LUND UNIVERSITY
School of Economics and Management

Master's Programme in Economic Development and Growth (MEDEG)

Are Children a Risky Investment?

How Cultural Risk Preferences Shape Fertility Responses to Macroeconomic Shocks

by

Chiara Giorgi - ch2334gi-s@student.lu.se

Abstract: This study investigates whether cultural risk preferences moderate the relationship between macroeconomic shocks — namely inflation, unemployment, and GDP contraction — and fertility across 76 developed and developing countries over the period 2000–2023. Drawing on the Becker–Tomes quantity–quality framework extended with a precautionary saving channel, I argue that more risk-averse societies treat the irreversible commitment of childbearing as a larger effective resource loss when faced with economic uncertainty, producing stronger fertility responses to adverse shocks. Using two-way fixed-effects panel regressions, I find that inflation and unemployment robustly reduce total fertility rates, while cultural risk tolerance significantly buffers the negative fertility response to labour-market instability. This moderating role is concentrated among middle- and high-income countries, where precautionary motives are most operative, and is robust across multiple specification checks. These findings provide strong evidence that pro-natalist policies targeting the sources of economic uncertainty could play a meaningful role in limiting the fertility costs of macroeconomic downturns.

Keywords: *Fertility, Cultural risk tolerance, Macroeconomic shocks*

Course code: EKHS42
Master's Thesis (15 credits ECTS)
June 2026
Supervisor: Martin Dribe
Examiner: Finn Hedefalk
Word Count: 15,982 words (excluding appendices and tables)

Acknowledgments

This thesis has been a deeply engaging project for me. I am especially grateful to my supervisor, Martin Dribe, for helping me refine ideas that were still taking shape. His ability to quickly grasp complex issues, together with his experience, greatly influenced my understanding of the material. Being able to work with focus also requires a sense of emotional stability and belonging. I am therefore very grateful to the friends who shared this year with me in Lund—at times a chaotic one—for providing the balance and support I needed. Finally, I would like to express my sincere appreciation to my parents, who, from Italy, have always supported me and encouraged me to continue my studies.

Contents

List of Tables

List of Figures

1 Introduction

Since the onset of the Great Recession in 2007, a growing body of empirical work has sought to identify the specific channels through which macroeconomic deterioration shapes reproductive behavior, with rising unemployment and labor-market precariousness emerging as the most consistently documented drivers of fertility decline across Europe (Goldstein et al., 2013). That this reversal coincided so closely with a macroeconomic shock is unlikely to be coincidental: it suggests that households, acting as economic agents, adjust their reproductive plans with remarkable sensitivity to changes in the broader economic environment, treating childbearing as a form of consumption that can be deferred when uncertainty rises. Indeed, there is a pervasive consensus among policymakers and researchers alike that birth rates are inextricably tied to the business cycle (Sobotka et al., 2011). This is not merely a matter of current income: since childbearing is an irreversible and long-horizon commitment, fertility decisions are governed above all by households' expectations about future economic stability (Adesete and Piano, 2026).

A separate but complementary strand of research, rooted in behavioral economics, emphasizes that objective economic conditions do not translate mechanically into economic behaviour: individual responses are fundamentally mediated by underlying preferences, and notably by attitudes toward risk (Falk et al., 2016). Risk preferences can be defined as the degree to which an individual is willing to accept uncertainty, also referred to as risk tolerance (Bellani and Arpino, 2021). In economics, risk is not just a background condition but a defining feature of decision-making, and has become one of the discipline's core tools for understanding how people navigate uncertain outcomes (Hudson et al., 2005). The coefficient of absolute risk aversion $r(x) = -u''(x)/u'(x)$, developed independently by Arrow and Pratt (1964), and commonly referred to as the Arrow-Pratt measure, provides the canonical measure of how strongly a decision maker wishes to avoid uncertainty at a given level of wealth. Intuitively, it captures the curvature of an individual's utility function: the more an individual's utility diminishes at the margin as wealth increases, the more they are willing to forgo expected returns in order to avoid risk. Its properties underpin virtually all subsequent empirical work on risk preferences.

While scholars have begun to explore how individual psychological characteristics shape demographic outcomes, this line of inquiry has developed largely in isolation from the macroeconomic literature reviewed above. On the one hand, some work has examined how cultural norms surrounding work and domestic life mediate the relationship between economic instability and family formation, shaping how individuals psychologically process and practically respond to precarity (Bernardi et al., 2008). On the other hand, a small but growing body of research has connected risk preferences directly to demographic behaviour: given the unpredictability surrounding the availability of suitable partners, the effectiveness of contraception, and the ability to conceive, women's attitudes toward risk are likely to influence the timing of both marriage and childbearing (Schmidt, 2008). Yet neither strand has examined how risk preferences specifically condition individual responses to macroeconomic shocks. When households face uncertainty about future income or employment prospects, it remains largely an open question whether and how risk preferences shape the resulting adjustments in reproductive plans.

Aim and Scope Addressing the gap outlined above is the primary aim of the present study. Using a panel dataset covering 76 developed and developing nations over the period 2000–2023, I investigate whether cross-country variation in risk-taking culture acts as a systematic moderator of the relationship between macroeconomic shocks and fertility, and specifically whether societies with a higher cultural propensity for risk-taking exhibit greater resilience in fertility rates when faced with adverse economic conditions. The inclusion of both developed and developing countries is deliberate, as it allows for a heterogeneity analysis of whether the moderating role of risk culture varies systematically across income levels. The temporal window is equally motivated: the period 2000–2023 is characterised by substantial heterogeneity in macroeconomic conditions — particularly the Global Financial Crisis and the COVID-19 pandemic — while also situating the risk preference data, measured in 2012, approximately at its midpoint, making them a plausible snapshot of the cultural environment prevailing throughout the observation window.

The policy stakes of such analysis are considerable. At current mortality levels, a total fertility rate (TFR) of 2.1 children per woman is required to ensure generational replacement (Fluchtmann et al., 2023), yet fertility in many countries now falls well below this threshold (Wu, 2024). But the relevance of understanding fertility dynamics is not confined to low-fertility contexts. According to the UNFPA (United Nations Population Fund, 2025), across all demographic settings, individuals frequently struggle to realise their preferred family size: many people have more children than they consider ideal, while many others have fewer than they would like. These mismatches between desired and actual fertility are widespread in both high- and low-fertility countries alike, underscoring that structural, economic, and social barriers to fulfilling reproductive aspirations are pervasive. If macroeconomic shocks are among the forces driving a wedge between reproductive intentions and outcomes — and if risk tolerance shapes how individuals navigate such shocks — then understanding this relationship has direct implications for policies aimed at helping people achieve their desired family size, regardless of whether a country sits above or below replacement level.

Outline of the Thesis The remainder of the thesis is structured as follows. Section 2 provides contextual background on cross-country variation in fertility patterns and risk attitudes, situating the analysis within a comparative perspective across developed and developing countries. Section 3 develops the theoretical framework, extending the Becker–Tomes quantity–quality model to incorporate macroeconomic uncertainty and risk preferences, and reviews the relevant empirical literature on macroeconomic shocks and fertility before stating the three main hypotheses. Section 4 describes the data, detailing variable selection, sources, and distributional diagnostics. Section 5 outlines the empirical methodology, including the baseline fixed-effects specification, the interaction model, and the heterogeneity analysis by income level. Section 6 presents and discusses the results — covering the baseline estimates, the moderating role of cultural risk tolerance, the stratified analysis by income group, and a series of robustness checks — and Section 7 concludes, drawing together the main findings and discussing their implications.

2 Context

This section provides the background for the analysis by placing fertility patterns and risk preferences in a comparative cross-country perspective. Since the effects of macroeconomic shocks on fertility are unlikely to be uniform, it is important to account for differences in development levels as well as in attitudes toward risk. In this regard, the 76 countries included in the analysis span a wide range of development levels, offering a particularly suitable setting to explore such heterogeneity. Their fertility patterns differ substantially across contexts, with trends evolving markedly over the period 2000–2023. Across OECD countries, the postponement of childbearing has become a defining demographic feature of the early twenty-first century, with the mean age of mothers at childbirth rising steadily between 2002 and 2019. Yet aggregate fertility trends have been far less monotone. After a prolonged decline beginning in the 1960s, the Total Fertility Rate (TFR) appeared to recover after 2002 and peaked in 2008, only to contract sharply in the wake of the global financial crisis (Fluchtmann et al., 2023). While the demographic impact of the recent economic downturn has been extensively documented in high-income contexts — particularly across Europe — there remains a notable scarcity of literature concerning less developed regions. This oversight is critical, as the mechanisms linking economic volatility to fertility in the Global South may diverge significantly from the patterns observed in the developed world (Sobotka et al., 2011).

On the other hand, in most low-income countries, fertility decline did not begin until the decades following the Second World War (Fernandez-Villaverde, 2001). Today, the West and Central Africa region records the highest levels of adolescent childbearing worldwide, followed by East and Southern Africa (United Nations Population Fund, 2025). Several mechanisms help explain why fertility remains comparatively high in many low-income settings. In these contexts, children may play an important economic role, functioning both as a form of insurance and as a source of intrinsic value (Bellani and Arpino, 2021). In Sub-Saharan Africa, structural conditions such as lower income, limited educational attainment, and higher child mortality continue to sustain a preference for larger families (Bongaarts, 2011). In this region specifically, although the fertility transition is underway, its pace has been comparatively gradual (Towriss and Timæus, 2018). In addition, evidence from specific country studies—such as Egypt—shows how social and familial expectations can further shape reproductive behaviour, sometimes encouraging earlier or more closely spaced childbearing than women themselves prefer (Abdel-Tawab et al., 2008).

Crucially, a similar cross-country heterogeneity extends to attitudes toward risk. Figure 1 illustrates the global distribution of average willingness to take risks, based on data from the Global Preference Survey (Falk et al., 2018) — the dataset used in this analysis. The survey covers 15 countries in the Americas, 25 in Europe, 22 in Asia and the Pacific, and 14 in Africa — 11 of which are in Sub-Saharan Africa — and accounts for roughly 90% of the world's population and global income. Countries are colored according to their mean risk-taking scores, ranging from dark blue (lowest willingness to take risks) to dark red (highest). Countries in white were not included in the survey and are therefore excluded from the present analysis.

The map highlights clear geographic patterns. Regional averages, expressed as standard deviations from the world individual mean, reveal that Sub-Saharan Africa stands out as the most risk-tolerant region (0.34), followed by North Africa and the Middle East (0.16) and Neo-Europe — the United States, Canada, and Australia — (0.15). South and East Asia sits slightly below the world average (0.10), while Western Europe (0.11), Eastern Europe (0.12), and South America (0.03) all show mild risk aversion relative to the global mean. These differences are not merely descriptive: among all 2,850 possible pairwise country comparisons, 78% are significant at the 1% level, confirming that the regional patterns reflect genuine and widespread differences in risk attitudes rather than sampling noise.

Figure 1: Cross-Country Variation in Willingness to Take Risks

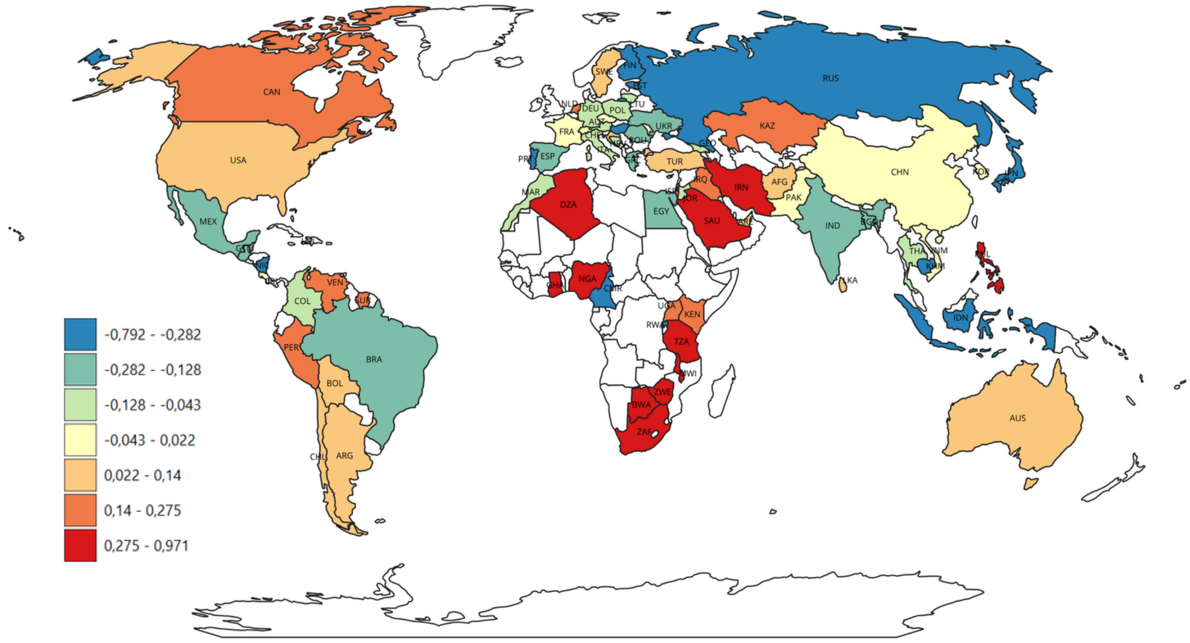

Note: The map displays average willingness to take risks by country, as measured on the original Global Preference Survey (GPS) scale (Falk et al., 2018). Countries in white are not included in the survey. The color scale ranges from dark blue (lowest willingness to take risks) to dark red (highest). Map boundaries are based on Natural Earth 110m cultural vectors (Natural Earth, 2023). The map was produced using QGIS.

This cross-country variation in fertility, income, and risk attitudes provides a rich setting to identify how macroeconomic shocks shape fertility outcomes, and to assess whether these effects differ systematically with countries' underlying risk preferences.

3 Theory and previous research

The intuition that economic conditions shape reproductive decisions has a long intellectual pedigree. As Smith (1776, p. 83) argued, wages must be sufficient not only to sustain the worker but also to allow him to raise a family, since without this capacity "the race of such workmen could not last beyond the first generation." Malthus (1798, p. 16) identified the mechanism: reason leads the prospective parent to ask whether he risks lowering his rank, being obliged to labour harder, or seeing his offspring in misery — considerations that "certainly do prevent a very great number in all civilised nations" from having children. Taken together, these arguments provide a foundation for viewing fertility decisions as fundamentally shaped by economic considerations, a premise that the formal framework developed below seeks to capture explicitly.

3.1 Theoretical Framework: The Quantity-Quality Trade-off Under Macroeconomic Uncertainty

Uncertainty permeates human experience, and individuals' responses to it depend on their perceived control, shaping how they seek stability and coherence (Casanova et al., 2024).Bernardi et al. (2008) identified two competing logics linking economic instability to reproductive behavior. The *insecurity narrative* holds that job and income instability depress fertility by threatening the financial foundations parenthood requires. The *uncertainty narrative* inverts this: biographical unpredictability can make parenthood attractive precisely because it imposes structure on an otherwise unmoored life course. Crucially, both narratives rest on the subjective meaning individuals attach to uncertainty, rather than on objective conditions alone — suggesting that cross-country variation in cultural risk attitudes could potentially moderate the aggregate fertility response to economic downturns. To conceptualize this relationship formally, I build on the "Quantity-Quality" trade-off model of Becker and Lewis (1973), extended by Becker and Tomes (1976) to incorporate child endowments.

3.1.1 Baseline Model

In Becker's model, each household is assumed to maximise a utility function[1] of the form:

$$U = U(n, w, y), \tag{1}$$

where n is the number of children, w is the quality of each child (assumed uniform across children), and y is the aggregate consumption of all other commodities. Child quality is produced according to the additive function:

$$w = e + q, \tag{2}$$

where $e \geq 0$ is the child's exogenous endowment — comprising, for instance, inherited ability, public investment in education, and luck — and $q \geq 0$ is the household's own contribution to quality. The household is subject to a budget constraint, defined as:

$$p_y y + p_q n q = I, \tag{3}$$

where p_y is the price of the consumption good, p_q is the unit cost of increasing q by one unit, and I is household income. Note that the total expenditure on children is $p_q n q$, reflecting the fact that a higher contribution q must be provided to *each* of the n children.

Maximising (1) subject to (3) and (2) yields the shadow prices of child quality and child quantity:

[1]The utility function U summarises the household's preferences over these three goods, with $\frac{\partial U}{\partial n} > 0$, $\frac{\partial U}{\partial w} > 0$, and $\frac{\partial U}{\partial y} > 0$, reflecting that more children, higher child quality, and greater consumption are each desirable.

$$\pi_w = np_q \qquad \text{and} \qquad \pi_n = qp_q. \tag{4}$$

These shadow prices embody the key non-linearity of the model: the cost of improving quality is proportional to the number of children already present, and the cost of having an additional child is proportional to the quality investment already committed. Substituting (4) and (2) into the budget constraint yields the household's *social income*:

$$\pi_y y + \pi_n n + \pi_w w = I + \pi_w e \equiv S, \tag{5}$$

where $\pi_y = p_y$ is the shadow price of consumption — equal to its market price since y enters neither the quality nor the quantity interaction — and $S \equiv I + \pi_w e$ is the household's social income. The term $\pi_w e$ reflects the implicit value of the child's exogenous endowment: since e raises child quality without requiring household expenditure, it effectively supplements available resources, so $S > I$ whenever $e > 0$.[2] Crucially, because the shadow prices in (4) are themselves functions of n and q, any change in one dimension of the fertility decision alters the cost of the other. An increase in n raises π_w, making quality more expensive; a rise in q raises π_n, making further children more expensive. This mutual cost-raising interaction is the mechanism through which the model can generate a low or even negative income elasticity of quantity alongside a high elasticity of quality (Becker and Lewis, 1973; Becker and Tomes, 1976).[3]

3.1.2 Incorporating Macroeconomic Shocks

The baseline model by Becker is extended here to incorporate macroeconomic shocks. Specifically, suppose a negative shock $\eta_t > 0$ reduces the household's realised income in period t. The budget constraint becomes:

$$p_y y + p_q nq = I - \eta_t, \tag{6}$$

and, since the endowment e is exogenous, social income is correspondingly reduced:

$$S_t = (I - \eta_t) + \pi_w e = S - \eta_t. \tag{7}$$

The fall in S_t induces a contraction in the demand for n, w, and y. However, the adjustment between quantity and quality is not symmetric. Because the endowment e is fixed, any desired

[2] Social income S can be thought of as the household's effective resource base: beyond monetary income I, the exogenous endowment e — comprising inherited ability, public schooling, or luck — contributes to child quality without requiring household expenditure, and is therefore valued at $\pi_w e$. A household whose children benefit from, say, high-quality public education needs to spend less to achieve any given level of w, and is in this sense richer than its income I alone would suggest. Social income S captures this by augmenting I with the implicit value of what the household receives for free.

[3] Income elasticity measures the percentage change in demand for a good in response to a one percent increase in income. For a normal good, this elasticity is positive — demand rises with income. The Q-Q model predicts that children are unusual in this respect: as household income rises, the mutual cost-raising interaction between n and q makes quantity relatively more expensive, inducing a substitution toward quality. The result is that richer households tend to have fewer children but invest substantially more in each one — a pattern documented empirically across a wide range of countries and time periods, and one of the most celebrated predictions of Becker's framework.

reduction in total quality w can only be achieved by reducing the household contribution q. By (4), a fall in q lowers the shadow price of quantity π_n, which partially offsets the income-driven fall in demand for n. Conversely, a reduction in n would lower the shadow price of quality π_w, encouraging greater investment in q. Figure 2 displays the mechanism.

Figure 2: The Quantity–Quality Trade-off Under a Negative Income Shock

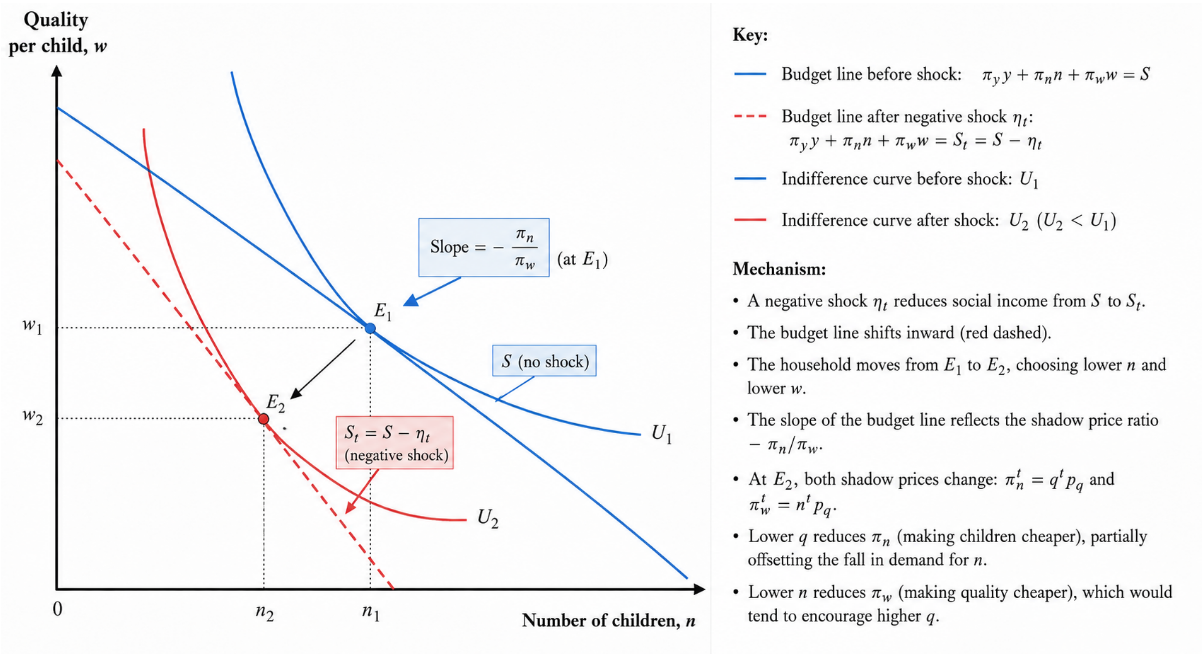

Note: The figure illustrates the household's optimal allocation between the number of children (n) and quality per child (w) under income uncertainty. The initial optimum is at E_1, where the baseline social budget line $\pi_y y + \pi_n n + \pi_w w = S$ (solid blue) is tangent to the initial indifference curve U_1, yielding the optimal choices n_1 and w_1. The slope of the budget line at this point reflects the shadow price ratio $-\pi_n/\pi_w$. A negative macroeconomic income shock $\eta_t > 0$ reduces social income to $S_t = S - \eta_t$, shifting the budget constraint inward to the red dashed line. The household moves to a lower welfare equilibrium at E_2 on indifference curve U_2 ($U_2 < U_1$). Due to the fixed endowment e, the reduction in quality is driven by a drop in the household's contribution q. This fall in q sequentially reduces the shadow price of quantity ($\pi_n^t = q^t p_q$), making children cheaper and partially offsetting the mechanical income-driven decline in n. Consequently, both quantity and quality fall asymmetrically ($n_2 < n_1$ and $w_2 < w_1$) based on endogenous shadow price adjustments.

