and for cash crop planting a 7.4% decline. In Mexico the interquartile range for smoking represents a 14.1% decline relative to the MXFLS2 baseline, for migration a 1.3% decline, and for self employment a 17% increase.

We can also quantify the magnitudes of lifetime changes in mean growth and growth volatility associated with changes in risky behavior. In Indonesia the first to fourth interquartile range represents a .32 percentage point decrease in mean growth (+.13pp $\rightarrow$ $-$.19pp) and a .27pp increase in the standard deviation of growth ($-$.61pp $\rightarrow$ $-$.34pp). In Mexico the first to fourth interquartile range represents a .22pp decrease in mean growth ($-$.01pp $\rightarrow$ $-$.23pp) and a .29pp increase in the standard deviation of growth ($-$.16pp $\rightarrow$ +.13pp).

D Robustness

We test the robustness of our main results to varying methodological choices in our analysis. First, in Appendix G we present the results of running our main analysis with *alternate sample compositions*. In particular, we (1) limit the analysis to individuals born after 1976 in Indonesia and 1940 in Mexico (for whom we have full lifetime macroeconomic histories); (2) include only individuals born after 1976 in Mexico (to compare with the parallel sample from Indonesia); and (3) exclude the "gamble averse" from our analysis in Indonesia (see subsubsection A.1). The results are qualitatively quite similar for each of these samples, though the sign of the linear mean term in Indonesia becomes positive when the gamble averse are excluded.

In Appendix H we present the results of our main analysis for *alternate specifications of measured risk aversion*. For both Indonesia and Mexico, we repeat the analysis with (1) a binarized measure of risk aversion (instead of using the 5 buckets of measured risk aversion, we set buckets 1 and 2 to be 0, and buckets 3, 4 and 5 to be 1); and (2) using an ordered

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

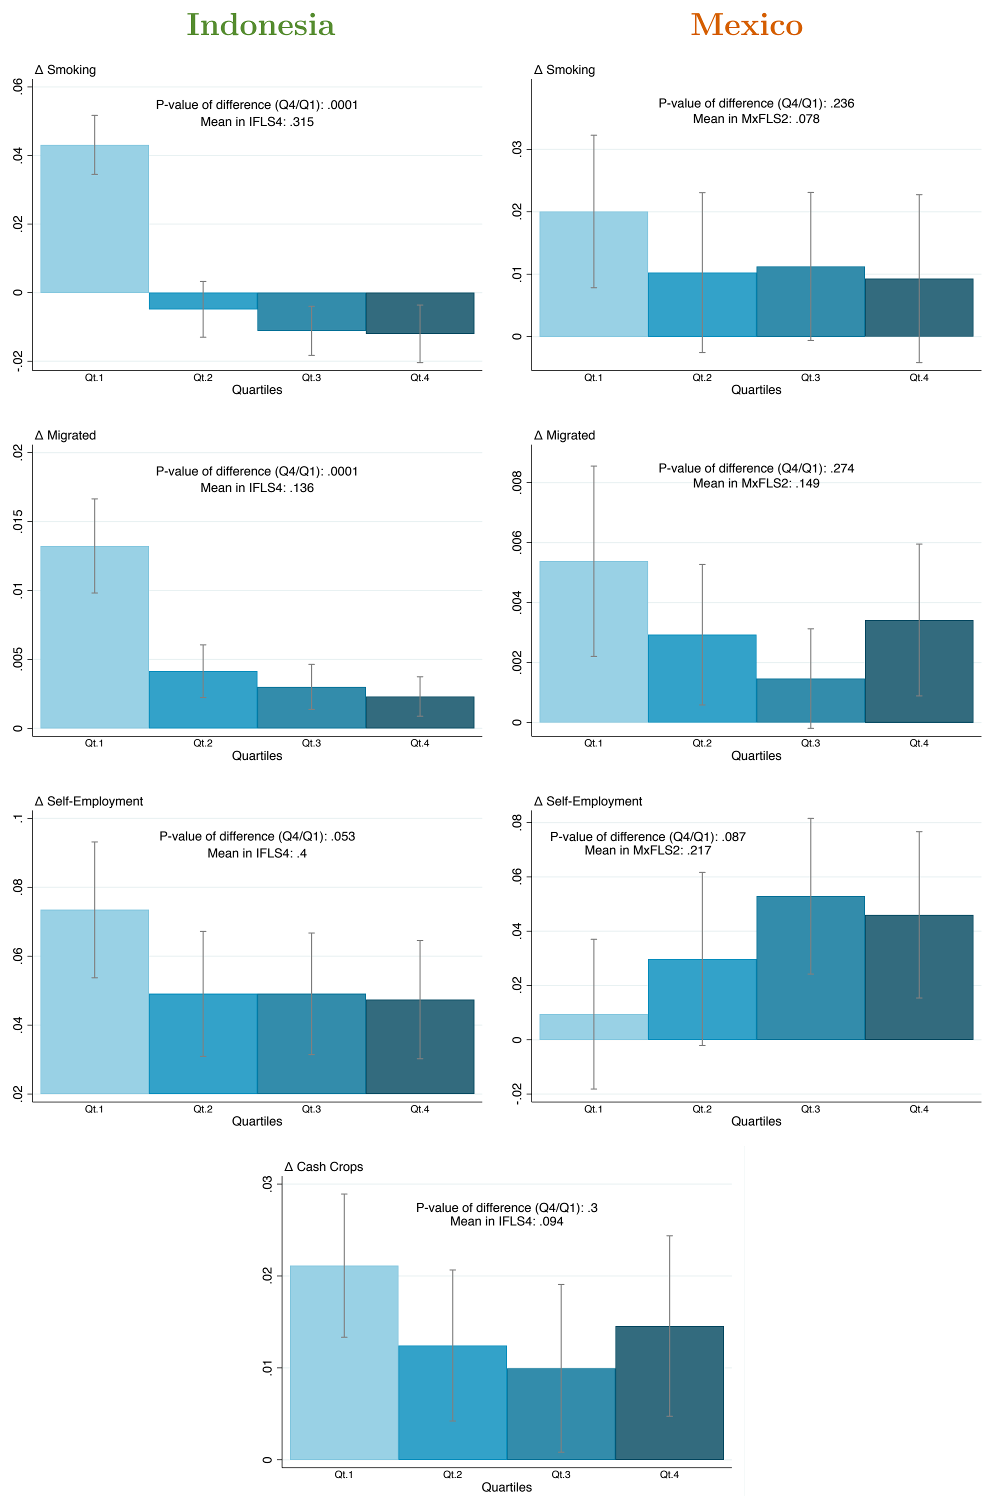

Notes: Bars represent quartiles of the predicted change in risk aversion distribution. Light blue is the bottom quartile of this distribution, representing the agents who are predicted to experience a decrease in risk aversion. Bottom panel is from the Indonesian data.

probit specification. The latter specification accounts explicitly for the ordinal nature of our risk aversion measure, though its results should be interpreted with care as the ordered probit with two way fixed effects estimator is known to be biased. For both specifications results are qualitatively quite similar to the baseline.

