Risk-Taking Adaptation to Macroeconomic Experiences*

Remy Levin
University of Connecticut

Daniela Vidart
University of Connecticut

First version: September 2019; **This version**: April 2024

Abstract

We study how lifetime experiences of macroeconomic volatility shape individual risk attitudes. We build a Bayesian model where risk aversion endogenously adapts to agents' beliefs about an exogenous income process. We combine panel data from Indonesia and Mexico containing elicited measures of risk aversion with state-level real GDP growth time series capturing individuals' lifetime macroeconomic experiences. In line with the model's predictions, we find that measured risk aversion increases with macroeconomic volatility, and that this is a first-order driver of risk attitudes. These results are robust to many alternate specifications and controls and extend to risk-taking behavior in other domains.

Keywords: Risk attitudes, experience effects, macroeconomic volatility, growth
JEL Codes: D81, D83, E32, E70, G50, O11

*Corresponding author: Remy Levin, remy.levin@uconn.edu. Vidart: daniela.vidart@uconn.edu. We are grateful for the generous and discerning feedback of numerous colleagues, including 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, Jeffrey Shrader, 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

1 Introduction

Recessions are a central feature of the economic history of every country. Although they are temporary, recessions can have long-lasting impacts on both aggregate (Barlevy (2007), Blanchard et al. (2015)) and individual (Ruhm (2000), Kahn (2010), Yagan (2019)) outcomes. One important individual outcome that recessions are known to affect is the propensity to take risks. In a 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 relationship between experienced recessions and decreased risk-taking has since been replicated in a number of studies (Sahm (2012), Appendino (2013), Dohmen et al. (2016), Malmendier and Steiny (2016), Ampudia and Ehrmann (2017), Guiso et al. (2018), Shigeoka (2019)).

The literature on macroeconomic experience effects for risk aversion has almost exclusively focused on the downturn aspect of recessions – that is, the fact that they make the economic environment worse on average.[1] This emphasis, however, ignores the uncertainty aspect of recessions – the fact that they increase macroeconomic volatility and the variance of the distribution of aggregate income shocks. Theoretical intuition from expected utility theory, where risk preference is thought of as the propensity to substitute between the mean and the variance of a gamble (Pratt (1964), Arrow (1970)), suggests that both these channels could be important drivers of changes in risk aversion. In this paper we show that lifetime experiences of macroeconomic volatility have first-order effects on individual risk attitudes. This implies that recessions, and macroeconomic experiences more broadly, shape individual risk aversion both by worsening or improving agents' economic environments and by stabilizing or destabilizing them.

[1] Malmendier and Nagel (2011) estimate the effects of average stock market returns over subjects' lives. Ampudia and Ehrmann (2017) replicate their analysis for Eurozone countries, and examine volatility as a robustness check. Sahm (2012) studies the role of the Index of Consumer Sentiment, the unemployment rate, and the real return of the S&P 500 in the month of measurement. Dohmen et al. (2016) study the effects of changes in average regional unemployment and average regional GDP growth. Malmendier and Steiny (2016) study the role of average inflation in shaping individuals' decisions to buy versus rent a home. Guiso et al. (2018) examine the effects of changes in expected income and stock returns and experimentally study the roles of salience and fear in mediating their effects. Notable exceptions are Appendino (2013), who examines experienced stock market volatility empirically, and Shigeoka (2019), who studies the effects of severe spells of elevated unemployment rates in early adulthood.

To this end, we employ a combination of theory and empirical analysis. We build a Bayesian model of risky choice where an agent's financial risk aversion endogenously adapts to their learning about the mean and variance of an exogenous income process. We then use large-scale panel 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 in both settings measured risk aversion decreases in experienced lifetime mean growth, increases in experienced growth volatility, and that the effects of growth volatility are first-order relative to those of mean growth. To our knowledge, ours is the first paper in the literature to provide robust empirical evidence that both lifetime experiences of average macroeconomic conditions and macroeconomic volatility directly shape individual attitudes towards risk, and the first to provide a model of experience effects for risk aversion due to exogenous income volatility.

We begin by presenting a theoretical framework that sheds light on the mechanisms through which lifetime macroeconomic experiences shape risk-taking and grounds our empirical analysis. In our model an expected utility maximizer is faced each period with two income risks: a *foreground* risk, which is a menu of risky lotteries from which the agent selects, analogous to the choice between investing in a safe bond or a risky stock; and a *background* risk (Ross (1981), Gollier and Pratt (1996)), analogous to an aggregate income shock, which is exogenous, unavoidable, and statistically independent of the foreground risk. The two risks are substitutes for the agent, so that the more exogenous risk exists in the environment, the less risk the agent wants to take on in their endogenous choice. The agent personally observes realizations of the background risk each period and updates their beliefs about its parameters as a Bayesian. Departing from the benchmark learning model in the economics literature, in which the mean of the distribution is the sole object of learning, we allow both the mean and the variance of the background risk to be unknown to the agent. As the agent's beliefs about the two moments of the background risk update, their choices over the foreground risk adapt in turn.

The central insight of our model is that macroeconomic experience effects for risk aversion can be thought of as a process of experiential Bayesian learning about background income risk. The most important empirical implication of this insight is that the response of individual risk-taking to new income shocks will be heterogeneous by agents' preexisting body of exogenous lifetime income experiences, and that this heterogeneity can be summarized by the additive effects of the new set of shocks on the mean and variance of the agent's set of previous experiences. In other words, when the agent observes a new exogenous income shock, they compare their set of experiences including the novel shock to their preexisting set of exogenous income experiences. If the shock increases the mean of their body of experiences, the agent will become more risk-seeking in their choices, because they now believe the economic environment is better than they had previously thought. If the shock increases the variance of their body of experiences, the agent becomes more risk averse, because they believe the environment is more unstable than they had previously thought. These two effects will be additive, meaning that the overall effect of a novel shock on the agent's foreground risk-taking will be the sum of the two moment effects.

Our main empirical analysis closely mirrors this theoretical result. We use a linear regression framework to estimate the effects of changes in the mean and variance of lifetime real GDP growth on changes in individual risk aversion. To conduct this analysis, we use microdata from the Indonesian Family Life Survey and the Mexican Family Life Survey. Both are high-quality panel surveys (total primary sample $N = 21{,}860$), each containing two elicited measures of risk aversion for the same subjects several years apart (IFLS: 2007 and 2014; MxFLS: 2006 and 2009). Risk aversion is measured in these surveys using staircase instruments 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, from which we construct binned measures of risk aversion in each wave. The main dependent variable in our analysis is within-person changes in these measures of risk aversion. To construct our growth experience variables we assign to each subject a combination of their national and birth-state time series of real GDP growth. We then calculate the means and standard deviations of these time series for each subject from birth to each survey wave. Our main independent variables are changes in subjects' experienced lifetime mean and experienced lifetime standard deviation of real GDP growth between the waves of the respective survey.

Our empirical approach offers several identification advantages for estimating the causal effects of lifetime growth experiences on individual risk aversion. The use of choices over objective income lotteries to measure risk aversion means that we do not have to worry about changes in beliefs about the payout distribution of the risky decision confounding observed changes in risk aversion. Our use of survey measures of risk aversion also means that the menu of lotteries from which subjects are choosing is exogenous to their own macroeconomic experiences, which is not the case for more-naturalistic data like real-world portfolio allocations. Since we use within-person changes in experiences and measured risk aversion, we can be confident that our results are not driven by time-invariant unobserved individual heterogeneity. Our analysis exploits identifying variation from several different sources, including cross-sectional and cohort variation within each of our two settings. Finally, it is worth noting that in both countries the two survey measurements of risk aversion were conducted both before and after the Great Recession, which allows us to leverage the variation arising from that substantial macroeconomic shock.[2]

We find that in both countries increases in the experienced lifetime mean of growth are significantly correlated with decreases in measured risk aversion, while increases in the experienced standard deviation of growth are significantly correlated with increases in measured risk aversion. The coefficients of experienced standard deviation of growth are comparable in magnitude to the coefficients of experienced mean growth (0.53 times as large in absolute terms in Indonesia, 1.87 times as large in Mexico), indicating that experienced macroeconomic volatility is a first-order driver of risk attitudes.

Our main results are estimated without the inclusion of any controls (aside from subnational inflation), because relevant covariates like income, wealth, and savings are potentially endogenous to risk aversion and thus their inclusion could threaten the causal interpretation of our results. Nevertheless, in order to evaluate whether the observed changes in measured risk aversion are driven by changes in subjects' personal economic circumstances or by other factors, we conduct a second analysis where we progressively add additional controls to the main regressions. Our controls include changes in subject income, assets, and savings between the waves of the survey, along with time-varying demographics and exposure to natural disasters and violence. 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 not driven by income effects or by other kinds of lifetime experiences.

We also explore whether the changes we observe in measured financial risk aversion induced by macroeconomic experiences correlate with changes in 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

[2] We also address several specific identification challenges unique to our setting, including issues arising from the existence of subjects who reject first-order stochastically dominant offers when presented with the risk aversion questions (whom we term "gamble averse"); changes in the risk aversion instrument between waves in Mexico; the possibility of confounding in the analysis due to endogenous migration; and questions regarding the mapping from our empirical measures of macroeconomic experiences to the theoretical construct of background income risk experiences in the model.

decreases (or smaller increases) in risk-taking behavior in the domains of investment in Indonesia, migration in Mexico, and smoking in Mexico.

Finally, we test the sensitivity of our baseline results to the most important methodological choices in the analysis. We repeat our main regressions using alternate specifications representing every possible combination of methodological choices along five dimensions: (1) the treatment of older cohorts in the construction of our macroeconomic experience variables; (2) the treatment of gamble averse subjects; (3) the construction of the measured risk buckets in the first wave of the MxFLS; (4) the treatment of subjects who report ever migrating in our sample; and (5) the use of macroeconomic conditions in subjects' state of birth versus subjects' state of residence. We find that the four coefficients of interest are quantitatively and qualitatively consistent with our baseline results across virtually all specifications.[3]

The remainder of the paper is organized as follows. Section 1.A situates this paper and its contributions in the literature. In Section 2 we present our model and its theoretical predictions. In Section 3 we describe our data and empirical methodology. In Section 4 we discuss the identification challenges in our setting and the steps we take to address them. We present the results of our empirical analysis in Section 5. In Section 6 we present additional robustness tests for our empirical results. Section 7 concludes.

1.A Related literature

This paper is related to several strands of the literature. Most closely, it is related to work studying the effects of economic experiences on risk aversion. On the empirical side, this literature includes papers examining the effects of macroeconomic experiences on risk attitudes and related outcomes (Malmendier and Nagel (2011), Sahm (2012), Appendino (2013), Dohmen et al. (2016), Malmendier and Steiny (2016), Ampudia and Ehrmann (2017), Guiso et al. (2018), Shigeoka (2019)), as well as work studying the effects of personal economic experiences of various kinds on related risky decisions, such as IPO bids (Kaustia and Knüpfer (2008)), savings decisions (Choi et al. (2009)), and consumption choices (Malmendier and Shen (2021)). Our empirical analysis makes two primary contributions to this literature. First, we examine the effects of both mean economic conditions and economic volatility when quantifying the effects of lifetime macroeconomic experiences on individual risk attitudes.

[3]We also conduct several additional robustness exercises to test the sensitivity of our results to more-minor methodological choices. These include using an ordered probit specification for the main analysis, binarizing the risk aversion measures, using progressively more conservative clustering levels, and using a non-linear temporal weighting derived from the methodology in Malmendier and Nagel (2011). The results for these exercises are very similar to the results from the baseline analysis.

This contrasts with most of the literature, which has focused only on one of these channels at a time. Second, our empirical analysis exploits substantially more sources of identifying variation in macroeconomic experiences, while also controlling for time-invariant unobserved individual heterogeneity. To our knowledge, ours is the first paper in the empirical experience effects literature to include data from multiple countries, subnational and cohort variation, and repeat observations of the outcome of interest for the same individuals.[4]

Our empirical analysis also contributes to the literature on experience effects for risk aversion in the developing world. This literature has centered on the effects of traumatic experiences, 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 et al. (2018)). To our knowledge, we provide the first empirical 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 developing country settings.

The literature on economic experience effects for risk aversion also includes several papers studying this phenomenon theoretically in the context of financial markets. Most notably, these include Schraeder (2016), who examines the role of learning from experience in shaping high frequency trading patterns such as overreaction, reversal, trading volume, and volatility; Ehling et al. (2018), who explore the effects of the young learning from experience for the risk premia and the risk-free rate; and Malmendier et al. (2020), who study the effects of heterogeneity between investors in both experience effects and recency bias for asset prices.

[4] Malmendier and Nagel (2011) use the national-level time series of stock market returns in the US and examine stock market participation and elicited risk aversion in a repeated cross-section specification. Sahm (2012) uses panel data from the US to study whether risk tolerance is stable and how it responds to national macroeconomic variation. Appendino (2013) follows the same empirical strategy as Malmendier and Nagel (2011) when studying the effects of volatility of stock market returns on risk-taking. Dohmen et al. (2016) uses panel data from both Germany and Ukraine to explore how risk aversion changed after the Great Recession, but do not exploit cohort nor subnational variation. Malmendier and Steiny (2016) 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 within-person changes or subnational macroeconomic data in their analyses. Guiso et al. (2018) study the evolution of risk attitudes of investing clients of an Italian bank before and after the 2008 financial crisis. Shigeoka (2019) exploits within-country variation in macroeconomic conditions in Japan, but does not include within-person changes or cross-country data. Kaustia and Knüpfer (2008) study the link between past IPO returns and future subscriptions in Finland. Choi et al. (2009) use data from 401(k) savings accounts in the US to show that individuals who experience more-rewarding outcomes save more. Malmendier and Shen (2021) exploit both panel data and within-country state-level variation to examine the effects of unemployment on consumption, but do so only for the United States.

Our model contributes to this literature in two ways. First, we study the role of experience effects in shaping the risk-taking behavior of individuals when markets for insurance are incomplete and exogenous risk cannot be avoided. As such, our theoretical results are broadly applicable to many settings outside financial markets where perfect arbitrage and risk-sharing are not reasonable assumptions. Second, we explore the consequences of agent learning over the variance of the body of economic experiences, whereas previous papers have only examined the effects of experiential learning over its mean.

Our model also contributes to the theoretical literature on the effects of additive background risk on risk aversion (Ross (1981), Pratt and Zeckhauser (1987), Kimball (1993), Gollier and Pratt (1996), Eeckhoudt et al. (1996)). Papers in this literature have generally modeled background risk statically or have explored comparative statics of changes in the moments of the background risk on foreground risk aversion. To our knowledge, ours is the first paper in the literature to combine a background risk framework with a dynamic model of Bayesian learning over the background risk and the first to use such a framework to study experience effects for risk aversion.

2 Model

Our overarching aim in this model is to understand the effects that personal lifetime experiences have on individual risk aversion. Achieving this goal is complicated by the fact that the concept of "personal experiences" is quite broad. For one thing, individuals have experiences in a variety of domains: experiences of the stock market, experiences of crime, experiences with the healthcare system, and so on. For another, some experiences are more personal than others, in the sense that some are idiosyncratic and some are aggregate. Nearly every American above a certain age has a clear memory of the September 11$^{\text{th}}$ attacks, but only some have the experience of fighting in the wars that followed. Although conceptually different, both types of experiences are likely to mark individuals psychologically and affect their long-run outcomes. This is apparent when we consider that works like memoirs and obituaries often weave together the idiosyncratic and the collective when recounting the important events that have shaped an individual's life.

A related but separate way to categorize experiences is by those that are under the control of the agent and those that are not. In decision-theoretic contexts this distinction is usually captured by contrasting endogenous outcomes with exogenous shocks. In general, we might expect that the more idiosyncratic an experience, the more likely it is to be endogenous, but this is often not the case: a surprise layoff or personal injury can be both idiosyncratic and exogenous. Aggregate experiences can likewise be endogenous: for instance, an individual might decide to migrate to a new country, thereby exposing themselves to a different history of recessions and booms.

In the model that follows, we focus on the effects that exogenous lifetime experiences of income have on risk aversion. We study experiences of income because our outcome of interest, financial risk aversion, is defined over income, and it is natural as a first step to consider experiences in the same domain. We study exogenous experiences for three reasons. First, their psychological effects are likely to be quite distinct from (and potentially more significant than) the experience effects of endogenous choices, as documented in a large body of work in the psychology and neuroscience literature.[5] Second, the set of exogenous shocks an individual experiences encompasses the historic times through which the individual lives, which has been the primary focus of the experience effects literature to date (Malmendier and Nagel (2011), Giuliano and Spilimbergo (2014)). This means that our theoretical results can directly shed light on previous findings in the literature. Third, by focusing on exogenous experiences we generate sharp predictions on the effects of plausibly exogenous macroeconomic experiences on risk aversion over independent lotteries, which we test in our subsequent empirical analysis.

Our model combines a background risk framework (Ross (1981), Gollier and Pratt (1996)) with a model of experiential Bayesian learning over the background risk. In our model, the agent is an expected utility maximizer faced with two income risks: the *foreground risk*, which is a menu of risky lotteries from which the agent chooses a single lottery each period; and the *background risk*, which is exogenous, uninsurable, and realizes each period after their choice over the foreground risk is complete.[6]

Our agent learns about the background risk through their own experiences. Departing from the benchmark Bayesian learning model in the economics literature, in which the mean of the distribution is the sole object of learning, we allow both the mean and the variance of the background risk are unknown to the agent. This allows perceptions of the variance

[5] This literature, which is centered around the Locus of Control construct (Rotter (1966), Ajzen (2002)), provides evidence that subjective feelings of control can affect the way that individuals process experiences (Bhanji and Delgado (2014)) and are an important determinant of individual choice (Stolz et al. (2020)), including in the domain of risky decision-making (Yi et al. (2018)).

[6] To fix ideas, think of a farmer in Indonesia who each year makes a single choice between planting a safe food crop or a risky cash crop. The agent is making this choice in the presence of background risk. For example, our Indonesian farmer may face background risk in the form of uncertain prices for their crops at the market, realizations of weather that occur after the planting, or macroeconomic conditions that affect their income from non-agricultural wage work.

to affect the agent's risky choice in the long run.[7] Our agent observes realizations of the background risk over their lifetime, and updates their beliefs about both moments accordingly. As their beliefs about the background risk change, their choices over the foreground risk adapt in turn.

A key element of our framework is that the agent lives in a world with incomplete insurance markets, where some income risks can be avoided while others cannot.[8] A necessary condition for a risk to be unavoidable is that the agent's behavior cannot affect their exposure to the risk. Mathematically, this condition is captured in the model by the assumption that the foreground and background risks are statistically independent, which means that the agent cannot explicitly hedge against the background risk through their choices over the foreground risk. Exposure to the background risk, however, can still affect choices over the foreground risk: under standard assumptions about their utility function the two sources of risk are substitutes for the agent, with more risk in the background leading to decreased risk-taking in the foreground.

The choice environment. Consider an agent born at time 0. In each period, indexed by $t \in \{1, 2, ..., T\}$, the agent receives a fixed wealth endowment w, and must choose an income lottery $\tilde{x}_t$ from a menu of lotteries $\mathcal{X}_t$. $\tilde{x}_t$, which we call the *foreground risk*, represents a risky choice that is under the agent's control, such as the choice between purchasing a risky stock or a risk-free bond, or the decision to plant either a safe food crop or a risky cash crop. Without loss of generality we assume that $\mathcal{X}_t$ is a menu of objective lotteries.[9]

In each period, the agent is also exposed to an exogenous and unavoidable *background* income risk $\tilde{y}_t$. Some parameters of $\tilde{y}_t$ are unknown to the agent, who holds beliefs over them that are updated each period as the agent observes the realization of $\tilde{y}_t$. Denote with $B_t(y)$ the CDF of the agent's beliefs distribution about the outcomes of $\tilde{y}_t$ at time t.

[7] In the mean-only learning model the variance shrinks monotonically over time, rendering its dynamic effects negligible in the long run. Our learning model allows the agent's beliefs about the value of the variance of the background risk to increase or decrease over time. This feature of our model distinguishes it from other models of history-dependent risk attitudes, such as those of Dillenberger and Rozen (2015) and Tserenjigmid (2021), in which risk aversion monotonically increases in the mean, a phenomenon known as the reinforcement effect. The central role of the variance in our model is consistent with the intuitive notion in expected utility theory of risk preference as a propensity to substitute between the mean and the variance of a gamble (Pratt (1964)).

[8] Missing and incomplete markets for insurance are a central feature of the economies of developing countries (Besley (1995)). They also feature prominently in developed countries, where many households face large and uninsured labor income risk, proprietary income risk, and real estate risk (Heaton and Lucas (2000)).

[9] Our results apply equally well to a setting where $\mathcal{X}_t$ is populated by subjective lotteries, so long as the agent's beliefs about the foreground and background risks satisfy the independence assumption below.

We assume that $\tilde{x}_t$ and $\tilde{y}_t$ are statistically independent for all t.[10] This ensures that the background risk is outside of the agent's control, because if $\tilde{x}_t$ and $\tilde{y}_t$ were not independent, the agent would be able to hedge against fluctuations in $\tilde{y}_t$ through their choice of $\tilde{x}_t$, and learn about the distribution of $\tilde{y}_t$ from the realizations of $\tilde{x}_t$.[11] We also assume that our agent exhibits hand-to-mouth consumption, so at the end of each period the realizations of $\tilde{y}_t$, $\tilde{x}_t$, and w are consumed.

Utility and risk. Our agent is a subjective expected utility maximizer. Following the background risk literature, we model the agent as possessing two utility functions. The *background utility function* u is the agent's period utility and takes $\tilde{y}_t$, $\tilde{x}_t$, and w as its arguments, additively:

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

The *foreground utility function* v_t[12] represents the agent's utility over the foreground risk alone, conditional on their expectations about the background risk:

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

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

[10] An alternative assumption we could make is that $\tilde{x}_t$ and $B_t(y)$ are independent for all t. Though all our results would still follow, the interpretation of the agent's choice process would be somewhat different: instead of the agent correctly perceiving the independence of the foreground and background risks, we would think of the agent as believing that the risks are independent, regardless of whether they are in reality.

[11] As noted above, the independence assumption formalizes the idea that the agent lives in a partial insurance world, where some exogenous risks cannot be insured away or avoided. Many real-world settings, however, may appear to exhibit correlation between $\tilde{x}_t$ and $\tilde{y}_t$. Consider for instance our Indonesian farmer facing background income risk from macroeconomic conditions, measured as regional GDP. Realizations of this regional GDP may correlate with the returns of planting a cash or food crop, because the farmer's economic activity is a component of local GDP. Note, though, that so long as the background risk can be decomposed into a component $\tilde{y}_t^C$ that is correlated with $\tilde{x}_t$ and a component $\tilde{y}_t^I$ that is independent, the model's insights apply to the effects of $\tilde{y}_t^I$ on choices of $\tilde{x}_t$. This is to say that our theoretical results are applicable to the effects of many real-world risks, even if they are not perfectly independent of the agent's risky choices, so long as they contain a substantial component that is exogenous from the agent's perspective.