The net direction of adjustment depends on the relative income and price elasticities and on the size of the endowment ratio e/q: the larger this ratio, the more the observed quantity elasticity is compressed relative to the quality elasticity (Becker and Tomes, 1976). Empirically, this model predicts that fertility, quantified as the number of children, is more sensitive to income shocks at higher income levels, where the endowment ratio e/q is smaller.

Another way through which a shock could operate is through prices rather than income directly. Inflation raises p_y, the cost of all other consumption, and plausibly also p_q, the unit cost of investing in child quality, as expenditures on nutrition, healthcare, and education all become more expensive. Through the shadow price mechanism in (4), a rise in p_q increases both $\pi_w = np_q$ and $\pi_n = qp_q$, making children more costly in both quantity and quality dimensions simultaneously. While the mechanism differs from a direct income shock η_t, the outcome is analogous: by eroding the household's real purchasing power, inflation effectively reduces the resources available for childbearing, mimicking a contraction in S_t.[4]

3.1.3 Macroeconomic shocks and risk preferences

Fertility decisions are often modeled as risk-free and fully controlled, but in reality they involve significant uncertainty. Women face three main costs: lost earnings from having children (especially in career-intensive jobs), the financial and practical costs of contraception, and the risk of infertility from delaying childbirth (Schmidt, 2008). The present analysis focuses on one specific and tractable source of uncertainty: shocks to future household income arising from macroeconomic fluctuations. These shocks add a broader, economy-wide layer of uncertainty on top of the individual risks outlined above, and constitute the main focus of the analysis.

To introduce risk preferences, I follow the precautionary saving logic of Leland (1968): when income is uncertain, risk-averse households behave as if their available resources were lower than the expected value of income, setting aside a *precautionary buffer* against future shocks. Because childbearing is a largely irreversible commitment — the costs of raising a child extend far beyond any single period — households treat an additional birth as they would any other durable and illiquid expenditure: a commitment that cannot be undone once made. A more risk-averse household will therefore require a larger compensation before proceeding with an additional birth under uncertainty.

I capture this by replacing the realised income $I - \eta_t$ in (6) with an *effective income* that discounts available resources by a risk premium $\phi(\rho, \sigma_\eta^2)$, which is increasing in both the household's risk aversion ρ and the variance of the shock σ_η^2:

$$p_y y + p_q n q = I - \eta_t - \phi(\rho, \sigma_\eta^2), \tag{8}$$

where η_t is the income loss that materialises during period t, and $\phi(\rho, \sigma_\eta^2) \geq 0$ is the precautionary wedge that risk-averse households impose on themselves *ex ante*, in response to uncertainty about future economic conditions. The two components are conceptually distinct: η_t is realised, while $\phi(\rho, \sigma_\eta^2)$ is anticipated. Social income therefore becomes:

$$S^* = S - \eta_t - \phi(\rho, \sigma_\eta^2) \tag{9}$$

Equation (9) identifies two distinct channels through which macroeconomic shocks shape fertility decisions, and — crucially — shows that they interact. The first is an *income channel*: a shock of size η_t directly reduces the resources available for childbearing. This channel operates regardless of risk preferences: it is a mechanical consequence of lower income.

The second is a *precautionary channel*: even setting aside the realised income loss, uncertainty about future economic conditions leads risk-averse households to behave as if their

[4] As will be discussed in the empirical section, this interpretation is consistent with Wu (2024), who documents that rising costs of living lead households to postpone or forgo additional children.

resources were lower than they currently are. This is captured by $\phi(\rho, \sigma_\eta^2)$, which is larger when households are more risk-averse (higher ρ) or when the shock is more volatile (higher σ_η^2). A household with higher ρ therefore suffers a larger effective income loss from the *same* shock — not because the shock is objectively larger, but because the precautionary response magnifies its impact. This interaction is the theoretical counterpart of the empirical specification in Section 5.2, where the fertility response to each macroeconomic shock is allowed to vary with the country's cultural level of risk tolerance.

The precautionary wedge $\phi(\rho, \sigma_\eta^2)$ has a precise microeconomic foundation in Pratt (1964), who shows that for a household with utility $u(\cdot)$ and wealth x, the maximum willingness to pay to avoid a zero-mean risk $\tilde{z}$ with variance $\sigma_{\tilde{z}}^2$ satisfies:

$$\pi(x, \tilde{z}) \approx \tfrac{1}{2}\, \sigma_{\tilde{z}}^2\, r(x), \tag{10}$$

where $r(x) \equiv -u''(x)/u'(x)$ is the Arrow–Pratt coefficient of absolute risk aversion. The term $\phi(\rho, \sigma_\eta^2)$ in (8) is the empirical counterpart of π: it is increasing in $r(x)$ — summarised by ρ — and in σ_η^2. Cross-country variation in willingness to take risks, documented by Falk et al. (2018) and discussed in Section 2, can therefore be read as variation in the average level of $r(x)$ across societies. I treat this as a time-invariant cultural parameter that moderates the strength of the precautionary channel — and hence the total fertility response to a given shock — in (9).[5]

A behavioural extension: loss aversion and reference dependence. The precautionary channel introduced above rests on expected utility theory, which assumes that households evaluate uncertain outcomes by weighting each possible consequence by its probability and choosing the option that maximises the expected value of a well-behaved utility function. This is a powerful framework, but it carries a demanding assumption: that people assess outcomes in absolute terms, anchored to no particular reference point. A substantial body of empirical work shows that this assumption systematically fails to describe how individuals actually behave when facing risky choices (Hudson et al., 2005).

Easterlin (1976) anticipated this failure in the fertility context by arguing that reproductive decisions are governed not by absolute income but by the tension between a couple's material aspirations and their available resources. When resources fall short of aspirations, couples become hesitant to have children, regardless of their objective income level. Kahneman and Tversky (1979) formalise this intuition in prospect theory, where people evaluate outcomes not in absolute terms but as deviations from a reference point. Outcomes below it are experienced as losses, outcomes above it as gains, and — crucially — losses loom larger than equivalent gains. This property, which they call *loss aversion*, is captured by a value function $v(\cdot)$ that is concave for gains, convex for losses, and steeper on the loss side, so that $|v(-x)| > v(x)$ for

[5] Treating ρ as time-invariant abstracts from evidence that risk aversion may vary with wealth or aggregate conditions. This is a modelling simplification; the cross-country variation documented by Falk et al. (2018) provides empirical support for treating cultural differences in $r(x)$ as relatively stable.

any $x > 0$. Easterlin's aspiration level maps directly onto this reference point: any shortfall of realised income below it is not merely a reduction in resources but a loss, felt more acutely than an equivalent gain would be welcome.

In the context of macroeconomic shocks, this suggests that the fertility response to income losses is asymmetric. A negative shock reduces fertility by more than an equivalent positive shock of the same size would increase it, because the former is experienced as a loss while the latter is not. Fertility is therefore expected to decline sharply during downturns but recover only partially and slowly during upturns, even when income returns to its pre-shock level. This asymmetry arises not from any change in objective circumstances but from the structure of preferences: households that have experienced a loss remain in a heightened state of sensitivity to further downside risk, which discourages the irreversible long-term commitment of having a child.[6] [7]

3.2 Previous research

When UNFPA grouped the determinants of fertility decisions into broad categories—such as health constraints, economic insecurity, uncertainty about the future, and lack of a supportive partner—it found that economic barriers were the most frequently cited, with more than half of respondents reporting that financial constraints limited their ability to have the number of children they desired (United Nations Population Fund, 2025). This emphasis on economic constraints aligns closely with the mechanisms formalised in the theoretical framework, where fertility responds to changes in household resources and the costs of child quantity and quality. Building on this insight, the empirical analysis translates these mechanisms into observable macroeconomic shocks. Inflation, unemployment, and GDP contractions correspond to the income and price shocks formalised in Equations (6)–(9): inflation raises the effective cost of both child quantity and quality while eroding real income, whereas unemployment and GDP contractions represent negative income shocks η_t. Existing research has examined how these macroeconomic shocks influence fertility, and a smaller but growing literature investigates how risk preferences shape reproductive behaviour. The next subsections review these strands of evidence and highlight the remaining gap: whether the fertility effects of macroeconomic shocks depend on underlying societal risk tolerance.

[6] This could explain why some studies find only a limited effect of actual inflation on fertility: households are not necessarily expected to react to the inflation rate itself, but rather to the gap between expected and realised economic conditions. While predictable inflation can be managed through savings and consumption smoothing, unexpected inflation disrupts long-term planning and can influence fertility decisions (Adesete and Piano, 2026).

[7] A full integration of prospect theory into the Becker–Tomes framework would require replacing the utility function with a reference-dependent value function, which lies beyond the scope of the present analysis; the argument here is therefore a theoretical interpretation rather than a structural derivation.

3.2.1 Macroeconomic shocks

Inflation As outlined in the theoretical background, inflation is expected to affect fertility primarily through its impact on income. By eroding real income, rising prices increase financial strain on households, potentially reducing their ability to afford childbearing and childrearing (Wu, 2024). Unexpected inflation can also disrupt fertility plans by heightening economic uncertainty: when prices rise beyond expectations, families face unclear prospects for wages, living costs, and borrowing conditions. This erosion of real purchasing power raises the perceived cost of raising children, prompting many households to delay or reconsider childbearing. Conversely, when inflation remains stable and predictable, families are less likely to alter their fertility decisions (Adesete and Piano, 2026).

Recent evidence further shows that inflation expectations themselves shape household behaviour in ways directly relevant for fertility. Coibion et al. (2021) demonstrate that experimentally induced increases in inflation expectations lead households to cut back on large durable purchases—including housing—while reallocating spending toward non-durables. Because housing acquisition and durable-goods investment are closely tied to household formation, this behavioural response provides a micro-level mechanism through which uncertainty about future economic conditions can delay fertility. Moreover, less-educated and lower-income households — those most vulnerable to real income erosion — tend to hold the most inaccurate and dispersed inflation beliefs, potentially amplifying uncertainty-driven postponement of childbearing among groups already facing higher economic constraints.

A complementary mechanism operates through household wealth. Evidence from the United States shows that increases in housing wealth and financial wealth can raise fertility by improving households' perceived long-term economic security. Affuso et al. (2022) find that both higher house prices and higher interest rates exert a positive effect on fertility, consistent with a wealth-effect interpretation in which households feel more financially able to support additional children when asset values or returns rise.

These findings reinforce the idea that inflation — by eroding real wealth or increasing uncertainty about future asset values — can indirectly depress fertility through its impact on household balance sheets.

Macroeconomic policy frameworks have been found to mitigate these uncertainty effects. Wu (2024) show that countries adopting inflation targeting (IT) experience a smaller decline in fertility rates compared to non-adopters. This suggests that, beyond achieving price stability, IT may help sustain higher fertility by anchoring inflation expectations and reducing the uncertainty that otherwise discourages long-term family planning.

Unemployment Of the major macroeconomic shocks studied in relation to fertility, unemployment has received the greatest attention in the literature. Unlike more abstract macroeconomic indicators, joblessness poses an immediate threat to the financial security of potential parents, often leading to a precautionary postponement of childbearing Sobotka et al. (2011). Its effects have been examined particularly in developed countries. Early evidence from Europe shows that unemployment has long been associated with delayed family formation: Adserà (2010) finds that the high and persistent unemployment of the 1980s and 1990s systematically postponed childbearing, including higher-order births, even after accounting for family-policy environments. This pattern persisted into the 2000s. Analysing the impact of the Great Recession across European countries between 2000 and 2010, Goldstein et al. (2013) document a robust negative relationship between rising unemployment and fertility, with the downturn's effects falling disproportionately on younger adults—suggesting that labour-market instability primarily discouraged the transition to first births. More recent cross-national evidence confirms the continued relevance of this mechanism. A 2023 OECD working paper examining labour-market conditions and family trends in 26 countries between 2002 and 2019 similarly finds a clear negative association between unemployment and total fertility rates (Fluchtmann et al., 2023). These studies show a consistent pattern across decades: higher unemployment depresses fertility in high-income settings, both by delaying entry into parenthood and by discouraging progression to additional births. Yet national aggregate unemployment figures may obscure important dynamics operating at the individual level, where the mechanisms linking economic conditions to fertility decisions can work quite differently. Nordic countries, with their rich register data, have allowed researchers to probe precisely this distinction. Hoem (2000) found that first-birth rates in Sweden between 1986 and 1997 moved in tandem with local employment levels — an intermediate unit of analysis between the national aggregate and the individual — with the response concentrated among women under 30, who leveraged their longer reproductive window to postpone childbearing during periods of local market volatility. Kristensen and Lappegård (2022) address the macro versus micro question most directly, finding that personal employment status outweighs aggregate municipal unemployment as a driver of fertility in Norway. The effects vary considerably: for childless individuals, personal unemployment is a persistent barrier to parenthood, while among existing parents its influence has weakened over time. There is also a gender dimension — unemployed fathers are less likely to have additional children, while unemployed mothers show the opposite pattern, possibly treating periods out of work as an opportunity for childbearing. Municipal unemployment, by contrast, has little bearing on first births but remains a negative predictor for higher-order ones, suggesting that aggregate economic conditions matter more for family expansion than for entry into parenthood.

Gross Domestic Product While previous sections examined the roles of inflation and unemployment, far fewer studies analyse fertility responses to aggregate indicators of economic performance (Sobotka et al., 2011). To situate short-run GDP fluctuations within a broader conceptual frame, it is useful to begin with the long-run relationship between economic development and fertility.

The demographic transition framework provides the historical backdrop against which cyclical shocks must be interpreted. First described by Thompson (1929) and later elaborated by Davis (1945), the theory of demographic transition captures the shift from regimes of high mortality and high fertility to low mortality and low fertility, driven by modernisation and rising living standards (Lesthaeghe, 2010).[8] This secular process established the foundational negative association between economic development and fertility that subsequent empirical work has sought to quantify. In the long run, Bampinas and Mavropoulos (2024) examine 144 years of data for eight advanced European economies and show that the GDP–fertility relationship is not simply monotonic but follows an inverse J-shape: fertility declines with income up to a certain level of development, beyond which the relationship reverses. When time-varying estimation methods are applied, the pattern becomes even more nuanced, following an N-shaped trajectory. These findings underscore that the GDP–fertility nexus evolves across historical and economic regimes. While essential for understanding broad demographic change, the demographic transition operates on a fundamentally different temporal scale from the present analysis. The transition describes structural, development-driven shifts unfolding over decades or generations, whereas this study focuses on short-run deviations from a country's own growth path. This distinction motivates the use of GDP contraction—defined as the negative of the annual change in real GDP per capita—rather than GDP levels as the relevant shock variable. GDP levels are dominated by persistent cross-country differences in development, whereas GDP contraction isolates transitory downward deviations, capturing precisely the cyclical dimension that the demographic transition literature leaves aside. Descriptive evidence supports the relevance of these short-run fluctuations. Sobotka et al. (2011) show that over the last three decades, declines in Total Fertility Rates were significantly more frequent during periods of economic contraction in 26 OECD countries. Even episodes of economic stagnation (GDP growth below 1%) were associated with falling fertility in two-thirds of observations. While the correlation between GDP and fertility often weakens in multivariate models—an outcome confirmed by both Sobotka et al. (2011) and Bampinas and Mavropoulos (2024)—this does not imply that economic cycles are irrelevant. Rather, it suggests that GDP alone may be too coarse to capture the specific channels through which recessions affect reproductive behaviour. Micro-level evidence from Sweden illustrates the role of income more directly. Using rich register data, Andersson (2000) show that women with low individual incomes or enrolled in education had systematically lower fertility, and that shifts in the composition of these groups contributed to the fertility decline of the 1990s. Yet individual characteristics alone were insufficient: macro-level factors—particularly social policy and aggregate female earnings—were also necessary to explain Sweden's fluctuating fertility.

[8] This demographic perspective fundamentally challenged the core assumptions of the Malthusian trap. Whereas Malthus predicted a positive relationship between income and fertility, empirical evidence from the demographic transition shows the opposite pattern: as income rises, fertility tends to decline across countries and over time (Fernandez-Villaverde, 2001). These historical and cross-national dynamics indicate that the Malthusian framework cannot account for modern demographic behaviour and performs poorly when confronted with empirical data (Ali, 2015).

3.2.2 Cultural Preferences and Willingness to Take Risks

While the literature reviewed so far highlights the importance of economic uncertainty, much less attention has been paid to the behavioural traits that shape how individuals respond to such uncertainty. Cross-country evidence reveals several systematic patterns in preference heterogeneity: women tend to exhibit greater impatience, lower tolerance for risk, and higher prosociality than men, and these relationships appear remarkably stable across national contexts. Risk-taking itself is strongly associated with meaningful behavioural and economic outcomes. Individuals reporting higher willingness to take risks are more likely to be self-employed or to engage in behaviours such as smoking, and at the aggregate level, societies with higher average risk-taking display greater entrepreneurial activity (Falk et al., 2018). These patterns suggest that cultural differences in risk tolerance may plausibly extend to fertility behaviour as well.

Direct evidence on this link remains limited, but the few existing studies point to a nuanced relationship. Schmidt (2008) shows that women with higher risk tolerance tend to postpone marriage, a finding consistent with both search-based models of partner choice and risk-pooling interpretations.[9] Risk preferences also shape fertility timing in ways that vary across marital status, education, and age. Among both married and unmarried women, greater willingness to take risks is associated with earlier childbearing at younger ages, likely reflecting lower propensity to engage in effective contraceptive behaviour. Yet this pattern reverses for college-educated, unmarried women nearing the end of their reproductive years: those who are more risk-tolerant are more inclined to delay motherhood relative to their more risk-averse peers, making them less likely to have a first birth (Schmidt, 2008).[10] More recent evidence from Italy reinforces the relevance of risk preferences for fertility behaviour. Bellani and Arpino (2021) find that individuals who are less tolerant of risk are more likely to have a first and a second child during the observation period than their more risk-tolerant counterparts, suggesting that risk aversion may encourage earlier or more certain family formation. Despite these insights, the existing literature remains fragmented: studies either examine how risk preferences shape fertility in general or how macroeconomic shocks affect fertility on average. What is missing is an analysis that brings these two strands together. In particular, no research has examined whether underlying cultural differences in risk tolerance moderate the fertility response to macroeconomic shocks at the cross-country level.

3.3 Hypothesis

Drawing on the theoretical framework and the empirical literature reviewed above, this study tests three main hypotheses:

[9] When models allow for heterogeneity in risk preferences, more risk-tolerant individuals set a higher reservation standard for an acceptable partner and are therefore willing to search longer before matching. From a risk-pooling perspective, individuals less concerned about income fluctuations have weaker incentives to marry early, leading to later union formation (Schmidt, 2008).

[10] Schmidt's risk measure—derived from hypothetical income gambles—captures a general risk-taking propensity that extends beyond purely financial behaviour.

H1: *Macroeconomic shocks negatively affect Total Fertility Rates.* Building on the quantity–quality trade-off model, adverse economic conditions are expected to reduce fertility through distinct but complementary mechanisms: inflation erodes real purchasing power and raises the shadow cost of childrearing, unemployment directly reduces household income, and GDP contraction signals broader declines in aggregate economic prospects. Each shock reduces effective resources available for childbearing, thereby depressing fertility, albeit potentially to differing degrees across shocks and contexts.

H2: *Cultural risk tolerance moderates the negative effect of macroeconomic shocks on fertility.* The theoretical framework identifies a precautionary channel through which the fertility response to a given shock is amplified in more risk-averse populations: a higher degree of risk aversion increases the precautionary wedge $\phi(\rho, \sigma_\eta^2)$, reducing effective resources beyond the mechanical income loss. Conversely, households in high-risk-tolerance societies are expected to attenuate their fertility response to transient shocks, as the precautionary motive is weaker. Accordingly, the interaction between each macroeconomic shock and the cultural risk-taking index is expected to carry a positive coefficient: since the baseline effect of shocks on fertility is negative, a positive interaction term implies that this negative effect is attenuated — that is, smaller in absolute value — in more risk-tolerant societies.

H3: *The moderating role of risk tolerance is concentrated among middle- and high-income countries.* The precautionary channel through which risk aversion amplifies the fertility response to uncertainty presupposes that households have sufficient resources to engage in precautionary behaviour and that labour-market instability constitutes a genuine threat to family planning. Both conditions are most plausible in middle- and high-income settings. In lower-income countries, fertility decisions are more heavily shaped by structural factors — such as limited contraceptive access, marriage market norms, and son preference — that operate largely independently of macroeconomic fluctuations and cultural risk preferences. The moderating role of risk tolerance is therefore expected to be statistically significant primarily among middle- and high-income countries, where the optimising logic underlying the precautionary channel is most likely to hold.

4 Data

The study uses a cross-country panel dataset covering 76 countries over the period 2000-2023. It is worth noting that each observation is measured at the country-level. Given the macro-level approach of the analysis, it is not possible to estimate whether these shocks might affect individuals differently.[11] The present section will discuss the variables chosen for the analysis, their sources and the data cleaning procedures.