In Appendix I we present the results of our main analysis using data on *macroeconomic conditions in subjects' province or state of residence*, rather than their province or state of birth. Here we use the migration histories for subjects and assign them growth time series based on their actual location in every given year. These data more closely match the intuitive notion of macroeconomic experiences, but, as discussed above, suffer from a potential identification problem due to endogenous migration. Results are again qualitatively similar to the main analysis, with the magnitudes of the mean coefficients in specification 4 in both countries increasing by 40%-80% and becoming more significant.

Malmendier and Nagel (2011) estimate a non-linear single parameter weighting function for the effects of mean stock market returns on later in life stock market participation and elicited risk aversion. Malmendier and Nagel (2016) extend this analysis to the context of inflation experiences. This *temporal weighting function* is meant to flexibly estimate higher relative weights on early, formative experiences or on recent experiences due to recency bias. In Appendix J we extend their method to the context of lifetime volatility experiences. Our results are qualitatively similar to those in our baseline model. We estimate substantial recency effects across all specifications using this method.

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

In Appendix L we conduct our main analysis using a *repeated cross-section specification* that drops the individual fixed effects from the regression. We perform this analysis both restricting to our primary sample and extending the sample to any individuals who respond to the risk instrument in one of the waves of the surveys. The msgnitude of the estimated coefficients drops considerably in both settings under this specification. While the estimates for Mexico are no longer significant, estimates for our Indonesian sample remain consistent and significant.

IV Conclusions

Our theoretical and empirical results paint a coherent picture of the ways in which the risk attitudes of individuals in the developing world are shaped by lifetime macroeconomic experiences. They suggest that novel shocks are evaluated relative to the agent's body of experiences, and that their effects on risk attitudes are driven by changes in agent beliefs about risk in their environment. Consequently, the effects of changes in both the perceived mean (whether conditions are getting better or worse) and the perceived variance (whether conditions are stabilizing or destabilizing) are first order. The overall effect of shocks on risk attitudes can be usefully thought of as the composition of these two moment effects. Measured shifts in risk attitudes translate into meaningful changes in risk-taking behavior in the domains of health, migration, investment, and occupational choice.

While our current set of results applies to macroeconomic experiences specifically, the notion that the dynamics of background risk drive the dynamics of individual risk attitudes is, potentially, more broadly applicable. A key open question given our results is whether other kinds of environmental risk affect risk attitudes in the same way that macroeconomic risk does. Paramount amongst these is climate change, both because of its increasing prominence as a source of aggregate risk going forward, and because understanding the relationship between climate change and risk attitudes holds the best promise for extending our analysis into the deep past, for which high fidelity economic data is hard to come by. Theoretically, addressing this question would involve relaxing the stationarity assumption in our model, to examine the consequences of learning over a source of background risk exhibiting both time-trends and cyclicality.

A second open question raised by our results is the role that agent choice over the environment plays in the relationship between background risk and risk attitudes. In our current analysis we abstracted from this choice by assuming that the agent has no control over the environment they are exposed to, and that all learning about the background risk is passive. In reality agents often have substantial control over their environments. Perhaps the main way that such control is exerted is through the choice to change their environment by migrating. To study the interaction between experienced background risk and migration decisions would necessarily involve building a richer theoretical framework that can accommodate both choice over the physical environment and agent choice over information. A promising avenue for doing so is to combine our current model with a multi-armed bandit framework where the agent can learn about both the mean and variance of each arm. This would capture the exploration-exploitation dynamic underlying the migration decision, and the ways in which this dynamic interacts with agent beliefs about risk in both the old and new environments.

Our analysis also leaves open a set of intriguing questions regarding the ways that risk-taking adaptation aggregates up to collective behavior. In our model we assume that personal lifetime experiences are a primary determinant of risk attitudes. This implies that different agents should respond differently to novel shocks that arise, based on their own previous body of experiences. Therefore, our results should have non-trivial implications for risk sharing, from which we abstract in our current framework. In particular, our analysis suggests that heterogeneous agent beliefs about background risk may themselves function as a driver of dynamic risk-sharing. Aggregate investment should therefore depend on the distribution of lifetime experiences of macroeconomic risk in the population.

References

Acemoglu, By Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2017. "Microeconomic Origins of Macroeconomic Tail Risks." *American Economic Review*, 107(1): 54–108.

Ampudia, Miguel, and Michael Ehrmann. 2017. "Macroeconomic experiences and risk taking of euro area households." *European Economic Review*, 91(C): 146–156.

Arrow, Kenneth J. 1970. "Aspects of the Theory of Risk Bearing." In *Essays in the Theory of Risk-Bearing.* Chapter 3.

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

Brown, Ryan, Verónica Montalva, Duncan Thomas, and Andrea Velásquez. 2019. "Impact of Violent Crime on Risk Aversion: Evidence from the Mexican Drug War." *Review of Economics and Statistics.*

Bryan, Gharad, Shyamal Chowdhury, and Ahmed Mushfiq Mobarak. 2014. "Under-investment in a Profitable Technology : The Case of Seasonal Migration in Bangladesh." *Econometrica*, 82(5): 1671–1748.

Caballé, Jordi, and Alexey Pomansky. 1996. "Mixed risk aversion." *Journal of Economic Theory*, 71(2): 485–513.

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

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

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

Deaton, Angus. 1991. "Saving and Liquidity Constraints." *Econometrica*, 59(5): 1221–1248.

Deck, Cary, and Harris Schlesinger. 2014. "Consistency of higher order risk preferences." *Econometrica*, 82(5): 1913–1943.

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

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

Dohmen, Thomas, Hartmut Lehmann, and Norberto Pignatti. 2016. "Time-Varying Individual Risk Attitudes over the Great Recession: A Comparison of Germany and Ukraine." *Journal of Comparative Economics*, 44(1): 182–200.

Ebert, Sebastian, and Daniel Wiesen. 2014. "Joint measurement of risk aversion, prudence, and temperance." *Journal of Risk and Uncertainty*, 48: 231–252.

Eeckhoudt, Louis, Christian Gollier, and Harris Schlesinger. 1996. "Changes in Background Risk and Risk Taking Behavior." *Econometrica*, 64(3): 683–689.

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

Frey, Renato, Andreas Pedroni, Rui Mata, Jorg Rieskamp, and Ralph Hertwig. 2017. "Risk preference shares the psychometric structure of major psychological traits." *Science Advances*, 1–14.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and Eugene H. Stanley. 2003. "A theory of power-law distributions in financial market fluctuations." *Nature*, 423(May): 267–270.

Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and Eugene H. Stanley. 2006. "Institutional Investors and stock market volatility." *Quartely Journal of Economics*, , (May): 461–504.

German-Soto, Vicente. 2005. "Generación del producto interno bruto Mexicano for entidad federativa, 1940-1992." *Trimestre Economico*, 72(3): 617–653.