[12] v_t is also referred to as the "derived" utility function in the background risk literature (Kihlstrom et al. (1981), Nachman (1982)). Note that v_t is not the agent's value function or indirect utility function. That would be $V_t(w) = \max_{\tilde{x}_t} \mathbb{E}_{\tilde{x}} v_t(w, \tilde{x}_t)$.

risk evolve over time.[13]

Our main variable of interest in the model is the agent's Arrow-Pratt coefficient of absolute risk aversion over the foreground utility function v_t:[14]

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

This transformation of the utility function is a broadly accepted measure of risk preference under complete insurance in expected utility theory. As demonstrated by Pratt (1964), $r(w)$ can be thought of as the approximate insurance premium (or willingness to pay) of an agent to completely avoid a small mean-zero risk at wealth level w. $r(\cdot)$, meanwhile, functions as a global measure of risk aversion.[15] Higher values of $r(w)$ have been linked to more risk-averse decisions in numerous settings, such as portfolio allocations, insurance choices, and entrepreneurship decisions.[16] Our main result details how $r_t(w)$ changes (at any w) in response to newly observed realizations of $\tilde{y}_t$, given the agent's preexisting body of background risk experiences.[17]

We assume the background utility function u is four-times differentiable and make four assumptions on the signs of its derivatives. Two of these assumptions are standard: monotonicity ($u' > 0$) and diminishing marginal utility ($u'' < 0$). The other two are somewhat less common and captured by the following definition:

Definition 2.1. (*Ross risk vulnerability*) Let $\tilde{y}$ have a bounded domain on the interval $[a, b]$. Then u is Ross risk vulnerable iff $\exists \lambda > 0$ such that for all $y_1, y_2 \in [a, b]$:

$$-\frac{u'''(w + y_1)}{u''(w + y_2)} \geq \lambda \geq -\frac{u''(w + y_1)}{u'(w + y_2)}$$

[13] Gollier and Pratt (1996) are the first to note that studying the effects of background risk on optimal risk-taking is equivalent to examining the effects of changes in risk preferences. We develop this idea by modeling the background risk as subject to learning, so that the agent's foreground risk preference adapts to new information about unavoidable background risks in their environment.

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

[15] To see this, consider two utility functions $v_1(w)$ and $v_2(w)$ with corresponding coefficients $r_{v_1}(w)$ and $r_{v_2}(w)$. Pratt (1964) shows that $r_{v_1}(w) > r_{v_2}(w)$ for all w if and only if $v_1 = G(v_2)$, where G is a globally concave function. In other words, $r(\cdot)$ ranks utility functions by the concave order and therefore by global risk aversion.

[16] See, for instance, Merton (1969), Ehrlich and Becker (1972), Kihlstrom and Laffont (1979). In our model a higher r_t would be associated with a higher probability of choosing a riskier $\tilde{x}_t$ from $\mathcal{X}_t$.

[17] This result can easily be extended to $r_t(\cdot)$ by assuming that definition Definition 2.1 holds for all w.

and

$$-\frac{u''''(w+y_1)}{u'''(w+y_2)} \geq \lambda \geq -\frac{u''(w+y_1)}{u'(w+y_2)}$$

Ross (1981) defines a stronger measure of absolute risk aversion than Arrow-Pratt, suitable for the incomplete insurance case. Where the Arrow-Pratt coefficient measures the local concavity of the utility function at the single, insurable state w, the Ross coefficient $r^R(w) = -u''(w+y_1)/u'(w+y_2)$ measures the concavity of the utility function across the insurable state $w+y_1$ and the uninsurable state $w+y_2$.[18] The Ross coefficient can be interpreted as the approximate insurance premium for an agent to avoid a small mean-zero risk in the presence of an additional, unavoidable risk. It is therefore the correct measure of risk aversion for our background utility function u.

Ross risk vulnerability (Eeckhoudt et al. (1996)) is a natural extension of Ross risk aversion into higher-order risk preferences on u. The first condition in the definition above corresponds to decreasing absolute risk aversion in the sense of Ross (i.e. $r^R(w)$ decreasing in w). The second condition corresponds to decreasing absolute prudence in the sense of Ross (i.e. $-u'''(w+y_1)/u''(w+y_2)$ decreases in w). This second condition implies that the agent's precautionary savings motive will decline with their wealth (Kimball (1990)). Collectively, the conditions in Ross risk vulnerability entail that the background risk and any risk in the foreground will be substitutes for the agent. In other words, the more exogenous risk exists in the environment, the less risk the agent will take on in their endogenous choices.

Learning. We combine the above background risk framework with a model of experiential learning over the background risk. We assume the agent is a Bayesian who uses personally observed iid realizations of the background risk to update their belief distribution $B_t(y)$. Our agent perceives the background risk to be a stationary Gaussian random variable with unknown mean and unknown variance. For analytical tractability, we further assume that the agent's prior over the moments of the background risk takes the form of the conjugate prior of the Gaussian with unknown mean and unknown variance likelihood, which is a normal-inverse-chi-squared distribution. We call this learning process formally in the

[18] This yields an ordering of utility functions by risk aversion for an agent facing a choice between prospects that *all* contain uncertainty, unlike the Arrow-Pratt ranking, which is appropriate for an agent trading off uncertainty in one prospect for certainty in another. Notably, Ross risk aversion ranks utility functions not by the general concave order, but by the additive concave order: for utility functions $u_1(w)$ and $u_2(w)$ $r^R_{u_1}(w) > r^R_{u_2}(w)$ if and only if $u_1 = \lambda u_2 + G$, where $\lambda > 0$ is a scalar and G is a monotonically decreasing and concave function. Ross risk aversion implies Arrow-Pratt risk aversion, but the reverse is not the case, which is why the former is a strengthening of the latter.

following definition:

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

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

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

   where M and Σ^2 are time-invariant scalars that are unknown to the agent.

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

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

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

Bayesian learning over the unknown mean of a stationary Gaussian with *known* variance is the benchmark learning model in the economics literature and has been used in a variety of empirical applications.[19] Normal mean-variance learning departs from this benchmark by allowing the agent to learn directly about the variance as well. This learning process has several appealing features for the modeling of experience effects of GDP growth on risk attitudes. Most importantly, it allows the agent's beliefs about the variance of the background risk to increase or decrease over time, unlike the mean-only model in which the variance of beliefs monotonically decreases to zero as new information arrives. Intuitively, this means that realizations of the background risk can meaningfully change the agent's beliefs about how stable the environment is as well as about how favorable it is. This is important because risk preference in expected utility theory can be thought of as an agent's propensity to substitute between the mean and the variance of a gamble (Pratt (1964)), which suggests that we should expect both moments to have first-order effects on the agent's risk preferences. Therefore, a model that strictly constrains the direction of change of the agent's beliefs about the variance is likely to neglect important dynamics of risky choice, as our empirical results confirm.

[19] See, for example, Foster and Rosenzweig (1995), Gibbons et al. (2005).

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

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

Timing. The timing of events for the first two periods in the model is shown in Figure 1. The agent enters each period with an income endowment and beliefs about the background risk. The agent then chooses the foreground risk from the menu available to them before the background risk realizes. Once the two risks realize, the agent updates their beliefs about the background risk. We assume that the agent updates about the two moments of the background risk sequentially.[20] At the end of the period the agent consumes whatever realizations they received that period and enters the next period with the appropriate prior and income endowment.

Figure 1: Timing of events in the model

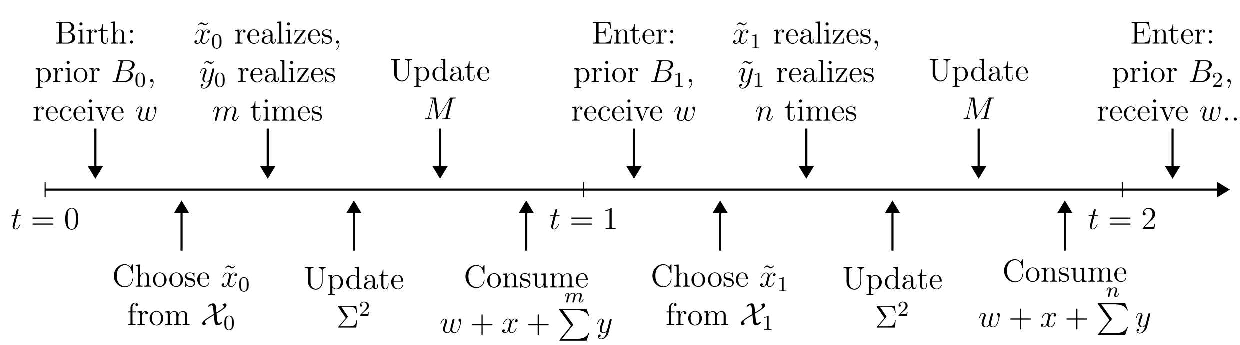

2.A Model result

We are now ready to state the main result of the model:

[20]The consequence of this simplifying assumption is that the change in the agent's risk preferences over the foreground risk is additive in the changes in their beliefs over the two moments of the background risk. This assumption entails minimal loss of generality, because no choices in the model are made *ex interim*.

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

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

Proof. See Appendix A. □

The goal of our model is to understand how endogenous risky choices will change when rational agents are faced with a new set of unavoidable income shocks, after having previously accumulated a large body of personal income shock experiences. Proposition 1 elucidates how this process takes place. Our Bayesian agent compares the cumulative body of data available to them about the unavoidable risk after the new shocks occur to their preexisting set of experiences. The agent's risk aversion will *decrease* if the new set of income shocks exhibits a higher mean than their body of experiences – that is, if the world in their experience has gotten better. The agent's risk aversion will *increase* if the new set of income shocks exhibits a higher variance than their body of experiences – that is, if in their experience the world has become less stable. These two effects will be additive, meaning that the overall effect of a novel shock on the agent's foreground risk preference will be the sum of the two moment effects.

The most important empirical implication of this result is that the response of individual risk-taking to new shocks will be heterogeneous by agents' preexisting body of macroeconomic experiences, and that this heterogeneity can be usefully summarized by the additive effects of the new set of shocks on the mean and variance of their set of previous experiences. We directly test this implication in the empirical analysis below, where our linear regression framework (with changes in measured risk aversion as the dependent variable and changes in the mean and standard deviation of lifetime aggregate income as the independent variables) closely mirrors the additive structure of Proposition 1. Our theory predicts that the marginal effects of mean increases will be positive, that the marginal effects of variance changes will be negative, and that the two marginal effects will be first-order, that is, of roughly the same order of magnitude.

An interesting corollary of our main result is that positive and negative income shocks have asymmetric effects on the agent's risk aversion. For a large negative shock, the mean and variance channels operate in the same direction, while for a large positive shock they operate in opposite directions. Therefore, the marginal effects of a large negative shock on risk aversion will be greater, in absolute terms, than the marginal effects of an equally large positive shock.[21] This result might explain why "depression babies" loom much larger in the popular imagination than "post-war-boom babies", a fact echoed by the empirical macroeconomic experience effects literature, which is almost exclusively focused on the effects of recessions (rather than booms) on risk aversion.[22]

3 Data and Methodology

Our general approach for estimating the effects of long-run experiences of exogenous income shocks on individual risk attitudes is to regress within-person changes in measured risk aversion, obtained from panel survey data from Indonesia and Mexico, on changes in the mean and standard deviation of subjects' lifetime state-level real GDP growth between the waves of the respective survey. Below we provide detailed descriptions of the data and methodology we employ for this purpose.

3.A Survey data

For the Indonesian analysis our source of micro data is the Indonesian Family Life Survey (IFLS) (Strauss et al. (2009), Strauss et al. (2016)). The IFLS is a longitudinal study administered by the RAND corporation in 13 states in Indonesia in five waves, starting in 1993.[23] For the Mexican analysis our source of micro data is the Mexican Family Life Survey (MxFLS), a longitudinal study administered in 16 states in three waves starting in 2002. The MxFLS was piloted by the RAND corporation, and is now managed by the Iberoamerican University (UIA) and the Center for Economic Research and Teaching (CIDE).

Both the IFLS and the MxFLS exhibit high recontact rates (>90%) and contain a wealth of economic and demographic information, allowing for a near-complete accounting of the balance sheet for subjects, including household income, assets, savings and borrowing. Both surveys also contain measures of risk aversion, which we discuss in the next section. We

[21] Conversely, the marginal effects of a small negative shock will be smaller in absolute terms than those of a small positive shock.

[22] Note that our empirical results below, which support the predictions of Proposition 1, also provide direct evidence for the asymmetric effects of positive and negative shocks on risk aversion. This is because by definition large negative shocks are those that decrease the mean and increase the variance, while large positive shocks are those that increase the mean and increase the variance, with the asymmetry following from the estimated coefficients of these moment effects.

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

restrict our attention to subjects who appear in both of the latest waves of the two surveys, and for whom responses to the risk aversion measure are recorded in both. This results in an initial sample (which we call the "full sample") of 17,183 subjects in Indonesia and 12,152 subjects in Mexico, with each subject observed twice. The primary sample we use for our analysis includes further refinements and is described in Section 3.D.

3.B Risk aversion measures

The two most recent waves of the IFLS (IFLS4 [2007–2008] and IFLS5 [2014]) and the MxFLS (MxFLS2 [2005–2006] and MxFLS3 [2009–2012]) include modules for measuring financial risk aversion, which we use to construct the main dependent variables in our analysis.[24]

In both surveys the risk aversion measurement modules employ staircase instruments. These have been shown to generate high-quality measures of risk preference with low subject response burden, which makes them ideal for field applications.[25] In a staircase risk aversion instrument, subjects are presented with a series of hypothetical, high-stakes choices between a safe lottery (often a sure amount of money) and a riskier lottery, which generally has a higher mean and a higher variance than the safe option. Risky lotteries are commonly in the form of fair coin flips with known probabilities. Based on the subject's choice in the first question, they are then asked one of two other questions with different amounts of money for the lotteries. If the subject previously chose the safe (risky) option, risk in the coin flip is reduced (increased) in their subsequent question. This process can then be repeated as many times as necessary to yield as fine a measure of risk aversion as desired. The result is an ordinal binned measure of risk aversion for each subject.

In IFLS4 and IFLS5 subjects answered between two and three questions each, which resulted in a measure with five bins. Each question offered the same sure amount of money (800,000 Indonesian rupiah), while the amounts of the risky lottery varied between questions. The same module structure and payment amounts were used in both waves of the survey. We code the resulting measure by numbering the bins, with higher numbers indicating a higher degree of risk aversion.[26] The dependent variable in our Indonesian analysis is within-person

[24]Risk module information was collected for individuals who were 15 years old or older at the time of the survey in both countries. Enumerators attempted to conduct an interview with all individuals age 12 and over living in sample households. The risk questions, among others, were only posed to individuals of 15 years of age or older.

[25]Notably, the instruments in the IFLS and MxFLS are similar to those employed in the experimentally validated global preference survey in Falk et al. (2018).

[26]The text of the questions used in the risk aversion measurement modules in both Indonesia and Mexico can be found in Appendix B. Our method for constructing the binned measures for the IFLS and MXFLS is

changes in the measured risk aversion bin between IFLS4 and IFLS5.

One complicating factor with the IFLS risk aversion module is that the first question in both waves offered subjects a choice between a sure payment amount and a coin flip over the same amount and a higher payment. A significant fraction of subjects in the sample chose the certain option, which is first-order stochastically dominated by the coin flip, even after being prompted to reconsider a second time. We describe our method for dealing with these "gamble averse" subjects in our analysis in Section 4.A.

In MxFLS2 subjects answered between two and five questions each, which resulted in a measure with seven bins. Most questions offered subjects a choice between a sure amount of money (1,000 Mexican pesos) and a riskier coin flip, with the amounts of the riskier coin flips generally changing between questions. Some questions in the MxFLS2 instrument offered subjects a choice between two coin flips, with the riskier lottery having a higher mean and a higher variance. We code the resulting measure by numbering the bins, with higher numbers indicating a higher degree of risk aversion. In MxFLS3 subjects answered between two and six questions each, resulting again in a measure with seven bins. All questions offered a choice between a sure amount of money (2,500 Mexican pesos in most of them), while the amounts of the risky lottery varied between questions. A gamble averse option was also offered in this instrument. We code the resulting measure by numbering the bins, with higher numbers indicating a higher degree of risk aversion and the top two bins indicating gamble aversion.[27] The dependent variable in our Mexican analysis is within-person changes in the measured risk aversion bin between MxFLS2 and MxFLS3.

Although the overarching methodology (staircase instrument) is consistent across the two waves of the MxFLS, the payment amounts and gamble aversion methodology were changed between the waves (with the MxFLS3 instrument aligning closely with the measurement module in the IFLS). This presents a complicating factor for using the MxFLS instrument to measure changes in risk aversion. We describe our strategy for addressing this measurement issue in Section 4.B.

3.C Macroeconomic experience variables

We use a combination of national-level and state-level data on real GDP growth in Indonesia and Mexico to construct measures capturing subject lifetime macroeconomic experiences,

presented graphically in Appendix C. Histograms of the distributions of the risk aversion measures in both surveys are in Figure D.1.

[27]Our construction of the risk aversion measure in the MxFLS follows the same methodology as Brown et al. (2019).

which are the primary independent variables in our analysis.

We obtain macroeconomic data in Indonesia 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. Yearly time series of national real GDP growth are available starting in 1961. State-level real GDP growth time series for Indonesia are available starting in 1977, and cover all 25 states (spread out over nine islands) reported as birth locations for subjects in our IFLS sample.[28] For Mexico, we obtain yearly national-level real GDP growth, starting in 1925, from the National Institute on Statistics and Geography (INEGI). We obtain yearly state-level growth data from German-Soto (2005), who constructs the time series using historical data from INEGI. These data are available starting in 1941, and cover all 32 states reported as birth locations in our MxFLS sample.[29]

We assign real GDP growth time series to subjects based on their birth year and location of birth.[30] Ideally, we would like to assign state-level time series from birth to all subjects in order to maximize relevant variation in our independent variables. However, a significant share of our sample (62.5% in Indonesia and 10.1% in Mexico) are born before state-level data are available. To address this issue, in our baseline specification we assign to all subjects the real GDP growth time series at the national level for the first 15 years of their life, and at their state of birth for the remaining years. This ensures that we exclude a smaller share of our sample from the baseline analysis (the 23.3% born before 1961 in Indonesia and the 0.2% born before 1925 in Mexico), while treating younger and older cohorts consistently and maintaining congruence in our methodology between the two settings.[31]

[28] One issue with the Indonesian state-level data is that 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 process of decentralization. As a consequence, many administrative units at all levels of the country split, with the number of states in particular increasing from 27 to 34 from 1993 to 2015. Because our analysis requires a consistent definition of states over time, we mapped back all state-level variables (including GDP, population, and inflation) to states as they were defined in 1993. This is possible because in all cases a larger state split into multiple states and in no cases did they recombine into novel states. Thus, every state in 2015 has exactly one corresponding ancestor state in 1993. To avoid confusion we refer to Indonesian states throughout by their names in 1993.

[29] To reduce the effects of possible bias of our results due to measurement error in state-level macroeconomic variables in both countries, we winsorize the state-level GDP measures at the 5–95 level in our main analysis. The winsorized data are presented in Figure 2. The equivalent unwinsorized data plots are available upon request.

[30] We use macroeconomic conditions in subjects' state of birth in our baseline specification to control for possible biasing of our results due to endogenous migration (see Section 4.C for a complete discussion of this issue). We also consider the robustness of our results to using time series from subjects' state of residence (which we observe at a six-month temporal resolution over their entire lifetimes). As shown in Section 6.A, our results are highly robust to using that alternative specification.

[31] To ensure robustness of our empirical results to this methodological choice, we also consider three

For each individual, we calculate the mean (A_{it}) and the standard deviation (V_{it}) from birth to year of measurement in the corresponding survey wave using their assigned growth time series.[32] For example, to calculate the IFLS4 macroeconomic experience variables for an individual born in East Java in 1981, we assign them the growth time series for Indonesia from 1981 to 1996, to which we append the growth series for East Java from 1997 to 2007. That same individual's IFLS5 macroeconomic experience variables are calculated from the national time series for 1981–1996, to which the state time series for 1997–2014 is appended. Let g_{is} be the growth rate assigned to person i in year s and b_i be their birth year. Then for year of measurement t these statistics are:

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

$$\text{where} \quad g_{is} = \begin{cases} g_{National,s} \text{ for } s \in [b_i, b_i + 15] \\ g_{State,s} \text{ for } s \in [b_i + 16, t] \end{cases} \quad \text{and } b_i \geq \begin{cases} 1961 & \text{if Country}_i = \text{Indonesia} \\ 1925 & \text{if Country}_i = \text{Mexico.} \end{cases}$$

Figure 2 presents the distributions of the raw macroeconomic data we use and the macroeconomic experience variables we construct. The top panels display the real GDP growth time series for all 25 Indonesian and 32 Mexican states in our data, as well as the national GDP time series. These time series exhibit substantial variation both in the cross section and over time. The bottom panels display the raw distributions of our main explanatory variables, ΔA_{it} and ΔV_{it}, plotted against each other in each country. These scatterplots demonstrate that substantial variation exists not only in macroeconomic conditions across states, but also in the dynamics of macroeconomic experiences at the cohort-state level.

variations of the growth series assignment rule: (1) we assign state-level growth to all older cohorts as if they were born in 1977 (Indonesia) or 1941 (Mexico); (2) we assign national-level growth for the first 15 years of life to cohorts born 1961–1976 in Indonesia and 1925–1940 in Mexico, while assigning only state-level time series to younger cohorts; and (3) we drop individuals born prior to the availability of state-level time series. Our results are highly robust to the variant of the assignment rule that we use, as shown in Section 6.A. Note that relative to our baseline assignment rule, variants (1) and (2) treat older and younger cohorts inconsistently in our analysis, while variant (3) decreases our sample size dramatically. Our baseline assignment rule can therefore be thought of as a happy medium between the imperatives of maintaining our sample size and treating our subjects consistently in the analysis.

[32]An alternative empirical approach we might have employed, rather than calculating these statistics, would be to regress our measure of risk-taking on growth in each life 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, creating a host of 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 its theoretical predictions in our data.