[11] Shocks may operate differently at the micro and macro level. For example, unemployment can—under certain household conditions such as sufficient savings—create incentives to have a child due to the additional time available to individuals. This limitation is widely acknowledged in the literature on macro-level fertility trends (see, for instance, Fluchtmann et al. (2023)). Nonetheless, the aim of this study is to provide a cross-country perspective, for which a macro-comparative approach is appropriate.

4.1 Variable Selection, Assumptions and Datasets

4.1.1 Macroeconomic Shocks and Total Fertility Rates

The analysis considers three macroeconomic shocks discussed in Section 3.2: headline consumer price inflation, unemployment as a percentage of the total labour force, and contraction of GDP per capita relative to the previous year. While data on inflation and unemployment were retrieved directly from the World Bank Database (Ha et al. (2023); World Bank (2026*d*)), the GDP contraction variable was derived from World Bank data on GDP per capita (World Bank (2024)). *GDP growth* was calculated as the annual percentage change in real GDP per capita, according to the following formula:

$$\text{GDP Growth}_t = \frac{\text{GDP per Capita}_t - \text{GDP per Capita}_{t-1}}{\text{GDP per Capita}_{t-1}} \tag{11}$$

This measure captures how much output per person has grown or shrunk relative to the previous year. To ensure directional consistency across all indicators, this measure was subsequently transformed into **GDP contraction** by reversing its sign (multiplying by -1). This adjustment ensures that higher values in every indicator consistently represent a worsening of the economic environment, effectively acting as a positive shock to macroeconomic uncertainty.

It is worth clarifying why GDP contraction – defined as the negative of the annual growth rate of real GDP per capita – is preferred over the level of GDP per capita as a measure of macroeconomic deterioration. The three shock variables considered in this analysis are not symmetric in their construction: inflation is itself a rate of change (the percentage increase in the price level), and unemployment is a flow measure expressed as a share of the labour force, both of which already capture current economic conditions rather than structural levels. GDP per capita, by contrast, is a stock variable whose absolute level is dominated by persistent cross-country differences in development rather than by transitory adverse shocks. Using it directly — even within a fixed-effects framework — would conflate long-run income trajectories with short-run deteriorations, and would exacerbate the multicollinearity concerns that also motivate its exclusion as a control variable. GDP contraction, being the negation of the annual log-difference in real GDP per capita, is instead intrinsically a first-difference measure: it captures sudden downward deviations from a country's growth path rather than its structural income level, and is therefore the most internally consistent operationalisation of a negative output shock within the present framework. In this respect, GDP contraction is conceptually the closest of the three indicators to a pure shock measure, since it is already expressed as a change rather than a level.

The main dependent variable is the natural logarithm of each country's Total Fertility Rate (TFR) in a given year. The TFR measures the average number of children a woman would have over her lifetime if she were to experience the age-specific fertility rates prevailing in a given

country and year. It is the most widely used summary measure of fertility in the demographic literature, as it is independent of the age structure of the population and therefore allows for meaningful comparisons across countries and over time.
Data on Total Fertility Rates are drawn from the World Bank Database (World Bank, 2026*a*), while data on Age-specific fertility rates used for a heterogeneity analysis (Section 5.1.1) are taken from the United Nations Population Division (United Nations, 2024).

4.1.2 Risk-Taking

Social scientists have long examined how individuals' attitudes toward risk shape their choices and behaviours, motivating the inclusion of experimental survey items designed to elicit respondents' risk preferences. These questions typically measure willingness to engage in a series of gambles over lifetime income, thereby generating quantitative indicators of individual risk tolerance (Schmidt, 2008). One prominent example of such an instrument is the Global Preference Survey (GPS), an experimentally validated dataset measuring economic preferences—including risk and time preferences, reciprocity, altruism, and trust—across 76 nations (Falk et al., 2016), which is used here as the source of cross-country data on cultural risk attitudes. The GPS is uniquely suited for cross-country comparative research, as it provides nationally representative samples covering approximately 90% of the global population and income.

With a median of 1,000 respondents per country (totaling roughly 80,000 observations), the survey allows for the construction of robust country-level averages. In this study, the risk-taking measure is utilized as a standardized numerical index. The data collection relied on strict standardization controls to maintain cross-national data integrity. Linguistic consistency was achieved through an iterative translation and back-translation protocol designed to preserve question meaning across different cultural contexts. Furthermore, financial variables within the survey were adjusted according to country-specific median household incomes, ensuring that the economic stakes of the questions remained uniform. Methodological bias was also controlled by conducting the majority of interviews via a single mode—face-to-face—though telephone surveys were utilized as an exception in 18 nations (Falk et al., 2018).

The survey combines two complementary items into a single index. The first is a qualitative self-assessment, asking respondents to rate their general willingness to take risks on a scale from 0 ("completely unwilling") to 10 ("very willing"). The second is a quantitative measure elicited through a staircase procedure consisting of five interdependent binary choices, in which respondents repeatedly choose between a lottery — a 50% chance of receiving an amount x or nothing — and a varying sure payment y. Respondents who prefer the lottery face a higher sure payment next; those who prefer the sure payment face a lower one.[12] The wording of both items is reproduced in Table 1. The final index is a weighted combination of both items, with weights

[12] The quantitative item more closely resembles incentivized laboratory measures by holding stakes, probabilities, and information conditions constant, while the qualitative item retains independent explanatory power for experimentally elicited behavior.

derived from an experimental validation study (Falk et al., 2018).[13] The precise branching structure of the elicitation procedure is illustrated in Falk et al. (2018, Appendix Figure 2).

Table 1: GPS Risk Preference Survey Items

Qualitative	Please tell me, in general, how willing or unwilling you are to take risks. Please use a scale from 0 to 10, where 0 means "completely unwilling to take risks" and a 10 means you are "very willing to take risks".
Quantitative	Please imagine the following situation. You can choose between a sure payment of a particular amount of money, or a draw, where you would have an equal chance of getting amount x or getting nothing. What would you prefer: a draw with a 50% chance of receiving amount x, and the same 50% chance of receiving nothing, or the amount of y as a sure payment? This question was repeated five times using a staircase procedure, where each answer determined the next value of y offered.

Note: Items are reproduced from Falk et al. (2018, Online Appendix). The qualitative item uses an 11-point scale. The quantitative item was elicited via an adaptive staircase procedure of five binary choices. The final risk preference index is a weighted combination of both items, with weights derived from an experimental validation study.

A key strength of the GPS risk module is, indeed, its use of both qualitative and quantitative items, rather than relying on a single survey question as is common in much of the literature. This multidimensional design is particularly valuable given evidence from Hudson et al. (2005) that risk-aversion measures vary substantially across elicitation methods, contexts, and gamble structures. Because no single question can fully capture this complexity, combining multiple formats reduces method-specific bias and yields a more robust measure of individual risk preferences.

Both the explanatory variables and the dependent variable are time-varying. The risk tolerance measure, by contrast, reflects a 2012 mid-period snapshot.[14] Following Bellani and Arpino (2021), who argue that risk preferences manifest during adulthood as stable, trait-like dispositions rather than situational fluctuations, I treat risk tolerance as time-invariant. This assumption also has an econometric justification: a time-varying measure of risk attitudes would likely become endogenous, since individuals may update their willingness to take risks in response to macroeconomic shocks. Treating national risk attitudes as a persistent cultural trait thus serves a dual purpose: it reflects the theoretical view that such dispositions are stable over time, while avoiding the endogeneity concerns that a time-varying measure would introduce.

4.1.3 Controls

Three control variables are considered: percentage of urban population, female percentage of school enrollment in tertiary education and infant mortality rate, all drawn from the World Bank Database (World Bank (2026*e*), World Bank (2026*c*), World Bank (2026*b*)).

[13] The final risk preference index assigns a weight of 0.527 to the qualitative self-assessment and 0.473 to the quantitative staircase measure (Falk et al., 2018).

[14] Using a single observation from the centre of the period minimises the potential measurement error that would arise from using a value at the very beginning or end of the timeframe.

Urbanization is included as a control for its consistent association with fertility levels, especially in low-income settings. This pattern aligns with the *evolutionary perspective*, which holds that high-density environments encourage "slow" life-history strategies—delayed reproduction, fewer births, and higher investment per child—because competition for resources intensifies (Lee, 2025). In particular, urbanization in contemporary Sub-Saharan Africa is closely linked to the ongoing fertility transition, mirroring the historical experience of Western economies as they moved from a (post-)Malthusian regime toward modern sustained growth (Flückiger and Ludwig, 2017).

Female enrollment in tertiary education is also included for several interrelated reasons. At the most fundamental level, education shapes the opportunity costs women face when considering childbearing. Higher-educated women invest substantially more in their human capital, making early or unintended pregnancies particularly costly.[15] Under conditions of effective fertility control, this calculus leads highly educated women to systematically postpone childbearing to later stages of life (Schmidt, 2008). Conversely, for women with lower levels of education, the foregone earnings associated with early childbearing are comparatively modest, and the incentive to delay parenthood is correspondingly weaker. This gradient in opportunity costs implies that the aggregate level of female educational attainment shapes the baseline fertility schedule of a population, independently of any macroeconomic fluctuation. This dynamic becomes especially salient in the context of economic downturns. Economic theory on the opportunity costs of motherhood suggests that highly educated women are particularly prone to treating childbearing as a high-stakes decision under uncertainty: since their career investments are larger, the perceived risk of permanent income loss during a recession is amplified, reinforcing the tendency to delay parenthood (Sobotka et al., 2011). Failing to account for cross-country differences in female educational attainment would therefore conflate variation in baseline fertility postponement with variation attributable to the macroeconomic shock itself — precisely the effect of interest in this study. Beyond its direct effect on the cost-benefit calculus of childbearing, female education also operates through its influence on marriage formation, which in turn affects fertility. Contrary to the economic-independence argument — which would predict that more financially autonomous women are less inclined to marry — micro-level evidence consistently shows that, once school enrollment is controlled for, higher-educated women actually display a greater propensity to marry (Oppenheimer et al., 1997). This suggests that education does not simply delay partnership formation but reshapes its quality and timing in ways that feed back into fertility decisions. Taken together, these mechanisms indicate that female tertiary enrollment captures a broad set of structural factors — opportunity costs, career investment, and partnership patterns — that must be held constant to isolate the effect of macroeconomic shocks on fertility behavior.

[15] Having children can disrupt career trajectories and significantly reduce lifetime income.

The inclusion of infant mortality as a regressor in the fertility equation is theoretically grounded in the demographic literature. Several competing theories postulate a causal link between these two variables (Chowdhury, 1988). The *demographic transition theory*,[16] already cited, and *choice theory* both imply a lagged causality running from infant mortality to fertility, while the *Ricardian theory* reverses this direction, arguing that higher fertility causes subsequent increases in mortality. The *modern economic theory of population*, instead, posits a feedback relationship between the two variables, treating both as jointly determined.[17] Empirical evidence across thirty-five developing countries lends the broadest support to demographic transition theory, with most countries exhibiting a positive and long-lagged causal effect of infant mortality on fertility, while evidence in favour of the Ricardian theory remains negligible (Chowdhury, 1988). The inclusion of infant mortality in the present model is therefore consistent with the dominant theoretical and empirical consensus.

In the present sample, this is particularly relevant for low-income countries, many of which remain in earlier stages of the demographic transition where infant mortality is still high and the mortality–fertility feedback is most active (Chowdhury, 1988). Failing to control for infant mortality would therefore risk attributing to macroeconomic shocks variation in fertility that instead reflects differences in child survival probabilities across countries and over time.

In the GPS data, risk preferences show no meaningful association with agricultural suitability, in contrast to other preferences such as trust or patience, which exhibit significant geographic correlations (Falk et al., 2018). Given this lack of empirical relevance — and the fact that geographical determinants of risk attitudes or fertility repsonses to macroeconomic shocks are not commonly emphasized in the existing literature — geographical variables were not included as controls in the analysis.

GDP per capita and female unemployment are not considered as controls due to possible issues of multicollinearity with the three macroeconomic shocks under study.[18] However, a heterogeneity analysis will consider whether the effects vary based on the countries level of GDP per capita.

[16] A defining feature of the transition is the temporal sequence in which mortality decline precedes fertility decline: as child survival probabilities rise, households gradually revise downward the number of births required to achieve a desired number of surviving children (Thompson, 1929; Davis, 1945; Lesthaeghe, 2010).

[17] According to demographic transition theory, declining infant mortality—driven by industrialisation, urbanisation, and improved medical practices—leads with a lag to declining fertility, as the need to compensate for child deaths gradually diminishes. Choice theory offers two related mechanisms: the *child replacement hypothesis*, whereby parents respond to a child's death by increasing subsequent births, and the *child survival hypothesis*, whereby the perceived probability of a child surviving to adulthood shapes the desired number of births. The Ricardian theory instead argues that rising wages above subsistence level improve health and life expectancy, raising the net reproduction rate and thus fertility, which in turn increases mortality with a lag. The modern economic theory of population distinguishes itself from choice theory in that both fertility and infant mortality are treated as choice variables, jointly determined by parental resource allocation (Chowdhury, 1988).

[18] Moreover, female unemployment may be a weaker indicator of economic constraints than overall unemployment, particularly in settings where a male-breadwinner model still structures household decision-making. In such contexts, men's employment prospects remain the primary determinant of perceived economic security, meaning that fluctuations in general unemployment rates better capture the shocks relevant for fertility behaviour.

4.2 Descriptive Statistics and Diagnostic Checks

4.2.1 Summary Statistics and Correlation Structure

This section presents the descriptive statistics for the main variables used in the analysis, examines the correlation structure among regressors, and reports diagnostic checks for multicollinearity.

Table A.1 reports the list of countries included in the analysis, along with their average levels of Total Fertility Rates, Inflation, Unemployment, GDP contraction across the 23 years and their risk taking level.

Table 2 presents the descriptive statistics for the main variables used in the analysis, based on a sample of 1,824 observations. The statistics reported refer to the raw, uncleaned sample prior to any data cleaning procedure. The average fertility rate stands at 2.46 births per woman (SD = 1.25), ranging from 0.72 to 7.57, reflecting substantial demographic heterogeneity across countries and over time. Headline consumer price inflation displays the most striking distributional features in the sample: while the mean is 56.95%, the standard deviation is exceptionally large (1,601.17), and the maximum reaches above 65,000%, pointing to the presence of extreme hyperinflationary episodes that heavily skew the distribution. The unemployment rate averages 7.84% of the total labour force (SD = 5.36), with values ranging from 0.12% to 34.01%, while negative GDP growth has a mean close to zero (−0.06) and a narrow range (−1.45 to 0.65), indicating that severe contractions are relatively rare events in the sample. Willingness to take risks is on average near zero (mean = 0.01, SD = 0.30). Urban population accounts for 61.32% of the total on average (SD = 20.77), spanning from 11.64% to 95.32% and capturing a wide spectrum of urbanisation levels. Female tertiary school enrolment averages 53.42% on a reduced subsample of 1,313 observations, with a notably wide range (0.18-171.79%).[19] The reduced sample size reflects missing data for a number of country-year cells; however, an inspection of the missing observations reveals no systematic pattern across countries, income groups, or time periods, suggesting that the data are missing at random and that the reduced subsample is unlikely to introduce any meaningful selection bias into the analysis. The infant mortality rate averages 20.61 deaths per 1,000 live births (SD = 20.92), ranging from 1.60 to 138.50, pointing to sharp cross-country disparities in health outcomes. Finally, GDP per capita averages approximately $14,000 but exhibits a very high standard deviation and a a cross-country income disparity exceeding $100,000, underscoring the large income inequalities present in the sample.

[19] Values exceeding 100% reflect the gross enrolment ratio methodology, which counts students regardless of age relative to the official age group.

Table 2: Summary Statistics

	Mean	SD	Min	Max	Count
Fertility rate, total (births per woman)	2.46	1.25	0.72	7.57	1824
Headline Consumer Price Inflation	56.95	1601.17	-72.73	65374.08	1824
Unemployment, total (% of total labor force)	7.84	5.36	0.12	34.01	1822
Negative GDP growth	-0.06	0.14	-1.45	0.65	1823
Will. to take risks	0.01	0.30	-0.79	0.97	1824
Urban population (% of total population)	61.32	20.77	11.64	95.32	1824
School enrollment, tertiary, female (% gross)	53.42	32.76	0.18	171.79	1313
Mortality rate, infant (per 1,000 live births)	20.61	20.92	1.60	138.50	1824
GDP per capita (current US dollars)	14264.64	17673.40	138.71	100623.55	1824

Note: The table reports summary statistics for the raw (uncleaned) dataset used in the analysis. Means, standard deviations, minimums, maximums, and observation counts are computed using the full available sample for each variable. The total number of observations ranges from 1,822 to 1,824, except for female tertiary school enrolment, which is available for a reduced subsample. Values for inflation and GDP per capita exhibit substantial dispersion due to the presence of extreme observations, including episodes of hyperinflation and large cross-country income differences.

Table A.2 in the Appendix reports the Pearson correlation matrix for the full set of covariates, where r denotes the pairwise correlation coefficient ranging from -1 (perfect negative linear association) to $+1$ (perfect positive linear association). Several patterns are worth noting. Among the macroeconomic shocks, GDP contraction shows no meaningful association with TFR ($r = -0.0154$), whereas Inflation is positively and significantly correlated with Fertility ($r = 0.106^{***}$), and Unemployment is negatively and significantly correlated with it ($r = -0.168^{***}$).

Risk taking is positively correlated with TFR ($r = 0.276^{***}$), and with both inflation and unemployment, while displaying no systematic correlation with GDP contraction.

Turning to the controls, female education is strongly negatively correlated with TFR ($r = -0.690^{***}$), consistent with the canonical quantity–quality trade-off in the fertility literature, and is also negatively associated with infant mortality ($r = -0.760^{***}$). Infant mortality itself is strongly positively correlated with TFR ($r = 0.862^{***}$), while urbanisation is negatively correlated with TFR and positively correlated with female education and GDP per capita.

Finally, GDP per capita is negatively correlated with TFR, in line with the cross-country stylised facts on demographic and economic development. These patterns indicate that the control variables chosen are substantively related to fertility and to broader structural development processes. Their inclusion in the regression models is therefore justified, as omitting them would risk confounding the estimated effects of macroeconomic shocks with underlying demographic and socioeconomic conditions. Importantly, while several regressors are correlated with one another, no pair of *independent* covariates exhibits a correlation that would raise serious multicollinearity concerns in the regression models.[20]

Table A.3 reports the Variance Inflation Factors (VIFs) for all regressors included in the

[20] As a rule of thumb, correlations above $|0.8|$ between *regressors* may be considered potentially problematic. The high correlation between TFR and infant mortality does not pose a multicollinearity issue, as TFR is the dependent variable.

baseline specification. Including this diagnostic alongside the correlation matrix is important because the latter only captures pairwise associations, whereas VIFs assess the combined linear dependence among regressors in the multivariate setting. All VIF values are well below conventional thresholds (typically 5 or 10), with a mean VIF of 1.87, indicating that multicollinearity is not a concern. This confirms that the estimated coefficients are not inflated by redundant information among the explanatory variables, and that the controls—although correlated with fertility and development indicators—contribute distinct explanatory power to the model.

4.2.2 Distributional Analysis and Model Specification Diagnostics

Macroeconomic shocks The preliminary analysis of the macroeconomic indicators revealed significant distributional challenges that necessitated specific data-cleaning and transformation procedures.

As shown in Figure A.1, the initial boxplots highlight the presence of extreme outliers, particularly within the inflation variable (HCPI),[21], values that were also visible in the descriptive statistics. To ensure model stability and prevent these extreme observations from exerting undue leverage, a threshold was applied to exclude any inflation values exceeding 100%. This led to the loss of 16 observations.[22]

Table 3 reports the new descriptive statistics, while Table A.4 shows the new country list.

Table 3: Summary Statistics Without Inflation Outliers

	Mean	SD	Min	Max	Count
Fertility rate, total (births per woman)	2.46	1.25	0.72	7.57	1808
Headline Consumer Price Inflation	6.26	9.11	-72.73	98.55	1808
Unemployment, total (% of total labor force)	7.85	5.38	0.12	34.01	1806
Negative GDP growth	-0.06	0.14	-1.45	0.65	1807
Will. to take risks	0.01	0.30	-0.79	0.97	1808
Urban population (% of total population)	61.25	20.70	11.64	95.32	1808
School enrollment, tertiary, female (% gross)	53.35	32.68	0.18	171.79	1312
Mortality rate, infant (per 1,000 live births)	20.48	20.87	1.60	138.50	1808
GDP per capita (current US dollars)	14362.20	17718.55	138.71	100623.55	1808

Note: The table reports summary statistics after excluding extreme inflation observations (HCPI >100%) to mitigate the influence of outliers. The trimming procedure resulted in a reduced sample of 1,808 observations (with minor variation across variables due to data availability). All other variables remain unchanged and are defined as in the main text.

Figure A.2 illustrates the Kernel density of the "cleaned" variables. The visual alignment of the shocks reveals different distributional patterns: while GDP contraction exhibits a slight left-skewness due to periods of robust economic growth, unemployment shows pronounced right-skewness, representing the "long tails" of severe economic crises. This visual evidence of

[21] Values for Venezuela occasionally reached levels exceeding 60, 000%, with the highest inflationary episode taking place in 2018

[22] The countries that in the dataset experienced episodes of hyperinflation were, beside Venezuela, Argentina and Zimbabwe.

non-normality is further confirmed by the histograms with normal overlays in Figure A.3. The mismatch between the empirical and theoretical distributions—especially the excess kurtosis in inflation and the differing skewness across indicators—motivates the use of **Z-score standardization**. By standardizing these variables, the analysis ensures that all macroeconomic shocks are centered and measured on a comparable scale.

Total Fertility Rates Figure A.4 displays the distribution of the Total Fertility Rate (TFR) in its raw and log-transformed forms. The left panel shows the raw TFR distribution, which is characterised by a pronounced right skew: the bulk of observations is concentrated between 1 and 3 births per woman, while a long right tail extends towards values of 6 and above, corresponding to high-fertility countries. The superimposed normal curve makes clear that the raw distribution departs substantially from normality, with the empirical density far exceeding the normal benchmark at low values and falling well below it in the upper tail.

The right panel presents the distribution of the log-transformed TFR. The log transformation substantially reduces the right skew and compresses the upper tail, yielding a distribution that is markedly more symmetric and closer to the normal benchmark, although a mild rightward asymmetry persists. The log TFR values are concentrated in the range of 0 to 1.5, which covers the vast majority of the sample.