Giné, Xavier, and Dean Yang. 2009. "Insurance, credit, and technology adoption: Field experimental evidencefrom Malawi." *Journal of Development Economics*, 89(1): 1–11.

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

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2018. "Time Varying Risk Aversion." *Journal of Financial Economics*, 128(3): 403–421.

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

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

Karlan, Dean, Robert Osei, Isaac Osei-Akoto, and Christopher Udry. 2014. “Agricultural Decisions after Relaxing Credit and Risk Constraints.” *The Quarterly Journal of Economics*, 129(2): 597–652.

Kelly, Bryan, and Hao Jiang. 2014. “Tail Risk and Asset Prices.” *Review of Financial Studies*, 27(10): 2842–2871.

Kimball, Miles. 1990. “Precautionary Saving in the Small and in the Large.” *Econometrica*, 58(1): 53–73.

Kimball, Miles S. 1993. “Standard Risk Aversion.” *Econometrica*, 61(3): 589–611.

Malmendier, Ulrike, and Leslie Sheng Shen. 2019. “Scarred Consumption.”

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

Malmendier, Ulrike, and Stefan Nagel. 2016. “Learning from inflation experiences.” *Quarterly Journal of Economics*, 131(1): 53–87.

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

Morris, Stephen, and Muhamet Yildiz. 2016. “Crises : Equilibrium Shifts and Large Shocks.”

Munshi, Kaivan. 2004. “Social learning in a heterogeneous population: Technology diffusion in the Indian Green Revolution.” *Journal of Development Economics*, 73(1): 185–213.

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

Noussair, Charles N, and Stefan T Trautmann. 2014. "Higher Order Risk Attitudes, Demographics, and Financial Decisions." *Review of Economic Studies*, 81(September 2013): 325–355.

Pratt, John W. 1964. "Risk Aversion in the Small and in the Large." *Econometrica*, 32(1/2): 122–136.

Pratt, John W, and Richard J Zeckhauser. 1987. "Proper Risk Aversion." *Econometrica*, 55(1): 143–154.

Rabin, Matthew. 2013. "Incorporating Limited Rationality into Economics." *Journal of Economic Literature*, 51(2): 528–543.

Rosenzweig, Mark R., and Hans P. Binswanger. 1993. "Wealth, Weather Risk and the Composition and Profitability of Agricultural Investments." *The Economic Journal*, 103(416): 56–78.

Rosenzweig, Mark R., and Kenneth I. Wolpin. 1993. "Credit Market Constraints, Consumption Smoothing, and the Accumulation of Durable Production Assets in Low-Income Countries: Investments in Bullocks in India." *Journal of Political Economy*, 101(2): 223.

Ross, Stephen A. 1981. "Some Stronger Measures of Risk Aversion in the Small and the Large with Applications." *Econometrica*, 49(3): 621–638.

Sahm, Claudia R. 2012. "How Much Does Risk Tolerance Change?" *Quarterly Journal of Finance*, 2(4).

Shigeoka, Hitoshi. 2019. "Long-Term Consequences of Growing up in a Recession."

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

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

Weber, Elke U, Ann-rené E Blais, and Nancy E Betz. 2002. "A Domain-specific Risk-attitude Scale: Measuring Risk Perceptions and Risk Behaviors." *Journal of Behavioral Decision Making*, 15(August): 263–290.

Weitzman, Martin L. 2007. "Subjective expectations and asset-return puzzles." *American Economic Review*, 97(4): 1102–1130.

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

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

Online Appendix

A Summary statistics

Table 3: Summary Statistics (sample mean)

	Indonesia		Mexico	
Sample:	Primary sample	Full sample	Primary sample	Full sample
Measured Risk aversion	3.55	3.52	2.43	2.41
Woman	0.55	0.53	0.58	0.59
Age	40.27	37.32	42.80	42.02
Married	0.89	0.81	0.66	0.65
Household Size	5.20	5.21	5.64	5.65
Comp. Elementary	0.41	0.35	0.51	0.50
Comp. Junior High	0.19	0.20	0.25	0.25
Comp. High School	0.28	0.32	0.13	0.13
Above High School	0.13	0.14	0.11	0.11
Self-employed	0.42	0.39	0.23	0.22
Currently smoke	0.32	0.32	0.08	0.08
Ever migrated	0.14	0.17	0.15	0.15
Income/month†	9.20	8.868	2,420	2,420
Consumption/month†	2.35	2.45	2,088	1,768
Savings†	8.09	8.76	9,677	9,522
Borrowing†	2.64	2.68	4,734	4,707
Observations	35,292	55,820	19,769	25,005

B Geographic distribution of survey samples in our data

Figure 4: Distribution of the Primary Sample in Indonesia by Province of Birth

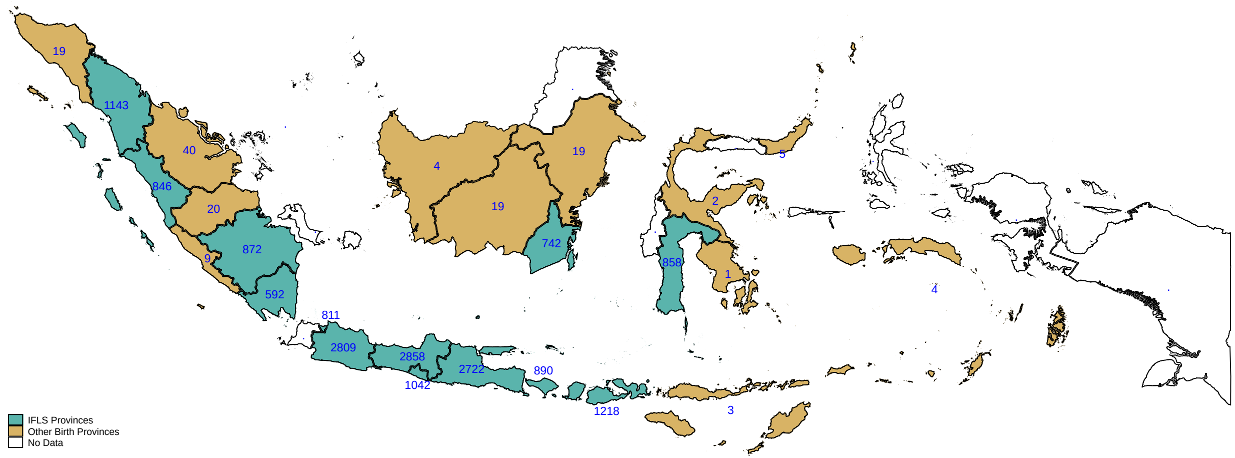

Notes: Provinces in blue are ones in which the IFLS has been deployed. Provinces in brown are non-IFLS provinces in which some subjects in our primary sample were born.