Figure 2: Variation in macroeconomic data

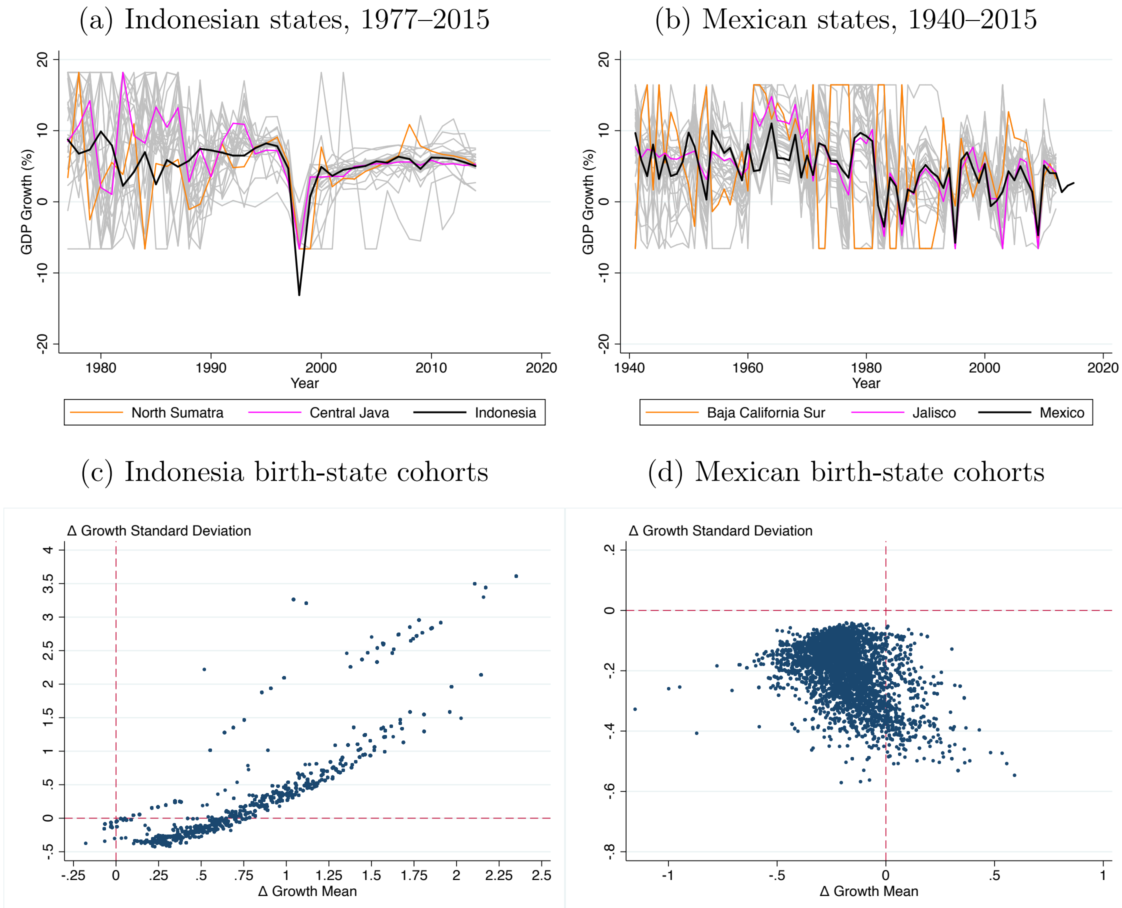

Notes: This figure displays our macroeconomic data in two ways. The top two panels display the real GDP growth time series for all 25 Indonesian and 32 Mexican states in our data (winsorized at the 5–95 level to reduce measurement error), as well as the national GDP time series. The bottom two panels display the raw distributions of our main explanatory variables ΔA_{it} and ΔV_{it} graphed against each other for the primary sample in each country. The primary sample is described in Section 3.D.

3.D Primary sample

After performing refinements on the risk aversion measure and macroeconomic experience variables, we obtain a primary sample of 11,636 subjects for Indonesia and 10,224 subjects for Mexico, with each subject appearing twice in our data. This primary sample comprises individuals who (1) responded to the risk aversion module in both waves of each survey; (2) were born after 1960 in Indonesia or 1925 in Mexico; and (3) either are not gamble averse in our main measure (Indonesia and Mexico), or are gamble averse and also exhibit the highest level of risk aversion in a secondary measure (Indonesia).[33]

[33] See Section 4.A for details on our treatment of the gamble averse.

Summary statistics for the full survey samples and the primary samples are available in Table D.1. This table shows that there are significant differences between the primary and full samples in Indonesia, stemming largely from the fact that individuals in the primary sample are younger, since this sample only includes subjects born after 1960. Differences between the full and primary samples are mostly small and not significant in Mexico.[34]

3.E Empirical specification

Our baseline empirical specification is a first-difference regression:

$$\Delta R_{it} = \alpha + \beta_1 \Delta A_{it} + \beta_2 \Delta V_{it} + \gamma Inflation_p + \epsilon_{it}, \quad (2)$$

where we regress within-person changes in measured risk aversion R_{it} on changes in the lifetime experienced real GDP growth mean A_{it}, and changes in the lifetime experienced real GDP growth standard deviation V_{it}. We include subnational inflation in all our regressions to reduce noise and address potential bias.[35] Standard errors ϵ_{it} are clustered at the state-of-birth by birth-year level, which is the level of treatment in our analysis.[36]

4 Identification

The aim of our empirical analysis is to cleanly identify parameters A and B in Proposition 1, which capture the causal marginal effects of long-run changes in the mean and variance of experienced background income risk on risk aversion over independent income lotteries. Our empirical setup has five key advantages for achieving this task, two stemming from the

[34]Maps of the geographic distributions of our primary samples in Indonesia and Mexico are available in Appendix E. The data covers 25 Indonesian states spread out over nine islands, including the five major ones, and all 32 states in Mexico.

[35]Because R_{it} is measured off of nominal hypothetical lotteries, national-level inflation can introduce noise into the analysis by changing the real value of the prizes offered in the risk elicitation task between waves. Inflation can also introduce bias into our estimates if it varies substantially at the state level, and if it correlates with state-level growth and with risk aversion. To address these concerns, we include a measure of subnational inflation for administrative unit p as a control variable in all specifications. This corresponds to the change in the CPI normalized to 100 during the year of the first wave (IFLS4 and MxFLS2) of the respective survey. In Indonesia p is state, whereas in Mexico due to data constraints p is region.

[36]Nevertheless, given the potential for serial correlation in the lifetime macroeconomic histories of individuals from the same state and adjacent years of birth, in Table O.1 we consider the robustness of our main results to using more conservative levels of clustering, namely state-of-birth by 5-, 10-, and 15-year birth-year bins clustering. Our results remain highly significant across both the Indonesia and Mexico specifications in all cases.

nature of our risk aversion measures and the other three from the identifying variation in our macroeconomic experience variables.

To see the advantages inherent in our risk aversion measures for the task at hand, consider an alternative approach using portfolio allocation behavior, a commonly employed measure of risky decision making. One major identification challenge with attributing changes in portfolio allocations to changes in underlying risk attitudes is that we would have to rule out the alternative explanation that behavioral changes are occurring due to changes in agent beliefs about the returns of different portfolio compositions. This would be impossible to do without additional data on agent beliefs. Since our measures of risk aversion are estimated from choices over objective lotteries with distributions that are known by subjects, we can be confident that the changes in choice behavior we observe are not driven by changes in subject beliefs about the distribution of payouts of the lotteries themselves.

The second major identification challenge with using choice data like portfolio allocations in our setting concerns the independence assumption between background and foreground risk in the model. This assumption codifies the separation between the agent's experiences and choices and is required for the behavioral predictions of Proposition 1. For this assumption to hold in practice, it must be the case that the choice menu from which risk aversion is estimated empirically is independent of subjects' own background income risk experiences. This assumption does not plausibly hold for portfolio allocations: it is quite likely that the menus of assets that subjects can allocate to their portfolios are endogenous to their lifetime experiences of income shocks (for instance, if luckier subjects, who are wealthier ex post, have access to assets with higher returns and fixed costs of investment). This assumption does hold, however, for our risk aversion measures, as the menu of lotteries from which subjects are choosing is by construction independent of their own lifetime experiences.[37]

While the use of the survey risk aversion instruments goes a long way towards ensuring that we are estimating a statistical relationship between personal experiences and underlying

[37]To be sure, using survey measures of risk aversion has some potential disadvantages relative to more-naturalistic choice data. One important concern is the high degree of noise often present in survey measures of risk preferences, as has been documented for similar measures in the Health and Retirement Survey (Barsky et al. (1997)). Measurement noise of this kind can lead to biased estimates when these measures are used as independent variables (Kimball et al. (2008)). In our setting, however, we use these survey measures as dependent variables, which means the primary consequence of the noise is likely to be attenuation in the statistical significance of our results, rather than bias in the estimates (Mccallum (1972), Griliches (1977)). It is also notable that our survey measures of risk aversion correlate in expected ways with measures of demographics and risk-taking behavior in the cross-section, suggesting relatively low noise, particularly in the Indonesian data (Table F.1). These relationships between measured risk aversion and risk-taking behavior also hold for several outcomes when considering changes over time, as shown in the results of Section 5.C.

risk attitudes, it does not guarantee that the correlations we estimate can be interpreted causally. Our empirical setup has three additional advantages on this front, stemming from the construction of our macroeconomic experience variables. First, we estimate within-person changes in experiences and measured risk aversion, which ensures that our results are not driven by time-invariant unobserved individual heterogeneity. Second, we exploit several different sources of variation in macroeconomic experiences. To our knowledge, ours is the first paper in the empirical experience effects literature to include data from multiple countries,[38] subnational and cohort variation in each, and repeat observations of the outcome of interest for the same individuals.[39] Finally, it is also worth noting that in both countries the two survey measurements of risk aversion were conducted both before and after the Great Recession (2007 and 2014 in Indonesia, 2006 and 2009 in Mexico), which means that we can leverage the variation arising from that substantial macroeconomic shock in both settings.

The remainder of this section addresses several additional identification issues unique to our setting. These involve specific measurement issues for risk aversion from our survey instruments (Section 4.A and Section 4.B), potential endogenous migration (Section 4.C), and the mapping between our empirical measures of real GDP growth experiences and the theoretical construct of background income risk experiences (Section 4.D).

4.A Treatment of the gamble averse

One important feature of our risk aversion data, particularly in the IFLS, is that a fraction of subjects in our sample are "gamble averse"—that is, when presented with a choice between a sure amount of money, \$X, and a fair coin toss where they may receive either \$X or a larger amount \$Y, subjects choose the option that delivers \$X with certainty, even after being prompted to reconsider.[40] This means that gamble averse subjects are rejecting first-order stochastically dominant lotteries, a choice that is not rationalizable under most simple models of risk preferences. Therefore, gamble aversion in our setting may represent an extreme form of risk aversion, or it may be driven by other factors such as subject misunderstanding,

[38]Notably, Indonesia and Mexico, while sharing a recent history of rapid and volatile economic change, offer a distinct contrast along many plausibly important dimensions, including geography, level of development, language, culture, religion, institutions, and other aspects of their histories. This aids in establishing the external validity of our results.

[39]See Section 1.A for a detailed review of the empirical strategies used in the literature to date.

[40]The gamble aversion (GA) questions were deployed as part of the risk aversion measurement module in both IFLS4 and IFLS5, as well as in MxFLS3 (see Appendix B and Appendix C). Of the full sample of 17,183 subjects in Indonesia, 42.1% were GA in IFLS4, 31% in IFLS5, 14.8% were GA in both waves, and 58.2% were GA in at least one wave. Of the full sample of 12,152 subjects in Mexico, 14.1% were GA in MxFLS3.

inattentiveness, or religiosity (due to the prohibition on gambling in Islam, for instance).

Since gamble aversion may not truly represent high levels of risk aversion for all subjects, including all the gamble averse in our analysis and treating them as highly risk averse could introduce bias into our results. On the other hand, simply excluding all gamble averse from the analysis could introduce significant bias as well, as there are strong theoretical and empirical grounds to believe that gamble aversion may represent a high level of risk aversion for some subjects in our sample.[41] Empirically, it is notable that gamble averse subjects in our data are older, lower income, and more likely to be female, all of which are well-established correlates of higher risk aversion (Falk et al. (2016)). Furthermore, a large share of non-gamble-averse subjects (27.9%) in the IFLS switch into gamble aversion, while a large percentage of the initially gamble averse (64.8%) switch out of gamble aversion between waves. For these switchers in particular, gamble aversion is more likely to be a true expression of risk aversion, because the alternative drivers of gamble aversion, such as subject misunderstanding or religiosity, are less likely to vary over time.[42]

To deal with this problem, we develop a set of three diagnostics to separate out the two subgroups of the gamble averse. First, we use behavior in a second risk elicitation task in the IFLS to detect subjects who are gamble averse but definitively not risk averse.[43] Second, we detect gamble averse individuals whose observable characteristics suggest that they are more likely to exhibit this behavior due to misunderstanding, inattentiveness, or religiosity.[44] Third, we flag subjects who are gamble averse in both waves of the IFLS, under

[41] Theoretically, Mu et al. (2020) show that in the presence of background risk and small-stakes risk aversion, any model of risk preferences would necessarily imply that agents will violate monotonicity with respect to first-order stochastic dominance. This result implies that in our setting, where background risk is large and important, we should in fact expect an expression of subjects' risk attitudes to involve a rejection of first-order stochastically dominant lotteries. This theoretical prediction of Mu et al. (2020) is borne out empirically in Gneezy et al. (2006), who document extensive gamble averse behavior in both field and laboratory experimental settings. It is quite unlikely that these results are driven by subject misunderstanding, inattentiveness, or religiosity. Rather, as the authors argue, rejection of first-order stochastically dominant lotteries (which they term uncertainty aversion) seems to be an important aspect of individual risk attitudes in numerous settings.

[42] Summary statistics for the gamble averse and the full transition matrix are available in Appendix G.

[43] Both IFLS4 and IFLS5 include a second risk preference elicitation measure—risk task B. This task has the same staircase structure as the task we use to construct our main risk aversion measure, but with substantially higher monetary stakes. Given the higher stakes, a large majority of subjects (77.5% in IFLS4, 70.7% in IFLS5) select the most risk averse option in risk task B. Some subjects who do not select the most risk averse option in this measure also choose the gamble averse option in our main measure. We surmise that these subjects cannot be gamble averse due to a high level of risk aversion.

[44] These include subjects who score in the lowest two buckets (1 or 0) of a Raven's matrices measure of cognitive ability; subjects who are assessed by enumerators to be in the lowest three categories (fair, not so good, very bad) of attentiveness and seriousness when answering the survey; and subjects who are in the highest bucket of a measure of self-reported religiosity (in Indonesia).

the assumption that they are more likely than switchers to exhibit this behavior due to reasons other than risk aversion.

Using these diagnostic criteria, we construct five subsamples of our IFLS subjects, ordered by decreasing likelihood of including gamble averse who are not risk averse (and by increasing likelihood of excluding gamble averse who are risk averse): (1) a sample including all gamble averse, coded as the highest level of risk aversion (full sample); (2) the full sample, minus gamble averse subjects who are risk-seeking in the second risk task; (3) the previous sample, minus subjects who are likely not gamble averse due to risk aversion based on observables; (4) the previous sample, minus subjects who are gamble averse in both waves of the IFLS; and (5) the full sample, minus all subjects who are gamble averse in any wave of the IFLS. For the MxFLS, we construct three analogous subsamples: (1) the full sample including all gamble averse; (2) the full sample minus subjects who are likely not gamble averse due to risk aversion based on observables; and (3) the full sample minus all subjects who are gamble averse in MxFLS3.

In our baseline specification, we use subsample 2 for the Indonesian analysis and subsample 3 for the Mexican analysis, because we believe these provide a good balance between the two sources of potential bias in each setting, while maintaining an adequate sample size. To ensure that our results are robust to these methodological choices, we examine the sensitivity of our main results to the choice of gamble averse subsample in Section 6. As the specification charts in this section indicate, our main results are highly robust to the treatment of the gamble averse. Across virtually all specifications, the signs of the coefficient estimates for the effects of experienced mean and variance in Mexico and experienced variance in Indonesia remain consistent, regardless of which gamble averse subsample we use.[45]

4.B Change in the risk aversion measure in the MxFLS

As noted in Section 3.B, although both waves of the MxFLS use a staircase instrument to measure risk aversion, the payment amounts and gamble aversion methodology were changed between MxFLS2 and MxFLS3.[46] We follow two strategies to deal with this methodological

[45] We have also examined the robustness of all of our main and robustness empirical results to the exclusion of all the gamble averse in Indonesia. Reassuringly, all results are qualitatively identical when excluding the gamble averse. These results are available upon request.

[46] Given the change in the questions asked between waves of the MxFLS, caution should be exercised when interpreting the results from Mexico. In particular, a conservative interpretation of the results is that they capture changes in individual risk aversion relative to overall changes in the subject population. In other words, a positive effect of "treatment" (here growth experiences) indicates that the level of risk aversion of more-treated individuals increases (or decreases less) relative to those less treated.

change between waves. First, in line with the method in Brown et al. (2019), for the main and subsequent analyses we collapse the 5th, 6th, and 7th measured risk aversion bins in MxFLS2 into one bin with level 5 of measured risk aversion. By doing this we obtain an equal number of bins of measured risk aversion in MxFLS2 and MxFLS3, and guarantee that the highest category of risk aversion covers the same range of relative risk aversion coefficient values across both MxFLS waves. To ensure that our results are robust to this methodological choice, we examine the sensitivity of our main results when we do not collapse the bins in this way in Section 6.A. As the specification charts in this section indicate, our main results are highly robust to different constructions of the measured risk buckets in MxFLS2.

Second, we structurally estimate a risk aversion parameter for subjects in MxFLS2 and MxFLS3 and repeat our empirical analysis using the estimated risk aversion coefficients. To do this, we assume that subjects are expected utility maximizers with constant relative risk aversion (CRRA) utility and estimate the range of values that their risk aversion coefficient can take, given their income and choices in the risk module in a given wave. This structural estimation exercise yields a comparable measure of risk aversion for each subject across both waves. We conduct the exercise assuming both narrow and broad bracketing of the lottery payouts, and using both point and set distance measures between the estimated parameters. We find that the results for both Indonesia and Mexico match our baseline results both quantitatively and quantitatively across all specifications of the structural estimates. Details on this analysis and results are presented in Appendix H.

4.C Accounting for endogenous migration

One source of potential bias in our analysis is endogenous migration by subjects. Since migration choices are affected by risk attitudes (Bryan et al. (2014)), and because it can affect the economic environment that individuals are exposed to, migration could generate reverse causality in our setting. In our baseline analysis, we use growth conditions in subjects' state of birth, rather than their state of residence, to control for the potential effects of endogenous migration. In Section 6.A we further explore the robustness of our results to this methodological choice by estimating alternative specifications of our main analysis using macroeconomic conditions in subjects' state of residence, and/or excluding all migrants from the analysis. Our empirical results are very similar under these alternative specifications.

In addition, we repeat our primary analysis for a subsample of individuals who migrated out of their state of birth when young (under the age of 17), using the state of residence to construct the lifetime macroeconomic series. Because early-life migration decisions are likely made by subjects' parents, we can both exploit the fact that state of residence may more accurately capture relevant macroeconomic variation while addressing concerns of endogenous migration. We find that the results in this analysis, which are presented in Table I.1, are quantitatively and quantitatively similar to those in our baseline analysis in both Indonesia and Mexico, though noisier due to the significant reduction in sample size.

4.D Real GDP growth as a measure of background risk

An important consideration for identification in our setting is whether our empirical measure of macroeconomic experiences, which is based on real GDP growth, is an appropriate measure for the background risk in our model. Since real GDP growth measures changes in aggregate income, and the exogenous risk facing agents in our model is a series of exogenous income shocks, theoretically the connection between these constructs is quite close. Nevertheless, it is conceivable that empirically, real GDP growth may not capture real changes in living standards, particularly in settings with low market integration and large shares of agricultural income, like the developing countries under consideration in this paper.

We adopt a three-pronged approach to showing that GDP growth is a good measure of background risk in our settings. First, we survey the literature on the relationship between GDP growth and living standards in developing countries. We find that this literature overwhelmingly indicates that GDP changes closely correlate with changes in living standards in both developed and developing countries, including for the poorest in these settings. For example, Dollar and Kraay (2002) find that the average incomes of the poorest fifth of society rise proportionately with GDP, and show that this finding holds across regions, time periods, and income levels. Similarly, Ravallion and Chen (1997) find that absolute poverty measures typically respond quite elastically to growth using household surveys for 67 developing economies between 1981 and 1994, and that in almost all cases, poverty fell with economic growth and rose with contractions. Kapsos (2006) finds that for every percentage point of additional GDP growth, total employment grew between 0.3 and 0.38 percentage points between 1991 and 2003, using data from 139 countries. Likewise, Ball et al. (2019) find that GDP growth correlates negatively with unemployment, with a 1% increase in GDP reducing unemployment by 0.2% in developing economies. These findings are echoed in many similar studies in developing and middle-income countries, such as Ravallion et al. (1999), Adams (2003), and Cravo and Schimanski (2019).

Second, we examine the empirical relationship between GDP growth and living standards for our Indonesia and Mexico samples specifically. Using data from the IFLS, we calculate the correlation between changes from IFLS4 to IFLS5 in four measures of living standards (log household income, unemployment, self-reported poverty, and self-reported hunger), and average yearly state-level real GDP growth in the years between IFLS4 and IFLS5. We find that subjects living in states with higher average GDP growth during this period report significantly larger increases in income, and significantly larger decreases in poverty and hunger, with nonsignificant changes in employment. We repeat this exercise for income and employment using the final two waves of the MxFLS, and find that subjects living in states with higher average yearly GDP growth between MxFLS2 and MxFLS3 report significantly larger decreases in unemployment, with nonsignificant effects on income. Since data on other dimensions of economic well-being is not readily available in the MxFLS, we supplement the Mexican analysis with data from the Mexican population and housing census, which is available in four waves from 1990-2010. We aggregate this data at the state by year level, and use the proportion of houses in the state with an earthen floor, a key input in multidimensional indices of poverty in the developing world (Alkire et al. (2018)), as a measure of living standards. We find that states with higher average yearly GDP growth exhibit significantly larger decreases in the proportion of homes with earthen floors. Taken as a whole, these results (described in detail in Appendix J) are strongly indicative that real GDP growth is a good measure of changes in living standards in both Indonesia and Mexico.