These distributional features provide a clear justification for using the log-transformed TFR as the dependent variable in the regression analysis. Beyond improving the approximation to normality, a desirable property for inference, the log specification also yields a convenient economic interpretation, as estimated coefficients can be read as approximate percentage changes in the fertility rate associated with a one-unit change in each regressor.

Willingness to take risks and TFR Figure 3 presents two complementary graphical analyses of the relationship between average willingness to take risks and the Total Fertility Rate (TFR) across the 76 countries considered in the study. Panel A displays a LOWESS (Locally Weighted Scatterplot Smoothing) curve, a non-parametric estimator that fits a smooth line through the data by running a series of local weighted regressions, without imposing any assumption on the functional form of the relationship. Panel B contrasts this non-parametric fit with a standard linear regression line, allowing the reader to assess how well a simple linear approximation captures the underlying pattern in the data.

Panel A reveals a generally positive association between risk-taking propensity and fertility, suggesting that countries characterised by higher average willingness to take risks tend to exhibit higher fertility rates. Theoretically, this is consistent with the view of childbearing as a long-term, high-commitment decision: in societies where risk-taking is culturally more accepted, individuals may be more inclined to start a family despite the inherent uncertainties of the future. However, the LOWESS curve also highlights a non-linear pattern: the relationship is relatively flat or slightly negative at low levels of risk-taking, before rising more steeply across the middle of the distribution, and flattening again toward the upper tail.

Panel B further illuminates this by juxtaposing the LOWESS curve with a linear fit. While the overall positive trend is confirmed, the non-parametric estimator suggests that the relationship is not perfectly linear across the full distribution. Rather than motivating a particular functional form for risk-taking, this preliminary evidence simply supports its inclusion in the model as a moderating variable: if risk attitudes shape fertility levels in the cross-section, they may equally condition how fertility responds to macroeconomic shocks, which is the hypothesis tested in the regression analysis that follows.

Figure 3: Relationship between Willingness to Take Risks and TFR

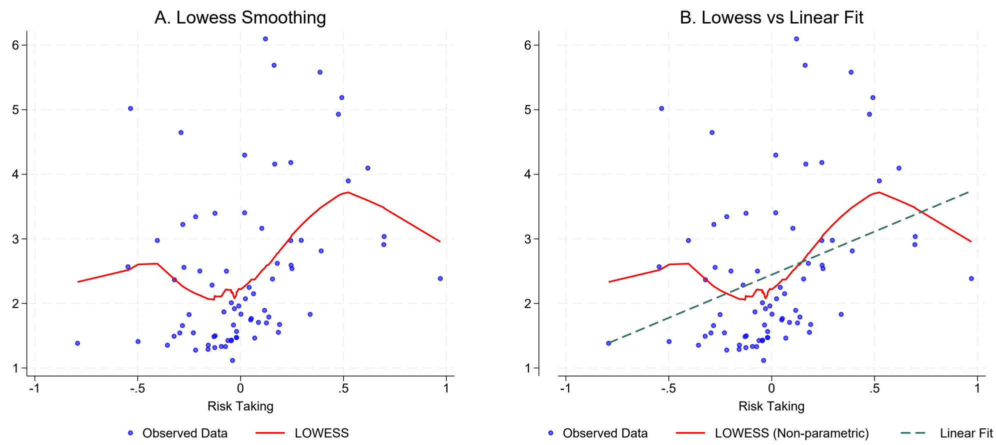

Note: Each observation represents a country average, obtained by collapsing the panel to a single mean value per country over the period 2000–2023. Risk-taking is time-invariant by construction (GPS, 2012 wave); TFR is averaged across all available years. The LOWESS bandwidth follows Stata's default (0.8).

5 Methodology

This section describes the empirical strategy used to examine the relationship between macroeconomic shocks, cultural risk tolerance, and fertility. The analysis proceeds in three steps. First, a baseline two-way fixed-effects model is estimated to identify the direct effects of each macroeconomic shock on the Total Fertility Rate. Second, an interaction model is introduced to test whether cultural risk tolerance moderates these effects. Third, both specifications are re-estimated separately by income group to assess whether the results mask systematic heterogeneity across development levels. All estimations were carried out in Stata 19.5.

5.1 Baseline

To examine the relationship between macroeconomic instability and fertility dynamics, I first estimate a baseline model that assesses the impact of lagged economic shocks. The specification is defined as follows:

$$\ln(TFR_{it}) = \beta_1 Shock_{i,t-1} + \mathbf{X}'_{it-1}\gamma + \alpha_i + \delta_t + (\alpha_i \times t) + \epsilon_{it} \tag{12}$$

where $\ln(TFR_{it})$ represents the natural logarithm of the Total Fertility Rate for country i in year t. The use of a log-level specification allows the coefficients to be interpreted as semi-elasticities. $Shock_{i,t-1}$ denotes the standardized macroeconomic shock (e.g., unemployment rate or GDP contraction) lagged by one year. Macroeconomic shocks are lagged to mitigate concerns about contemporaneous reverse causality and to ensure that the shocks can plausibly affect observed fertility outcomes, given that fertility decisions are typically made several months in advance. The inclusion of lagged variables of shocks to measure their effects on fertility choices is considered a standard practice in the literature (e.g. Fluchtmann et al. (2023), Wu (2024), Goldstein et al. (2013)).

$\mathbf{X}'_{it-1}$ is a lagged vector of time-varying controls, while α_i and δ_t represent country and year fixed effects, respectively, and ϵ_{it} is the idiosyncratic error term, assumed to be uncorrelated with the regressors conditional on the fixed effects.

The two-way fixed-effects (TWFE) specification with country and year fixed effects, plus country-specific linear time trends $(\alpha_i \times t)$ is also an established practice in the fertility-macroeconomics literature (e.g., Goldstein et al., 2013; Fluchtmann et al., 2023; Wu, 2024). This approach accounts for time-invariant country-level unobservables, global temporal shocks, and nation-specific fertility paths such as long-term demographic transitions or structural institutional shifts. Identification is therefore driven entirely by within-country deviations of each shock from a country's own historical trajectory — which also addresses the concern raised by Sobotka et al. (2011), who argues that fertility dynamics are best understood as a function of changes in economic conditions rather than levels. The log-level specification of the dependent variable further reinforces this, as coefficients capture proportional changes in the TFR associated with fluctuations in the shock variables rather than cross-sectional differences in fertility levels.

5.1.1 Age-Specific Fertility Rates

To further unpack potential heterogeneity across cohorts, the baseline model is also estimated separately for age-specific fertility rates (ASFRs), defined as the number of births to women in a given age group relative to the total female population in that group:

$$ASFR_{i,a,t} = \frac{B_{i,a,t}}{N_{i,a,t}} \tag{13}$$

where $B_{i,a,t}$ denotes the number of births to women aged a in country i during year t, and $N_{i,a,t}$ the corresponding female population. The specification follows Equation 12 in all other respects, with $\ln(ASFR_{i,a,t})$ as the dependent variable estimated separately for each age group.

5.2 Risk Taking Interaction

To test the core hypothesis regarding the moderating role of cultural traits, I introduce an interaction term between the lagged shock and the standardized, time-invariant risk-taking index ($Risk_i$):

$$\ln(TFR_{it}) = \beta_1 Shock_{i,t-1} + \beta_2(Shock_{i,t-1} \times Risk_i) + \mathbf{X}'_{it-1}\gamma + \alpha_i + \delta_t + (\alpha_i \times t) + \epsilon_{it} \quad (14)$$

In this specification, both *Shock* and $Risk$ are standardized to have a mean of zero and a standard deviation of one, facilitating the interpretation of the interaction term. Since the specification includes country fixed effects, the main effect of the risk-taking variable cannot be separately identified, as it is time-invariant and therefore perfectly collinear with the country dummies. The coefficient of interest is β_2: if cultural risk tolerance buffers against economic uncertainty, we expect $\beta_1 < 0$, confirming that shocks depress fertility, and $\beta_2 > 0$, indicating that this negative effect is attenuated in societies with higher average risk tolerance. A positive and statistically significant β_2 would imply that the fertility cost of an economic shock is smaller in absolute terms in high-risk-tolerance countries.

5.3 Heterogeneity Analysis by Income Level

To investigate whether the baseline results mask systematic heterogeneity across development levels, the sample is partitioned into three groups based on each country's time-averaged GDP per capita. Specifically, for each country i the average real GDP per capita over the full sample period is computed as $\overline{\text{GDP}}_i = T_i^{-1} \sum_t \text{GDPperCapita}_{it}$. The distribution of these country-level averages is then divided at the 33rd and 67th percentiles, yielding three balanced terciles: *Low GDP p.c.*, *Medium GDP p.c.*, and *High GDP p.c.*, each containing approximately one third of the 76 countries in the sample. The list of countries by income group is reported in Table B.1.

All baseline and interaction specifications described in Sections 12 and 14 are re-estimated separately within each sub-sample, retaining country fixed effects, country-specific linear trends, year fixed effects, and Driscoll–Kraay standard errors with two lags. This stratification allows for an assessment of whether the precautionary channel — and the moderating role of cultural risk tolerance — operates differently across development contexts, as hypothesised in H3. Results are reported in Section 6.3.

5.4 Standard Errors

Panel data are appealing in economic analysis since they align with theories of individual behavior and provide detailed individual-level information (Bhargava et al., 1982). However, panel data estimators are subject to two forms of error-term misspecification that, if left uncorrected, render ordinary least squares (OLS) and within-group fixed-effects standard errors inconsistent: *serial autocorrelation* within cross-sectional units and *cross-sectional dependence* across units (Bhargava et al., 1982; Driscoll and Kraay, 1998; Hoechle, 2007).

Diagnosing serial autocorrelation Serial autocorrelation in panel data arises when the error terms of a given unit are correlated across time periods, as a result of possibly failing to include all relevant variables (Bhargava et al., 1982). To detect first-order serial correlation in the residuals of the fixed-effects estimator, I employ the Modified Durbin–Watson statistic by (Bhargava et al., 1982) and the Baltagi–Wu Locally Best Invariant (LBI) test (Baltagi and Wu, 1999). Under the null hypothesis of no serial correlation, both statistics are expected to be approximately equal to 2; values markedly below 2 provide evidence of positive first-order autocorrelation. As reported in Table A.5, all five model specifications yield Durbin–Watson statistics in the range [0.12, 0.20] and LBI statistics in the range [0.32, 0.51], far below the benchmark value of 2. The null hypothesis of no serial autocorrelation is therefore strongly rejected across all specifications.

Diagnosing cross-sectional dependence Cross-sectional, or "spatial", dependence is a problematic feature of many panel data sets in which the cross-sectional units are not randomly sampled (Driscoll and Kraay, 1998). It emerges when unobserved common factors affect multiple units simultaneously, inducing contemporaneous correlation in the error terms — a feature that is especially relevant in macroeconomic panels where countries are subject to shared global shocks (Pesaran, 2021). I test for cross-sectional dependence using the Pesaran (2021) CD test, which is valid for both large-N and large-T panels and is robust to structural breaks and non-stationary series. Because the CD test requires a minimum number of overlapping time periods for each pair of cross-sectional units ($T_{ij} > 3$) to ensure the validity of the standard normal approximation (Pesaran, 2021), it could only be implemented for the baseline model without controls, which is the only specification with sufficient temporal overlap across units. [23] The Pesaran test rejects the null of cross-sectional independence ($CD = -2.482, p = 0.0131$), indicating significant interdependence among the countries in the panel.

Driscoll–Kraay standard errors Given the simultaneous presence of serial autocorrelation and cross-sectional dependence, standard heteroskedasticity-robust or cluster-robust standard errors are insufficient. As Cameron and Miller (2015) make explicit, the cluster-robust variance estimator is derived under the assumption that errors are uncorrelated across clusters; when clustering on the panel unit to address within-unit serial correlation, independence of errors across units is therefore imposed by assumption, leaving cross-sectional dependence unaddressed. I therefore employ the nonparametric covariance estimator proposed by Driscoll and

[23] The remaining specifications exhibit a degree of panel imbalance such that Stata's `xtcd` command cannot compute the test reliably and returns an error.

Kraay (1998), commonly referred to as Driscoll–Kraay (DK) standard errors.[24] This estimator is robust to heteroskedasticity, serial autocorrelation up to some lag order, and general forms of cross-sectional dependence, without requiring the panel to be balanced or imposing parametric assumptions on the dependence structure. The estimator has been shown to perform well in panels with the time dimension typical of macroeconomic applications (Hoechle, 2007).[25]

6 Results

6.1 Baseline model

The baseline regression results for Equation 12, reported in Table 4, reveal a significant negative relationship between macroeconomic distress and fertility rates across both the parsimonious and fully controlled specifications. Given that the dependent variable is the natural logarithm of the Total Fertility Rate and the independent variables are lagged and z-standardized, the estimated coefficients represent semi-elasticities relative to a one-standard-deviation change in the previous year. Inflation and unemployment emerge as robust predictors of fertility decline; in the full model with controls (Panel B, Column 4), a one-standard-deviation increase in lagged inflation and unemployment is associated with a 1.95% and 4.39% reduction in the TFR, respectively. Interestingly, while GDP contraction exhibits a small positive coefficient when controls are introduced, its magnitude remains secondary to the primary drivers of economic uncertainty. Regarding the demographic controls, urbanization maintains a consistent negative pressure on fertility, while infant mortality remains positively correlated. Conversely, female education does not reach statistical significance in these baseline models.

[24] All models are estimated using the `xtscc` command in Stata, which implements the Driscoll–Kraay estimator with a Newey–West-type lag selection (Hoechle, 2007).

[25] The performance improves with larger time dimension T (Hoechle, 2007).

Table 4: Baseline Model

	Dependent variable: ln(Total Fertility Rate)			
	(1)	(2)	(3)	(4)
	Inflation	Unemployment	GDP contraction	Full model
Panel A: No Controls				
Headline Consumer Price inflation	-0.00778*			-0.0110**
	(0.00420)			(0.00405)
Unemployment		-0.0616***		-0.0640***
		(0.0161)		(0.0163)
GDP contraction			0.00627	0.00941*
			(0.00584)	(0.00526)
Within R^2	0.301	0.354	0.300	0.362
Observations	1729	1728	1728	1727
Panel B: Controls				
Headline Consumer Price Inflation	-0.0178***			-0.0195***
	(0.00398)			(0.00425)
Unemployment		-0.0401**		-0.0439***
		(0.0145)		(0.0139)
GDP contraction			0.00863*	0.0117**
			(0.00479)	(0.00446)
Urbanization	-0.130***	-0.103***	-0.123***	-0.105***
	(0.0298)	(0.0289)	(0.0273)	(0.0264)
Female Education	0.00507	0.000966	0.00121	0.00149
	(0.0118)	(0.0144)	(0.0138)	(0.0108)
Infant Mortality	0.0922***	0.0911***	0.0932***	0.0878***
	(0.0127)	(0.0150)	(0.0124)	(0.0142)
Observations	1252	1252	1252	1252
Within R^2	0.364	0.380	0.358	0.398

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is the natural logarithm of the Total Fertility Rate. All independent variables (macroeconomic shocks and controls) are z-standardized and lagged by one period ($t - 1$). All specifications include country fixed effects, country-specific linear trends, and year fixed effects.

6.1.1 Age-Specific Heterogeneity in Fertility Responses

To unpack the mechanisms driving the aggregate results in Table 4, Figure 4 decomposes the semi-elasticities by age group, as defined by Equation 13, revealing the specific demographic cohorts responsible for the observed TFR fluctuations in response to different macroeconomic shocks. In each panel, confidence intervals are widest in the last age group, likely reflecting the smaller number of births recorded among women aged 40–44, which reduces estimation precision.

Inflation (Panel A) displays a negative gradient that intensifies from the age group 25–29, although the effect is seemingly absent for the first two age groups. These findings contrast with Adesete and Piano (2026), who find in the US that younger cohorts are the most responsive to *unexpected* inflation, with women aged 25–44 specifically showing no statistically significant response. This divergence likely stems from two factors. First, whereas unexpected shocks may trigger immediate postponement among younger women, headline inflation exerts a more sustained structural pressure on the budgets of established, older households. This pattern is consistent with the budget erosion mechanism embedded in the quantity–quality framework of Becker and Lewis (1973) and Becker and Tomes (1976): as inflation raises both p_y and p_q, the shadow prices of quantity and quality in Equation 4 increase proportionally to the number of children already present, implying that established households with existing childcare commitments face a compounding cost increase that younger, childless households do not yet bear. Second, the global scope of this study introduces substantial heterogeneity in institutional contexts — in particular in the generosity of social safety nets and the structure of childcare costs — which may condition the age profile of the fertility response to inflation in ways that a single cross-country estimate cannot fully capture.

For unemployment (Panel B), semi-elasticities are negative across all age groups, with the effect most pronounced among women aged 30–34. While the broader negative response is consistent with the postponement mechanism documented in the European literature (Goldstein et al., 2013; Hoem, 2000), the concentration of the effect in the 30–34 bracket — rather than among the youngest cohorts as these studies find — may reflect the greater prevalence of higher-order birth decisions in this age group, which are more discretionary and therefore more sensitive to labour-market conditions. However, given the cross-country heterogeneity of the sample, this pattern should be interpreted with caution.

The most striking pattern emerges in Panel C: GDP contractions are associated with positive and increasing semi-elasticities across the age distribution. This counter-cyclical pattern may reflect a catching-up mechanism: a temporary contraction reduces the opportunity cost of time out of the labour market, potentially making it a relatively favourable moment to proceed with a previously postponed birth. This effect is plausibly stronger among older women, for whom the biological cost of further delay is highest (Schmidt, 2008).[26]

[26]According to Schmidt (2008), over the life cycle, two opposing forces shape the net cost of childbearing. On the one hand, the opportunity cost of having a child—measured in terms of foregone lifetime earnings—declines with age, as women have fewer remaining working years. On the other hand, the expected cost of remaining childless increases over time due to biological constraints. Together, these forces generate a U-shaped profile of net childbearing costs across the reproductive years. Women at older ages lie on the right-hand side of this curve, where the cost of further delay rises sharply. Bridging this framework with the present empirical evidence, a GDP contraction likely suppresses current labor market opportunities, temporarily lowering the opportunity cost of childbearing. For this older cohort, the macroeconomic shock thus tips the balance, prompting them to proceed with a birth rather than risking further postponement.

Figure 4: Semi-elasticity of Economic Shocks on Age-Specific Fertility Rates

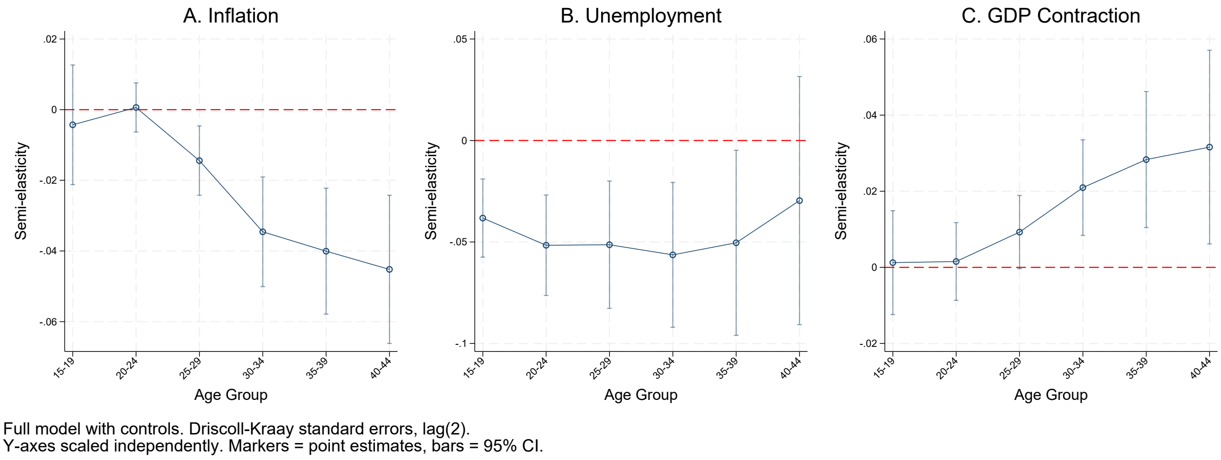

Note: Each panel reports semi-elasticities of age-specific fertility rates with respect to different macroeconomic shocks: unexpected inflation (A), unemployment (B), and GDP contraction (C). Estimates are obtained from fixed-effects regressions using Driscoll–Kraay standard errors with two lags, controlling for urbanization, female education, and infant mortality, as well as country trends and year effects. Markers represent point estimates and vertical bars denote 95% confidence intervals. Y-axes are scaled independently.

6.2 Interaction Models

The interaction models (Equation 14) are presented in Tables 5, 6, and 7.

While inflation remains a robust negative predictor of fertility across all specifications — a one-standard-deviation increase corresponding to a 1.97% reduction in the TFR in the full model — the interaction with risk-taking is only marginally significant in the parsimonious specification (Table 5, Column 1) and loses significance once controls are introduced, indicating that the buffering role of cultural risk tolerance against inflationary shocks is not robust at the aggregate level.

Table 5: Inflation: Interaction Model

	Dependent variable: ln(Total Fertility Rate)		
	(1)	(2)	(3)
	Base Interaction	With Controls	Full Model
Headline Consumer Price Inflation	-0.0108**	-0.0180***	-0.0197***
	(0.00484)	(0.00449)	(0.00468)
Inflation × Risk Taking	0.0124*	0.00719	0.00567
	(0.00675)	(0.00627)	(0.00608)
Unemployment			-0.0439***
			(0.0139)
GDP contraction			0.0114**
			(0.00456)
Urbanization		-0.130***	-0.105***
		(0.0291)	(0.0259)
Female Education		0.00494	0.00145
		(0.0115)	(0.0106)
Infant Mortality		0.0917***	0.0874***
		(0.0132)	(0.0146)
Observations	1729	1252	1252
Within R^2	0.304	0.365	0.399

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: All explanatory variables - comprising macroeconomic shocks (unemployment, GDP contraction, and headline consumer price inflation) and control variables (urbanization, female education, and infant mortality) - are standardized (z-scored to mean zero and unit standard deviation) and included with a one-year lag to mitigate reverse causality concerns. Risk-taking is standardized but time-invariant and therefore not lagged; because the specification includes time fixed effects, its main effect is absorbed and not separately identified, so only its interaction with time-varying shock variables is included. All models also include country fixed effects and country-specific linear trends to account for unobserved heterogeneity and differential trajectories across countries.