Figure 5: Distribution of the Primary Sample in Mexico by Province of Birth

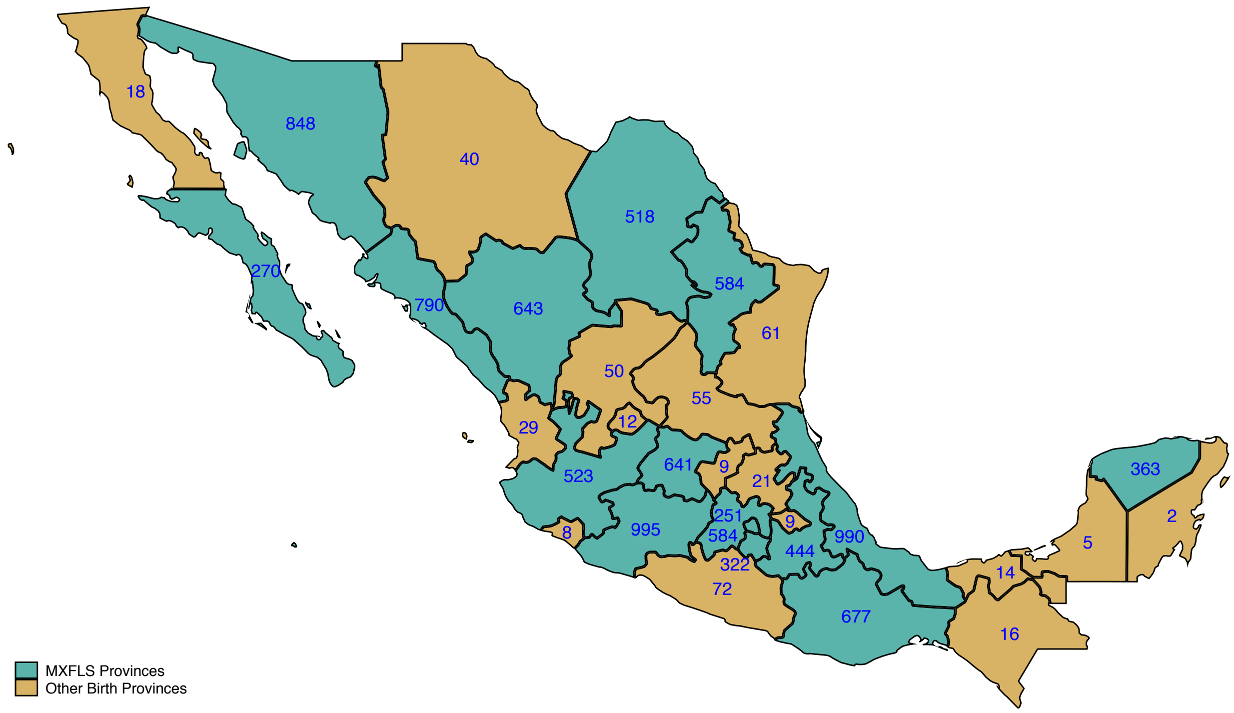

Notes: States in blue are ones in which the MXFLS has been deployed. States in brown are non-MXFLS states in which some subjects in our primary sample were born.

C Construction of Risk aversion measures

Figure 6: Construction of risk aversion measure in IFLS2 and IFLS3

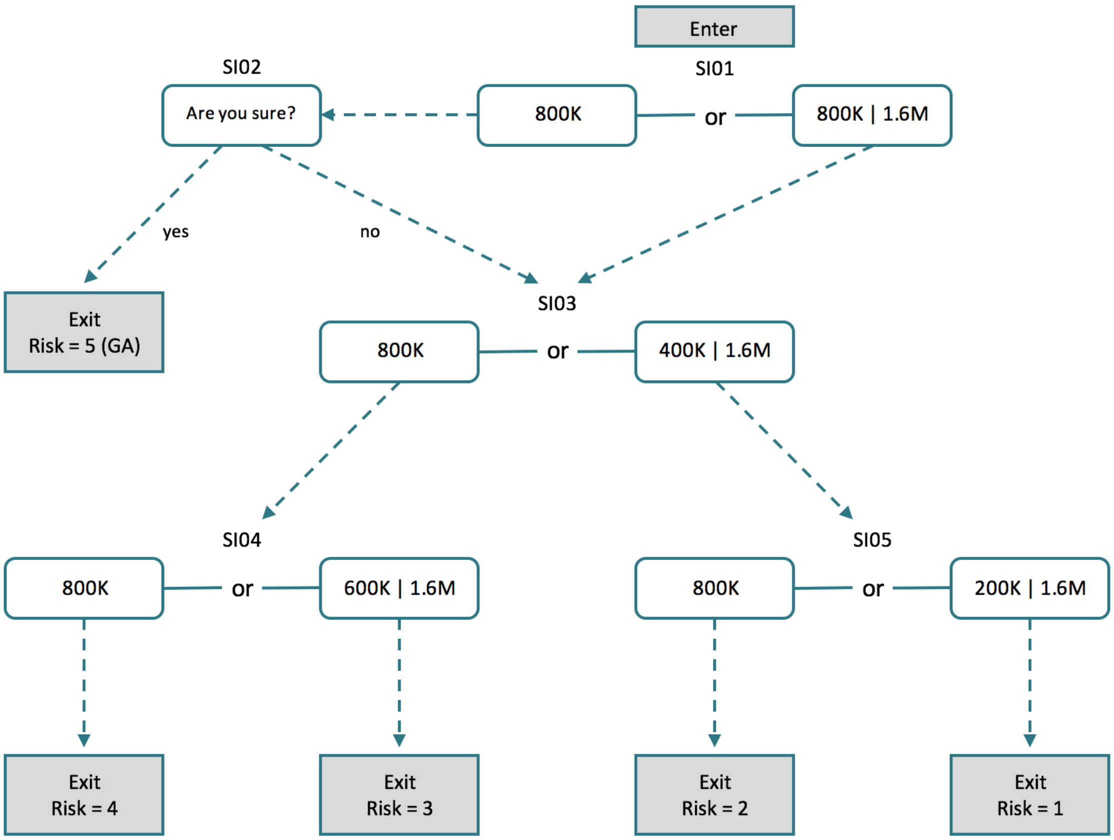

Notes: Higher numbers for "Risk" indicate a higher rate of measured risk aversion. Values are in Indonesian Rupiah.