Third, we study the heterogeneity of our main results by sector of employment. If GDP growth is a good measure of background income risk for subjects, one pattern we might expect to see is that the effects of growth experiences on measured risk aversion will be larger for subjects employed in sectors of the economy that are more strongly correlated with GDP growth. We find a pattern of results that is consistent with this idea: in both Indonesia and Mexico, subjects who are employed in sectors of the economy that are known in the literature to be more procyclical in developing countries, namely agriculture, social services, and finance, exhibit stronger responses to lifetime experiences of GDP growth. We present the full results of this analysis in Appendix K.

5 Empirical Results

5.A Main results

Our main empirical results for the primary subject sample defined in Section 3.D and employing the empirical specification in Section 3.E are presented in Table 1.

Table 1: Main results

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)
Indonesia			
Δ Growth Mean	-0.23***		-0.85***
	(0.06)		(0.12)
Δ Growth Std. Dev.		-0.01	0.45***
		(0.04)	(0.08)
Observations	11636	11636	11636
Mexico			
Δ Growth Mean	-1.04***		-0.86***
	(0.20)		(0.20)
Δ Growth Std. Dev.		2.02***	1.61***
		(0.41)	(0.42)
Observations	10224	10224	10224

Notes: *Measured Risk Aversion*: 1–5 (Indonesia and Mexico), with 5 being the highest measured risk aversion. State (Indonesia) and regional (Mexico) inflation included in all regressions. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

In column 1, we find that the estimated effect of mean growth in both countries is negative and highly significant.[47] In column 2 we show the results of regressing changes in measured risk aversion on changes in the standard deviation of growth. The estimated effect of changes in volatility is positive and highly significant in Mexico but not significant in Indonesia. Once we regress changes in measured risk aversion on both mean and volatility in column 3, however, we find that the estimated effects of increases in mean growth are negative, the effects of increases in the standard deviation of growth are positive, and that the effects of both are highly significant in both settings. These results are consistent with Proposition 1, the main result of our model.

It is also notable that the magnitude of the standard deviation coefficient is 0.53 (Indonesia) to 1.87 (Mexico) times as large in absolute terms as the mean coefficient in column 3, where their magnitudes are most directly comparable. This implies that the marginal effects on measured risk aversion of experienced long-run changes in variance are first-order relative to experienced changes in the mean, in line with the theoretical result in Proposition 1.

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

5.B Additional controls

Our main results are estimated without the inclusion of any 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 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 potentially 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 driving the observed changes or whether other kinds of experiences, which may be correlated with growth dynamics, are in fact playing a central role.

We provide evidence on these points in Table 2, where we progressively incorporate additional controls to the specification for the last column in Table 1. These include time-varying demographics, namely marital status, educational attainment, household size, and household size squared; changes in income, assets, and savings;[48] and self-reported exposure to natural disasters, self-reported exposure to violence, and a measure of municipal homicide rate in Mexico.[49] Full details on the controls are available in Table L.1.

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.[50]

[48] All monetary variables are at the household level and inflation-adjusted to local currency in the first wave of the survey. As such, in Indonesia monetary variables are in terms of millions of rupiah of 2007 and in Mexico monetary variables are in terms of 2007 pesos. To make these adjustments, we use consumer price index information for the states (Indonesia) or regions (Mexico) of residence of individuals in the two countries.

[49] Our measure of the municipal-level homicide rate is built by Brown et al. (2019). It is based on the municipality of residence in MxFLS2 and captures the homicide rate in the 12 months prior to being interviewed. See Brown et al. (2019) for details. We thank the authors of that paper for kindly sharing their data.

[50] We have also extended the above analysis by adding further asset and income controls in order to better address potential unobserved wealth concerns. Specifically, we have considered the following asset controls: (1) the second, third, and fourth orders of assets; (2) quintile dummies of assets; (3) liquid and illiquid assets (liquid assets comprise savings, receivables, jewelry, livestock, poultry, and cash crops in Indonesia; and savings, livestock, and poultry in Mexico. Illiquid assets comprise the difference between assets and liquid assets.); (4) the second, third, and fourth orders of liquid and illiquid assets; (5) housing and land assets

Table 2: Additional controls

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Indonesia							
Δ Growth Mean	-0.85***	-0.89***	-0.89***	-0.89***	-0.89***	-0.88***	-0.89***
	(0.12)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)	(0.13)
Δ Growth Std. Dev.	0.45***	0.47***	0.47***	0.47***	0.47***	0.46***	0.47***
	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)	(0.08)
Observations	11,636	11,282	11,281	11,281	11,281	11,281	11,281
Mexico							
Δ Growth Mean	-0.86***	-0.91***	-0.91***	-0.91***	-0.91***	-0.80***	-0.78***
	(0.20)	(0.20)	(0.20)	(0.20)	(0.20)	(0.21)	(0.21)
Δ Growth Std. Dev.	1.61***	1.69***	1.70***	1.69***	1.69***	1.87***	1.80***
	(0.42)	(0.43)	(0.43)	(0.43)	(0.43)	(0.44)	(0.44)
Observations	10,224	10,050	10,050	10,050	10,050	9,627	9,627
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 Net yearly savings					X	X	X
Δ Violence						X	X
Δ Natural disasters							X

Notes: *Measured Risk Aversion*: 1–5 (Indonesia and Mexico), with 5 being the highest measured risk aversion. State (Indonesia) or regional (Mexico) inflation included in all regressions. Demographics include marital status, household size, household size squared, and educational attainment. All monetary variables are at the household level and inflation-adjusted to local currency in the first wave of the survey (millions of rupiah of 2007 in Indonesia and pesos of 2005 in Mexico). Violence variables from self-reported exposure only for Indonesia, and self-reported exposure and municipal homicide rate built by Brown et al. (2019) for Mexico. Natural disasters variables from self-reported exposure in both settings. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

5.C Changes in risk-taking behavior in other domains

In this section we explore whether the changes we observe in measured financial risk aversion induced by macroeconomic experiences correspond with changes in risk-taking behavior in other domains. Our objective with this analysis is to better understand the behavioral mechanisms driving changes in risky behavior in a variety of behavioral domains, and how

and other assets; and (6) the second, third, and fourth orders of housing and land assets and other assets. In addition, we have considered the following income controls: (1) the second, third, and fourth orders of income; (2) quintile dummies of income; (3) labor income and other income; and (4) the second, third, and fourth orders of labor income and other income. These additional controls are similar to those used by Malmendier and Shen (2021) to address unobserved wealth concerns. We find that in both settings our results barely change when we include these additional controls. These results are available upon request.

these changes are similar to or different from those in the domain of financial risk aversion. While an active debate exists in the literature on the domain specificity of economic risk preferences when considered statically (Blais and Weber (2006)), to our knowledge these results are the first presented in the literature on domain specificity of the dynamics of risk-taking.

To conduct this analysis, we construct a variable measuring predicted change in risk aversion ($\widehat{\Delta R_{it}}$) using our preferred specification (column 3 of Table 1), and examine its correlation with changes in risk-taking behaviors for our subjects. We focus on behaviors commonly examined in relation to risk-taking in the literature and for which we have data: smoking, having ever migrated across state lines, self-employment status, and, in Indonesia, whether subjects report that their land is planted in at least one cash crop, a measure of risky agricultural investment.[51]

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

We draw four conclusions from these results. First, the choice to plant cash crops is an analogous behavioral measure to financial risk aversion, and responds strongly and consistently to macroeconomic experiences in our Indonesian data. Subjects who become more risk tolerant significantly increase their plantings of cash crops, and subjects who become more risk averse significantly reduce (or do not change) their planting behavior. Differences in behavioral changes between these two groups are highly statistically significant. These results are reassuring for the interpretation of our main results, because investment behavior is closely related in theory to financial risk aversion.

Second, we find consistent evidence across both countries that rates of migration increase in concert with increased financial risk tolerance. In addition, we find that within these individuals, those with a larger increase in risk tolerance have a larger increase in rates of migration. These findings are consistent with the interpretation in the literature of migration

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

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

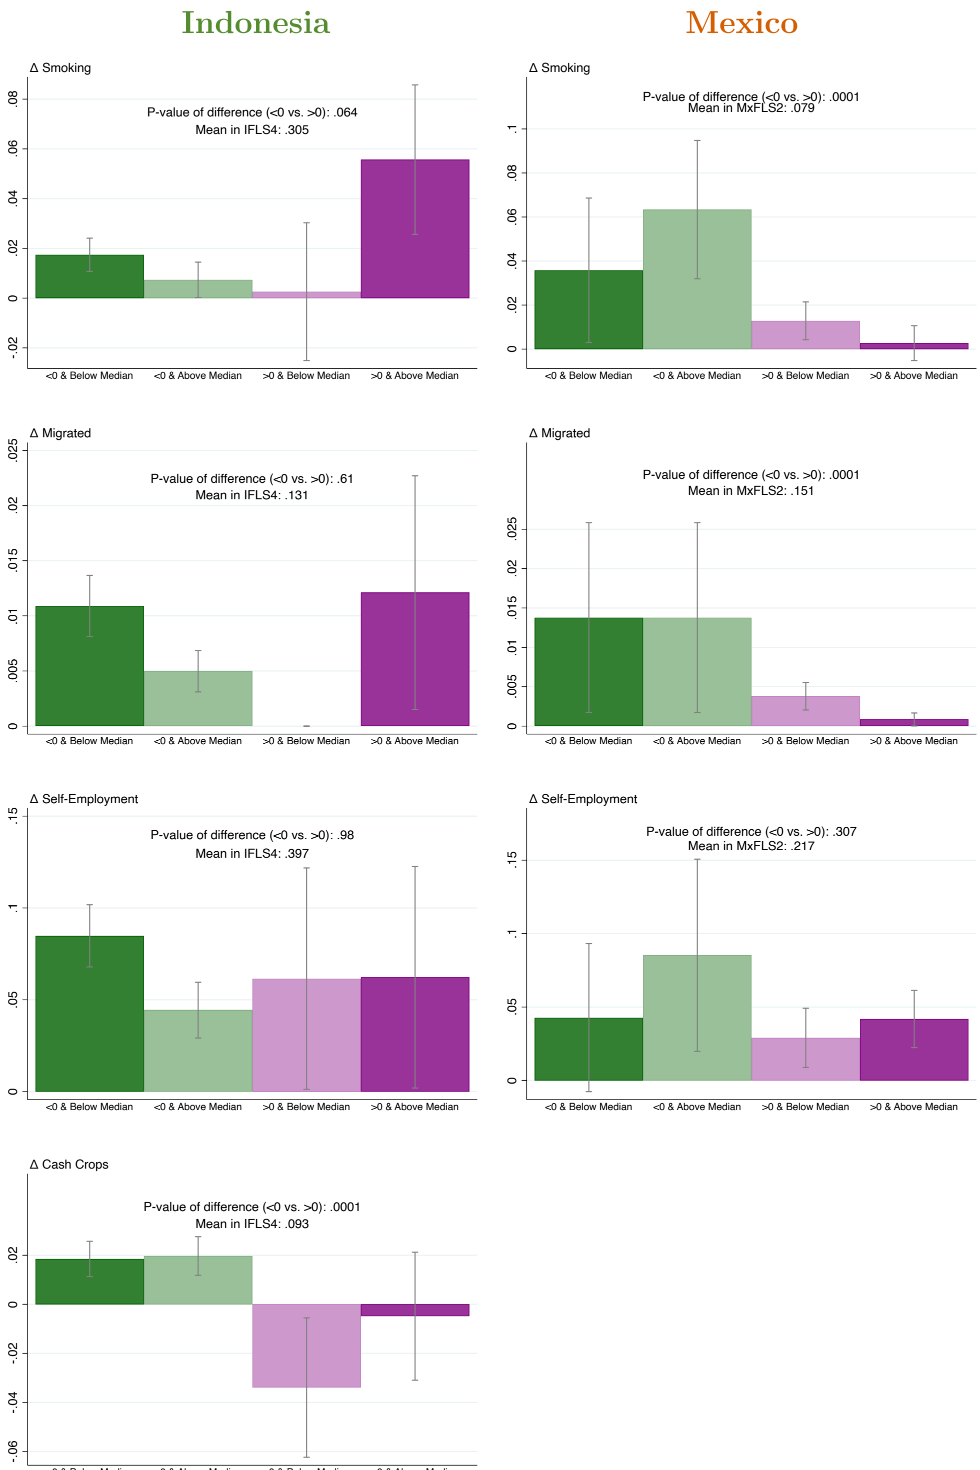

Notes: Bars represent results of mean group changes of the indicated risky behavior for different bins of the predicted change in risk aversion distribution. Green bars represent subjects who are predicted to become less risk averse, while purple bars represent subjects who are predicted to become more risk averse. Darker-colored bars within each of these groups represent subjects whose change in predicted risk aversion is larger (in absolute terms) than the median among individuals in the group, while lighter-color bars represent subjects whose change in predicted risk aversion is smaller (in absolute terms) than the median among individuals in the group. 95% confidence intervals are indicated for each group. P-values of differences are from a two-sided t-test of difference between individuals with positive and negative predicted changes in risk aversion. These results are for subjects in the primary sample in Indonesia, described in Section 3.D.

as a risky choice (Bryan et al. (2014)). However, in the Indonesian context we observe a non-linear relationship between migration rates and predicted changes in financial risk aversion, with individuals who are predicted to become the most risk averse reporting the largest increases in cross-state migration. These results suggest that migration may not be a straightforwardly driven by risk attitudes.

Third, we find mixed results across both settings in the domain of smoking. In Mexico we find that rates of smoking increase with increased risk tolerance. In Indonesia, though rates of smoking increase with risk tolerance, the largest increases in smoking rates occur for those individuals who are predicted to become the most risk averse. This latter finding is particularly interesting, because Indonesia is one of the countries with the highest rates of smoking in the world (World Health Organisation (2019)). More research on the relationship between risk aversion and smoking as a risky health behavior is warranted.

Finally, we find across-the-board increases in self-employment across all bins in both settings, with no significant differences in behavioral change in this domain by predicted changes in financial risk tolerance. This finding bolsters a set of recent findings in the literature that distinguish between measures of self-employment and entrepreneurship in developing countries, and that argue that the former do not capture the latter, given high rates of subsistence self-employment in these settings (Herreño and Ocampo (2021)).

6 Robustness

6.A Sensitivity to alternative specifications

As a primary test of the robustness of our baseline analysis, we consider the sensitivity of our main results to alternate specifications along five dimensions: (1) the treatment of older cohorts in the construction of our macroeconomic experience variables;[52] (2) the treatment of the gamble averse for the purposes of our risk aversion measures;[53] (3) the construction

[52] As discussed in Section 3.C, in our baseline specification we assign to all subjects the real GDP growth time series at the national level for the first 15 years of their life, and at their state of birth for the remaining years. Here we consider three alternative assignment rules: (i) assigning state-level growth to all older cohorts as if they were born in 1977 (Indonesia) or 1941 (Mexico); (ii) assigning national-level growth for the first 15 years of life to cohorts born 1961–1976 in Indonesia and 1925–1940 in Mexico, while assigning only state-level time series to younger cohorts; and (iii) dropping individuals born prior to 1977 (Indonesia) and 1941 (Mexico) from the analysis. Details on the construction of growth experience variables for each of these variations are available in Appendix M.

[53] As discussed in Section 4.A, in our baseline specification in Indonesia we include all subjects who are not gamble averse, or who are gamble averse and in the highest risk aversion bucket in a second measure.

Figure 4: Specification charts and associated sample sizes for **Indonesia**

(a) Δ Growth Mean (b) Δ Growth Std. Dev.

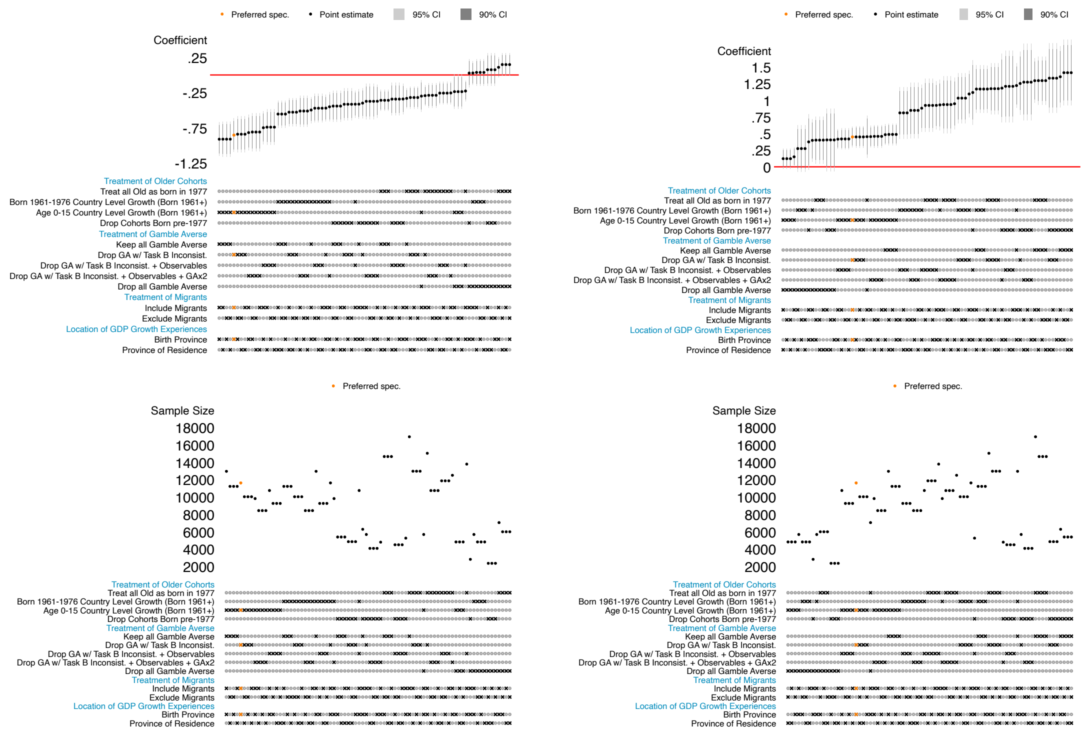

Notes: The top two figures display the effects of changes in experienced mean and standard deviation of growth, respectively, on measured risk aversion in Indonesia for different specifications. 90% and 95% confidence intervals from standard errors clustered at the birth-year by state-of-birth level plotted. The bottom two figures display the associated sample sizes in each of these specifications.

of the measured risk buckets in MxFLS2;[54] (4) the treatment of subjects who report ever migrating in our sample; and (5) the use of macroeconomic conditions in subjects' state of birth versus subjects' state of residence.[55] We present the estimated coefficients from column

In Mexico we include all subjects who are not gamble averse in MxFLS3. Here we consider four alternative treatments of the gamble averse in Indonesia, ordered by increasing restrictiveness: (i) including all the gamble averse; (ii) excluding the gamble averse who are risk-seeking in a second task, as well as those who may be gamble averse but not due to risk aversion based on observables; (iii) the previous sample, plus excluding those who are gamble averse in both waves of the IFLS; and (iv) our primary sample, excluding all subjects who are gamble averse in either wave of the IFLS. For Mexico, we consider two alternative treatments of the gamble averse: (i) including all subjects who are gamble averse in MxFLS3 and (ii) excluding only those gamble averse who may not be gamble averse due to risk aversion based on observables.

[54] As discussed in Section 4.B, in our baseline specification in Mexico we collapse the 5th, 6th, and 7th measured risk aversion bins in MxFLS2 into one bin, with level 5 of measured risk aversion. Here we consider the results when we do not collapse these bins.

[55] As discussed in Section 4.C, in our baseline specification we do not exclude migrants. Here we present alternative specifications excluding all those who who report migrating across state lines at least once for a

3 of Table 1 for every combination of these alternative specifications, as well as the associated sample sizes, in the specification charts in Figure 4 and Figure 5.

Figure 5: Specification charts and associated sample sizes for **Mexico**

(a) Δ Growth Mean (b) Δ Growth Std. Dev.

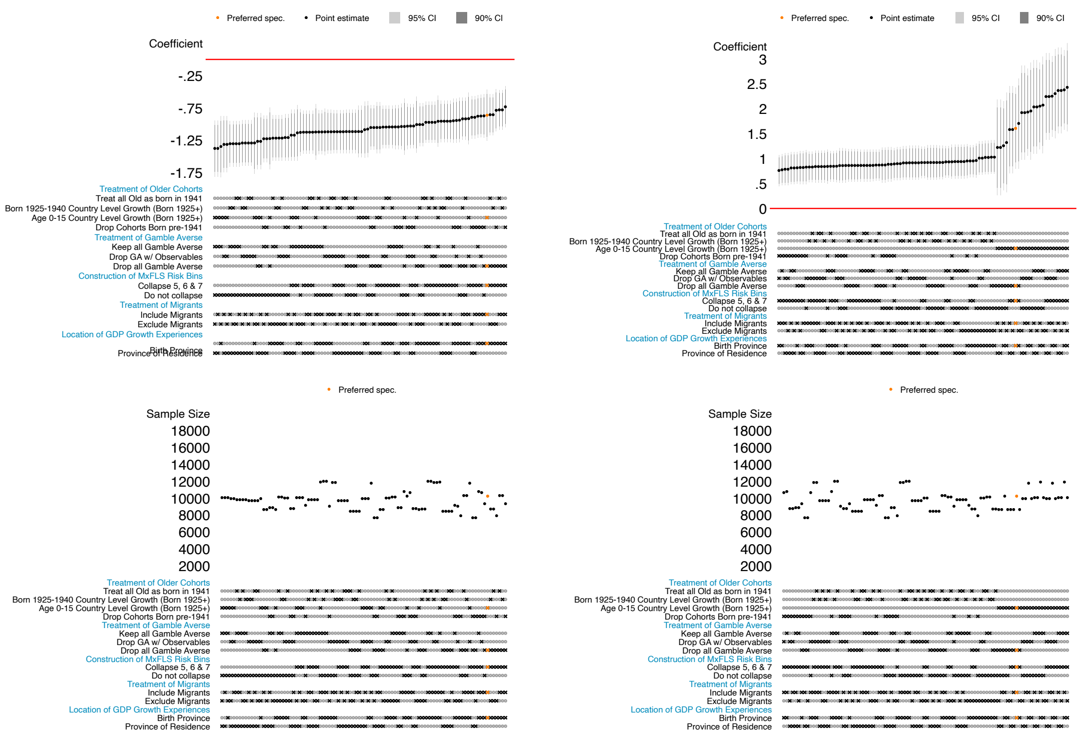

Notes: The top two figures display the effects of changes in experienced mean and standard deviation of growth, respectively, on measured risk aversion in Mexico for different specifications. 90% and 95% confidence intervals from standard errors clustered at the birth-year by state-of-birth level plotted. The bottom two figures display the associated sample sizes in each of these specifications.