The most compelling evidence for the moderating role of risk preferences appears in the interaction with unemployment (Table 6). In the full model (3), the direct effect of unemployment is strongly negative: a change in one standard deviation of unemployment corresponds to a 3.97% decrease in the TFR. However, the interaction term is positive and significant (0.0203**), indicating that the negative effect of a one-standard-deviation increase in unemployment on TFR is attenuated by 2.03 percentage points for each one-standard-deviation increase in cultural risk tolerance.[27]

[27] To illustrate the economic magnitude: in a country with average risk tolerance, a one-standard-deviation increase in unemployment is associated with a 3.97% reduction in TFR. In a country one standard deviation *above* the mean in risk tolerance, the same unemployment shock is associated with a reduction of approximately 1.94% ($-3.97 + 2.03$). In a country one standard deviation *below* the mean, the reduction would be approximately 6.00% ($-3.97 - 2.03$), suggesting that low risk-tolerance societies bear a substantially heavier fertility cost from labour-market instability.

Table 6: Unemployment: Interaction Model

	Dependent variable: ln(Total Fertility Rate)		
	(1)	(2)	(3)
	Base Interaction	With Controls	Full Model
Unemployment	-0.0612***	-0.0357**	-0.0397***
	(0.0165)	(0.0141)	(0.0136)
Unemployment × Risk Taking	0.0105	0.0215**	0.0203**
	(0.00896)	(0.00889)	(0.00872)
Headline Consumer Price Inflation			-0.0190***
			(0.00421)
GDP contraction			0.0119**
			(0.00438)
Urbanization		-0.102***	-0.104***
		(0.0311)	(0.0284)
Female Education		0.00237	0.00271
		(0.0146)	(0.0112)
Infant Mortality		0.0920***	0.0887***
		(0.0155)	(0.0148)
Observations	1728	1252	1252
Within R^2	0.355	0.384	0.402

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: All explanatory variables - comprising macroeconomic shocks (unemployment, GDP contraction, and headline consumer price inflation) and control variables (urbanization, female education, and infant mortality) - are standardized (z-scored to mean zero and unit standard deviation) and included with a one-year lag to mitigate reverse causality concerns. Risk-taking is standardized but time-invariant and therefore not lagged; because the specification includes time fixed effects, its main effect is absorbed and not separately identified, so only its interaction with time-varying shock variables is included. All models also include country fixed effects and country-specific linear trends to account for unobserved heterogeneity and differential trajectories across countries.

GDP contractions (Table 7) display a positive direct coefficient in the full model (0.0134**), consistent with the counter-cyclical catching-up mechanism discussed in Section 6.1.1: a temporary contraction may reduce the opportunity cost of time out of the labour market, making it a relatively favourable moment to proceed with a previously postponed birth. The negative interaction with risk-taking (-0.00594^{*}) could be consistent with this interpretation: in more risk-tolerant societies, individuals may be less likely to exit the labour market during downturns, attenuating the opportunity cost reduction that drives the counter-cyclical response. Consequently, the catching-up mechanism operates most strongly precisely where risk tolerance is lowest.

Table 7: GDP contraction: Interaction Model

	Dependent variable: ln(Total Fertility Rate)		
	(1)	(2)	(3)
	Base Interaction	With Controls	Full Model
GDP contraction	0.00964	0.00995*	0.0134**
	(0.00735)	(0.00564)	(0.00544)
GDP Contraction × Risk Taking	-0.00649*	-0.00495	-0.00594*
	(0.00330)	(0.00304)	(0.00318)
Headline Consumer Price Inflation			-0.0189***
			(0.00429)
Unemployment			-0.0451***
			(0.0141)
Urbanization		-0.124***	-0.105***
		(0.0278)	(0.0266)
Female Education		0.00147	0.00166
		(0.0137)	(0.0108)
Infant Mortality		0.0914***	0.0856***
		(0.0119)	(0.0137)
Observations	1728	1252	1252
Within R^2	0.303	0.360	0.401

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: All explanatory variables - comprising macroeconomic shocks (unemployment, GDP contraction, and headline consumer price inflation) and control variables (urbanization, female education, and infant mortality) - are standardized (z-scored to mean zero and unit standard deviation) and included with a one-year lag to mitigate reverse causality concerns. Risk-taking is standardized but time-invariant and therefore not lagged; because the specification includes time fixed effects, its main effect is absorbed and not separately identified, so only its interaction with time-varying shock variables is included. All models also include country fixed effects and country-specific linear trends to account for unobserved heterogeneity and differential trajectories across countries.

Overall, the results suggest that populations with a higher willingness to take risks are significantly less deterred by labor market instability when making fertility decisions. For Inflation and GDP contraction, in contrast, risk-taking does not appear to meaningfully mitigate the shocks once controls are included.

Marginal Effects Plots Figure 5 illustrates how the marginal effect of each shock on ln(TFR) varies across the distribution of cultural risk tolerance.

Panel A displays a negative marginal effect of inflation on fertility at low levels of risk tolerance, which moves toward zero as risk tolerance increases. The confidence intervals are wider at the lower end of the risk-taking distribution, where the negative effect is largest, introducing some uncertainty around the precise magnitude of the effect in low risk-tolerance societies. This is consistent with the fragility of the interaction term in Table 5.

Panel B provides the clearest visual evidence of a buffering role: the marginal effect of unemployment on fertility is negative and statistically significant at low levels of risk tolerance, but diminishes progressively as risk tolerance increases, becoming statistically indistinguishable from zero at the upper end of the distribution (+2 SD). This pattern visually confirms the robust interaction result reported in Table 6.

Panel C reflects the counter-cyclical pattern for GDP contraction: the marginal effect is positive and significant at low levels of risk tolerance, weakening monotonically and becoming indistinguishable from zero around +1 SD, suggesting that the catching-up mechanism is concentrated in low risk-tolerance societies.

Figure 5: Marginal Effects of Economic Shocks on ln(TFR)

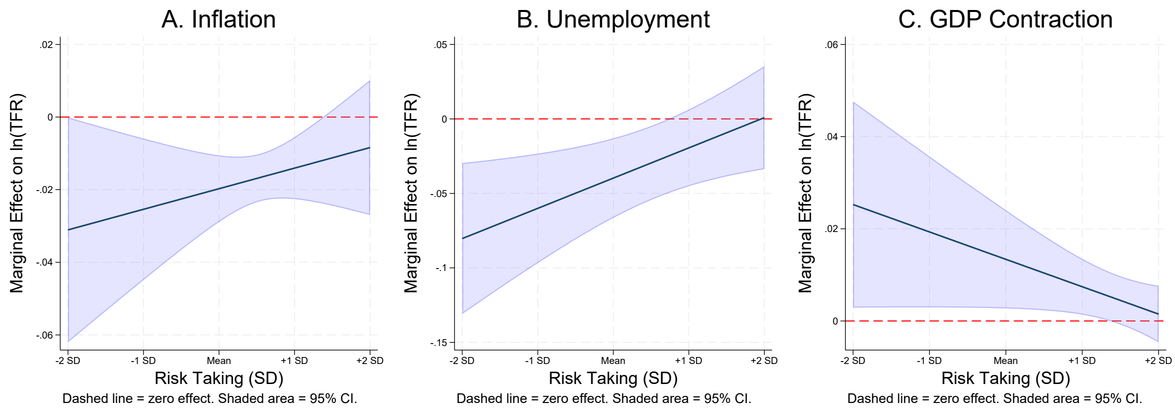

Note: Y-axes are scaled independently across panels to maximize readability. Full model estimates with Driscoll-Kraay standard errors, lag(2).

6.3 Heterogeneity Analysis by Income Level

The stratified analysis presented here investigates whether the aggregate results mask systematic heterogeneity across development levels, and whether the moderating role of cultural risk tolerance varies along the income distribution.

6.3.1 Baseline results by income group

Without controls In Table B.2, the only predictor that holds consistently across all three income groups is unemployment, which enters with a negative and statistically significant coefficient in each tercile. Its magnitude is largest for low-income countries (-8.84%, $p < 0.01$), declining to approximately -6.5% for medium-income and -5.5% for high-income countries. Inflation is statistically significant only for the middle-income group (-0.90%, $p < 0.05$), and GDP contraction is insignificant across all three panels in this parsimonious specification.

Once controls are introduced In Table B.3, the picture changes substantially, and in instructive ways.

For **low-income countries**, the unemployment effect disappears entirely upon inclusion of controls — and in fact reverses sign — suggesting that the raw correlation was driven by structural confounders rather than a genuine precautionary response to labour-market instability. The most notable shift in this group concerns female education, which enters with a positive and statistically significant coefficient. This finding is consistent with the demographic literature on a U-shaped relationship between female education and fertility. GDP contraction acquires a marginally significant positive coefficient in this group, consistent with the counter-cyclical pattern documented in the aggregate results, and discussed further in the interaction analysis below.

For **medium-income countries**, results are the most consistent and economically meaningful. Inflation reduces the TFR by approximately 2.11% per standard deviation ($p < 0.01$), and unemployment by 6.5–6.7% ($p < 0.01$). The significance of inflation in this group — absent for both low- and high-income countries — may not be coincidental. The medium-income tercile includes countries such as Argentina, Iran, Turkey, and Serbia, economies with historically volatile price levels where inflationary pressure is a structural feature of the macroeconomic environment rather than a transitory shock. In such contexts, households may be particularly sensitive to price instability as a signal of broader economic uncertainty, reinforcing the fertility-reducing channel identified in the theoretical framework. Female education enters negatively and significantly in this group, consistent with the standard opportunity-cost mechanism: once development crosses a threshold at which women face genuine career trade-offs, higher education depresses fertility. This reversal of sign relative to the low-income group is precisely what the U-shaped education–fertility hypothesis would predict.

For **high-income countries**, unemployment is the only robust predictor, with a semi-elasticity of approximately −4.95% to −5.41% ($p < 0.01$). Neither inflation nor GDP contraction reaches significance. This pattern closely mirrors the European literature on the Great Recession, which identifies labour-market conditions as the primary macroeconomic channel through which fertility responds to economic downturns in wealthy settings (Goldstein et al., 2013; Fluchtmann et al., 2023).

6.3.2 Interaction models by income group

Tables B.4, B.5, and B.6 present the interaction results between each macroeconomic shock and the cultural risk-taking index, stratified by income group.

Inflation and risk tolerance Table B.4 presents the biggest puzzle to interpret. For low-income countries, neither the direct effect of inflation nor its interaction with risk tolerance reaches significance in any specification, consistent with the general weakness of the inflation channel in this group. For medium-income countries, the interaction is large and significant in the parsimonious specification ($\beta = 0.0221$, $p < 0.01$), suggesting that risk-seeking cultures partly buffer the negative fertility response to inflation; however, this effect is fully absorbed once structural controls are added, pointing to fragility rather than robustness. The most striking

— and most difficult to interpret — finding concerns high-income countries, where the interaction coefficient is negative and significant with controls ($\beta = -0.0348$, $p < 0.05$), implying that higher risk tolerance *amplifies* rather than attenuates the negative fertility response to inflation in wealthy countries. However, the significance disappears in the full specification.

Unemployment and risk tolerance Table B.5 exhibits the clearest income gradient. For low-income countries, the interaction term is not statistically significant in any specification, and the direct effect of unemployment reverses sign once controls are included — reinforcing the interpretation that labour-market instability operates through qualitatively different mechanisms, or is simply less operative, at the lower end of the income distribution. For medium- and high-income countries, the interaction is positive and significant with controls ($\beta \approx 0.023$–0.026 and $\beta \approx 0.030$ respectively, $p < 0.05$), closely mirroring the aggregate result. This confirms that the buffering role of cultural risk tolerance against unemployment shocks is concentrated precisely where precautionary motives are most plausibly the operative mechanism — that is, in settings where labour-market instability constitutes a genuine long-horizon threat to household financial planning.

GDP contraction and risk tolerance Table B.6 tells a story that is specific to low-income countries. In this group, GDP contractions are associated with a positive and significant direct effect on fertility ($\beta \approx 0.010$, $p < 0.05$) that persists across all three specifications — consistent with the counter-cyclical catching-up pattern noted in the aggregate analysis. The interaction term is negative and significant ($\beta \approx -0.006$, $p < 0.01$ in the base model), implying that the counter-cyclical response is attenuated in countries with higher risk tolerance. One possible interpretation is that in low-income settings, where labour-market attachment may be weaker and opportunity costs of childbearing lower, economic downturns reduce the cost of time out of the workforce sufficiently to encourage births — a mechanism that is dampened in more risk-tolerant societies where individuals are more likely to remain attached to the labour market even during contractions. For medium- and high-income countries, neither the direct effect of GDP contraction nor its interaction with risk tolerance reaches significance once controls are included, confirming that this counter-cyclical dynamic is specific to the lower end of the income distribution.

Figure 6 summarises the stratified interaction results visually. Panel A (Inflation) reveals a clear asymmetry across income groups: the direct effect of inflation is concentrated in medium-income countries and is absent elsewhere. Panel B (Unemployment) displays the most orderly income gradient of the three shocks. The direct effect is negative and significant for medium and high income economies, that also experience while a positive buffering interaction with willingness to take risks. The confidence intervals for the low-income interaction are also noticeably wider, reflecting the greater heterogeneity within this tercile. Panel C (GDP

Contraction) tells a qualitatively different story: the direct effect is positive and significant only for low-income countries—consistent with the counter-cyclical pattern documented in the aggregate analysis—and the interaction term is negative, indicating that higher cultural risk tolerance attenuates this counter-cyclical response. For medium- and high-income countries, both the direct effect and the interaction are close to zero and statistically indistinguishable from it, confirming that the counter-cyclical mechanism is specific to the lower end of the income distribution.

Figure 6: Heterogeneity in the Effects of Economic Shocks on Fertility by Income Group

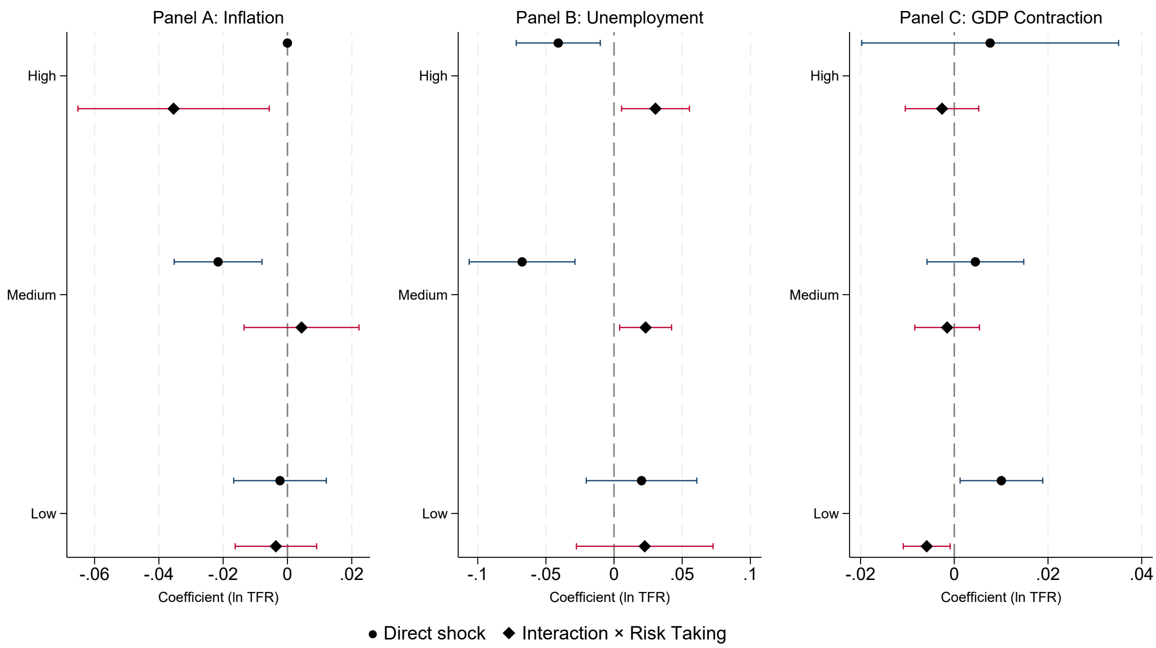

Note: Each panel reports coefficients from a separate fixed-effects regression (Driscoll–Kraay standard errors, lag = 2) estimated on the relevant income tercile. The dependent variable is the natural log of the total fertility rate (ln TFR). All shock variables are lagged by one period and standardised to unit variance. The solid circles (●) denote the direct semi-elasticity of each macroeconomic shock; the solid diamonds (♦) denote the coefficient on the interaction between the shock and the cultural risk-taking index. Full controls include urbanisation, female education, and infant mortality, as well as country-specific linear trends and year fixed effects. Horizontal bars represent 95% confidence intervals. A coefficient estimate crossing the dashed vertical line at zero is not statistically significant at the 5% level.

Taken together, both the regression results and their graphical representation reveal that the aggregate relationship between macroeconomic shocks and fertility conceals fundamentally different mechanisms across development levels. While the unemployment channel appears robust across all income groups in the parsimonious baseline, it disappears for low-income countries once structural controls are introduced, and its moderation by cultural risk tolerance is confined to medium- and high-income contexts. In contrast, the counter-cyclical fertility response to GDP contraction is specific to the poorest tercile, whereas inflation plays a more prominent role in middle-income economies, particularly those with histories of price volatility. As illustrated in the figure, neither the sign nor the magnitude of these effects is uniform across development levels: both the dominant transmission channels and the moderating influence of cultural risk tolerance vary systematically with income.

6.4 Robustness Checks

The baseline and interaction results were subjected to three robustness checks: excluding peak years of major crisis episodes, replacing linear with quadratic country-specific time trends, and purifying the risk-taking index of its correlation with patience. Each check re-examines both the direct effects of macroeconomic shocks on the Total Fertility Rate and the moderating role of cultural risk tolerance. The latter is assessed exclusively for the unemployment interaction, which was the only term to achieve robust statistical significance in the main analysis reported in Section 5.2.[28] Across all three checks, the main conclusions hold.

Excluding Crisis Episodes A key concern is that the baseline results may be disproportionately driven by two exceptional macroeconomic episodes—the Global Financial Crisis and the COVID-19 pandemic—rather than capturing the systematic relationship between ordinary business-cycle fluctuations and fertility.[29] Both crises generated simultaneous, large-scale shocks to unemployment, inflation, and output across most of the 76 countries in the sample, and their extraordinary nature may have triggered fertility responses that differ qualitatively from those observed in normal downturns. To assess whether the main findings are robust to the exclusion of these episodes, I re-estimate the baseline specification sequentially dropping the peak year of the GFC (2008), the peak year of the COVID-19 shock (2021), and both years simultaneously.[30]

The results, reported in Table C.1, confirm the robustness of the baseline findings. Across all four columns and both panels, inflation and unemployment retain their negative sign and remain statistically significant, with magnitudes that are remarkably stable relative to the full-sample estimates. The positive coefficient on GDP contraction likewise persists throughout. The stability of the coefficients across the three exclusion schemes suggests that the main results are not unduly influenced by the GFC or the pandemic, and instead reflect a more general relationship between macroeconomic conditions and reproductive behavior.

Country-Specific Quadratic Trends The baseline specification controls for country-specific linear time trends to absorb gradual, nation-level shifts in fertility trajectories such as demographic transitions or structural institutional changes. A potential concern is that some of these trajectories may follow a non-linear path — for instance, countries that first experienced fertility

[28] Re-testing the interactions between risk tolerance and inflation or GDP contraction is not warranted here: since neither reached significance in the baseline specification, confirming or rejecting their robustness would carry limited interpretive value.

[29] Adesete and Piano (2026), in their analysis of unexpected inflation rates in the U.S., also exclude the COVID-19 period due to concerns about potential distortions.

[30] Both crises extended across multiple years — the GFC broadly covers 2007–2009 and the pandemic shock 2020–2021. I restrict the exclusion to a single peak year per episode to limit the loss of observations in a panel spanning only 23 years; dropping the full crisis windows would reduce the sample by approximately 10%, with a non-trivial cost to statistical power.

decline and then partial recovery — which a linear trend would fail to capture. If such curvature is systematically correlated with macroeconomic conditions or cultural risk attitudes, the baseline estimates could be confounded by residual trend heterogeneity. To address this, I augment the specification with country-specific quadratic trends, including both $\alpha_i \times t$ and $\alpha_i \times t^2$ for each country i.

The results, reported in Table C.3, are reassuringly stable. The direct effects of unemployment and inflation remain negative and statistically significant across all three specifications, with magnitudes closely mirroring the baseline. Most importantly, the interaction between unemployment and cultural risk tolerance retains its positive sign and statistical significance ($\hat{\beta} = 0.0231^{**}$), virtually unchanged relative to the main estimate of 0.0203^{**}. This evidence validates the protective role of risk tolerance against labor-market shocks, proving the relationship is robust to potential specification errors regarding non-linear national fertility trajectories.

Risk Tolerance Purified A potential concern with the cultural risk-taking index is that it may not capture risk tolerance in isolation. Falk et al. (2018) document a positive correlation between risk tolerance and patience within the Global Preference Survey. If more patient societies tend also to be more risk-tolerant, and if patience independently shapes fertility responses to economic shocks — for instance because more forward-looking households plan further ahead and are therefore less reactive to transitory deteriorations — then the interaction term estimated in Section 6.2 could be capturing a blend of both cultural traits rather than risk culture alone.

To address this concern, I residualise the risk-taking index on the GPS patience measure. Concretely, I collapse the panel to a single cross-section of country-level means, regress risk-taking on patience, and retain the residuals as a *purified* risk-taking measure that is, by construction, orthogonal to patience. This purified index is then standardised and used in place of the original risk-taking variable to re-estimate the unemployment interaction model.