Figure 7: Construction of risk aversion measure in MXFLS2

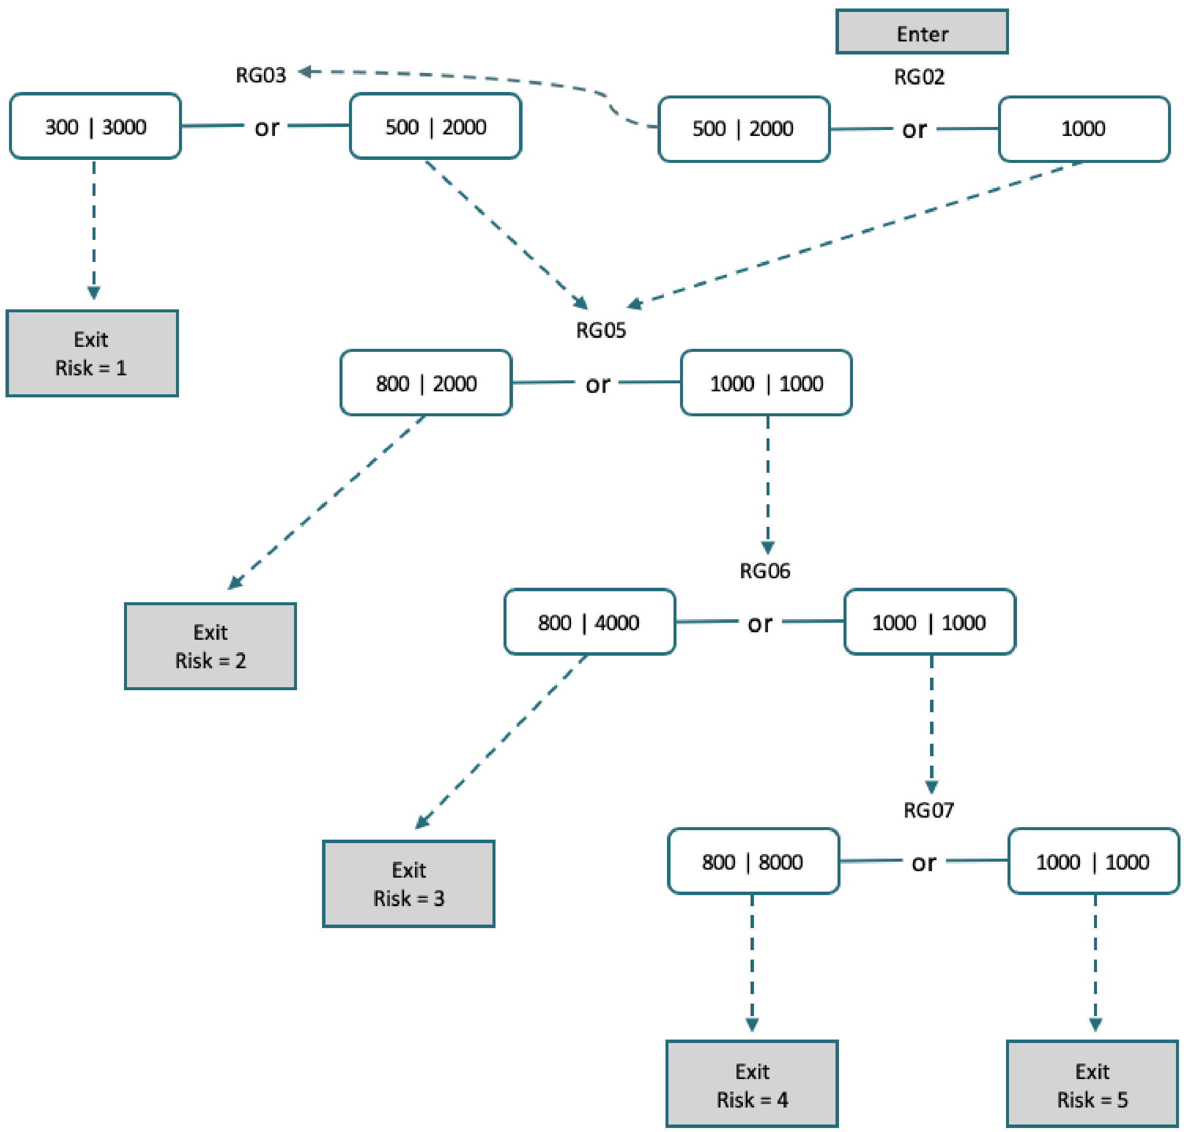

Notes: Higher numbers for "Risk" indicate a higher rate of measured risk aversion. Values are in Mexican Pesos.