We find that the results across virtually all specifications are quantitatively and qualitatively consistent with our baseline results. For three of our four coefficients (standard deviation of growth in Indonesia, mean growth in Mexico, and standard deviations of growth in Mexico) the sign of the effect never changes from baseline, and results remain significant at conventional levels, regardless of specification. For the mean growth coefficient in Indonesia, the sign and significance are the same as in the baseline in almost all specifications, with the exception of a small number of specifications that attenuate the coefficient to zero. These

period of at least six months. In addition, in our baseline specification we use growth time series in subjects' state of birth to control for endogenous migration. Here we present alternative specifications using growth time series in subjects' state of residence.

specifications, however, require discarding around two-thirds of our primary sample, and are thus likely to introduce significant bias into our analysis. These findings imply that our results are highly robust to alternative methodological choices.

6.B Additional robustness

We perform several additional tests of the robustness of our results. In Appendix N we present the results of our main analysis for alternate specifications of measured risk aversion. For both Indonesia and Mexico, we repeat the analysis using (1) an ordered probit specification;[56], and (2) 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). We present the results in Table N.1, Table N.2. In all specifications results are very similar to the baseline. In Appendix O we then repeat the analysis using more conservative levels of clustering, namely state-of-birth by 5-, 10-, and 15-year birth-year bins clustering. We present the results in Table O.1, and find that our results remain highly significant across both the Indonesia and Mexico specifications in all cases.

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 P we extend their method to the context of lifetime volatility experiences. The results are qualitatively and quantitatively similar to those in the baseline model. We find evidence of recency bias across all specifications using this method, though this is especially marked in Indonesia.

7 Conclusions

In this paper we investigated how lifetime experiences of macroeconomic volatility, in addition to lifetime experiences of mean macroeconomic conditions, shape individual attitudes towards risk. To achieve this goal we built a Bayesian model of learning over background income risk that grounded our empirical analysis and shed light on the mechanisms through which exogenous long-run income experiences shape risk-taking. We then used data from

[56]The ordered probit specification accounts explicitly for the ordinal nature of our risk aversion measure, though it may potentially introduce bias due to its non-linear nature.

Indonesia and Mexico to link within-person changes in measured risk aversion to lifetime changes in state-level real GDP growth mean and variance. In line with our model's predictions, we found that in both countries increases in the experienced lifetime mean of growth are significantly correlated with decreases in measured risk aversion, while increases in the experienced variance of growth are significantly correlated with increases in measured risk aversion. These results are robust to a variety of alternate specifications and controls and extend to changes in risk-taking behavior in other domains. Our results, both theoretical and empirical, suggest that experienced macroeconomic volatility is a first-order driver of individual risk attitudes.

While our current empirical analysis focuses on developing countries, a fruitful area for future research would be to extend our analysis examining the effects of experienced growth mean and variance to developed country contexts. It is unclear, *a priori*, whether we might expect our methodology to yield stronger or weaker results in those settings. On the one hand, our theory suggests that the more robust insurance markets that exist in developed countries may ameliorate the impact of macroeconomic fluctuations on risk-taking. On the other hand, real GDP growth may be more strongly correlated with personal income experiences in settings with greater financial and market integration. Parsing out the importance of these two competing mechanisms would allow us to further generalize our empirical and theoretical findings to other settings.

Another interesting question left open by our work is whether experienced volatility stemming from other types of exogenous experiences shapes individual risk attitudes in similar ways to macroeconomic volatility. Are the destabilizing effects of climate change, pandemics, and terrorist attacks similar in their behavioral implications to the effects of recessions? If not, how do they differ? The experience effects literature, particularly in developing countries, contains some intriguing hints on this front, but more research is clearly warranted in this area. A more comprehensive account of experience effects for risk aversion would need to grapple with how to integrate multiple sources of environmental risk into a single framework and with understanding how these different sources of experienced risk might interact in shaping individual risk attitudes.

References

Adams, R. H. (2003). Economic Growth, Inequality, and Poverty: Findings from a New Data Set. *Policy Research Working Papers*. (Cited on page 29.)

Ajzen, I. (2002). Perceived behavioral control, self-efficacy, locus of control, and the theory of planned behavior. *Journal of Applied Social Psychology*, 32(4):665–683. (Cited on page 9.)

Alkire, S., Kanagaratman, U., and Suppa, N. (2018). The Global Multidimensional Poverty Index (MPI): 2018 Revision. *OPHI MPI Methodological Notes 46*, 31(7):1800–1806. (Cited on page 30.)

Ampudia, M. and Ehrmann, M. (2017). Macroeconomic experiences and risk taking of euro area households. *European Economic Review*, 91(C):146–156. ISBN: 2014-10. (Cited on pages 2, 6, and 7.)

Appendino, M. (2013). Lifetime Experience of Volatility: A New Determinant of Household Demand for Stocks. *Working Paper*, page 46. (Cited on pages 2, 6, and 7.)

Arrow, K. J. (1970). Aspects of the Theory of Risk Bearing. In *Essays in the Theory of Risk-Bearing*. (Cited on page 2.)

Bahadir, B. and Gumus, I. (2016). *Credit decomposition and business cycles in emerging market economies*, volume 103. Elsevier B.V. (Cited on page 63.)

Ball, L., Furceri, D., Leigh, D., and Loungani, P. (2019). Does One Law Fit All? Cross-Country Evidence on Okun's Law. *Open Economies Review*, 30(5):841–874. (Cited on page 29.)

Barlevy, G. (2007). On the Cyclicality of Research and Development. *American Economic Review*, 97(4):1131–1164. (Cited on page 2.)

Barsky, R. B., Juster, F. T., Kimball, M. S., and Shapiro, M. D. (1997). Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study. *The Quarterly Journal of Economics*, 112(2):537–579. ISBN: 00335533. (Cited on page 24.)

Bernanke, B. and Gertler, M. (1990). Financial Fragility and Economic Performance. *The Quarterly Journal of Economics*, 105(1):87–114. (Cited on page 63.)

Besley, T. (1995). Nonmarket Institutions for Credit and Risk Sharing in Low-Income Countries. *Journal of Economic Perspectives*, 9(3):115–127. (Cited on page 10.)

Bhanji, J. P. and Delgado, M. R. (2014). Perceived control influences neural responses to setbacks and promotes persistence. *Neuron*, 83(6):1369–1375. (Cited on page 9.)

Blais, A.-r. and Weber, E. U. (2006). A Domain-Specific Risk-Taking (DOSPERT) scale for adult populations. *Judgement and Decision Making*, 1(1):33–47. ISBN: 1930-2975. (Cited on page 34.)

Blanchard, O., Cerutti, E., and Summers, L. (2015). Inflation and Activity – Two Explorations and their Monetary Policy Implications. *IMF Working Paper*. (Cited on page 2.)

Brown, P., Daigneault, A., Tjernström, E., and Zou, W. (2018). Natural Disasters, Social Protection, and Risk Perceptions. *World Development*, 104:310–325. (Cited on page 7.)

Brown, R., Montalva, V., Thomas, D., and Velásquez, A. (2019). Impact of Violent Crime on Risk Aversion: Evidence from the Mexican Drug War. *Review of Economics and Statistics*. (Cited on pages 7, 19, 28, 32, 33, 50, and 65.)

Bryan, G., Chowdhury, S., and Mobarak, A. M. (2014). Under-investment in a Profitable Technology : The Case of Seasonal Migration in Bangladesh. *Econometrica*, 82(5):1671–1748. ISBN: 9781604138795. (Cited on pages 28 and 36.)

Callen, M., Isaqzadeh, M., Long, J. D., and Sprenger, C. (2014). Violence and Risk Preference : Experimental Evidence from Afghanistan. *American Economic Review*, 104(1):123–148. ISBN: 9550091031. (Cited on page 7.)

Cameron, L. and Shah, M. (2015). Risk-taking behaviour in the wake of natural disasters. *Journal of Human Resources*, 50(2):484–515. (Cited on page 7.)

Choi, J. J., Laibson, D., Madrian, B. C., and Metrick, A. (2009). Reinforcement Learning and Savings Behavior. *The Journal of Finance*, 64(6):2515–2534. Publisher: [American Finance Association, Wiley]. (Cited on pages 6 and 7.)

Chávez Martín del Campo, J. C., Rodríguez Vargas, R., and Fonseca Hernández, F. d. J. (2010). Vacas gordas y vacas flacas: La polítca Fiscal y el Balance Estructural en México, 1990-2009. *Estudios Económicos*, 25(2):309–336. (Cited on page 63.)

Cravo, T. and Schimanski, C. (2019). Who Suffers During Recessions in Brazil? *IDB Working Paper Series*. (Cited on page 29.)

Degroot, M. H. (1970). *Optimal Statistical Decisions*. McGraw-Hill. (Cited on page 46.)

Dillenberger, D. and Rozen, K. (2015). History-dependent risk attitude. *Journal of Economic Theory*, 157:445–477. Publisher: Elsevier Inc. (Cited on page 10.)

Dohmen, T., Lehmann, H., and Pignatti, N. (2016). Time-Varying Individual Risk Attitudes over the Great Recession: A Comparison of Germany and Ukraine. *Journal of Comparative Economics*, 44(1):182–200. (Cited on pages 2, 6, and 7.)

Dollar, D. and Kraay, A. (2002). Growth is Good for the Poor. *Journal of Economic Growth*, 7(3):195–225. (Cited on page 29.)

Eeckhoudt, L., Gollier, C., and Schlesinger, H. (1996). Changes in Background Risk and Risk Taking Behavior. *Econometrica*, 64(3):683–689. (Cited on pages 8, 13, and 48.)

Ehling, P., Graniero, A., and Heyerdahl-Larsen, C. (2018). Asset prices and portfolio choice with learning from experience. *Review of Economic Studies*, 85(3):1752–1780. (Cited on page 7.)

Ehrlich, I. and Becker, G. S. (1972). Market Insurance, Self-Insurance, and Self-Protection. *Journal of Political Economy*, 80(4):623–649. (Cited on page 12.)

Falk, A., Becker, A., Dohmen, T., Enke, B., and Huffman, D. (2018). Global Evidence on Economic Preferences. *Quarterly Journal of Economics*, 133(4):1645–1692. (Cited on page 18.)

Falk, A., Becker, A., Dohmen, T., Enke, B., Huffman, D., and Sunde, U. (2016). The Nature and Predictive Power of Preferences: Global Evidence. *Quartely Journal of Economics*. (Cited on page 26.)

Foster, A. D. and Rosenzweig, M. R. (1995). Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture. *Journal of Political Economy*, 103(6):1176–1209. (Cited on page 14.)

German-Soto, V. (2005). Generación del producto interno bruto Mexicano for entidad federativa, 1940-1992. *Trimestre Economico*, 72(3):617–653. (Cited on page 20.)

Gibbons, R., Katz, L. F., Lemieux, T., and Parent, D. (2005). Comparative advantage, learning, and sectoral wage determination. *Journal of Labor Economics*, 23(4):681–723. (Cited on page 14.)

Giuliano, P. and Spilimbergo, A. (2014). Growing up in a recession. *Review of Economic Studies*, 81(2):787–817. arXiv: 1011.1669v3 ISBN: 0034-6527. (Cited on page 9.)

Glen, J. and Mondragón-Vélez, C. (2011). Business cycle effects on commercial bank loan portfolio performance in developing economies. *Review of Development Finance*, 1(2):150–165. (Cited on page 63.)

Gneezy, U., List, J. A., and Wu, G. (2006). The uncertainty effect: When a risky prospect is valued less than its worst possible outcome. *Quarterly Journal of Economics*, 121(4):1283–1309. (Cited on page 26.)

Gollier, C. and Pratt, J. W. (1996). Risk Vulnerability and the Tempering Effect of Background Risk. *Econometrica*, 64(5):1109–1123. (Cited on pages 3, 8, 9, and 12.)

Griliches, Z. (1977). Estimating the Returns to Schooling: Some Econometric Problems. *Econometrica*, 45(1):1–22. Publisher: [Wiley, Econometric Society]. (Cited on page 24.)

Guiso, L., Sapienza, P., and Zingales, L. (2018). Time Varying Risk Aversion. *Journal of Financial Economics*, 128(3):403–421. (Cited on pages 2, 6, and 7.)

Hanaoka, C., Shigeoka, H., and Watanabe, Y. (2018). Do Risk Preferences Change? Evidence from the Great Eastern Japan Earthquake. *American Economic Journal: Applied Economics*, 10(2)(2):298–330. ISBN: w21400. (Cited on page 7.)

Heaton, J. and Lucas, D. (2000). Portfolio choice in the presence of background risk. *The Economic Journal*, 110(JANUARY):1–26. (Cited on page 10.)

Herreño, J. and Ocampo, S. (2021). The macroeconomic consequences of subsistence self-employment. Technical report. ISBN: 9781316798553. (Cited on page 36.)

Im, T., Cho, W., and Porumbescu, G. (2011). An Empirical Analysis of the Relation Between Social Spending and Economic Growth in Developing Countries and OECD Members. *Asia Pacific Journal of Public Administration*, 33(1):37–55. (Cited on page 63.)

Jakiela, P. and Ozier, O. (2019). The Impact of Violence on Individual Risk Preferences: Evidence from a natural experiment. *Review of Economics and Statistics*, 101(3). ISBN: 1389-4978\r1573-7780. (Cited on page 7.)

Jordà, Ò., Schularick, M., and Taylor, A. M. (2013). When credit bites back. *Journal of Money, Credit and Banking*, 45(2):3–28. (Cited on page 63.)

Jordà, Ò., Schularick, M., and Taylor, A. M. (2016). Macrofinancial history and the new business cycle facts. *NBER Macroeconomics Annual*, 31(1):213–263. (Cited on page 63.)

Kahn, L. B. (2010). The long-term labor market consequences of graduating from college in a bad economy. *Labour Economics*, 17(2):303–316. Publisher: Elsevier B.V. ISBN: 09275371. (Cited on page 2.)

Kapsos, S. (2006). Labor markets in Asia: Issues and perspectives. In Felipe, J. and Hasan, R., editors, *Labor Markets in Asia: Issues and Perspectives*, pages 143–201. Palgrave Macmillan, London. (Cited on page 29.)

Kaustia, M. and Knüpfer, S. (2008). Do Investors Overweight Personal Experience? Evidence from IPO Subscriptions. *The Journal of Finance*, 63(6):2679–2702. Publisher: [American Finance Association, Wiley]. (Cited on pages 6 and 7.)

Kihlstrom, R. E. and Laffont, J.-J. (1979). A General Equilibrium Entrepreneurial Theory of Firm Formation Based on Risk Aversion. *Journal of Political Economy*, 87(4):719–748. (Cited on page 12.)

Kihlstrom, R. E., Romer, D., and Williams, S. (1981). Risk Aversion with Random Initial Wealth. *Econometrica*, 49(4):911–920. (Cited on page 11.)

Kimball, M. (1990). Precautionary Saving in the Small and in the Large. *Econometrica*, 58(1):53–73. (Cited on page 13.)

Kimball, M. S. (1993). Standard Risk Aversion. *Econometrica*, 61(3):589–611. (Cited on page 8.)

Kimball, M. S., Sahm, C. R., and Shapiro, M. D. (2008). Imputing Risk Tolerance From Survey Responses. *Journal of the American Statistical Association*. (Cited on page 24.)

Koren, M. and Tenreyro, S. (2007). Volatility and Development. *Quarterly Journal of Economics*, 122(1):243–287. (Cited on page 63.)

Malmendier, U. and Nagel, S. (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking? *The Quarterly Journal of Economics*, 126(1):373–416. ISBN: 00335533. (Cited on pages 2, 6, 7, 9, 21, 31, 39, and 69.)

Malmendier, U. and Nagel, S. (2016). Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87. (Cited on page 39.)

Malmendier, U., Pouzo, D., and Vanasco, V. (2020). Investor experiences and financial market dynamics. *Journal of Financial Economics*, 136(3):597–622. arXiv: 1612.09553 Publisher: Elsevier B.V. (Cited on page 7.)

Malmendier, U. and Shen, L. S. (2021). Scarred Consumption. *Working Paper*. (Cited on pages 6, 7, and 33.)

Malmendier, U. and Steiny, A. (2016). Rent or Buy? The Role of Lifetime Experiences of Macroeconomic Shocks within and across Countries. Technical report. (Cited on pages 2, 6, and 7.)

Mccallum, B. T. (1972). Relative Asymptotic Bias from Errors of Omission and Measurement. *Econometrica*, 40(4):757–758. (Cited on page 24.)

Merton, R. C. (1969). Lifetime Portfolio Selection under Uncertainty : The Continuous-Time Case. *The Review of Economics and Statistics*, 51(3):247–257. (Cited on page 12.)

Mu, X., Pomatto, L., Strack, P., and Tamuz, O. (2020). Background risk and small-stakes risk aversion. (1994):1–25. arXiv: 2010.08033. (Cited on page 26.)

Murphy, K. P. (2007). Conjugate Bayesian Analysis of the Gaussian Distribution. Technical report. ISSN: <null>. (Cited on page 46.)

Nachman, D. C. (1982). Preservation of " More Risk Averse " under Expectations. *Journal of Economic Theory*, 368:361–368. (Cited on page 11.)

Pratt, J. W. (1964). Risk Aversion in the Small and in the Large. *Econometrica*, 32(1/2):122–136. (Cited on pages 2, 10, 12, and 14.)

Pratt, J. W. and Zeckhauser, R. J. (1987). Proper Risk Aversion. *Econometrica*, 55(1):143–154. (Cited on page 8.)

Ravallion, M. and Chen, S. (1997). What can new survey data tell us about recent changes in developing countries? *The World Bank Economic Review*, 11(2):357–382. (Cited on page 29.)

Ravallion, M., Squire, L., and Bruno, M. (1999). *Equity and Growth in Developing Countries: Old and New Perspectives on the Policy Issues*. The World Bank. (Cited on page 29.)

Ross, S. A. (1981). Some Stronger Measures of Risk Aversion in the Small and the Large with Applications. *Econometrica*, 49(3):621–638. (Cited on pages 3, 8, 9, and 13.)

Rotter, J. B. (1966). Generalized Expectancied for Internal Versus External Control of Reinforcement. *Psychological Monographs: General and Applied*, 80(1). arXiv: 1011.1669v3 ISBN: 0096-9753. (Cited on page 9.)

Ruhm, C. J. (2000). Are Recessions Good for Your Health? *The Quarterly Journal of Economics*, 115(2):617–650. Publisher: Oxford University Press. (Cited on page 2.)

Sahm, C. R. (2012). How Much Does Risk Tolerance Change? *Quarterly Journal of Finance*, 2(4). (Cited on pages 2, 6, and 7.)

Schraeder, S. (2016). Information processing and non-Bayesian learning in financial markets. *Review of Finance*, 20(2):823–853. (Cited on page 7.)

Shigeoka, H. (2019). Long-Term Consequences of Growing up in a Recession. Technical report. (Cited on pages 2, 6, and 7.)

Stolz, D. S., Müller-Pinzler, L., Krach, S., and Paulus, F. M. (2020). Internal control beliefs shape positive affect and associated neural dynamics during outcome valuation. *Nature Communications*, 11(1):1–13. Publisher: Springer US. (Cited on page 9.)

Strauss, J., Beegle, K., Sikoki, B., Dwiyanto, A., Herawati, Y., and Witoelar, F. (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. (Cited on page 17.)

Strauss, J., Witoelar, F., and Bondan, S. (2016). The Fifth Wave of the Indonesia Family Life Survey: Overview and Field Report. Technical report. Volume: 1 Issue: March. (Cited on page 17.)

Tserenjigmid, G. (2021). History-Dependent Risk Aversion, the Reinforcement Effect, and Dynamic Monotonicity. Technical report. (Cited on page 10.)

World Health Organisation (2019). WHO global report on trends in prevalence of tobacco use 2000-2025, third edition. Technical report. Publication Title: World Health Organisation ISBN: 9789240000032. (Cited on page 36.)

Yagan, D. (2019). Employment hysteresis from the great recession. *Journal of Political Economy*, 127(5). (Cited on page 2.)

Yi, W., Mei, S., Li, Q., Liu, X., and Zheng, Y. (2018). How choice influences risk processing: An ERP study. *Biological Psychology*, 138(9):223–230. (Cited on page 9.)