The results are reported in Table C.4. The interaction coefficient obtained with the purified measure ($\hat{\beta} = 0.0224^{***}$) is virtually identical in magnitude to — and slightly larger and more precisely estimated than — the baseline estimate ($\hat{\beta} = 0.0203^{**}$), while all other coefficients remain stable across the two specifications. This finding indicates that the moderating role of cultural risk tolerance on the unemployment–fertility relationship is not an artefact of its correlation with patience: the buffering effect survives, and if anything strengthens, once the patience component is removed from the risk-taking index.

7 Discussion and Conclusion

7.1 Discussion of Results

This study has investigated how three macroeconomic shocks — inflation, unemployment, and GDP contraction — affect fertility across 76 countries over the period 2000–2023, and whether country-level cultural risk tolerance moderates these effects. The analysis was guided by three hypotheses derived from an extended Becker–Tomes quantity–quality framework that incorporates a precautionary saving channel: (1) that macroeconomic shocks negatively affect total fertility rates; (2) that cultural risk tolerance buffers these effects; and (3) that both the baseline effects and their moderation vary systematically across income groups.

Hypothesis 1: The Direct Effects of Macroeconomic Shocks The baseline results offer robust support for the first hypothesis. Across both parsimonious and fully controlled two-way fixed-effects specifications, inflation and unemployment emerge as consistent, statistically significant negative predictors of the Total Fertility Rate. These findings are directly interpretable through the theoretical framework developed in Section 3.1. Inflation operates through the shadow price mechanism formalised in Equation (4): as consumer prices and the unit cost of child quality p_q rise simultaneously, the shadow prices of both child quantity and quality increase, compressing the household's effective resource base in the same way as a direct income shock — consistent with the interpretation in Wu (2024) that rising costs of living lead households to postpone or forgo additional births. Unemployment, by contrast, operates primarily through the income channel of Equation (6): job loss reduces realised household income $I - \eta_t$, shifting the social budget constraint inward and inducing a contraction in the demand for both the number and quality of children, as illustrated in Figure 2. The negative effect of unemployment on age-specific fertility rates is most pronounced among women aged 30–34, who face a concentration of higher-order birth decisions that are discretionary and therefore more sensitive to labour-market conditions. The coefficient on GDP contraction is positive and small in the levels specification and, as it will be discussed, it is likely driven by low-income countries.

Hypothesis 2: The Moderating Role of Cultural Risk Tolerance The interaction results provide targeted support for the second hypothesis: cultural risk tolerance acts as a buffer against macroeconomic shocks on fertility, but this moderating role is specific to labour-market uncertainty, rather than a general feature of all macroeconomic shocks. Its interpretation follows directly from the precautionary channel in Equation (9). The buffering effect of risk tolerance on unemployment is robust across all specifications and all three robustness checks, and its interpretation is grounded in the precautionary channel formalised in Equation (9). The key intuition is that labour-market instability is precisely the kind of shock that activates precautionary behaviour. Unemployment generates a direct, personalised threat to household income: losing a job — or facing a heightened probability of doing so — immediately raises the perceived variance of future resources σ_η^2. GDP contraction, by contrast, is a more diffuse aggregate signal: it may or may not translate into a concrete deterioration in any given household's labour-market position, and its fertility effects appear to operate through the opportunity-cost channel rather than through precautionary restraint, as discussed in the stratified analysis below.

Inflation bypasses the precautionary channel altogether by operating through the price structure of childbearing regardless of how households feel about uncertainty. Unemployment thus stands apart as the shock whose fertility impact is most directly mediated by the subjective weight households place on downside risk. Where that weight is culturally high, even a moderate rise in unemployment pushes the effective social income S^* below the threshold at which households are willing to commit to an additional birth. In more risk-tolerant societies, the same labour-market deterioration produces a smaller fertility response, because households are less prone to amplify the shock through precautionary restraint. The marginal effects plot in Figure 5 Panel B makes this gradient visible: the fertility cost of unemployment falls monotonically as cultural risk tolerance rises. The buffering effect of risk tolerance on unemployment is robust across three robustness checks.

Hypothesis 3: Heterogeneity by Income Group The stratified analysis broadly supports the third hypothesis, while revealing mechanisms that are more differentiated than a simple precautionary story would predict. For *high-income countries*, unemployment is the only robust predictor of fertility, closely mirroring the European literature on the Great Recession, and the interaction with cultural risk tolerance is positive and significant — confirming that the precautionary channel formalised in Equation (9) is operative precisely where households have sufficient resources to engage in forward-looking reproductive planning and where labour-market instability constitutes a genuine long-horizon threat to household welfare.

For *middle-income countries*, both inflation and unemployment significantly depress fertility, and the interaction with risk tolerance moderates the unemployment effect. The salience of inflation in this group is not coincidental: the middle-income tercile includes economies such as Argentina, Iran, Turkey, and Serbia — countries historically characterized by high price volatility. The erasure of the inflation–risk-tolerance interaction once structural controls are added suggests that, consistent with the aggregate result, the cost-raising mechanism of inflation operates uniformly through the price structure of childbearing rather than through precautionary behaviour.

For *low-income countries*, the results are the most structurally distinct from the rest of the sample. Unemployment loses significance — and reverses sign — once controls are added. One plausible explanation for this muted reaction, connecting to Schmidt (2008), lies in the structure of fertility timing in these settings. Where access to modern contraception is limited — due to financial barriers, supply constraints, or restrictive social norms — fertility is less tightly planned and more strongly shaped by constrained contraceptive choice (Bongaarts, 2011).

The most notable and theoretically revealing finding for this group concerns GDP contraction. Unlike in the middle- and high-income terciles, GDP contractions in low-income countries are associated with a positive and significant direct effect on fertility that persists across all three specifications. This counter-cyclical pattern is consistent with a catching-up mechanism: a temporary contraction reduces the opportunity cost of time out of the labour market, making it a relatively favourable moment to proceed with births that had been postponed. This result should, however, be situated within the broader demographic context of these countries. Many low-income nations in the sample remain in earlier stages of the fertility transition, where aggregate fertility has not yet undergone the sustained secular decline driven by rising income, female educational attainment, and the deepening of women's labour-market integration that characterises more advanced economies (Lesthaeghe, 2010). In these settings, the Becker–Tomes endowment ratio e/q is larger on average: children's exogenous endowments constitute a more significant share of total child quality. As formalised in Equation (7), and as Becker and Tomes (1976) demonstrate, a larger ratio e/q compresses the quantity elasticity relative to the quality elasticity, making the number of children less sensitive to income shocks. In this context, a temporary GDP contraction does not trigger the strong precautionary postponement that characterises households in high-income settings; instead, by reducing the opportunity cost of time outside the labour market, it may temporarily lower the effective price of childbearing in ways that are fertility-enhancing at the margin. The negative and significant interaction between GDP contraction and risk tolerance in low-income countries implies that the counter-cyclical fertility response is attenuated in more risk-tolerant societies. One possible interpretation is that in low-income settings, where labour-market attachment is generally weaker and the opportunity costs of childbearing lower, economic downturns reduce the cost of time out of the workforce sufficiently to encourage births. In more risk-tolerant societies, however, individuals may be more likely to remain attached to the labour market even during contractions, dampening the opportunity-cost reduction that drives the counter-cyclical response.

7.2 Limitations and Future Directions

Several limitations warrant acknowledgement. The validity of the risk-taking index rests on the assumption that cultural risk preferences measured in 2012 provide a stable representation of underlying preferences across the full 2000–2023 window. While Bellani and Arpino (2021) argue that risk preferences manifest as stable, trait-like dispositions in adulthood, and while the cross-country differences documented by Falk et al. (2018) are both large and statistically robust, Hudson et al. (2005) show that even within the same individuals different elicitation methods and framing can generate inconsistent estimates. The GPS index partially addresses this concern through its combination of qualitative and quantitative items, but the possibility that the measure partly reflects the design of the elicitation instrument rather than a purely underlying preference cannot be fully dismissed.

A further behavioural mechanism not explicitly captured in the empirical specification concerns the indirect pathway from risk preferences through marriage timing to fertility. More risk-tolerant individuals tend to set lower reservation thresholds for an acceptable partner and match earlier (Schmidt, 2008), meaning that part of the fertility response identified here may operate through postponed partnership formation among the risk-averse, rather than through fertility behaviour conditional on partnership status. Disentangling these pathways would require cross-country data on marriage timing and union formation.

Finally, the analysis operates at the country level and cannot identify whether macroeconomic shocks affect individuals differently depending on their own employment status, wealth, or risk preferences. A natural extension would combine the cultural-level risk measure used here with individual-level panel data to estimate heterogeneous treatment effects across the within-country distribution of risk preferences.

7.3 Policy Implications

The findings carry clear implications for pro-natalist policy. The behavioural framework developed in Section 3.1.3 offers a precise diagnosis of why temporary fertility incentives such as baby bonuses tend to produce modest and short-lived effects. Households operating in the loss domain — where realised income has fallen below the pre-shock reference point — discount temporary transfers because such transfers do not remove the underlying source of uncertainty. The perceived gain from a one-time payment is insufficient to offset the persistent precautionary wedge $\phi(\rho, \sigma_\eta^2)$ that governs the decision to undertake an irreversible long-horizon commitment. Effective pro-natalist policy must therefore operate on the source of uncertainty itself rather than on its symptoms: by stabilising labour markets and anchoring inflation expectations, policymakers can reduce the effective size of $\phi(\rho, \sigma_\eta^2)$ for the entire distribution of households, not only those with high risk tolerance. Monetary frameworks that reduce inflation uncertainty — such as inflation targeting, as documented by Wu (2024) — and active labour-market policies that shorten unemployment spells are therefore more structurally suited to limiting the fertility costs of macroeconomic downturns than direct financial transfers. For low-income countries specifically, where the precautionary channel is less operative but access to reproductive health services is more constrained, policies that expand contraceptive availability and strengthen family planning programmes are likely to be the most effective tools for helping people realise their desired family size (Bongaarts, 2011; United Nations Population Fund, 2025).

The study demonstrates that childbearing is indeed perceived and treated as a risky, irreversible investment. The magnitude of the fertility cost of economic shocks is not determined solely by the size of the shock itself but is systematically shaped by the cultural lens through which households interpret and respond to uncertainty. Understanding this cultural dimension is essential for both interpreting contemporary demographic trends and designing policies that support reproductive autonomy across the full range of macroeconomic environments.

References

Abdel-Tawab, Nahla G., Sarah Loza and Abeer Zaki (2008), Helping egyptian women achieve optimal birth spacing intervals through fostering linkages between family planning and maternal/child health services, Frontiers final report, Population Council, Washington, DC.
URL: *https://knowledgecommons.popcouncil.org/departments*$_{s}bsr - rh/$389

Adesete, Ahmed A. and Clara E. Piano (2026), 'Unexpected inflation and fertility rates in the us'. Working Paper.

Adserà, Alícia (2010), 'Where are the babies? labor market conditions and fertility in europe', *European Journal of Population* **26**(3), 211–242.

Affuso, Ermanno, Khandokar Istiak and James Swofford (2022), 'Interest rates, house prices, fertility, and the macroeconomy', *Journal of Risk and Financial Management* **15**(9), 403.
URL: *https://doi.org/10.3390/jrfm15090403*

Ali, S. (2015), 'An empirical analysis of population growth on economic development: The case study of bangladesh', *International Journal of Economics, Finance and Management Sciences* **3**(3), 252–258.

Andersson, Gunnar (2000), 'The impact of labour-force participation on childbearing behaviour: Pro-cyclical fertility in sweden during the 1980s and the 1990s', *European Journal of Population* **16**(4), 293–333.

Baltagi, Badi H. and Ping X. Wu (1999), 'Unequally spaced panel data regressions with AR(1) disturbances', *Econometric Theory* **15**(6), 814–823.

Bampinas, Georgios and Georgios Mavropoulos (2024), 'Asymmetric effects between economic development and fertility: What do 140 years of data tell us?', *The Journal of Economic Asymmetries* **30**, e00368.

Becker, Gary S. and H. Gregg Lewis (1973), 'On the interaction between the quantity and quality of children', *Journal of Political Economy* **81**(2), S279–S288.

Becker, Gary S. and Nigel Tomes (1976), Child endowments, and the quantity and quality of children, Working Paper 123, National Bureau of Economic Research.

Bellani, Daniela and Bruno Arpino (2021), 'Risk tolerance and fertility: Evidence from a lottery question in italy', *Journal of Marriage and Family* **84**(2), 457–475.
URL: *https://onlinelibrary.wiley.com/doi/10.1111/jomf.12820*

Bernardi, Laura, Andreas Klärner and Holger von der Lippe (2008), 'Job insecurity and the timing of parenthood: A comparison between eastern and western germany', *European Journal of Population / Revue européenne de Démographie* **24**(3), 287–313.

Bhargava, Alok, Luisa Franzini and Wiji Narendranathan (1982), 'Serial correlation and the fixed effects model', *Review of Economic Studies* **49**(4), 533–549.

Bongaarts, John (2011), 'Can family planning programs reduce high desired family size in sub-saharan africa?', *International Perspectives on Sexual and Reproductive Health* **37**(4), 209–216.
URL: *https://www.jstor.org/stable/41336478*

Cameron, A. Colin and Douglas L. Miller (2015), 'A practitioner's guide to cluster-robust inference', *Journal of Human Resources* **50**(2), 317–372.
URL: *https://www.jstor.org/stable/24735989*

Casanova, Mariana, Lara Pacheco and Joaquim Luís Coimbra (2024), 'New forms of uncertainty in the individualized society: Adaptation and validation of the uncertainty response scale (urs, greco roger, 2001) to portuguese population and the creation of a scale on the perception of uncertainty in the social context and its psychological consequences', *University of Porto* . Adaptation and validation study of the Uncertainty Response Scale (URS) and development of a new scale on social-contextual uncertainty perception.

Chowdhury, Abdur R. (1988), 'The infant mortality–fertility debate: Some international evidence', *Southern Economic Journal* **54**(3), 666–674.

Coibion, Olivier, Dimitris Georgarakos, Yuriy Gorodnichenko and Maarten van Rooij (2021), 'How does consumption respond to news about inflation? Field evidence from a randomized control trial', *The Review of Economic Studies* **88**(5), 2136–2168.

Davis, Kingsley (1945), 'The world demographic transition', *Annals of the American Academy of Political and Social Science* **237**, 1–11.
URL: *https://www.jstor.org/stable/i242844*

Driscoll, John C. and Aart C. Kraay (1998), 'Consistent covariance matrix estimation with spatially dependent panel data', *Review of Economics and Statistics* **80**(4), 549–560.

Easterlin, Richard A. (1976), 'The conflict between aspirations and resources', *Population and Development Review* **2**(3/4), 417–425.
URL: *https://www.jstor.org/stable/1971619*

Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman and Uwe Sunde (2018), 'Global evidence on economic preferences', *Quarterly Journal of Economics* **133**(4), 1645–1692.

Falk, Armin, Anke Becker, Thomas J. Dohmen, David Huffman and Uwe Sunde (2016), The preference survey module: A validated instrument for measuring risk, time, and social preferences, IZA Discussion Paper 9674, IZA Institute of Labor Economics.

Fernandez-Villaverde, Jesus (2001), Was malthus right? economic growth and population dynamics, Pier working paper 01-046, Penn Institute for Economic Research, University of Pennsylvania, Philadelphia.

Fluchtmann, Jonas, Violetta van Veen and Willem Adema (2023), Fertility, employment and family policy: A cross-country panel analysis, OECD Social, Employment and Migration Working Papers 299, OECD, Paris.
URL: *https://dx.doi.org/10.1787/326844f0-en*

Flückiger, Matthias and Markus Ludwig (2017), 'Urbanization, fertility and child education in sub-saharan africa', *Economics Letters* **157**, 97–102.
URL: *https://doi.org/10.1016/j.econlet.2017.05.024*

Goldstein, Joshua R., Michaela Kreyenfeld, Aiva Jasilioniene and Deniz Karaman Örsal (2013), 'Fertility reactions to the "great recession" in europe: Recent evidence from order-specific data', *Demographic Research* **29**, 85–104.

Ha, Jongrim, M. Ayhan Kose and Franziska Ohnsorge (2023), 'One-stop source: A global database of inflation', World Bank Dataset. Policy Research Working Paper Series.
URL: *https://openknowledge.worldbank.org/handle/10986/35997*

Hoechle, Daniel (2007), 'Robust standard errors for panel regressions with cross-sectional dependence', *Stata Journal* **7**(3), 281–312.

Hoem, Britta (2000), 'Entry into motherhood in sweden: the influence of economic factors on the rise and fall in fertility, 1986-1997', *Demographic Research* **2**(4).

Hudson, Darren, Keith Coble and Jayson Lusk (2005), 'Consistency of risk premium measures', *Agricultural Economics* **33**(1), 41–49.

Kahneman, Daniel and Amos Tversky (1979), 'Prospect theory: An analysis of decision under risk', *Econometrica* **47**(2), 263–291.

Kristensen, Axel Peter and Trude Lappegård (2022), 'Unemployment and fertility: The relationship between individual and aggregated unemployment and fertility during 1994–2014 in norway', *Demographic Research* **46**(35), 1037–1064.

Lee, Sanghoon (2025), 'Fertility and urbanization', *Cities* **163**, 106002.
URL: *https://doi.org/10.1016/j.cities.2025.106002*

Leland, Hayne E. (1968), 'Saving and uncertainty: The precautionary demand for saving', *The Quarterly Journal of Economics* **82**(3), 465–473.
URL: *http://www.jstor.org/stable/1879518*

Lesthaeghe, Ron (2010), 'The unfolding story of the second demographic transition', *Population and Development Review* **36**, 211–251.

Malthus, Thomas Robert (1798), *An Essay on the Principle of Population*, ElecBook, London. Text of the first edition, published 1798. This digital edition 1997.

Natural Earth (2023), 'Admin 0 – countries, 1:110m cultural vectors', https://www.naturalearthdata.com/downloads/110m-cultural-vectors/. Accessed: 2026-04-29.

Oppenheimer, Valerie Kincade, Matthijs Kalmijn and Nelson Lim (1997), 'Men's career development and marriage timing during a period of rising inequality', *Demography* **34**(3), 311–330.
URL: *https://www.jstor.org/stable/3038286*

Pesaran, M. Hashem (2021), 'General diagnostic tests for cross section dependence in panels', *Empirical Economics* **60**(1), 13–50. First circulated as CESifo Working Paper No. 1229, 2004.

Pratt, John W. (1964), 'Risk aversion in the small and in the large', *Econometrica* **32**(1/2), 122–136.
URL: *https://www.jstor.org/stable/1913738*

Schmidt, Lucie (2008), 'Risk preferences and the timing of marriage and childbearing', *Demography* **45**(2), 439–460.

Smith, Adam (1776), *An Inquiry into the Nature and Causes of the Wealth of Nations*, Vol. 1, 1 edn, W. Strahan and T. Cadell, London.

Sobotka, Tomáš, Vegard Skirbekk and Dimiter Philipov (2011), 'Economic recession and fertility in the developed world', *Population and Development Review* **37**(2), 267–306.

Thompson, Warren S. (1929), 'Population', *American Journal of Sociology* **34**(6), 957–975.

Towriss, C. A. and I. M. Timæus (2018), 'Contraceptive use and lengthening birth intervals in rural and urban eastern africa', *Demographic Research* **38**(64), 2027–2052.

United Nations (2024), 'World population prospects 2024, online edition. file FERT/02: Age-specific fertility rates', Dataset. POP/DB/WPP/Rev.2024/FERT/F02.
URL: *https://population.un.org/wpp/*

United Nations Population Fund (2025), 'State of world population 2025'.
URL: *https://www.unfpa.org/publications*

World Bank (2024), 'Gdp per capita (current US$)', World Development Indicators. Accessed: May 6, 2026.
URL: *https://data.worldbank.org/indicator/NY.GDP.PCAP.CD*

World Bank (2026*a*), 'Fertility rate, total (births per woman)', World Development Indicators. Source: United Nations Population Division. Accessed: May 6, 2026.
URL: *https://data.worldbank.org/indicator/SP.DYN.TFRT.IN*

World Bank (2026*b*), 'Mortality rate, infant (per 1,000 live births)', World Development Indicators. Source: UN Inter-agency Group for Child Mortality Estimation. Accessed: May 6, 2026.
URL: *https://data.worldbank.org/indicator/SP.DYN.IMRT.IN*

World Bank (2026*c*), 'School enrollment, tertiary, female (% gross)', World Development Indicators. Source: UNESCO Institute for Statistics (UIS). Accessed: March 19, 2026.
URL: *https://data.worldbank.org/indicator/SE.TER.ENRR.FE*

World Bank (2026*d*), 'Unemployment, total (% of total labor force) (modeled ILO estimate)', World Development Indicators. Source: ILO Modelled Estimates database (ILOEST). Accessed: January 17, 2026.
URL: *https://data.worldbank.org/indicator/SL.UEM.TOTL.ZS*

World Bank (2026*e*), 'Urban population (% of total population)', World Development Indicators. Source: United Nations Population Division. Accessed: May 6, 2026.
URL: *https://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS*

Wu, Tong (2024), 'Inflation and reproduction: the impact of inflation targeting on fertility in emerging economies', *China Population and Development Studies* **8**(4), 337–355.
URL: *https://doi.org/10.1007/s42379-024-00167-2*

Appendix A. AI Statement, Data Analysis, and Diagnostics

AI statement

The ideas developed in this research are entirely my own and stem from the demographic and behavioral finance courses I completed during my two-year master's program, as well as from my personal interest in identifying a coherent conceptual common ground. Generative AI tools — specifically Claude, Gemini, and Copilot — were used primarily for language refinement, most extensively in the introduction and conclusion to ensure clarity of expression, and occasionally to improve specific phrasings elsewhere in the thesis. I wrote the Stata and LaTeX code myself, drawing on similarities with analyses I had previously conducted; only in cases of substantial technical difficulty, after unsuccessful independent attempts and online searches, did I consult AI tools for code syntax. All interpretations of results, analytical reasoning, and substantive insights are exclusively my own.