Figure 8: Construction of risk aversion measure in MXFLS3

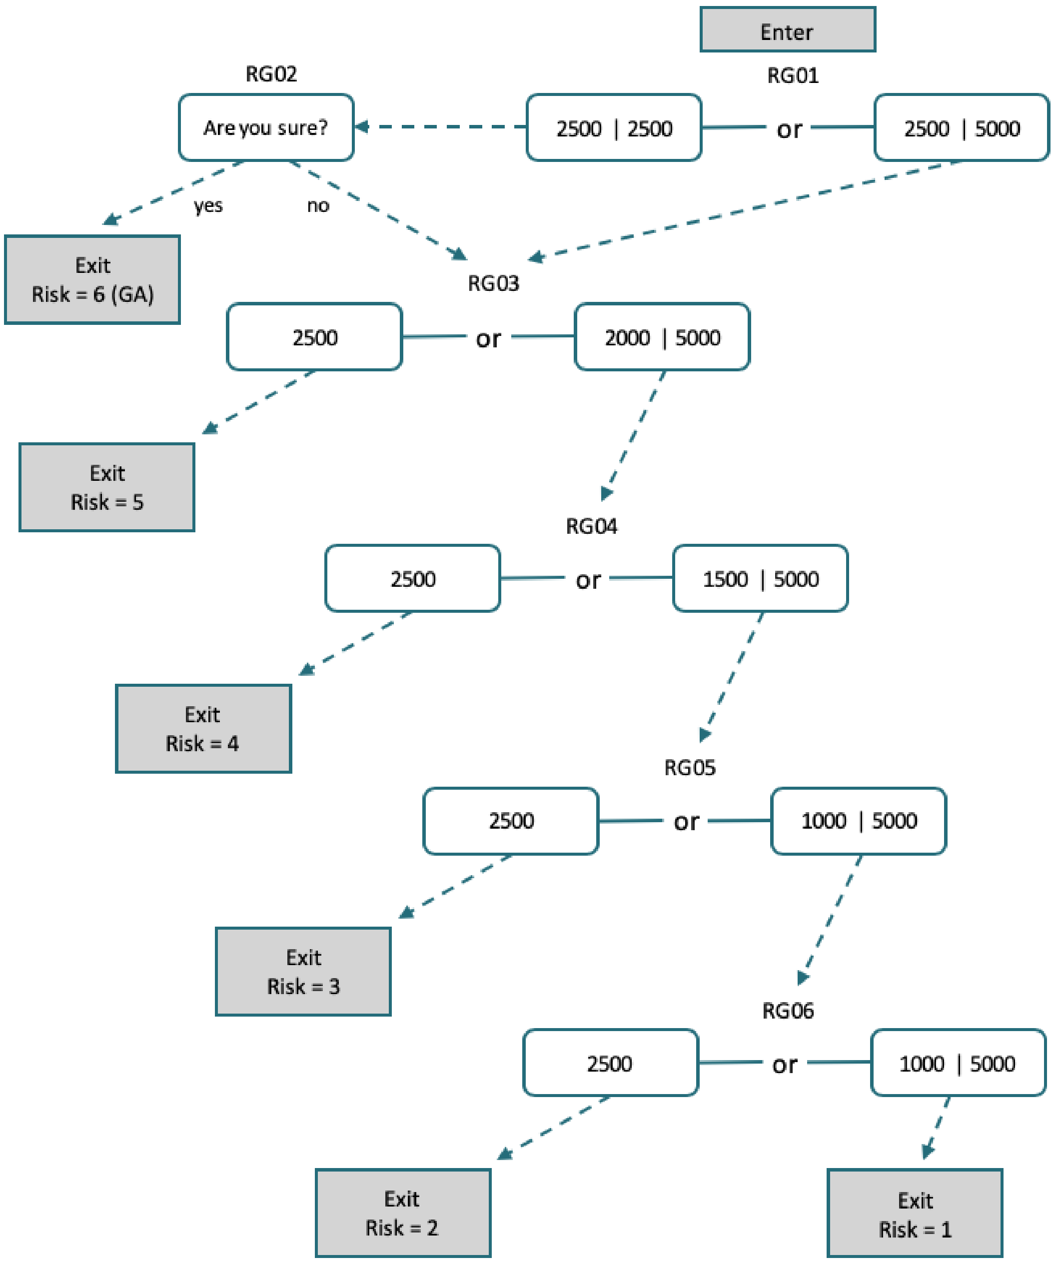

Notes: Higher numbers for "Risk" indicate a higher rate of measured risk aversion. Values are in Mexican Pesos.

D Sample distribution for risk aversion measures

Figure 9: Histogram of Measured Risk Aversion buckets across IFLS4 and IFLS5

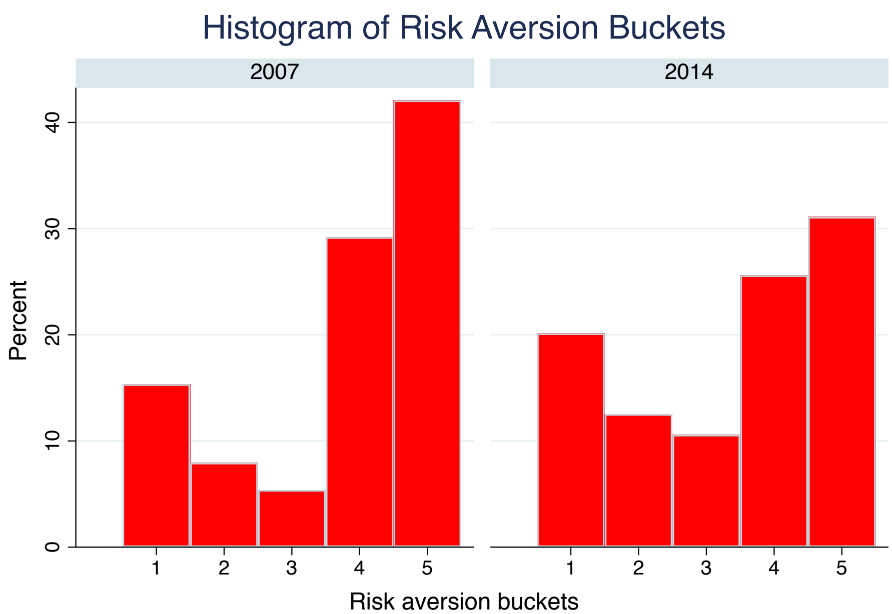

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Distributions for individuals in main regressions: present in both 2007 and 2014 surveys.

Figure 10: Histogram of Measured Risk Aversion buckets across MXFLS2 and MXFLS3

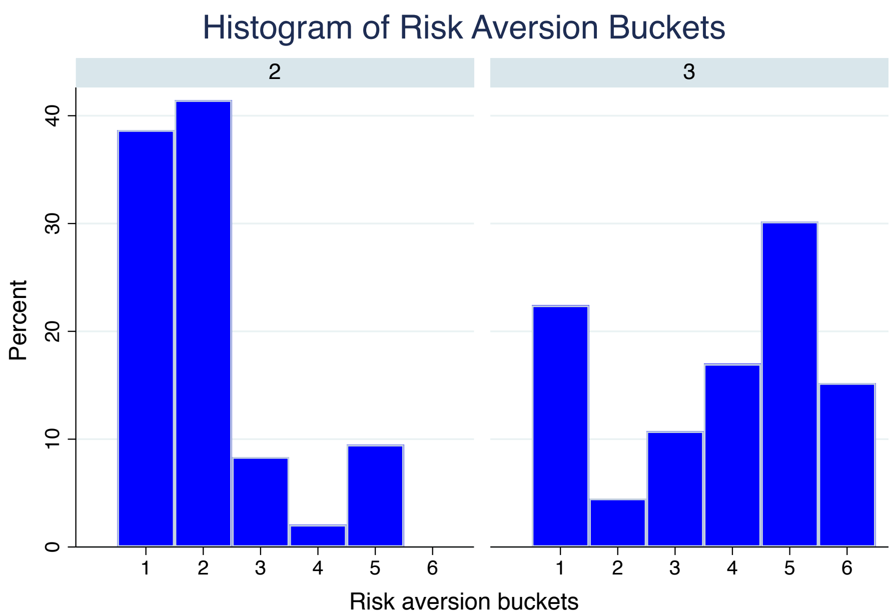

Notes: *Measured Risk Aversion*: 1-6, 6 highest measured risk aversion. Distributions for individuals in main regressions: present in both 2005 and 2009 surveys.

E Correlates of risk aversion measures in the cross-section

Table 4: Correlates of risk preference measures

	Indonesia		Mexico	
Dep. Var:	Measured Risk Aversion		Measured Risk Aversion	
Sample:	X-Sec	Panel	X-Sec	Panel
Self-employed	-0.11***	-0.10***	0.01	0.03
	(0.018)	(0.021)	(0.03)	(0.04)
Migrated	-0.10***	-0.08**	0.02	0.03
	(0.023)	(0.033)	(0.034)	(0.039)
Income	1.52e-06***	1.75e-06***	0.07	0.08*
	(3.39e-07)	(3.86e-07)	(0.05)	(0.04)
Consumption	-0.015***	-0.018***	-0.16	-0.05
	(0.004)	(0.005)	(0.24)	(0.26)
Total assets	-3.47e-05	-3.39e-05	0.009	0.01
	(3.04e-05)	(3.82e-05)	(0.01)	(0.01)
Borrowing	-0.001**	-0.001**	-0.07	-0.1
	(0.0004)	(0.0005)	(0.21)	(0.23)
Savings	-0.0003	-0.0002	0.06	0.18
	(0.0002)	(0.0003)	(0.32)	(0.33)
Smoker	0.09***	0.07*	-0.17***	-0.15***
	(0.0 30)	(0.038)	(0.05)	(0.06)
Cigs/day	-0.06***	-0.04**	0.001	0.001
	(0.02)	(0.02)	(0.0007)	(0.0008)
Woman	0.28***	0.26***	0.04	0.02
	(0.023)	(0.028)	(0.03)	(0.03)
Age	-0.015***	-0.014***	-0.012**	-0.012**
	(0.004)	(0.005)	(0.005)	(0.006)
Age^2	0.002***	0.002***	0.0001**	0.0001**
	(4.25e-05)	(5.64e-05)	(5.55e-05)	(6.22e-05)
Observations	35,848	23,995	11,740	9,335
R-squared	0.052	0.055	0.18	0.168