Online Appendix

A Proof of Proposition 1

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

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

Proof. By assumption, the agent's prior mean M and variance Σ^2 are the relevant parameters of the $NI\chi^{-2}(\mu_0, \kappa_0, \sigma_0^2, \nu_0)$ prior distribution: $M_0 = \mu_0$ and $\Sigma_0^2 = \sigma_0^2$. Given mean-variance normal learning, since the agent's prior is conjugate, Degroot (1970) shows[57] that after observing data set $\mathcal{D}_t$ containing k observations, the agent's posterior will be a $NI\chi^{-2}(\mu_t, \kappa_t, \sigma_t^2, \nu_t)$ where

$$\mu_t = \mu_0 + \frac{k}{\kappa_0 + k}(\bar{y}_t - \mu_0) = M|\mathcal{D}_t = M_t$$

$$\kappa_t = \kappa_0 + k$$

$$\nu_t = \nu_0 + k$$

$$\sigma_t^2 = \frac{1}{\nu_t}\left[\nu_0\sigma_0^2 + ks_t^2 + \frac{k\kappa_0}{\kappa_0 + k}(\bar{y}_t - \mu_0)^2\right] = \Sigma^2|\mathcal{D}_t = \Sigma_t^2.$$

Consider the change in the agent's posterior mean from period 1 to period 2, with $\mathcal{D}_1$ containing m observations, and $\mathcal{D}_2$ containing $m+n$ observations:

$$M_2 - M_1 = \mu_0 + \frac{m+n}{\kappa_0 + m + n}(\bar{y}_2 - \mu_0) - \mu_0 - \frac{m}{\kappa_0 + m}(\bar{y}_1 - \mu_0)$$

Add and subtract $\frac{m+n}{\kappa_0+m+n}\bar{y}_1$ and collect terms to yield:

$$= \frac{m+n}{\kappa_0 + m + n}(\bar{y}_2 - \bar{y}_1) + \frac{n\kappa_0}{(\kappa_0 + m + n)(\kappa_0 + m)}(\bar{y}_1 - \mu_0).$$

[57] Degroot (1970) [pg.169] proves these facts 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 second term above, which is a function for the agent's prior, is $O(m^{-2})$, whereas the first term is constant in m. Therefore, for large m:

$$M_2 - M_1 \approx \bar{y}_2 - \bar{y}_1.$$

Consider next the change in the agent's posterior variance from period 1 to period 2:

$$\begin{aligned}\Sigma_2^2 - \Sigma_1^2 =& \frac{1}{\nu_0 + m + n}\left[\nu_0\sigma_0^2 + (m+n)s_2^2 + \frac{(m+n)\kappa_0}{\kappa_0 + m + n}(\bar{y}_2 - \mu_0)^2\right] \\ &- \frac{1}{\nu_0 + m}\left[\nu_0\sigma_0^2 + ms_1^2 + \frac{m\kappa_0}{\kappa_0 + m}(\bar{y}_1 - \mu_0)^2\right]\end{aligned}$$

Add and subtract $\frac{m+n}{\nu_0+m+n}s_1^2$ and collect terms to yield:

$$\begin{aligned}=& \frac{m+n}{\nu_0 + m + n}(s_2^2 - s_1^2) + \frac{n\nu_0}{(\nu_0 + m + n)(\nu_0 + m)}s_1^2 - \frac{n\nu_0}{(\nu_0 + m + n)(\nu_0 + m)}\sigma_0^2 \\ &+ \frac{(m+n)\kappa_0}{(\nu_0 + m + n)(\kappa_0 + m + n)}(\bar{y}_2 - \mu_0)^2 - \frac{m\kappa_0}{(\nu_0 + m)(\kappa_0 + m)}(\bar{y}_1 - \mu_0)^2\end{aligned}$$

Note that the first term is constant in m, the second and third are $O(m^{-2})$, and the final two are $O(m^{-1})$. Therefore, for large m:

$$\Sigma_2^2 - \Sigma_1^2 \approx s_2^2 - s_1^2.$$

The relevant beliefs distribution over the background risk for the agent at time t is the posterior predictive $p(y|\mathcal{D}_t)$, the agent's best forecast of the realization of the background risk given their previous body of experiences, and given the uncertainty about M and Σ_2. For a normal mean-variance Bayesian learner with a conjugate prior, the posterior predictive is a location-scale student's-t distribution with ν_t degrees of freedom: $p(y|\mathcal{D}_t) = t_{\nu_t}(M_t, \frac{(1+\kappa_t)}{\kappa_t}\Sigma_t^2)$. The mean and variance of this distribution are $E[y|\mathcal{D}_t] = M_t$ and $Var(y|\mathcal{D}_t) = \frac{1+\kappa_t}{\kappa_t}\frac{\nu_t}{\nu_t - 2}\Sigma_t^2$. Consider the changes in the mean and variance of the agent's posterior predictive from period 1 to period 2. For the mean, this is $M_2 - M_1 \approx \bar{y}_2 - \bar{y}_1$ by above. For the variance, it is

$$Var(y|\mathcal{D}_2) - Var(y|\mathcal{D}_1) = \frac{(1 + \kappa_0 + m + n)(\nu_0 + m + n)}{(\kappa_0 + m + n)(\nu_0 + m + n - 2)}\Sigma_2^2 - \frac{(1 + \kappa_0 + m)(\nu_0 + m)}{(\kappa_0 + m)(\nu_0 + m - 2)}\Sigma_1^2.$$

Add and subtract $\frac{(1+\kappa_0+m+n)(\nu_0+m+n)}{(\kappa_0+m+n)(\nu_0+m+n-2)}\Sigma_1^2$ to get

$$\begin{aligned}&\frac{(1+\kappa_0+m+n)(\nu_0+m+n)}{(\kappa_0+m+n)(\nu_0+m+n-2)}(\Sigma_2^2-\Sigma_1^2)\\&+\bigg(\frac{(1+\kappa_0+m+n)(\nu_0+m+n)(\kappa_0+m)(\nu_0+m-2)}{(\kappa_0+m+n)(\nu_0+m+n-2)(\kappa_0+m)(\nu_0+m-2)}\\&-\frac{(1+\kappa_0+m)(\nu_0+m)(\kappa_0+m+n)(\nu_0+m+n-2)}{(\kappa_0+m+n)(\nu_0+m+n-2)(\kappa_0+m)(\nu_0+m-2)}\bigg)\Sigma_1^2.\end{aligned}$$

Note that the first term is constant in m, while the second term in parentheses is $O(m^{-1})$. Therefore, for large m, it follows that[58]

$$Var(y|\mathcal{D}_2)-Var(y|\mathcal{D}_1)\approx\Sigma_2^2-\Sigma_1^2\approx s_2^2-s_1^2.$$

Eeckhoudt et al. (1996) show that under the assumption of Ross Risk Vulnerability, both first-order and second-order stochastic dominance deteriorations in an objective background risk will result in an increase in the agent's risk aversion over the foreground risk. This result can be trivially extended to a subjective background risk, with changes in agent beliefs about the background risk taking the place of actual changes in an objective distribution. Since an increase in the mean is a first-order stochastic dominance improvement, this entails that an increase in the observed sample mean (considered separately from changes in the sample variance) would result in a decrease in the agent's foreground risk aversion. In other words, for $A>0$:

$$r_2(w)-r_1(w)\approx-A(\bar{y}_2-\bar{y}_1).$$

Likewise, because a mean-preserving spread is a second-order stochastic dominance deterioration, an increase in the observed sample variance (considered separately from changes in the sample mean) would result in an increase in the agent's foreground risk aversion. In other words, for $B>0$:

$$r_2(w)-r_1(w)\approx B(s_2^2-s_1^2).$$

It remains to be shown that these two effects are additive when considered jointly. This follows from our assumption of sequential updating over the moments. To see this, consider the agent's risk aversion parameter in the *ex-interim* period after updating over the variance

[58] Alternatively, one can prove this statement by noting that as m gets large $\nu_t\to\infty$ and $\frac{1+\kappa_t}{\kappa_t}\to 1$, so $t_{\nu_t}(M_t,\frac{(1+\kappa_t)}{\kappa_t}\Sigma_t^2)\xrightarrow{d}\mathcal{N}(M_t,\Sigma_t^2)$.

of the background risk, but before updating over the mean. Label this parameter $r_I(w)$. It follows that

$$r_I(w) - r_1(w) \approx B(s_2^2 - s_1^2)$$

while

$$r_2(w) - r_I(w) \approx -A(\bar{y}_2 - \bar{y}_1).$$

Adding the two equations together completes the proof. $\square$

B Wording of Risk Aversion Questions

Figure B.1: Text of risk aversion questions in IFLS4 and IFLS5

SECTION SI:RISK AND TIME PREFERENCES

RANDOM_SI : A

SI01.	Suppose you are offered two ways to earn some money. With option 1, you are guaranteed Rp 800 thousand per month. With option 2, you have an equal chance of ither the same income, Rp 800 thousand per month, or, if you are lucky, Rp 1.6 million. per month, which is more. Which option will you choose?	1. Rp 800 thousand per month 2. Rp 1.6 million or Rp 800 thousand per month→**SI03** 8. DON'T KNOW
SI02.	Are you sure? In option 2 you will get at least Rp 800 thousand per month and you may get Rp 1.6 million per month. In option 1 you will always get Rp 800 thousand per month.	1. Still picks option 1 →**SI11** 2. Switches to option 2 8. DON'T KNOW
SI03.	Now, in option 2 you have an equal chance of receiving either Rp. 1.6 million per month or Rp. 400 thousand per month, depending on how lucky you are. Option 1 guarantees you an income of Rp 800 thousand per month. Which option will you choose?	1. Rp 800 thousand 2. Rp 1.6 million or Rp 400 thousand→**SI05** 8. DON'T KNOW
SI04.	Now, in option 2 you have an equal chance of receiving either Rp 1.6 million per month or Rp 600 thousand per month, depending on how lucky you are. Option 1 guarantees you an income of Rp 800 thousand per month. Which option will you choose?	1. Rp 800 thousand 2. Rp 1.6 million or Rp 600 thousand 8. DON'T KNOW →**SI11**
SI05.	Now, in option 2 you have an equal chance of receiving either Rp 1.6 million per month or Rp 200 thousand per month, depending on how lucky you are. Option 1 guarantees you an income of Rp 800 thousand per month. Which option will you choose?	1. Rp 800 thousand 2. Rp 1.6 million or Rp 200 thousand 8. DON'T KNOW →**SI11**

Notes: Text of risk questions taken directly from questionnaire deployed in IFLS4 and IFLS5. Original survey done in Indonesian; English translation provided with data. Hypothetical lottery values are in Indonesian rupiah.

Figure B.2: Text of risk aversion questions in MxFLS2

RISK (SECCIÓN RG)

Now imagine a game of random chance. In a bag there is a blue chip and a yellow chip and an amount of money is written on each of them. (INTERVIEWER: SHOW THE SLIDES). If you stick your hand inside the bag and take out the yellow chip, we would pay you what is written on the yellow chip, if you take out the blue chip, we will pay what is written on the blue chip. Now you reach inside the bag, but you do not know yet what chip you will get.

RG01.	Before we continue, what color chip do you have the highest probability of getting? **(INTERVIEWER: After writing down the answer, explain the correct answer)** 1. Blue 2. Yellow 3. Same probability 8. DK	1 2 3 8
RG02.	**(INTERVIEWER: show slide RG02, indicate and read the quantities for each game of chance)** Now imagine you can choose between the two bags shown on the slide: 1. In bag 1, if you get the blue chip **or** the yellow chip, you receive $1,000 2. In bag 2, if you get the blue chip you receive $500 or $2,000 if you get the yellow chip Which one of the bags do you choose? 8. DK	1 → **RG05** 2 8 → **RG05**
RG03.	**(INTERVIEWER: show slide RG03, indicate and read the quantities for each game of chance)** Now imagine you can choose between the two bags shown on the slide: 1. In bag 1, if you get the blue chip you receive $500 or $2,000 if you get the yellow chip 2. In bag 2, if you get the blue chip you receive $300 or $3,000 if you get the yellow chip Which one of the bags do you choose? 8. DK	1 → **RG05** 2 8 → **RG05**
RG04.	**(INTERVIEWER: show slide RG04, indicate and read the quantities for each game of chance)** Now imagine you can choose between the two bags shown on the slide: 1. In bag 1, if you get the blue chip you receive $100 or $4,000 if you take out the yellow chip 2. In bag 2, if you get the blue chip you receive $100 or $7,000 if you take out the yellow chip Which one of the bags do you choose? 8. DK	1 → **RG08** 2 → **RG08** 8
RG05.	**(INTERVIEWER: show slide RG05, indicate and read the quantities for each game of chance)** Now imagine you can choose between the two bags shown on the slide: 1. In bag 1, if you get the blue chip you receive $1,000 or $1,000 if you get the yellow chip 2. In bag 2, if you get the blue chip you receive $800 or $2,000 if you get the yellow chip Which one of the bags do you choose? 8. DK	1 2 → **RG08** 8 → **RG08**
RG06.	**(INTERVIEWER: show slide RG06, indicate and read the quantities for each game of chance)** Now imagine you can choose between the two bags shown on the slide: 1. In bag 1, if you get the blue chip you receive $1,000 or $1,000 if you get the yellow chip 2. In bag 2, if you get the blue chip you receive $800 or $4,000 if you get the yellow chip Which one of the bags do you choose? 8. DK	1 2 → **RG08** 8
RG07.	**(INTERVIEWER: show slide RG07 and read the quantities for each game of chance)** Now imagine you can choose between the two bags shown on the slide: 1. In bag 1, if you get the blue chip you receive $1,000 or $1,000 if you get the yellow chip 2. In bag 2, if you get the blue chip you receive $800 or $8,000 if you get the yellow chip Which one do you choose? 8. DK	1 2 8

Notes: Text of risk questions taken directly from questionnaire deployed in MxFLS2. Original survey done in Spanish; English translation provided with data. Hypothetical lottery values are in Mexican pesos. Please note that there is a typo in option 2 in question RG04, which should say "In Bag 2, if you get the blue chip you lose $100, or get $7000 if you take out the yellow chip." This is evidenced when the English translation is compared to the Spanish questionnaire, and is also noted by Brown et al. (2019).

Figure B.3: Text of risk aversion questions in MxFLS3

RISK (SECTION RG) **MxFLS 2009**

RG01.	**(INTERVIEWER: SHOW HIM/HER SLIDE RG01)** Imagine you can choose between two bags. Once you have chosen one of the bags, you will put your hand inside the bag and without looking you will pick a ball which will show the amount of money you have won. Bag No. 1 has a ball that is worth $2,500. Bag No. 2 has two balls: one is worth $2,500 (same as Bag No. 1) and the other ball is worth $5,000. Which one of the two bags do you choose? **(INTERVIEWER: SHOW AND READ THE AMOUNT WRITTEN ON EACH BALL)**	1. $2,500 or $2,500 2. $2,500 or $5,000→ **RG03** 8. DK
RG02.	**(INTERVIEWER: SHOW HIM/HER SLIDE RG02)** Are you sure? You are going to pick only one ball from the bag you choose. Things would not change if we put another ball that is worth $2,500 into Bag No. 1. Now Bag No. 1 has two balls worth $2,500 each, as shown in the image. If you choose Bag No. 1 you will win $2,500. If you choose Bag No. 2, you will win at least $2,500 and probably, you will win $5,000, depending on your luck. Which one of the two bags do you choose? **(INTERVIEWER: SHOW AND READ THE AMOUNT WRITTEN ON EACH BALL).**	1. He/She still chooses Bag No. 1 → **RG07** 2. He/She changes to Bag No. 2 8. DK
RG03.	**(INTERVIEWER: SHOW HIM/HER SLIDE RG03)** Now, imagine you can choose between the following two bags: Bag No. 1 guarantees that you will win $2,500. Bag No. 2 has a ball that is worth $2,000 and another ball that is worth $5,000. Which one of the two bags do you choose? **(INTERVIEWER: SHOW AND READ THE AMOUNT WRITTEN ON EACH BALL)**	1. $2,500→ **RG08** 2. $5,000 or $2,000 8. DK
RG04.	**(INTERVIEWER: SHOW HIM/HER SLIDE RG04)** And if now you could choose between: Bag No. 1, which again guarantees that you will win $2,500; or Bag No. 2 which has a ball that is worth $1,500 and another ball that is worth $5,000. Which one of the two bags do you choose? **(INTERVIEWER: SHOW AND READ THE AMOUNT WRITTEN ON EACH BALL)**	1. $2,500→ **RG08** 2. $5,000 or $1,500 8. DK
RG05.	**(INTERVIEWER: SHOW HIM/HER SLIDE RG05)** Now, suppose you can choose between: Bag No. 1, which guarantees $2,500; or Bag No. 2, which now has a ball that is worth $1,000 and another ball that is worth $5,000. Which one of the two bags do you choose? **(INTERVIEWER: SHOW AND READ THE AMOUNT WRITTEN ON EACH BALL)**	1. $2,500→ **RG08** 2. $5,000 or $1,000 8. DK
RG06.	**(INTERVIEWER: SHOW HIM/HER SLIDE RG06)** Now, suppose you can choose between: Bag No. 1, which guarantees $2,500; or Bag No. 2, which has a ball that is worth $500 and another ball that is worth $5,000. Which one of the two bags do you choose? **(INTERVIEWER: SHOW AND READ THE AMOUNT WRITTEN ON EACH BALL)**	1. $2,500→ **RG08** 2. $5,000 or $500→ **RG08** 8. DK → **RG08**
RG07.	**(INTERVIEWER: SHOW HIM/HER SLIDE RG07)** Finally, suppose you can choose between: Bag No. 1, which guarantees $2,000; or Bag No. 2, which has a ball that is worth $5,000 and another ball that is worth $2,500. Which one of the two bags do you choose? **(INTERVIEWER: SHOW AND READ THE AMOUNT WRITTEN ON EACH BALL)**	1. $2,000 2. $5,000 or $2,500 8. DK

Notes: Text of risk questions taken directly from questionnaire deployed in MxFLS3. Original survey done in Spanish; English translation provided with data. Hypothetical lottery values are in Mexican pesos.

C Construction of Risk Aversion Measures

Figure C.1: Construction of risk aversion measure in IFLS4 and IFLS5

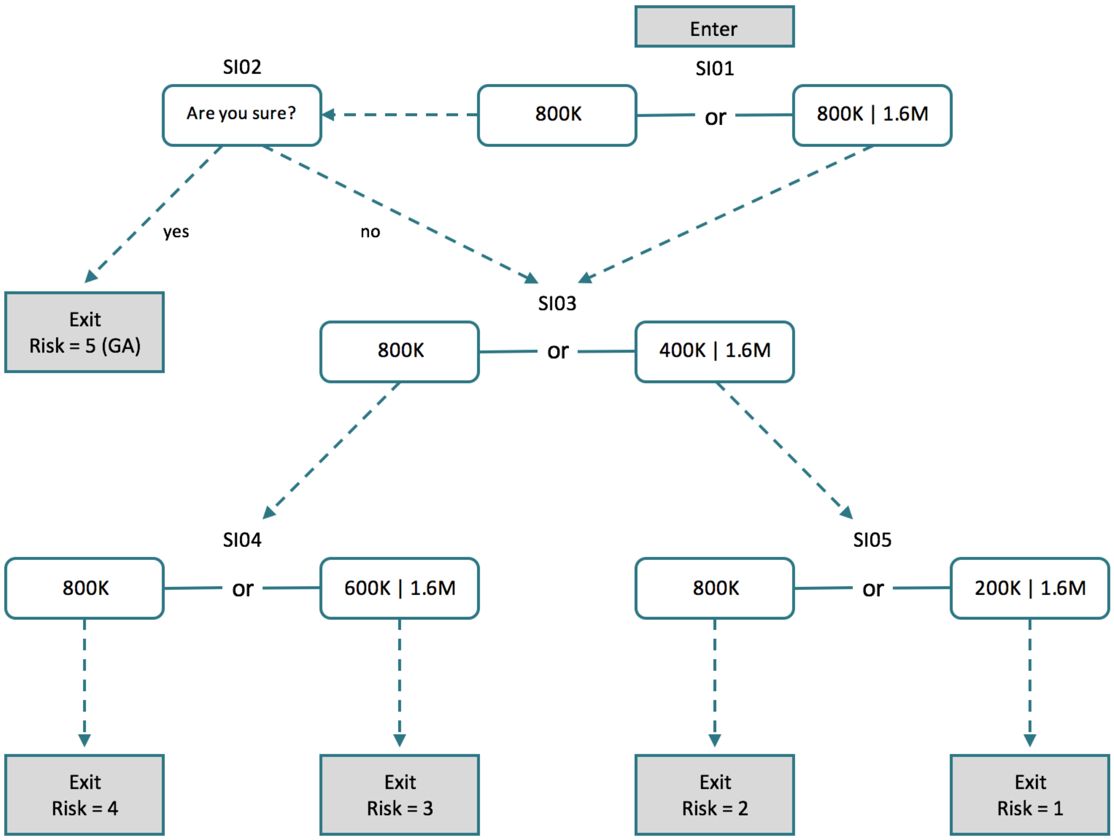

Notes: Hypothetical lottery values are in Indonesian rupiah. All lotteries offer a 50% probability of winning each of the two prizes. Higher values for "Risk" indicate a higher level of measured risk aversion (Risk = 1 indicates the most risk-seeking choice).

Figure C.2: Construction of risk aversion measure in MxFLS2

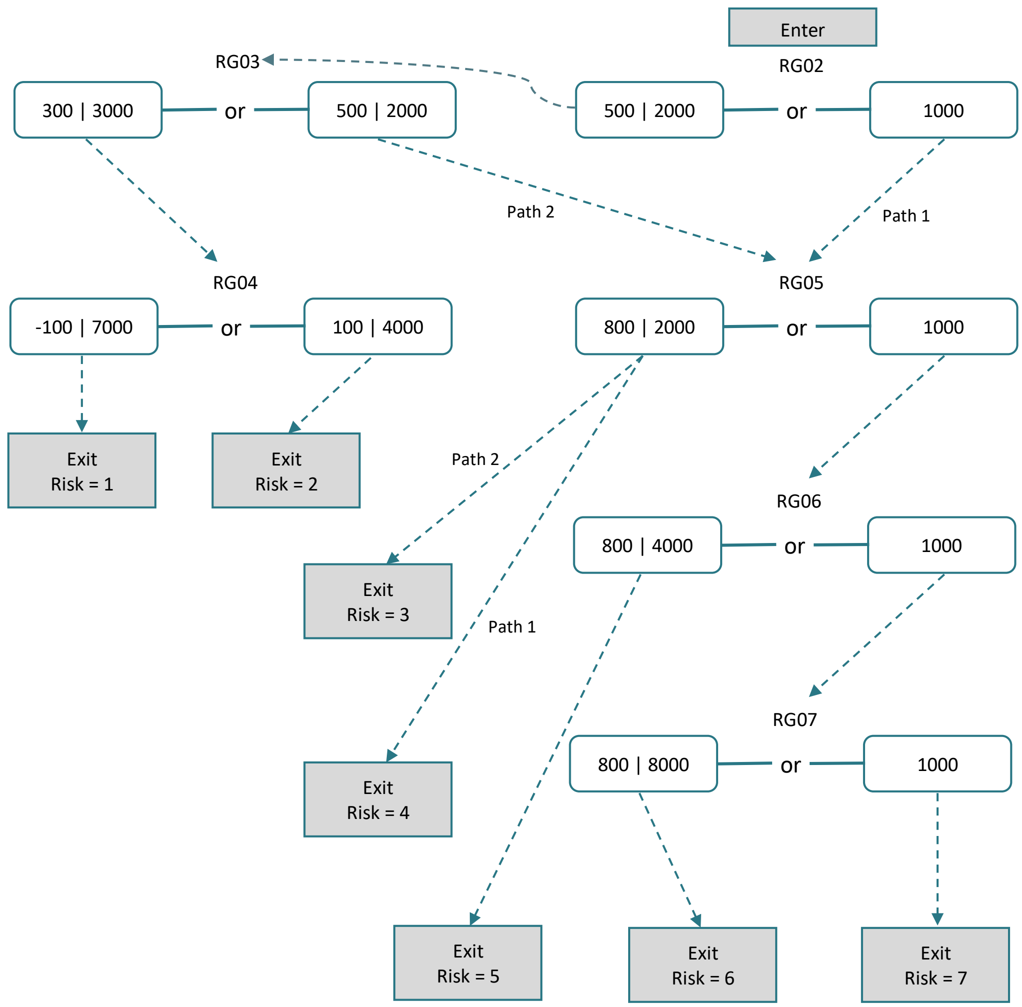

Notes: Hypothetical lottery values are in Mexican pesos. All lotteries offer a 50% probability of winning each of the two prizes. The safe option in each stage is depicted on the right side of the stage. Higher values for "Risk" indicate a higher level of measured risk aversion (Risk = 1 indicates the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