Table A.1: List of Countries

	TFR	Inflation	Unemployment	GDP contraction	Risk taking
AFG	6.10	6.82	9.28	-0.05	0.12
ARE	1.71	3.12	2.61	-0.03	0.09
ARG	2.22	26.70	10.00	-0.06	0.04
AUS	1.79	2.86	5.37	-0.06	0.14
AUT	1.43	2.46	5.17	-0.03	-0.06
BGD	2.50	6.35	4.28	-0.08	-0.20
BIH	1.32	2.40	23.92	-0.09	-0.13
BOL	3.16	4.07	3.12	-0.07	0.10
BRA	1.83	6.35	9.95	-0.06	-0.25
BWA	3.04	6.74	19.49	-0.04	0.70
CAN	1.55	2.22	6.99	-0.04	0.18
CHE	1.47	0.61	4.21	-0.04	-0.02
CHL	1.70	3.73	8.58	-0.06	0.13
CHN	1.57	2.09	4.48	-0.12	-0.02
CMR	5.02	2.60	4.30	-0.04	-0.54
COL	2.01	5.33	11.64	-0.06	-0.05
CRI	1.83	6.33	8.27	-0.07	0.00
CZE	1.47	3.23	5.50	-0.07	-0.02
DEU	1.43	1.91	6.32	-0.03	-0.04
DZA	2.81	4.69	14.48	-0.06	0.39
EGY	3.22	10.36	10.08	-0.05	-0.28
ESP	1.29	2.35	15.54	-0.04	-0.16
EST	1.54	4.14	8.53	-0.09	-0.30
FIN	1.66	1.84	8.48	-0.03	-0.28
FRA	1.92	1.74	8.89	-0.03	-0.03
GBR	1.75	2.45	5.48	-0.03	0.05
GEO	1.87	5.38	14.68	-0.12	-0.08
GHA	4.10	15.93	5.15	-0.10	0.62
GRC	1.35	2.18	15.31	-0.03	-0.16
GTM	3.34	5.47	2.77	-0.06	-0.22
HRV	1.46	2.68	11.89	-0.07	0.07
HTI	3.40	15.07	13.93	-0.06	0.02
HUN	1.41	5.27	6.66	-0.07	-0.50
IDN	2.37	5.97	5.36	-0.09	-0.32
IND	2.56	6.23	7.29	-0.08	-0.28
IRN	1.83	21.25	11.30	-0.06	0.34
IRQ	4.16	10.81	10.61	-0.09	0.17
ISR	2.97	1.68	7.64	-0.04	0.24
ITA	1.33	2.10	9.31	-0.03	-0.09
JOR	3.40	3.04	14.97	-0.05	-0.12
JPN	1.35	0.32	3.89	0.00	-0.36
KAZ	2.59	8.72	6.55	-0.12	0.25
KEN	4.18	7.62	3.45	-0.07	0.24
KHM	2.97	3.86	0.61	-0.09	-0.40
KOR	1.12	2.50	3.34	-0.05	-0.04
LKA	2.15	9.34	5.67	-0.07	0.06
LTU	1.42	3.45	10.05	-0.10	-0.05
MAR	2.50	1.90	10.08	-0.04	-0.07
MDA	1.67	9.61	3.17	-0.13	-0.03
MEX	2.28	4.76	3.90	-0.04	-0.14
MWI	4.93	13.86	5.04	-0.06	0.48
NGA	5.58	13.11	4.10	-0.08	0.39
NIC	2.57	7.20	5.83	-0.04	-0.55
NLD	1.67	2.31	4.55	-0.04	0.19
PAK	4.30	9.10	2.37	-0.06	0.02
PER	2.38	3.11	4.19	-0.06	0.15
PHL	2.98	4.32	3.27	-0.06	0.29
POL	1.33	3.56	9.77	-0.07	-0.07
PRT	1.38	2.17	8.57	-0.04	-0.79
ROU	1.55	9.31	6.34	-0.12	-0.23
RUS	1.49	9.82	6.31	-0.12	-0.32
RWA	4.65	6.66	12.00	-0.06	-0.29
SAU	2.91	2.28	5.64	-0.06	0.70
SRB	1.49	12.52	15.63	-0.10	-0.13
SUR	2.62	20.88	8.70	-0.06	0.18
SWE	1.77	1.79	7.07	-0.03	0.05
THA	1.50	1.99	1.11	-0.06	-0.12
TUR	2.07	20.17	10.63	-0.06	0.02
TZA	5.19	6.39	2.80	-0.05	0.49
UGA	5.69	6.22	3.07	-0.07	0.16
UKR	1.28	12.60	8.65	-0.11	-0.22
USA	1.89	2.56	5.76	-0.04	0.12
VEN	2.39	3790.16	8.44	-0.03	0.25
VNM	1.96	5.79	1.80	-0.11	-0.01
ZAF	2.39	5.28	25.51	-0.04	0.97
ZWE	3.86	76.34	6.03	-0.10	0.52

Note: Countries are identified using the ISO 3166-1 alpha-3 codes. The table reports the mean values (2000–2023) of total fertility rate (TFR), inflation, unemployment, GDP contraction, and willingness to take risks (the latter being time-invariant).

Table A.2: Correlation Matrix

	TFR	HCPI	Unempl.	GDP contraction	Risk Taking	Urbanization	Female Education	Infant Mortality	GDP per Capita
TFR	1								
HCPI	0.106***	1							
Unempl.	-0.168***	0.00964	1						
GDP contraction	-0.0154	-0.0831**	0.0683*	1					
Risk Taking	0.276***	0.109***	0.120***	0.0264	1				
Urbanization	-0.537***	-0.0600*	0.184***	0.118***	0.0594*	1			
Female Education	-0.690***	-0.0568*	0.134***	0.0928***	-0.161***	0.699***	1		
Infant Mortality	0.862***	0.141***	-0.139***	-0.0772**	0.147***	-0.673***	-0.760***	1	
GDP per Capita	-0.410***	-0.247***	-0.123***	0.131***	0.0128	0.577***	0.585***	-0.530***	1

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: This table reports pairwise Pearson correlation coefficients, ranging from 1 to +1, for all variables used in the analysis, based on the listwise-complete sample. TFR denotes the total fertility rate; HCPI the headline consumer price index; GDP contraction is defined as the negative of real GDP growth; Risk Taking the measure from the Global Preference Survey; all other variables are defined in Section 3. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table A.3: Variance Inflation Factors (VIFs)

Variable	VIF	1/VIF
Inflation	1.130	0.887
Unemployment	1.180	0.845
GDP contraction	1.030	0.971
Risk Taking	1.160	0.861
Urbanization	2.530	0.395
Female Education	3.150	0.317
Infant Mortality	2.720	0.367
GDP per Capita	2.040	0.490
Mean	VIF	1.870

Note: VIFs are computed from the baseline regression with TFR as the dependent variable. The Variance Inflation Factor (VIF) measures how much the variance of an estimated coefficient is increased due to multicollinearity. Values below 5 are generally considered acceptable, and values close to 1 indicate negligible collinearity. Mean VIF = 1.87.

Figure A.1: Boxplot of Economic Shocks

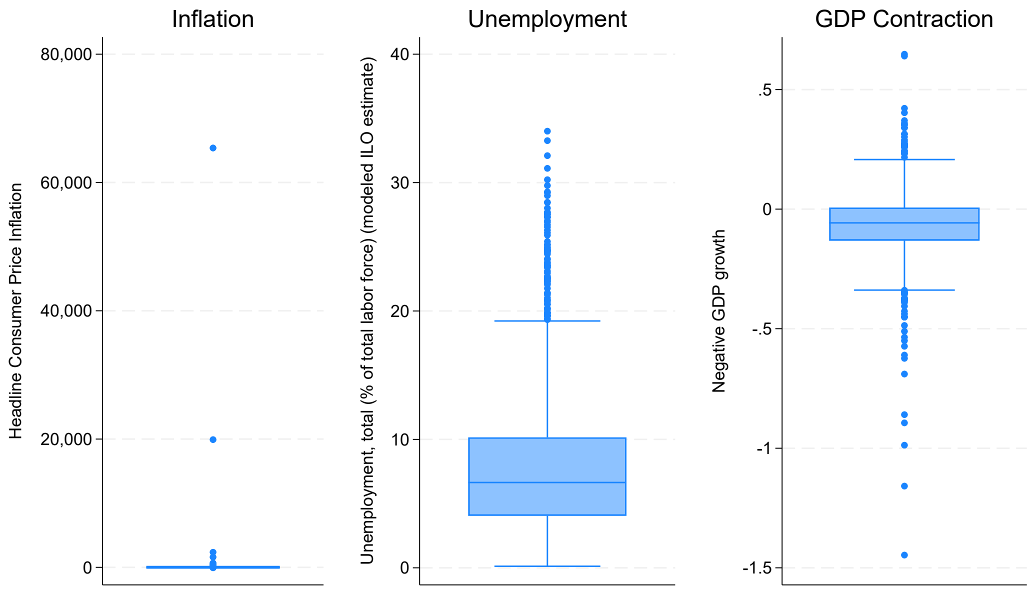

Note: The figure presents boxplots of the three macroeconomic shock variables — headline consumer price inflation, unemployment, and negative GDP growth — based on raw data. Each box displays the interquartile range (IQR), with the internal line denoting the median and whiskers extending to 1.5 times the IQR; observations beyond this range are plotted individually as outliers. The inflation panel reveals substantial right skew driven by extreme hyperinflation episodes, with values reaching up to approximately 65,000 percent. The unemployment panel likewise exhibits considerable right skew and dispersion, reflecting cross-country heterogeneity in labor market conditions. The GDP contraction panel is more tightly distributed around zero, though notable negative outliers are present, corresponding to severe recession episodes. The extreme inflation outliers visible in the left panel prompted the exclusion of observations with inflation rates exceeding 100 percent in subsequent analyses, so as to improve cross-country comparability and the robustness of estimation. Boxplots are produced using Stata's default `graph box` command.

Table A.4: List of Countries After Inflation Outlier Removal

	TFR	Inflation	Unemployment	GDP contraction	Risk taking
AFG	6.10	6.82	9.28	-0.05	0.12
ARE	1.71	3.12	2.61	-0.03	0.09
ARG	2.25	22.05	10.17	-0.06	0.04
AUS	1.79	2.86	5.37	-0.06	0.14
AUT	1.43	2.46	5.17	-0.03	-0.06
BGD	2.50	6.35	4.28	-0.08	-0.20
BIH	1.32	2.40	23.92	-0.09	-0.13
BOL	3.16	4.07	3.12	-0.07	0.10
BRA	1.83	6.35	9.95	-0.06	-0.25
BWA	3.04	6.74	19.49	-0.04	0.70
CAN	1.55	2.22	6.99	-0.04	0.18
CHE	1.47	0.61	4.21	-0.04	-0.02
CHL	1.70	3.73	8.58	-0.06	0.13
CHN	1.57	2.09	4.48	-0.12	-0.02
CMR	5.02	2.60	4.30	-0.04	-0.54
COL	2.01	5.33	11.64	-0.06	-0.05
CRI	1.83	6.33	8.27	-0.07	0.00
CZE	1.47	3.23	5.50	-0.07	-0.02
DEU	1.43	1.91	6.32	-0.03	-0.04
DZA	2.81	4.69	14.48	-0.06	0.39
EGY	3.22	10.36	10.08	-0.05	-0.28
ESP	1.29	2.35	15.54	-0.04	-0.16
EST	1.54	4.14	8.53	-0.09	-0.30
FIN	1.66	1.84	8.48	-0.03	-0.28
FRA	1.92	1.74	8.89	-0.03	-0.03
GBR	1.75	2.45	5.48	-0.03	0.05
GEO	1.87	5.38	14.68	-0.12	-0.08
GHA	4.10	15.93	5.15	-0.10	0.62
GRC	1.35	2.18	15.31	-0.03	-0.16
GTM	3.34	5.47	2.77	-0.06	-0.22
HRV	1.46	2.68	11.89	-0.07	0.07
HTI	3.40	15.07	13.93	-0.06	0.02
HUN	1.41	5.27	6.66	-0.07	-0.50
IDN	2.37	5.97	5.36	-0.09	-0.32
IND	2.56	6.23	7.29	-0.08	-0.28
IRN	1.83	21.25	11.30	-0.06	0.34
IRQ	4.16	10.81	10.61	-0.09	0.17
ISR	2.97	1.68	7.64	-0.04	0.24
ITA	1.33	2.10	9.31	-0.03	-0.09
JOR	3.40	3.04	14.97	-0.05	-0.12
JPN	1.35	0.32	3.89	0.00	-0.36
KAZ	2.59	8.72	6.55	-0.12	0.25
KEN	4.18	7.62	3.45	-0.07	0.24
KHM	2.97	3.86	0.61	-0.09	-0.40
KOR	1.12	2.50	3.34	-0.05	-0.04
LKA	2.15	9.34	5.67	-0.07	0.06
LTU	1.42	3.45	10.05	-0.10	-0.05
MAR	2.50	1.90	10.08	-0.04	-0.07
MDA	1.67	9.61	3.17	-0.13	-0.03
MEX	2.28	4.76	3.90	-0.04	-0.14
MWI	4.93	13.86	5.04	-0.06	0.48
NGA	5.58	13.11	4.10	-0.08	0.39
NIC	2.57	7.20	5.83	-0.04	-0.55
NLD	1.67	2.31	4.55	-0.04	0.19
PAK	4.30	9.10	2.37	-0.06	0.02
PER	2.38	3.11	4.19	-0.06	0.15
PHL	2.98	4.32	3.27	-0.06	0.29
POL	1.33	3.56	9.77	-0.07	-0.07
PRT	1.38	2.17	8.57	-0.04	-0.79
ROU	1.55	9.31	6.34	-0.12	-0.23
RUS	1.49	9.82	6.31	-0.12	-0.32
RWA	4.65	6.66	12.00	-0.06	-0.29
SAU	2.91	2.28	5.64	-0.06	0.70
SRB	1.49	12.52	15.63	-0.10	-0.13
SUR	2.62	20.88	8.70	-0.06	0.18
SWE	1.77	1.79	7.07	-0.03	0.05
THA	1.50	1.99	1.11	-0.06	-0.12
TUR	2.07	20.17	10.63	-0.06	0.02
TZA	5.19	6.39	2.80	-0.05	0.49
UGA	5.69	6.22	3.07	-0.07	0.16
UKR	1.28	12.60	8.65	-0.11	-0.22
USA	1.89	2.56	5.76	-0.04	0.12
VEN	2.54	25.85	9.93	-0.06	0.25
VNM	1.96	5.79	1.80	-0.11	-0.01
ZAF	2.39	5.28	25.51	-0.04	0.97
ZWE	3.90	-1.28	5.56	-0.16	0.52

Note: Countries are identified using the ISO 3166-1 alpha-3 codes. The table reports the mean values (2000–2023) of total fertility rate (TFR), inflation, unemployment, GDP contraction, and willingness to take risks (the latter being time-invariant) after excluding extreme inflation observations (HCPI >100%).

Figure A.2: Kernel Distribution of Economic Shocks

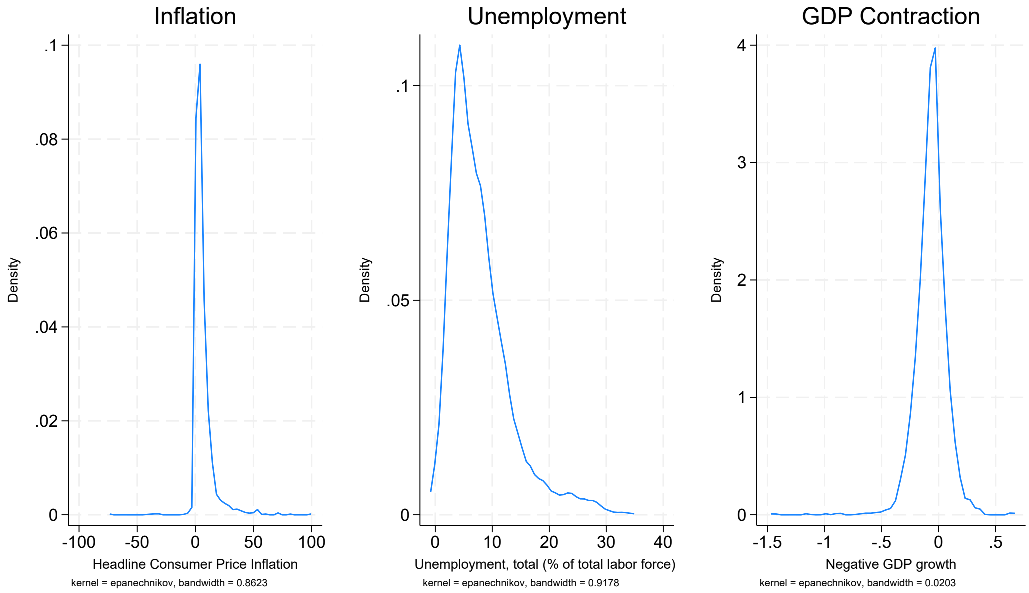

Note: The figure displays kernel density estimates for three economic shock variables. Densities are estimated using the Epanechnikov kernel function. Bandwidths are selected by Stata's default optimal bandwidth rule (minimizing mean integrated squared error), yielding bandwidths of 0.8623 for headline consumer price inflation, 0.9178 for unemployment (% of total labor force), and 0.0203 for negative GDP growth. The inflation distribution (excluding hyperinflation) is centered near zero but exhibits moderate right skew and fat tails, indicating occasional spikes in inflation. The unemployment distribution is right-skewed, consistent with the non-negative nature of unemployment rates and the prevalence of low-unemployment observations. The GDP contraction distribution is tightly concentrated near zero, indicating that large contractions are rare relative to moderate fluctuations.

Figure A.3: Histograms with Normal Overlay of Economic Shocks

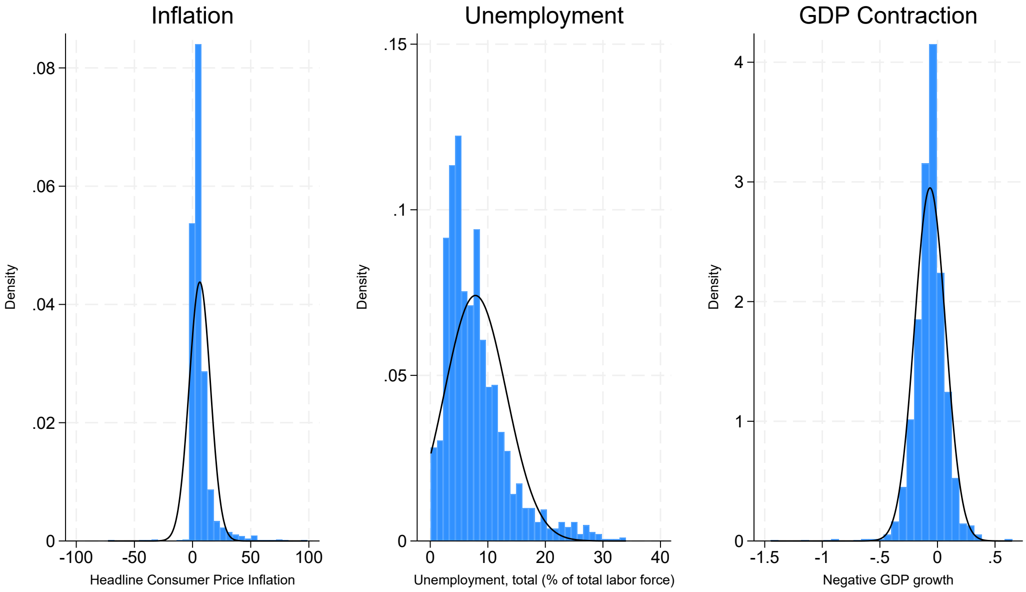

Note: Histograms are produced after removing inflation observations exceeding 100% (16 observations dropped). All three variables are in their raw, unstandardised form. Inflation is measured as headline consumer price inflation (%); unemployment as a percentage of the total labour force; GDP contraction as negative annual GDP per capita growth. The normal overlay is computed using the sample mean and standard deviation of each variable.

Figure A.4: Raw vs Log Distribution of Total Fertility Rate

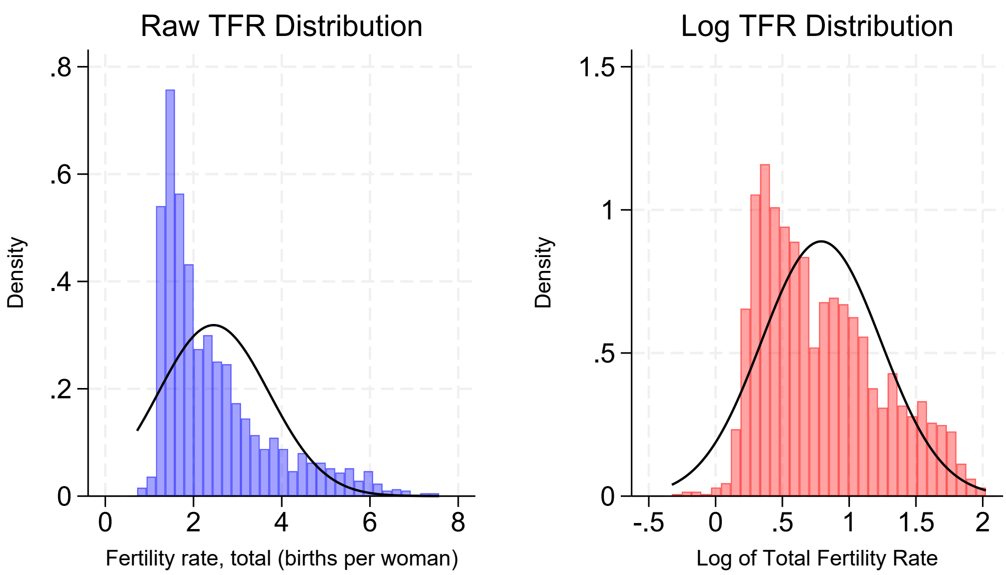

Note: This figure displays the frequency distribution of the Total Fertility Rate (TFR) in its raw form (left) and after a natural logarithmic transformation (right). The red and blue shaded areas represent the observed data, while the solid black curves represent the theoretical normal distribution. The log-transformation is employed to reduce right-skewness and ensure the dependent variable is more suitable for linear regression analysis, facilitating the interpretation of coefficients as semi-elasticities.