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

F Details of additional controls

Table 5: Description of controls included in Table 2

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

G Results for alternate samples

A Restricting sample by birth year

Table 6: Alternate samples

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia (born post-1976)				
Δ Growth Mean	-0.42***		-0.37***	0.99*
	(0.12)		(0.10)	(0.56)
Δ Growth Mean^2				-0.14**
				(0.06)
Δ Growth Volatility		1.48***	1.41***	1.19***
		(0.25)	(0.22)	(0.25)
Observations	6374	6374	6374	6374
Mexico (born post-1940)				
Δ Growth Mean	-0.97***		-0.94***	-1.67***
	(0.20)		(0.19)	(0.47)
Δ Growth Mean^2				0.10*
				(0.06)
Δ Growth Volatility		0.85***	0.83***	0.81***
		(0.17)	(0.17)	(0.17)
Observations	7420	7420	7420	7420
Mexico (born post-1976)				
Δ Growth Mean	-0.84***		-0.80***	-1.73
	(0.31)		(0.30)	(1.06)
Δ Growth Mean^2				0.18
				(0.19)
Δ Growth Volatility		0.49**	0.44*	0.44*
		(0.24)	(0.24)	(0.24)
Observations	2284	2284	2284	2284

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

B Excluding Gamble Averse Individuals in Indonesia

Table 7: Excluding Gamble Averse Individuals

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia (excluding gamble averse)				
Δ Growth Mean	0.10		0.13*	1.50***
	(0.07)		(0.07)	(0.50)
Δ Growth Mean^2				-0.13***
				(0.05)
Δ Growth Volatility		0.48***	0.52***	0.31*
		(0.16)	(0.16)	(0.18)
Observations	7193	7193	7193	7193

Notes: *Measured Risk Aversion*: 1-4, 4 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

H Alternate Specifications of Measured Risk Aversion

Table 8: Ordered Probit with two-way fixed effects

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.18*		-0.16*	0.13
	(0.04)		(0.04)	(0.23)
Δ Growth Mean^2				-0.03
				(0.02)
Δ Growth Volatility		0.71*	0.69*	0.65*
		(0.09)	(0.08)	(0.09)
Observations	16083	16083	16083	16083
Mexico				
Δ Growth Mean	-0.49***		-0.47***	-0.78***
	(0.10)		(0.10)	(0.23)
Δ Growth Mean^2				0.04
				(0.03)
Δ Growth Volatility		0.44***	0.42***	0.42***
		(0.09)	(0.09)	(0.09)
Observations	8046	8046	8046	8046

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 9: Binarized measure of risk aversion

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.08*		-0.06*	0.22
	(0.02)		(0.02)	(0.13)
Δ Growth Mean^2				-0.03**
				(0.01)
Δ Growth Volatility		0.33*	0.32*	0.29*
		(0.05)	(0.05)	(0.05)
Observations	16083	16083	16083	16083
Mexico				
Δ Growth Mean	-0.32***		-0.31***	-0.58***
	(0.06)		(0.06)	(0.13)
Δ Growth Mean^2				0.04**
				(0.02)
Δ Growth Volatility		0.26***	0.25***	0.24***
		(0.05)	(0.05)	(0.05)
Observations	8046	8046	8046	8046

Notes: *Binarized Measured Risk Aversion*: Measured Risk Aversion buckets 1 and 2 are set to 0, and buckets 3, 4, and 5 to 1. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

I Results for province of residence macroeconomic conditions

Table 10: Province of Residence Macroeconomic Experiences

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.36***		-0.31***	0.77*
	(0.08)		(0.07)	(0.42)
Δ Growth Mean^2				-0.10***
				(0.04)
Δ Growth Volatility		1.26***	1.19***	1.06***
		(0.17)	(0.16)	(0.17)
Observations	17394	17394	17394	17394
Mexico				
Δ Growth Mean	-1.25***		-1.26***	-2.39***
	(0.19)		(0.18)	(0.41)
Δ Growth Mean^2				0.15***
				(0.05)
Δ Growth Volatility		0.78***	0.80***	0.82***
		(0.17)	(0.17)	(0.17)
Observations	7971	7971	7971	7971

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

J Results using non-linear weighting based on Malmendier and Nagel (2011)

Malmendier and Nagel (2011) estimate a non-linear single parameter weighting function for the effects of mean stock market returns on later in life stock market participation and elicited risk aversion. We extend this method to the context of lifetime experiences of growth volatility. For individual i measured at time t, with experienced growth g_t occurring s years before t, our weighting function is:

$$w_{it}(s,\lambda) = \frac{(age_{it}-s)^{\lambda}}{\sum_{s=1}^{age_{it}-1}(age_{it}-s)^{\lambda}}.$$

This weighting function yields a set of monotonic weights for experiences that always add up to unity, regardless of the age of the individual at measurement. Figure 11 illustrates this weighting scheme for different values of λ for a 30 year old subject. For all ages, higher values of λ mean placing more relative weight on recent experiences, $\lambda = 0$ implies a flat weighting scheme like the one used in our baseline analysis, and negative values of λ indicate placing more relative weight on early life experiences.

Using this weighting scheme we construct a measure of average experienced lifetime growth as follows:

$$A_{it}(\lambda) = \sum_{s=1}^{age_{it}-1} w_{it}(s,\lambda) g_{t-s}.$$

We also construct an analogous measure of experienced lifetime growth volatility using the weighted standard deviations of experienced growth:

$$V_{it}(\lambda) = \sqrt{\frac{\sum_{s=1}^{age_{it}-1} w_{it}(s,\lambda)(g_{t-s}-A_{it}(\lambda))^2}{\frac{age_{it}-2}{age_{it}-1}\sum_{s=1}^{age_{it}-1} w_{it}(s,\lambda)}}$$

We estimate the marginal effects of these weighted macroeconomic experiences variables on measured risk aversion by estimating Equation 18 using non-linear least squares. This allows us to simultaneously estimate the marginal coefficients of mean growth, squared mean growth and growth volatility (β_1, β_2 and β_3), and the value of the non-linear weighting parameter λ:

$$\Delta R_{it} = \alpha + \beta_1 \Delta A_{it}(\lambda) + \beta_2 \Delta \left(A(\lambda)_{it}\right)^2 + \beta_3 \Delta V_{it}(\lambda) + \gamma Inflation_p + \epsilon_{it}. \quad (18)$$

Since non-linear estimation methods are known to be sensitive to initial seed values, we choose the initial value of λ that maximizes the likelihood in a linear specification of our model. Specifically, we build a fine grid of λ values, and generate our macroeconomic variables of interest using each of these values. We then use maximum likelihood estimation on these linearized versions of the model considering for each value of λ, and choose as the initial seed for the nonlinear estimation the value of λ that maximizes the likelihood. Results of this estimation are below.