Figure C.3: Construction of risk aversion measure in MxFLS3

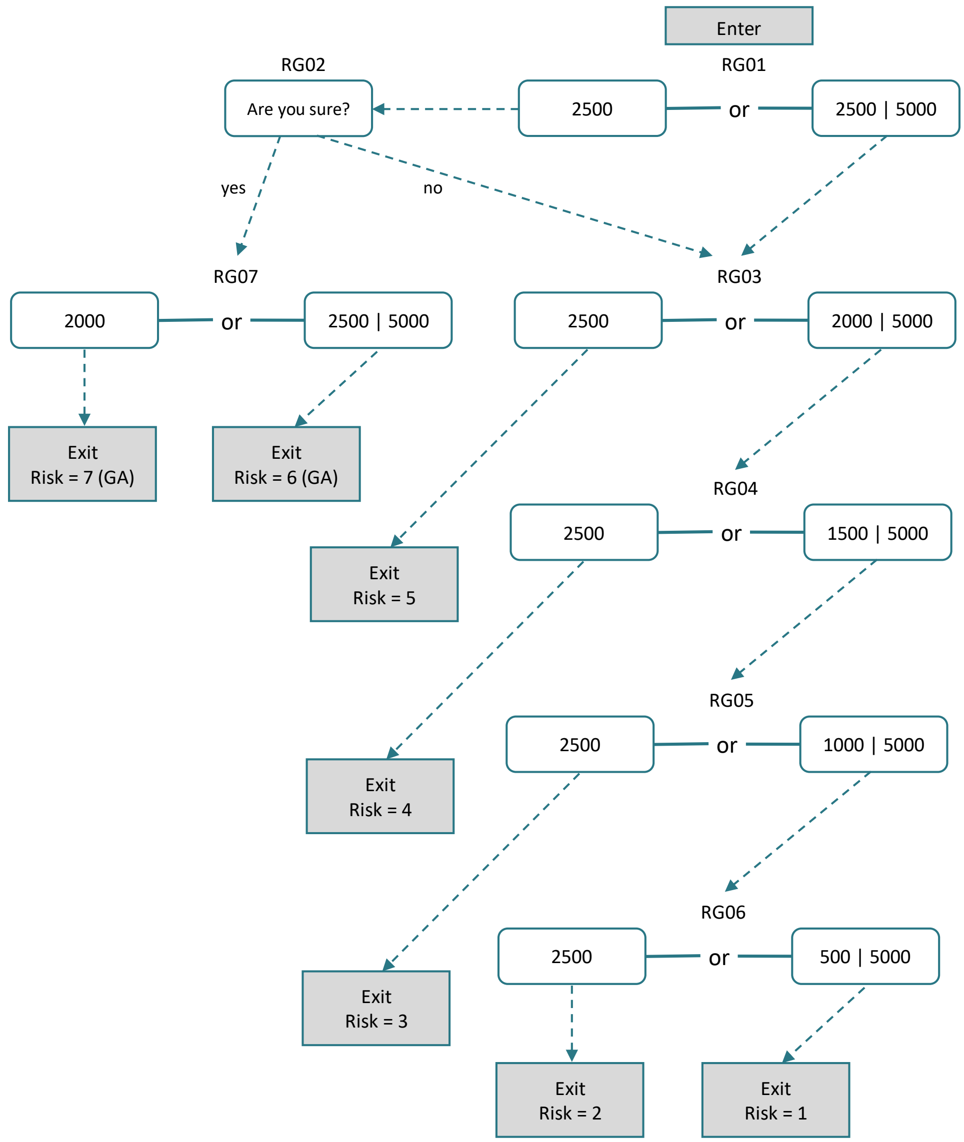

Notes: Hypothetical lottery values are in Mexican pesos. All lotteries offer a 50% probability of winning each of the two prizes. The safe option in each stage is depicted on the left side of the stage. Higher values for "Risk" indicate a higher level of measured risk aversion (Risk = 1 indicates the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

D Distribution of Measured Risk Aversion and Summary Statistics

Figure D.1: Histograms of measured risk aversion buckets across IFLS4 and IFLS5, and across MxFLS4 and MxFLS5

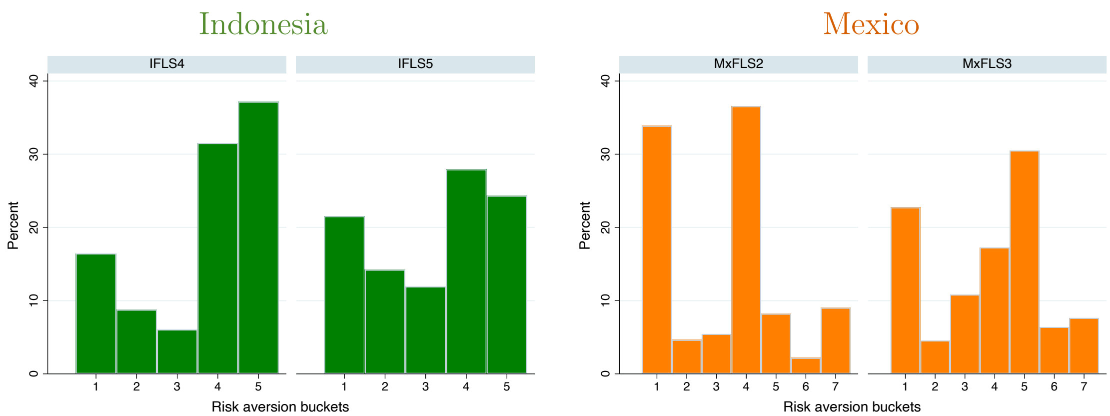

Notes: *Measured Risk Aversion*: 1–5 (Indonesia), 1–7 (Mexico), with 5 or 7 being the highest measured risk aversion. These distributions are for subjects in the primary sample, described in Section 3.D.

Table D.1: Summary statistics

	Indonesia					Mexico				
Sample:	Primary sample		Full sample*		Significance of	Primary sample		Full sample*		Significance of
	Mean	St.Dev	Mean	St.Dev	mean diffs.	Mean	St.Dev	Mean	St.Dev	mean diffs.
Measured risk aversion	3.42	1.49	3.56	1.5	***	3.28	1.78	3.51	1.92	***
Woman	0.56	0.50	0.56	0.50	***	0.59	0.49	0.59	0.49	
Age	34.27	9.35	40.05	13.80	***	42.04	16.13	42.9	16.68	
Married	0.85	0.35	0.89	0.32	***	0.66	0.47	0.66	0.47	***
Household size	4.96	3.00	5.18	3.06	***	5.64	2.7	5.65	2.72	
Comp. elementary	0.32	0.47	0.41	0.49	***	0.51	0.5	0.52	0.5	
Comp. middle school	0.21	0.41	0.19	0.39	***	0.25	0.43	0.25	0.43	
Comp. high school	0.33	0.47	0.28	0.45	**	0.13	0.34	0.13	0.33	
Above high school	0.15	0.35	0.13	0.33		0.11	0.32	0.11	0.31	**
Ever migrated	0.13	0.34	0.14	0.35	***	0.15	0.36	0.15	0.36	
Income	108.5	5590	112.2	5665		52,404	272,835	51,098	252,705	
Assets	132.80	236.2	139	246.1	***	281,288	1.21e+06	275,923	1.17e+06	
Net yearly savings	3.96	35.98	3.55	33.9		-2,014	65,142	-2,165	61,981	
Observations	23,272		34,366			20,448		24,304		
Individuals	11,636		17,183			10,224		12,152		

Notes: Income, assets, and net yearly savings are at the household level and inflation-adjusted to local currency in the first wave of the survey (millions of rupiah of 2007 in Indonesia and pesos of 2005 in Mexico). The full sample comprises individuals who have risk information in both waves of the survey. The primary sample is described in Section 3.D. For the significance of mean difference column: *** p<0.01, ** p<0.05, * p<0.1.

E Geographic Distribution of Survey Samples in our Data

Figure E.1: Distribution of the primary sample in Indonesia by state of birth

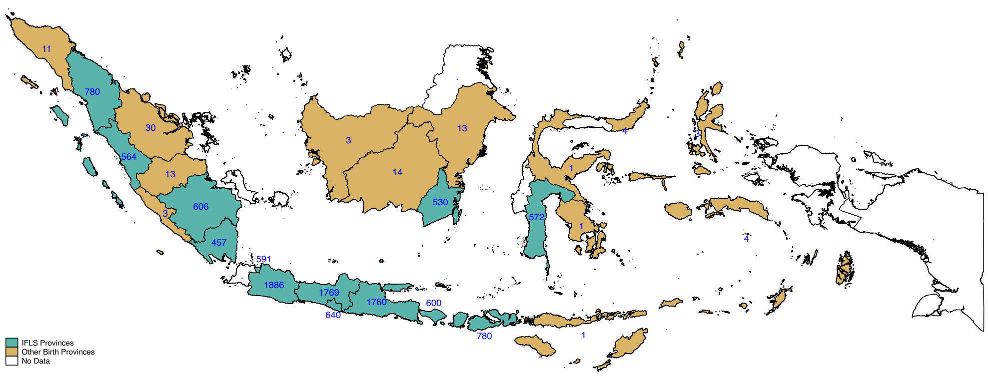

Notes: States in blue are those in which the IFLS has been deployed. States in brown are non-IFLS states in which some subjects in our primary sample were born. The primary sample is described in Section 3.D.

Figure E.2: Distribution of the primary sample in Mexico by state of birth

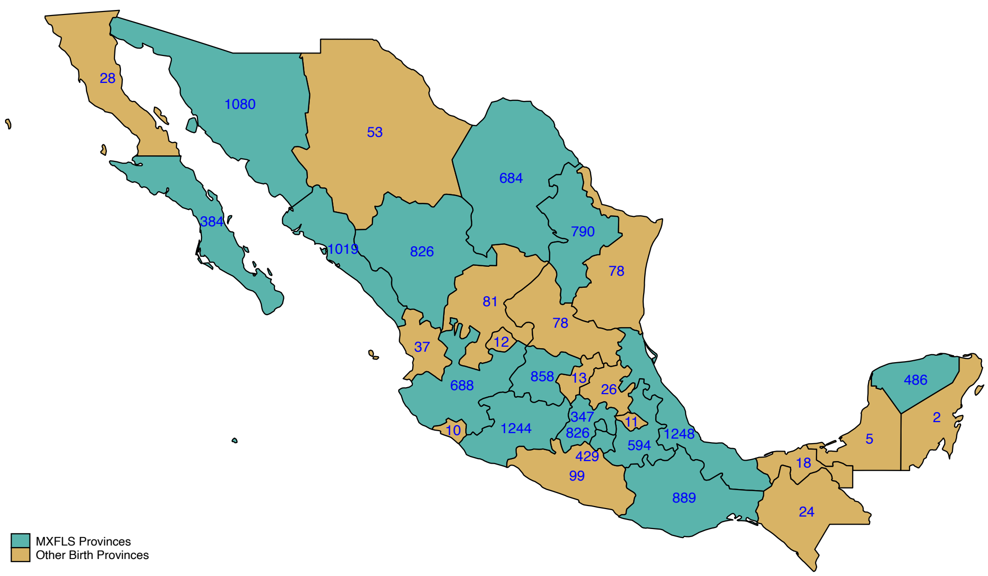

Notes: States in blue are those in which the MxFLS has been deployed. States in brown are non-MxFLS states in which some subjects in our primary sample were born. The primary sample is described in Section 3.D.

F Correlates of Risk Aversion Measures in the Cross-Section

Table F.1: Correlates of risk aversion measures

	Indonesia	Mexico
Dep. Var:	Measured Risk Aversion	Measured Risk Aversion
Self-employed	-0.11***	0.03
	(0.03)	(0.04)
Ever migrated	-0.12***	0.0003
	(0.03)	(0.04)
Income	1.75e-06***	6.19e-08**
	(4.67e-07)	(2.47e-08)
Assets	-0.0002***	2.33e-08***
	(0.0001)	(6.18e-09)
Yearly savings	0.0003	1.72e-07
	(0.0003)	(3.08e-07)
Yearly borrowing	-0.002***	-2.75e-07
	(0.001)	(2.68e-07)
Consumption	-0.00001	-2.44e-07
	(0.00004)	(3.09e-07)
Currently smoke	0.03	-0.20***
	(0.04)	(0.06)
Cigarettes/day	-0.001	0.002*
	(0.003)	(0.001)
Woman	0.31***	0.02
	(0.03)	(0.03)
Age	-0.02*	-0.02***
	(0.01)	(0.01)
Age^2	0.0003*	0.0002***
	(0.0002)	(0.0001)
Observations	17,158	10,608

Notes: Coefficients from regressions of dependent variables on all covariates. Standard errors clustered at the birth-year by state of birth in parentheses. Observations are at the individual by year level. Controls: Time fixed effects, state fixed effects, household size, marital status, education dummies, and religiosity dummies (the last only for Indonesia). Monetary variables are at the household level and inflation-adjusted to local currency in the first wave of the survey (millions of rupiah of 2007 in Indonesia and pesos of 2005 in Mexico). These results are for subjects in the primary sample, described in Section 3.D. *** p<0.01, ** p<0.05, * p<0.1.

G Information on Gamble Averse

Table G.1: Transition matrix in risk aversion measure between two waves of IFLS

	Risk Aversion Bucket in IFLS5					
Risk Aversion Bucket in IFLS4	1	2	3	4	5(GA)	Total
1	26.04%	11.75%	9.44%	23.62%	29.15%	100%
2	23.44%	13.90%	10.77%	25.69%	26.20%	100%
3	20.91%	16.06%	14.33%	25.43%	23.28%	100%
4	18.66%	12.54%	11.62%	28.59%	28.59%	100%
5(GA)	18.22%	12.31%	9.99%	24.24%	35.23%	100%
Total	20.11%	12.62%	10.68%	25.60%	30.99%	

Notes: This table presents the transition matrix between buckets of the risk aversion measure in our Indonesian analysis for the full sample, which comprises individuals who have risk aversion information in both waves of the IFLS.

Table G.2: Summary statistics by never, once, and twice gamble averse in IFLS

	Never GA	Once GA	Twice GA
Avg. age	39.02	40.86	42.51
Prop. female	0.51	0.59	0.63
Prop. Muslim	0.89	0.9	0.91
Raven's score	5.25	4.76	4.44
Prop. with comp. elementary	0.32	0.44	0.54
Prop. with comp. middle school	0.18	0.19	0.19
Prop with comp. high school	0.32	0.26	0.20
Prop. with above high school	0.17	0.10	0.07
Avg. income/month	12.28	8.10	4.77

Notes: Across both waves in full sample, comprising individuals who have risk aversion information in both waves of the survey. Income at household level in millions of rupiah.

Table G.3: Summary statistics by gamble aversion in MxFLS3

	Not GA	GA
Avg. age	44.19	47.42
Prop. female	0.59	0.57
Raven's score	5.53	4.90
Prop. with comp. elementary	0.5	0.62
Prop. with comp. middle school	0.24	0.22
Prop with comp. high school	0.13	0.10
Prop. with above high school	0.13	0.07
Avg. income/month	3993.2	3782.2

Notes: Gamble averse or not in MxFLS3 for full sample, comprising individuals who have risk aversion information in both waves of the survey. Income at household level in pesos.

H Structural Estimation of Risk Aversion Parameter

We test the robustness of the main results when we retrieve specific value ranges for the risk aversion parameter using a structural estimation approach, as described in Section 4.B. We assume subjects are expected utility maximizers and hold a CRRA utility function (parameterized as $\frac{C^{1-\theta}-1}{1-\theta}$), and find the range of values of the risk aversion parameter for each individual in each survey wave that is consistent with their choices in the baseline binned measure, given their level of household income.[59] In particular, we assume that the gamble choices presented in each survey in the IFLS and MxFLS are evaluated by the individual using the utility function above, where C represents the consumption level that would follow after adding the gamble amounts to (1) per-person household income (i.e. assuming broad bracketing), or (2) zero (i.e. assuming narrow bracketing). Figure H.1 presents a graphical representation of the mapping from the binned risk aversion measures to the estimated risk aversion parameter ranges obtained for subjects in this exercise. Notice that because we rely on ordinal choices to compute these structural parameters, the lower limit of the range estimated for subjects with the lowest reported bucket of risk aversion tends to negative infinity, while the upper limit of the range estimated for subjects with the highest reported bucket of risk aversion tends to positive infinity.

We calculate the magnitude of the within-person change in this risk aversion interval between the two waves of the surveys using two distinct approaches. First, we use a midpoint approach that assigns the midpoint in the interval to each individual. In particular, for individuals who are not in the end buckets, we assign their risk aversion parameter to be the midpoint between the upper and lower limits in each survey. We then compute the average distance between the lower or upper limit and this mid-interval measure among all individuals in these mid-range risk buckets. We then use this average distance to assign a point estimate of risk aversion for individuals in the end buckets. In particular, for individuals in the highest reported bucket of risk aversion, we assign their risk aversion parameter to be the lower bound plus this average distance. Conversely, for individuals in the lowest reported bucket of risk aversion, we assign their risk aversion parameter to be the upper bound minus this average distance. This approach preserves the scale of the risk aversion

[59] Note that for this exercise we do not collapse the 5th, 6th, and 7th measured risk aversion bins in MxFLS2 into one bin as in the baseline analysis, in order to correctly consider the different gamble choices made by individuals. In addition, we set the gamble option yielding a value of -100 in MxFLS2 (see Figure B.2 for details on the risk questions in this survey, and Figure C.2 for details on the construction of the risk measures used.) equal to zero to ensure positive consumption levels consistent with CRRA utility.

parameters estimated and thus allows us to interpret the magnitudes more easily. However, it requires us to make strong assumptions about the range of the risk aversion parameter for subjects in the end buckets.

Figure H.1: Risk aversion parameter range predicted by structural estimation exercise

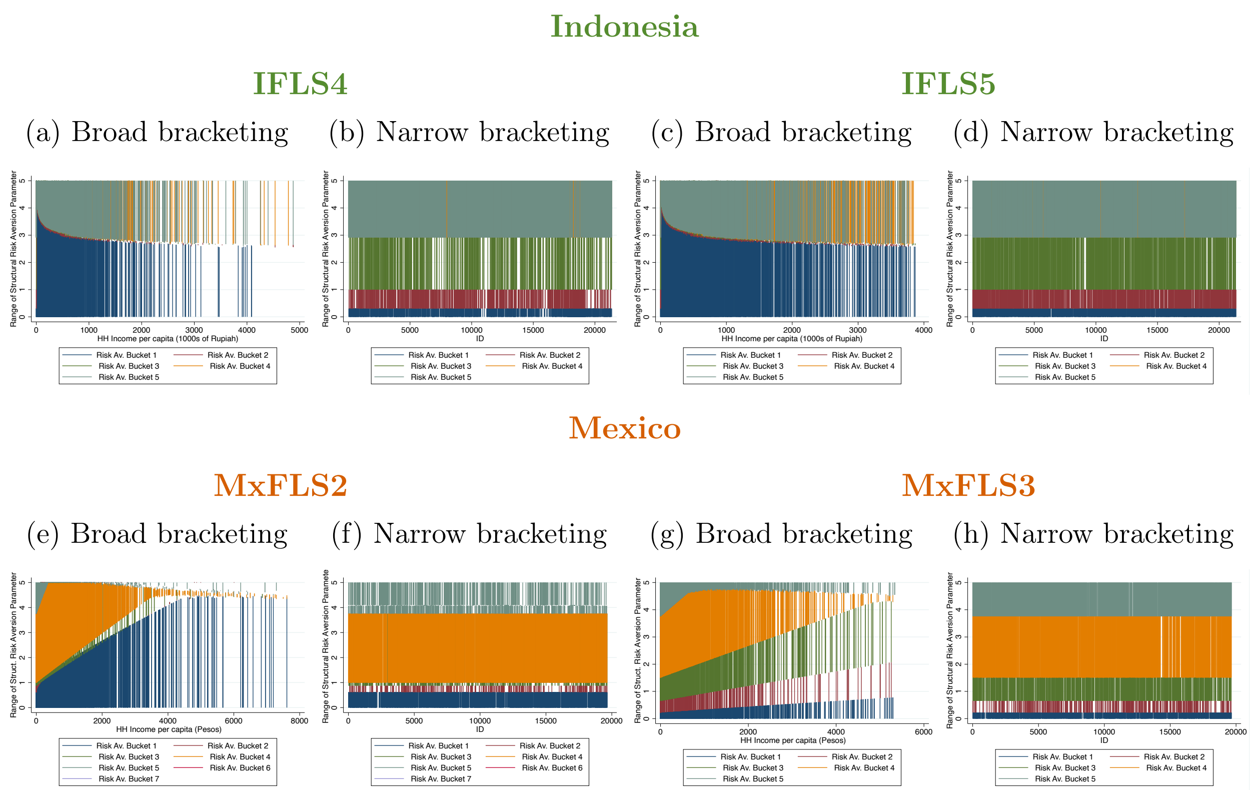

Notes: The y-axes in these figures denote the risk aversion parameter ranges estimated from our structural estimation framework. Risk parameter range values for each task are calculated using a structural estimation approach that imposes a CRRA utility function and finds the values of the risk aversion parameter consistent with utility maximization, the individual's risk aversion bucket, and household income (plotted along the x-axis in (a), (c), (e), and (g)), or zero (plotted along the x-axis in (b), (d), (f), and (h)). Given that the x-variable has no variation under narrow bracketing, in those figures the x-axes present individual ID numbers. Since the lower and upper bounds tend to infinity or negative infinity for some individuals, we have truncated the ranges to fall above 0 and below 5 for display purposes. In addition, we have excluded individuals for whom the income variable is above the 99th percentile. These results are for subjects in the primary sample, described in Section 3.D.

Second, we use a Hausdorff metric (a standard topological measure of distance between sets) to directly calculate the distance between the ranges corresponding to each individual between the two surveys. To deal with the complication of unboundedness for individuals in the end buckets, we attach normalized values at every point using the CDF of a normal distribution whose mean and standard deviation correspond to the mean and standard deviation of the mid-interval measure described above among all individuals in mid-range risk buckets. This approach does not preserve the magnitudes of the original risk aversion parameters, but parsimoniously maintains the unbounded nature of the risk aversion intervals for individuals in the end buckets.