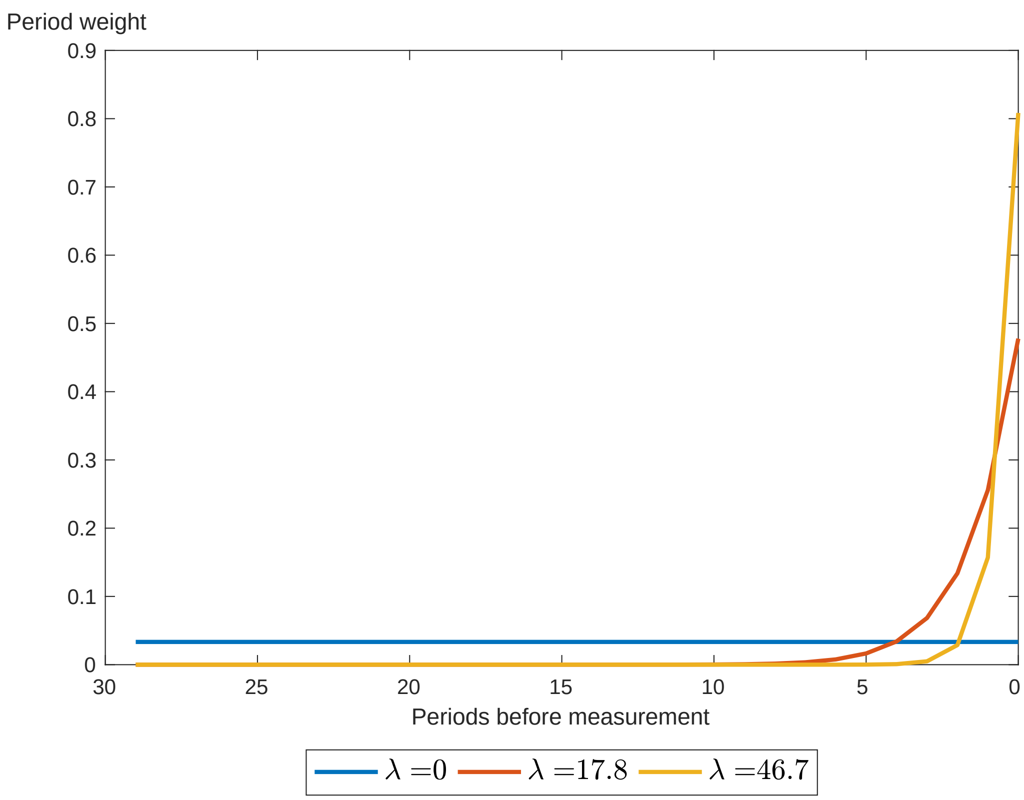

Figure 11: Relative weights placed on years of growth for individual aged 30 at different levels of λ

Table 11: Non-linear temporal λ weighting

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.35***		0.06**	0.49***
	(0.04)		(0.03)	(0.11)
Δ Growth Mean^2				-0.04***
				(0.009)
Δ Growth Volatility		0.89***	0.93***	0.97***
		(0.09)	(0.09)	(0.09)
λ	5.1***	46.1***	46.7***	46.8***
	(0.20)	(3.5)	(13.1)	(2.7)
Observations	17299	17299	17299	17299
Mexico				
Δ Growth Mean	-0.12		-0.32***	-0.06**
	(0.08)		(0.08)	(0.02)
Δ Growth Mean^2				0.03***
				(0.003)
Δ Growth Volatility		0.30***	0.32***	0.02
		(0.10)	(0.09)	(0.02)
λ	-0.2***	0.1***	-0.1***	17.78***
	(0.03)	(0.01)	(0.01)	(2.05)
Observations	8187	8187	8187	8187

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Regressions estimated via NLLS. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

K Results with birth-province/state level clustering

Table 12: Birth-province/state level clustering

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.35		-0.30	0.42
	(0.52)		(0.39)	(0.82)
Δ Growth Mean^2				-0.07
				(0.65)
Δ Growth Volatility		1.36	1.30	1.21
		(0.25)	(0.14)	(0.18)
Observations	17302	17302	17302	17302
Mexico				
Δ Growth Mean	-1.02		-0.97	-1.69**
	(0.19)		(0.12)	(0.05)
Δ Growth Mean^2				0.10
				(0.31)
Δ Growth Volatility		0.91*	0.87	0.86
		(0.10)	(0.13)	(0.14)
Observations	8187	8187	8187	8187

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the province level using the wild bootstrap procedure in Cameron, Gelbach and Miller (2008). P-values from 1000 repetitions in parenthesis. *** p <0.01, ** p <0.05, * p <0.1.

L Results for repeated cross-section specifications

The empirical specification for these regressions, including age, time, and province of birth fixed effects, is

$$R_{it} = \alpha + \alpha_t + \beta_{Age}Age_{it} + \beta_{Prov}Province_i + \beta_1 A_{it} + \beta_2 A_{it}^2 + \beta_3 V_{it} + \gamma ProvInf_{it} + \delta X_{it} + \epsilon_{it}$$

Where X_{it} in this case are controls for gender and ethnicity.

Table 13: Repeated Cross-Section (with panel sample only)

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.07*		-0.09*	-0.14
	(0.03)		(0.03)	(0.18)
Δ Growth Mean^2				0.00
				(0.02)
Δ Growth Volatility		0.04	0.07*	0.07*
		(0.03)	(0.03)	(0.03)
Observations	34851	34851	34851	34851
Mexico				
Δ Growth Mean	-0.00		-0.01	0.04
	(0.03)		(0.03)	(0.07)
Δ Growth Mean^2				-0.01
				(0.01)
Δ Growth Volatility		0.01	0.01	0.01
		(0.02)	(0.02)	(0.02)
Observations	18015	18015	18015	18015

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 14: Repeated Cross-Section

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)
Indonesia				
Δ Growth Mean	-0.05*		-0.05**	-0.15
	(0.02)		(0.02)	(0.11)
Δ Growth Mean^2				0.01
				(0.01)
Δ Growth Volatility		0.05**	0.05**	0.05**
		(0.02)	(0.02)	(0.02)
Observations	55111	55111	55111	55111
Mexico				
Δ Growth Mean	-0.00		-0.01	0.08
	(0.02)		(0.02)	(0.06)
Δ Growth Mean^2				-0.01
				(0.01)
Δ Growth Volatility		0.02	0.02	0.02
		(0.02)	(0.02)	(0.02)
Observations	20976	20976	20976	20976

Notes: *Measured Risk Aversion*: 1-5, 5 highest measured risk aversion. Province (Indonesia) and regional (Mexico) inflation included in all regressions. Standard errors clustered at the cohort by province/state of birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.