Table H.1: Main results with structurally estimated risk aversion parameter

Dep. Var: Δ Struct. Risk Av.	Mid-Interval Approach		Hausdorff Metric + C.D.F Approach	
	Broad br.	Narrow br.	Broad br.	Narrow br.
Indonesia				
Δ Growth Mean	-0.04*	-0.79***	-0.15***	-0.24***
	(0.02)	(0.13)	(0.03)	(0.04)
Δ Growth Volatility	0.02	0.43***	0.07***	0.13***
	(0.01)	(0.08)	(0.02)	(0.02)
N	11458	11636	11458	11636
Mexico				
Δ Growth Mean	-1.30***	-1.16***	-0.27***	-0.26***
	(0.27)	(0.25)	(0.06)	(0.07)
Δ Growth Volatility	2.17***	1.87***	0.53***	0.59***
	(0.58)	(0.51)	(0.13)	(0.13)
N	9811	10224	9811	10224

Notes: *Struct. Risk Aversion*: Values for the risk aversion parameter calculated using our structural estimation approaches. State (Indonesia) and regional (Mexico) inflation included in all regressions. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

We perform our baseline regression on these two metrics, and present the results in Table H.1. We find that the results for both Indonesia and Mexico remain qualitatively the same, with changes in average experienced GDP growth reducing the structurally estimated risk aversion parameter and changes in the standard deviation of experienced GDP growth raising it. The size of the results also roughly matches up with the baseline analysis, with the effect of volatility being about twice as large as the effect of the mean in Mexico and half as large as the effect of the mean in Indonesia.

I Results for Young Migrants

Table I.1: Limiting to individuals who migrated when young and using state of residence to build macroeconomic experiences

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)
Indonesia			
Δ Growth Mean	-0.37***		-0.76**
	(0.14)		(0.29)
Δ Growth Volatility		-0.14	0.27
		(0.09)	(0.19)
Observations	1197	1197	1197
Mexico			
Δ Growth Mean	-1.17**		-0.82
	(0.57)		(0.59)
Δ Growth Volatility		3.82***	3.44***
		(1.08)	(1.12)
Observations	1025	1025	1025

Notes: *Measured Risk Aversion*: 1–5 (Indonesia and Mexico), with 5 being the highest measured risk aversion. State (Indonesia) and regional (Mexico) inflation included in all regressions. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

J Link between Real GDP Growth and Living Standards

In this appendix and specifically in Table J.1 we report the results of our analysis examining the relationship between state-level real GDP growth and measures of living standards in Indonesia and Mexico (see Section 4.D for discussion of these results).

We use data from the IFLS and MxFLS to examine this relationship at the individual level by regressing the change in measures of living standards between the relevant waves of each survey on average yearly state-level real GDP growth in this period. Given that the MxFLS data has less information on living standards than the IFLS, we further investigate the relevance of GDP on individual's living standards in Mexico using census data. Specifically, we use data from the population and housing censuses of 1990, 2000, and 2010 and the population and dwelling counts of 2005. We aggregate this data at the state by year level and regress the change in living standards from wave to wave of the surveys on the average yearly real GDP growth in this period and some demographic controls. Specifically, we run:

$$\Delta W_{pt} = \alpha + \beta_1 GDPGrowth_{p,t-1:t} + \gamma \Delta X_{pt} + \alpha_t + \alpha_p + \epsilon_{pt}.$$

ΔW_{pt} denotes the change in living standards in each state p between the different waves of data. As a measure of living standards, we use the share of subjects who report living in a dwelling with an earthen floor, a key input into multidimensional poverty indices in the developing world. $GDPGrowth_{p,t-1:t}$ denotes the average annual growth of real GDP in each state between wave $t-1$ and wave t of the survey. We also include additional time-varying controls, which are represented by X_{it}: change in average age, change in proportion of women, and change in population. α_t and α_p denote time and state fixed effects, respectively.

Table J.1: Relationship of changes in state-level real GDP growth to changes in measured living standards

	Δ Log HH Income	Δ Unemployed	Δ Poverty	Δ Hunger	Δ Share w/ Earth Floor
Indonesia					
Average Annual Real GDP Growth	0.07***	0.0003	-0.03***	-0.03***	
	(0.02)	(0.001)	(0.01)	(0.01)	
Observations	9,261	9,430	9,426	9,426	
Mexico					
Average Annual Real GDP Growth	0.002	-0.005***			-.004***
	(0.02)	(0.002)			(.001)
Observations	6,521	9,587			64

Notes: The first four columns in this table are from IFLS and MxFLS data. Average Annual Real GDP Growth is the average yearly state-level growth in the years between IFLS4 and IFLS5, and MxFLS2 and MxFLS3, respectively. The last column is constructed using data from the Mexican population and housing censuses conducted in 1990, 2000, 2005, and 2010, with Average Annual Real GDP Growth as the average yearly state-level growth in the years between the current and last waves. Log HH Income: Log of Income at the household level, inflation-adjusted to local currency in the first wave of the survey (millions of rupiah of 2007 in Indonesia and pesos of 2005 in Mexico). Unemployed: 1 for a individuals without a job and whose main activity in the previous week was searching for one. Poverty: subjective indicator for less-than-adequate living standard. Hunger: subjective indicator for less-than-adequate food consumption. Share w/earth floor: share of subjects in state reporting that dwelling has earthen floor, by state. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

K Heterogeneity of Effects by Employment Sector

Our main empirical results show that lifetime experiences of GDP growth shape measured risk aversion. Our theory suggests that these results stem from the fact that GDP fluctuations capture changes in background income risk. If this interpretation is correct, we might expect our results to be stronger for individuals who are employed in sectors where conditions are more highly correlated with GDP growth. We test this in this section.

Using our data, we split employed subjects in our primary sample by their sector of employment in the first wave of the respective survey (starting with the primary sample defined in Section 3.D and filtering on employment status yields a sample of 8,213 subjects in Indonesia and 5,101 in Mexico). We consider nine sectors, harmonized across the two countries: manufacturing, agriculture, mining, utilities, construction, wholesale and retail, transportation, finance, and social services. We conduct the heterogeneity analysis by repeating our baseline analysis while adding variables coding the interaction between our lifetime growth experience variables and employment sector dummies. The specification is thus:

$$\Delta R_{it} = \alpha + \beta_1 \Delta A_{it} + \beta_2 \Delta V_{it} + \beta_{1,s} \Delta A_{it} \times D_{is} + \beta_{2,s} \Delta V_{it} \times D_{is} + \phi_s D_{is} + \gamma \Delta Inflation_p + \epsilon_{it},$$

where D_{is} is a dummy variable taking a value of 1 if individual i is employed in industry s in the first wave of the respective survey.

To interpret these results we must determine which sectors are more sensitive to GDP growth in our setting. Since all sectors are procyclical, a higher correlation with GDP is typically captured by measuring sectoral volatility, with higher volatility indicating greater procyclicality. In the most definitive study on sectoral volatility in developing countries to date, Koren and Tenreyro (2007) find that "agriculture and mining and quarrying tend to be more volatile than all the manufacturing sectors, and the services sectors tend to be less volatile than manufacturing". Many studies have shown that the financial sector is highly volatile and procyclical in developed countries (Bernanke and Gertler (1990), Jordà et al. (2013), Jordà et al. (2016)), while a smaller set of studies echoes these findings for developing countries (Glen and Mondragón-Vélez (2011), Bahadir and Gumus (2016)). In addition, Im et al. (2011) report that social spending is strongly positively correlated with GDP growth in a group of lower middle income countries that includes Indonesia, while Chávez Martín del Campo et al. (2010) report the same for Mexico. Based on these findings, we expect our results to be strongest for those employed in agriculture, mining, finance, and social services; intermediate for those employed in manufacturing; and weakest for those employed in retail and wholesale and transportation, the two non-financial service sectors in our data.

Our empirical findings, which we report in Figure K.1, largely bear out this predicted pattern. Relative to those employed in manufacturing, subjects employed in agriculture, social services, and finance exhibit significantly stronger effects of growth experiences on risk aversion in both countries. Subjects employed in wholesale and retail and in transportation exhibit indistinguishable changes relative to those in manufacturing in Indonesia and weaker effects in Mexico.

Figure K.1: Heterogeneity of main results by sector of employment

Indonesia

(a) Effect of mean growth by industry

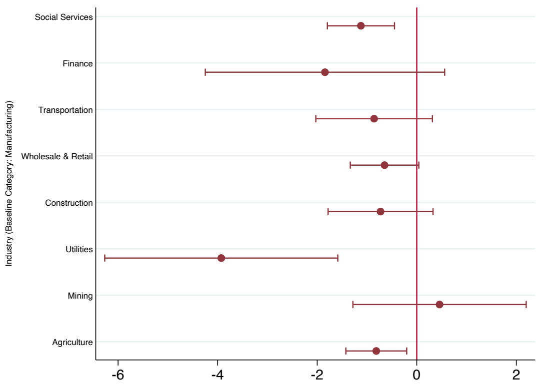

(b) Effect of growth std. dev. by industry

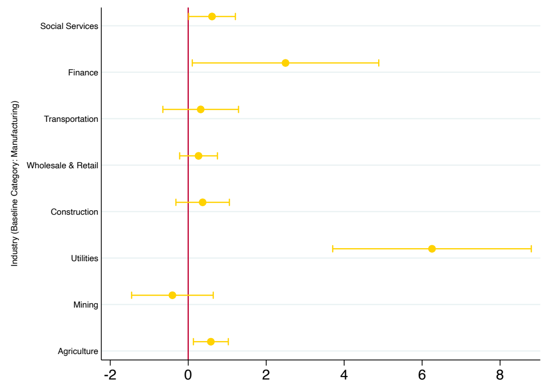

Mexico

(c) Effect of mean growth by industry

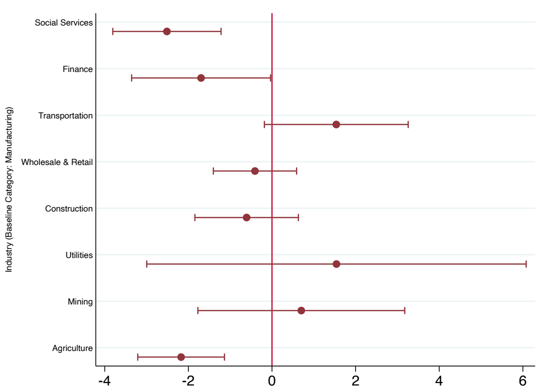

(d) Effect of growth std. dev. by industry

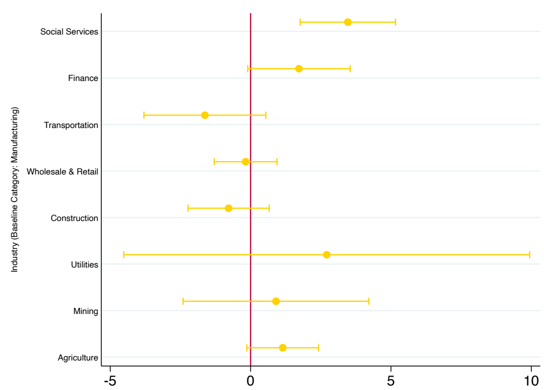

Notes: Graphs display the heterogeneity in the effects of experienced average and volatility of growth on measured risk aversion by industry of employment in the first wave of the survey (2007 in Indonesia and 2005 in Mexico), relative to the excluded sector of manufacturing. Note that stronger sectoral effects are indicated by negative coefficients in the case of mean growth and by positive coefficients in the case of growth standard deviation. These results are for subjects in the primary sample, described in Section 3.D. 90% confidence intervals from standard errors clustered at the birth-year by state-of-birth level plotted.

L Details of Additional Controls in Table 2

Table L.1: Description of controls included

Category	Variables Included
Demographics (**Indonesia** and **Mexico**)	Married Household size Household size squared Educational attainment
Income, Assets and Net yearly savings (**Indonesia** and **Mexico**)	Total household income Total value of household assets Net household 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**)	Feels safe at home Fear of assault during the day Fear of assault at night No. of times robbed, assaulted, kidnapped No. of family/friends robbed, assaulted, kidnapped in last 12 months Homicide rate in last 12 months in municipality of residence in MxFLS2 (built by Brown et al. (2019))
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

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

M Alternative Treatment of Older Cohorts

M.A Treat old cohorts as born in 1977 in Indonesia or 1941 in Mexico

For this variation, we build the mean (A_{it}) and the standard deviation (V_{it}) of subjects' time series from birth to year of measurement in the corresponding survey as follows. Let g_{is} be the growth rate assigned to person i in year s, d_i the year their series construction starts, and b_i their birth year. Then for year of measurement t these statistics are:

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

$$\text{where} \quad g_{is} = \Big\{ g_{State,s}, \quad d_i = \begin{cases} b_i & \text{if } b_i \geq B \\ B & \text{if } b_i < B, \end{cases} \text{and } B = \begin{cases} 1977 & \text{if Country}_i = \text{Indonesia} \\ 1941 & \text{if Country}_i = \text{Mexico.} \end{cases}$$

M.B Attach country-level growth for the first 15 years only to subjects born in 1925–1940 in Mexico and 1961–1976 in Indonesia

For this variation, we build the mean (A_{it}) and the standard deviation (V_{it}) of subjects' time series from birth to year of measurement in the corresponding survey as follows. Let g_{is} be the growth rate assigned to person i in year s and b_i their birth year. Then for year of measurement t these statistics are:

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

$$\text{where for Indonesia: } g_{is} = \begin{cases} g_{National,s} \text{ for } s \in [b_i, b_i + 15] \text{ if } b_i \leq 1976 \\ g_{State,s} \text{ for } s \in [b_i + 16, t] \text{ if } b_i \leq 1976 \\ g_{State,s} \text{ for } s \in [b_i, t] \text{ if } b_i > 1976 \end{cases}$$

$$\text{and for Mexico: } g_{is} = \begin{cases} g_{National,s} \text{ for } s \in [b_i, b_i + 15] \text{ if } b_i \leq 1940 \\ g_{State,s} \text{ for } s \in [b_i + 16, t] \text{ if } b_i \leq 1940 \\ g_{State,s} \text{ for } s \in [b_i, t] \text{ if } b_i > 1940, \end{cases}$$

$$\text{and } b_i \geq \begin{cases} 1961 & \text{if Country}_i = \text{Indonesia} \\ 1925 & \text{if Country}_i = \text{Mexico.} \end{cases}$$

M.C Drop cohorts born pre- 1977 in Indonesia and 1941 in Mexico

For this variation, we build the mean (A_{it}) and the standard deviation (V_{it}) of subjects' time series from birth to year of measurement in the corresponding survey as follows. Let g_{is} be the growth rate in the state of birth of person i in year s and b_i their birth year. Then for year of measurement t these statistics are:

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

$$\text{where} \quad g_{is} = \begin{cases} g_{State,s} \end{cases} \quad b_i \geq \begin{cases} 1961 & \text{if Country}_i = \text{Indonesia} \\ 1925 & \text{if Country}_i = \text{Mexico.} \end{cases}$$

N Alternate Specifications of Measured Risk Aversion

Table N.1: Ordered probit

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)
Indonesia			
Δ Growth Mean	-0.12*		-0.44*
	(0.03)		(0.06)
Δ Growth Volatility		-0.00	0.23*
		(0.02)	(0.04)
Observations	11636	11636	11636
Mexico			
Δ Growth Mean	-0.46***		-0.39***
	(0.09)		(0.09)
Δ Growth Volatility		0.87***	0.69***
		(0.18)	(0.18)
Observations	10224	10224	10224

Notes: *Measured Risk Aversion*: 1–5 (Indonesia and Mexico), with 5 being the highest measured risk aversion. State (Indonesia) and regional (Mexico) inflation included in all regressions. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table N.2: Binarized measure of risk aversion

Dep. Var: Δ Binarized Meas. Risk Av.	(1)	(2)	(3)
Indonesia			
Δ Growth Mean	-0.05***		-0.22***
	(0.02)		(0.04)
Δ Growth Volatility		0.00	0.12***
		(0.01)	(0.02)
Observations	11636	11636	11636
Mexico			
Δ Growth Mean	-0.29***		-0.24***
	(0.06)		(0.06)
Δ Growth Volatility		0.58***	0.46***
		(0.12)	(0.13)
Observations	10224	10224	10224

Notes: *Binarized Measured Risk Aversion*: Measured Risk Aversion buckets 1 and 2 are set to 0, and buckets 3, 4, and 5 to 1. State (Indonesia) and regional (Mexico) inflation included in all regressions. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

O Alternate Clustering

Table O.1: Different clustering levels

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)
Indonesia			
Δ Growth Mean	-0.85***	-0.85***	-0.85***
	(0.18)	(0.23)	(0.27)
Δ Growth Std. Dev.	0.45***	0.45***	0.45***
	(0.11)	(0.11)	(0.11)
Observations	11636	11636	11636
Cluster	5-year YOB bins by POB	10-year YOB bins by POB	15-year YOB bins by POB
Mexico			
Δ Growth Mean	-0.86***	-0.86***	-0.86**
	(0.25)	(0.29)	(0.34)
Δ Growth Std. Dev.	1.61***	1.61**	1.61**
	(0.56)	(0.65)	(0.70)
Observations	10224	10224	10224
Cluster	5-year YOB bins by POB	10-year YOB bins by POB	15-year YOB bins by POB

Notes: *Measured Risk Aversion*: 1–5 (Indonesia and Mexico), with 5 being the highest measured risk aversion. State (Indonesia) and regional (Mexico) inflation included in all regressions. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at different levels in parentheses. The clustering level in each column is: (1) 5-year birth-year bins by province/state-of-birth; (2) 10-year birth-year bins by province/state-of-birth; and (3) 15-year birth-year bins by province/state-of-birth. *** p<0.01, ** p<0.05, * p<0.1.

P Non-Linear Weighting à la 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:

$$w_{it}(s, \lambda) = \frac{(age_{it} - s)^{\lambda}}{\sum_{s=0}^{age_{it}} (age_{it} - s)^{\lambda}}.$$

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 P.1 illustrates this weighting scheme for the different values of λ we obtain in our analysis below and $\lambda = 0$ for a 30 year old subject. For all ages, higher values of λ mean relatively more weight is placed on recent experiences, $\lambda = 0$ implies a flat weighting scheme like the one used in our baseline analysis, and negative values of λ indicate relatively more weight is placed on early life experiences.

Figure P.1: Relative weights placed on years of growth for an individual of age 30 at different levels of λ

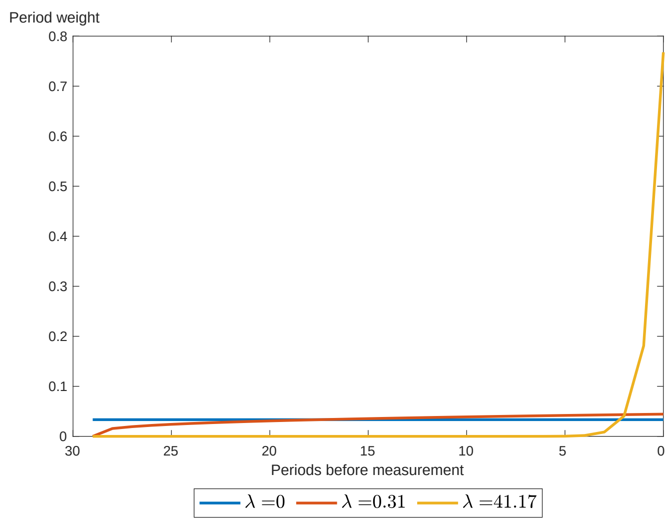

Using this weighting scheme, we construct a measure of average experienced lifetime growth, along with a measure of experienced lifetime growth volatility that uses the weighted standard deviations of experienced growth:

$$A_{it}(\lambda) = \sum_{s=0}^{age_{it}} w_{it}(s,\lambda) g_{t-s} \quad \text{and} \quad V_{it}(\lambda) = \sqrt{\frac{\sum_{s=0}^{age_{it}} w_{it}(s,\lambda)(g_{t-s} - A_{it}(\lambda))^2}{\frac{age_{it}-1}{age_{it}} \sum_{s=1}^{age_{it}} w_{it}(s,\lambda)}}.$$

We estimate the marginal effects of these weighted macroeconomic experiences variables on measured risk aversion by estimating Equation (3) using non-linear least squares (NLLS).[60] This allows us to simultaneously estimate the marginal coefficients of mean growth and growth volatility (β_1 and β_2) and the value of the non-linear weighting parameter λ:

$$\Delta R_{it} = \alpha_{FD} + \beta_1 \Delta A_{it} + \beta_2 \Delta V_{it} + \gamma_1 \Delta Inflation_p + \gamma_2 \Delta X_{it} + \epsilon_{it}. \tag{3}$$

We present the results from this exercise in Table P.1. We find that all results remain broadly consistent. The effects of volatility of growth on measured risk aversion remain positive and highly significant for all specifications in both samples. The effects of average growth on measured risk aversion remain negative across all specifications, and are highly significant across almost all specifications. In Indonesia, the effect of mean growth becomes nonsignificant in the specification that includes both average and volatility of growth. This is likely driven by the value of the weighting parameter in this case (41.17), which places a disproportionate weight in recent experiences. Notice that when the effect of mean growth is considered in isolation (column 1), the value of this weighting parameter is much lower (3.68), and the marginal effect of mean growth is negative and statistically significant. This indicates that recency matters more for the effects of volatility of growth rather than for the effects of mean growth. This implies that the push towards a higher positive weighting parameter is driven by the importance of recent experiences of macroeconomic growth volatility on measured risk preferences, and that this weighting scheme decreases the importance of experienced average growth on measured risk.

[60] Since non-linear estimation methods are known to be sensitive to initial seed values, we choose the initial value of λ by maximizing 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 estimate linearized versions of the model (analogous to Equation (2)) using these macroeconomic variables via maximum likelihood. We choose the initial seed for the non-linear estimation to match the value of λ that maximizes the likelihood across all values of λ.

Table P.1: Non-linear temporal λ weighting

Dep. Var: Δ Meas. Risk Av.	(1)	(2)	(3)
Indonesia			
Δ Growth Mean	-0.10***		-0.02
	(0.03)		(0.02)
Δ Growth Volatility		0.76***	0.72***
		(0.13)	(0.14)
λ	3.68***	43.00***	41.17***
	(0.56)	(5.41)	(5.51)
Observations	11633	11633	11633
Mexico			
Δ Growth Mean	-0.65***		-0.75***
	(0.18)		(0.18)
Δ Growth Volatility		0.99***	1.04***
		(0.29)	(0.28)
λ	0.30***	0.41***	0.31***
	(0.01)	(0.02)	(0.01)
N	10223	10223	10223

Notes: *Measured Risk Aversion*: 1–5 (Indonesia and Mexico), with 5 being the highest measured risk aversion. Regressions estimated via NLLS. State (Indonesia) and regional (Mexico) inflation included in all regressions. These results are for subjects in the primary sample, described in Section 3.D. Standard errors clustered at the birth-year by state-of-birth level in parentheses. *** p<0.01, ** p<0.05, * p<0.1.