Table A.5: Diagnostic Tests for Serial Autocorrelation and Cross-Sectional Dependence

	Serial Autocorrelation		*Cross-Sectional Dependence*		
Full Models	Modified Bhargava et al. Durbin–Watson	Baltagi–Wu LBI	Pesaran's test of cross sectional independence	Pr	Average absolute value of the off-diagonal elements
Baseline (No Controls)	0.1224	0.3159	-2.482	0.0131	0.608
Baseline (With Controls)	0.1970	0.5080		—	
Interaction with Inflation	0.1983	0.5113		—	
Interaction with Unemployment	0.1970	0.5131		—	
Interaction with GDP Contraction	0.2017	0.5139		—	
Expected under H_0	≈ 2.00	≈ 2.00		> 0.05	

Note: The Modified Bhargava et al. Durbin–Watson and Baltagi–Wu LBI tests assess first-order serial autocorrelation; values substantially below 2 indicate positive autocorrelation. The Pesaran (2021) CD test for cross-sectional dependence was only feasible for the baseline model with no controls, which is the only sufficiently balanced panel specification. The Pesaran CD test rejects the null of cross-sectional independence ($|2.482| > 1.96$, $p < 0.05$). The average absolute pairwise correlation of 0.608 is also a critical metric in this output. It indicates that, on average, the residuals between countries are correlated at 61%. This level of correlation proves that these nations are deeply integrated and subject to the same global or regional shocks. Together, these results justify the use of Driscoll–Kraay standard errors, which are robust to heteroskedasticity, serial autocorrelation, and cross-sectional dependence.

Appendix B. Heterogeneity by Income level

Table B.1: List of Countries by Income Level

	TFR	Inflation	Unemployment	GDP Contraction	Risk Taking
Panel A: Low GDP p.c.					
AFG	6.10	6.82	9.28	-0.05	0.12
BGD	2.50	6.35	4.28	-0.08	-0.20
BOL	3.16	4.07	3.12	-0.07	0.10
CMR	5.02	2.60	4.30	-0.04	-0.54
EGY	3.22	10.36	10.08	-0.05	-0.28
GHA	4.10	15.93	5.15	-0.10	0.62
GTM	3.34	5.47	2.77	-0.06	-0.22
HTI	3.40	15.07	13.93	-0.06	0.02
IDN	2.37	5.97	5.36	-0.09	-0.32
IND	2.56	6.23	7.29	-0.08	-0.28
KEN	4.18	7.62	3.45	-0.07	0.24
KHM	2.97	3.86	0.61	-0.09	-0.40
LKA	2.15	9.34	5.67	-0.07	0.06
MAR	2.50	1.90	10.08	-0.04	-0.07
MDA	1.67	9.61	3.17	-0.13	-0.03
MWI	4.93	13.86	5.04	-0.06	0.48
NGA	5.58	13.11	4.10	-0.08	0.39
NIC	2.57	7.20	5.83	-0.04	-0.55
PAK	4.30	9.10	2.37	-0.06	0.02
PHL	2.98	4.32	3.27	-0.06	0.29
RWA	4.65	6.66	12.00	-0.06	-0.29
TZA	5.19	6.39	2.80	-0.05	0.49
UGA	5.69	6.22	3.07	-0.07	0.16
UKR	1.28	12.60	8.65	-0.11	-0.22
VNM	1.96	5.79	1.80	-0.11	-0.01
ZWE	3.90	-1.28	5.56	-0.16	0.52
Panel B: Medium GDP p.c.					
ARG	2.25	22.05	10.17	-0.06	0.04
BIH	1.32	2.40	23.92	-0.09	-0.13
BRA	1.83	6.35	9.95	-0.06	-0.25
BWA	3.04	6.74	19.49	-0.04	0.70
CHL	1.70	3.73	8.58	-0.06	0.13
CHN	1.57	2.09	4.48	-0.12	-0.02
COL	2.01	5.33	11.64	-0.06	-0.05
CRI	1.83	6.33	8.27	-0.07	0.00
DZA	2.81	4.69	14.48	-0.06	0.39
GEO	1.87	5.38	14.68	-0.12	-0.08
IRN	1.83	21.25	11.30	-0.06	0.34
IRQ	4.16	10.81	10.61	-0.09	0.17
JOR	3.40	3.04	14.97	-0.05	-0.12
KAZ	2.59	8.72	6.55	-0.12	0.25
MEX	2.28	4.76	3.90	-0.04	-0.14
PER	2.38	3.11	4.19	-0.06	0.15
POL	1.33	3.56	9.77	-0.07	-0.07
ROU	1.55	9.31	6.34	-0.12	-0.23
RUS	1.49	9.82	6.31	-0.12	-0.32
SRB	1.49	12.52	15.63	-0.10	-0.13
SUR	2.62	20.88	8.70	-0.06	0.18
THA	1.50	1.99	1.11	-0.06	-0.12
TUR	2.07	20.17	10.63	-0.06	0.02
VEN	2.54	25.85	9.93	-0.06	0.25
ZAF	2.39	5.28	25.51	-0.04	0.97
Panel C: High GDP p.c.					
ARE	1.71	3.12	2.61	-0.03	0.09
AUS	1.79	2.86	5.37	-0.06	0.14
AUT	1.43	2.46	5.17	-0.03	-0.06
CAN	1.55	2.22	6.99	-0.04	0.18
CHE	1.47	0.61	4.21	-0.04	-0.02
CZE	1.47	3.23	5.50	-0.07	-0.02
DEU	1.43	1.91	6.32	-0.03	-0.04
ESP	1.29	2.35	15.54	-0.04	-0.16
EST	1.54	4.14	8.53	-0.09	-0.30
FIN	1.66	1.84	8.48	-0.03	-0.28
FRA	1.92	1.74	8.89	-0.03	-0.03
GBR	1.75	2.45	5.48	-0.03	0.05
GRC	1.35	2.18	15.31	-0.03	-0.16
HRV	1.46	2.68	11.89	-0.07	0.07
HUN	1.41	5.27	6.66	-0.07	-0.50
ISR	2.97	1.68	7.64	-0.04	0.24
ITA	1.33	2.10	9.31	-0.03	-0.09
JPN	1.35	0.32	3.89	0.00	-0.36
KOR	1.12	2.50	3.34	-0.05	-0.04
LTU	1.42	3.45	10.05	-0.10	-0.05
NLD	1.67	2.31	4.55	-0.04	0.19
PRT	1.38	2.17	8.57	-0.04	-0.79
SAU	2.91	2.28	5.64	-0.06	0.70
SWE	1.77	1.79	7.07	-0.03	0.05
USA	1.89	2.56	5.76	-0.04	0.12

Note: Countries are identified using the ISO 3166-1 alpha-3 codes. The grouping into low-, medium-, and high-income categories is based on the distribution of countries' average GDP per capita over the period 2000–2023. Specifically, the 33rd and 67th percentiles of the cross-country distribution of average GDP per capita were computed, and these percentile cut-offs were used to assign countries to the three groups. The table reports the mean values (2000–2023) of total fertility rate (TFR), inflation, unemployment, GDP contraction, and willingness to take risks (the latter being time-invariant).

Table B.2: Regression Results — No controls (Heterogeneity by Income level)

	Dependent variable: ln(Total Fertility Rate)			
	(1)	(2)	(3)	(4)
	Inflation	Unemployment	GDP contraction	Full model
Panel A: Low GDP p.c.				
Headline Consumer Price Inflation	-0.00320			0.000690
	(0.00585)			(0.00549)
Unemployment		-0.0884***		-0.0898***
		(0.0262)		(0.0275)
GDP contraction			0.00269	0.00430
			(0.00300)	(0.00323)
Observations	589	588	588	587
Within R^2	0.661	0.687	0.659	0.687
Panel B: Medium GDP p.c.				
Headline Consumer Price Inflation	-0.00896**			-0.0168***
	(0.00412)			(0.00491)
Unemployment		-0.0651***		-0.0706***
		(0.0166)		(0.0173)
GDP contraction			0.00707	0.0117**
			(0.00587)	(0.00477)
Observations	588	588	588	588
Within R^2	0.212	0.306	0.210	0.330
Panel C: High GDP p.c.				
Headline Consumer Price Inflation	0.0177			-0.0155
	(0.0416)			(0.0420)
Unemployment		-0.0551***		-0.0576***
		(0.0141)		(0.0135)
GDP contraction			-0.00589	0.00562
			(0.0200)	(0.0217)
Observations	552	552	552	552
Within R^2	0.207	0.263	0.207	0.264

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). Low GDP p.c..

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the baseline specification, estimated separately for each income tercile, without controls. Countries are assigned to terciles based on their time-averaged GDP per capita over the period 2000–2023, with cutpoints at the 33rd and 67th percentiles of the cross-country distribution. All macroeconomic shock variables are z-standardised and lagged by one period ($t-1$). All specifications include country fixed effects, country-specific linear trends, and year fixed effects.

Table B.3: Regression Results with Controls (Heterogeneity by Income level)

	Dependent variable: ln(Total Fertility Rate)			
	(1)	(2)	(3)	(4)
	Inflation	Unemployment	GDP contraction	Full model
Panel A: Low GDP p.c.				
Headline Consumer Price Inflation	-0.000701			-0.00265
	(0.00844)			(0.00761)
Unemployment		0.0129		0.0126
		(0.0271)		(0.0257)
GDP contraction			0.00678*	0.00694*
			(0.00380)	(0.00352)
Urbanization	-0.0337	-0.0291	-0.0318	-0.0287
	(0.0382)	(0.0436)	(0.0372)	(0.0417)
Female Education	0.212***	0.218***	0.208***	0.213***
	(0.0325)	(0.0301)	(0.0330)	(0.0307)
Infant Mortality	0.0307	0.0314	0.0317	0.0327
	(0.0271)	(0.0258)	(0.0268)	(0.0261)
Observations	374	374	374	374
Within R^2	0.710	0.710	0.712	0.712
Panel B: Medium GDP p.c.				
Headline Consumer Price Inflation	-0.0211***			-0.0242***
	(0.00639)			(0.00607)
Unemployment		-0.0620***		-0.0667***
		(0.0208)		(0.0197)
GDP contraction			-0.00221	0.00426
			(0.00578)	(0.00504)
Urbanization	-0.231***	-0.171***	-0.224***	-0.175***
	(0.0324)	(0.0229)	(0.0332)	(0.0210)
Female Education	-0.0197*	-0.0395***	-0.0255**	-0.0339***
	(0.0109)	(0.0134)	(0.0121)	(0.0108)
Infant Mortality	-0.267***	-0.331***	-0.262***	-0.342***
	(0.0446)	(0.0356)	(0.0465)	(0.0383)
Observations	448	448	448	448
Within R^2	0.354	0.400	0.331	0.432
Panel C: High GDP p.c.				
Headline Consumer Price Inflation	-0.0109			-0.0378
	(0.0274)			(0.0293)
Unemployment		-0.0495***		-0.0541***
		(0.0135)		(0.0129)
GDP contraction			-0.00196	0.00707
			(0.0129)	(0.0138)
Urbanization	-0.155*	-0.0720	-0.152*	-0.0737
	(0.0824)	(0.0656)	(0.0816)	(0.0702)
Female Education	-0.0123	0.00813	-0.0116	0.00735
	(0.0141)	(0.0173)	(0.0141)	(0.0172)
Infant Mortality	0.271**	0.339***	0.276**	0.327***
	(0.101)	(0.103)	(0.106)	(0.0948)
Observations	430	430	430	430
Within R^2	0.415	0.462	0.414	0.468

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). Low GDP p.c..

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the baseline specification (Equation 12) estimated separately for each income tercile, with controls. Countries are assigned to terciles based on their time-averaged GDP per capita over the period 2000–2023, with cutpoints at the 33rd and 67th percentiles of the cross-country distribution. All macroeconomic shock variables and controls (urbanisation, female tertiary enrolment, and infant mortality) are z-standardised and lagged by one period ($t-1$). All specifications include country fixed effects, country-specific linear trends, and year fixed effects.

Table B.4: Interaction Inflation x Risk Taking (Heterogeneity by Income level)

	Dependent variable: ln(Total Fertility Rate)		
	(1)	(2)	(3)
	Base	With Controls	Full Model
Panel A: Low GDP p.c.			
Headline Consumer Price Inflation	-0.00537	-0.000784	-0.00272
	(0.00895)	(0.00811)	(0.00733)
Inflation × Risk Taking	0.00436	-0.00340	-0.00337
	(0.00709)	(0.00663)	(0.00635)
Unemployment			0.0120
			(0.0257)
GDP contraction			0.00698*
			(0.00344)
Urbanization		-0.0324	-0.0276
		(0.0371)	(0.0407)
Female Education		0.213***	0.214***
		(0.0325)	(0.0307)
Infant Mortality		0.0306	0.0326
		(0.0270)	(0.0259)
Observations	589	374	374
Within R^2	0.661	0.710	0.713
Panel B: Medium GDP p.c.			
Headline Consumer Price Inflation	-0.0122*	-0.0214***	-0.0248***
	(0.00635)	(0.00720)	(0.00669)
Inflation × Risk Taking	0.0221***	0.00318	0.00524
	(0.00418)	(0.00957)	(0.00806)
Unemployment			-0.0667***
			(0.0197)
GDP contraction			0.00358
			(0.00498)
Urbanization		-0.230***	-0.172***
		(0.0327)	(0.0207)
Female Education		-0.0198*	-0.0340***
		(0.0108)	(0.0108)
Infant Mortality		-0.266***	-0.340***
		(0.0470)	(0.0369)
Observations	588	448	448
Within R^2	0.223	0.355	0.433
Panel C: High GDP p.c.			
Headline Consumer Price Inflation	0.0123	-0.0211	-0.0445*
	(0.0437)	(0.0228)	(0.0232)
Inflation × Risk Taking	-0.0249	-0.0348**	-0.0258
	(0.0297)	(0.0144)	(0.0190)
Unemployment			-0.0522***
			(0.0132)
GDP contraction			0.00533
			(0.0146)
Urbanization		-0.142*	-0.0665
		(0.0781)	(0.0690)
Female Education		-0.0181	0.00247
		(0.0143)	(0.0182)
Infant Mortality		0.236***	0.299***
		(0.0774)	(0.0835)
Observations	552	430	430
Within R^2	0.211	0.423	0.472

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). Low GDP p.c..

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the interaction model between headline consumer price inflation and the cultural risk-taking index, estimated separately for each income tercile. Countries are assigned to terciles based on their time-averaged GDP per capita over the period 2000–2023, with cutpoints at the 33rd and 67th percentiles of the cross-country distribution. All macroeconomic shock variables and controls are z-standardised and lagged by one period ($t-1$). Risk-taking is standardised but time-invariant and therefore not lagged; its main effect is absorbed by the country fixed effects and not separately identified.

Table B.5: Interaction Unemployment x Risk Taking (Heterogeneity br Income level)

	Dependent variable: ln(Total Fertility Rate)		
	(1)	(2)	(3)
	Base	With Controls	Full Model
Panel A: Low GDP p.c.			
Unemployment	-0.0872***	0.0209	0.0202
	(0.0270)	(0.0219)	(0.0207)
Unemployment × Risk Taking	0.0244	0.0238	0.0225
	(0.0145)	(0.0240)	(0.0256)
Headline Consumer Price Inflation			-0.00220
			(0.00778)
GDP contraction			0.00679*
			(0.00351)
Urbanization		-0.0267	-0.0263
		(0.0444)	(0.0425)
Female Education		0.213***	0.208***
		(0.0285)	(0.0287)
Infant Mortality		0.0327	0.0340
		(0.0252)	(0.0254)
Observations	588	374	374
Within R^2	0.689	0.711	0.713
Panel B: Medium GDP p.c.			
Unemployment	-0.0700***	-0.0627***	-0.0674***
	(0.0194)	(0.0210)	(0.0198)
Unemployment × Risk Taking	0.0164	0.0261**	0.0232**
	(0.0117)	(0.0108)	(0.00971)
Headline Consumer Price Inflation			-0.0232***
			(0.00579)
GDP contraction			0.00554
			(0.00485)
Urbanization		-0.163***	-0.166***
		(0.0247)	(0.0209)
Female Education		-0.0393***	-0.0342***
		(0.0134)	(0.0107)
Infant Mortality		-0.336***	-0.346***
		(0.0395)	(0.0406)
Observations	588	448	448
Within R^2	0.312	0.409	0.439
Panel C: High GDP p.c.			
Unemployment	-0.0489**	-0.0367**	-0.0409**
	(0.0175)	(0.0156)	(0.0157)
Unemployment × Risk Taking	0.0111	0.0295**	0.0304**
	(0.0118)	(0.0114)	(0.0127)
Headline Consumer Price Inflation			-0.0414
			(0.0282)
GDP contraction			0.00529
			(0.0137)
Urbanization		-0.0975	-0.100
		(0.0683)	(0.0722)
Female Education		0.0109	0.0101
		(0.0176)	(0.0176)
Infant Mortality		0.375***	0.362***
		(0.102)	(0.0938)
Observations	552	430	430
Within R^2	0.265	0.471	0.477

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). Low GDP p.c..

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the interaction model between unemployment and the cultural risk-taking index, estimated separately for each income tercile. Countries are assigned to terciles based on their time-averaged GDP per capita over the period 2000–2023, with cutpoints at the 33rd and 67th percentiles of the cross-country distribution. All macroeconomic shock variables and controls are z-standardised and lagged by one period ($t-1$). Risk-taking is standardised but time-invariant and therefore not lagged; its main effect is absorbed by the country fixed effects and not separately identified.

Table B.6: Interaction GDP Contraction x Risk Taking (Heterogeneity by Income level)

	Dependent variable: ln(Total Fertility Rate)		
	(1)	(2)	(3)
	Base	With Controls	Full Model
Panel A: Low GDP p.c.			
GDP contraction	0.0106**	0.00983**	0.0101**
	(0.00439)	(0.00463)	(0.00450)
GDP Contraction × Risk Taking	-0.00926***	-0.00587**	-0.00588**
	(0.00245)	(0.00256)	(0.00255)
Headline Consumer Price Inflation			-0.00321
			(0.00702)
Unemployment			0.0115
			(0.0259)
Urbanization		-0.0367	-0.0342
		(0.0378)	(0.0427)
Female Education		0.201***	0.206***
		(0.0324)	(0.0297)
Infant Mortality		0.0314	0.0324
		(0.0264)	(0.0258)
Observations	588	374	374
Within R^2	0.666	0.715	0.715
Panel B: Medium GDP p.c.			
GDP contraction	0.00805	-0.00239	0.00450
	(0.00630)	(0.00598)	(0.00527)
GDP Contraction × Risk Taking	-0.00380*	0.00139	-0.00153
	(0.00221)	(0.00317)	(0.00352)
Headline Consumer Price Inflation			-0.0241***
			(0.00606)
Unemployment			-0.0671***
			(0.0198)
Urbanization		-0.224***	-0.174***
		(0.0333)	(0.0206)
Female Education		-0.0255**	-0.0341***
		(0.0122)	(0.0109)
Infant Mortality		-0.261***	-0.343***
		(0.0466)	(0.0384)
Observations	588	448	448
Within R^2	0.211	0.331	0.432
Panel C: High GDP p.c.			
GDP contraction	-0.00574	-0.00213	0.00765
	(0.0199)	(0.0132)	(0.0140)
GDP Contraction × Risk Taking	-0.00345	0.000968	-0.00262
	(0.00631)	(0.00361)	(0.00400)
Headline Consumer Price Inflation			-0.0366
			(0.0301)
Unemployment			-0.0548***
			(0.0135)
Urbanization		-0.152*	-0.0751
		(0.0812)	(0.0701)
Female Education		-0.0119	0.00861
		(0.0141)	(0.0176)
Infant Mortality		0.274**	0.332***
		(0.105)	(0.0964)
Observations	552	430	430
Within R^2	0.208	0.414	0.468

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). Low GDP p.c..

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the interaction model between GDP contraction and the cultural risk-taking index, estimated separately for each income tercile. Countries are assigned to terciles based on their time-averaged GDP per capita over the period 2000–2023, with cutpoints at the 33rd and 67th percentiles of the cross-country distribution. All macroeconomic shock variables and controls are z-standardised and lagged by one period ($t-1$). Risk-taking is standardised but time-invariant and therefore not lagged; its main effect is absorbed by the country fixed effects and not separately identified.

Appendix C. Robustness checks

Table C.1: Robustness: Excluding Crisis Episodes

	Dependent variable: ln(Total Fertility Rate)			
	(1)	(2)	(3)	(4)
	Full Sample	Excl. GFC	Excl. COVID	Excl. Both
Panel A: No Controls				
Headline Consumer Price Inflation	-0.0110**	-0.0109**	-0.00984**	-0.00974**
	(0.00405)	(0.00418)	(0.00399)	(0.00412)
Unemployment	-0.0640***	-0.0641***	-0.0582***	-0.0582***
	(0.0163)	(0.0162)	(0.0158)	(0.0158)
GDP contraction	0.00941*	0.00917	0.00941*	0.00923
	(0.00526)	(0.00569)	(0.00518)	(0.00559)
Observations	1727	1652	1653	1578
Within R^2	0.362	0.357	0.349	0.344
Panel B: Controls				
Headline Consumer Price Inflation	-0.0195***	-0.0195***	-0.0178***	-0.0178***
	(0.00425)	(0.00416)	(0.00418)	(0.00408)
Unemployment	-0.0439***	-0.0423***	-0.0387***	-0.0370**
	(0.0139)	(0.0137)	(0.0134)	(0.0132)
GDP contraction	0.0117**	0.0111**	0.0116**	0.0110**
	(0.00446)	(0.00470)	(0.00442)	(0.00464)
Urbanization	-0.105***	-0.109***	-0.107***	-0.111***
	(0.0264)	(0.0262)	(0.0285)	(0.0283)
Female Education	0.00149	0.000911	0.00425	0.00356
	(0.0108)	(0.0106)	(0.0111)	(0.0109)
Infant Mortality	0.0878***	0.0878***	0.0895***	0.0896***
	(0.0142)	(0.0143)	(0.0152)	(0.0152)
Observations	1252	1204	1189	1141
Within R^2	0.398	0.393	0.393	0.389

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). GFC: 2008. COVID: 2021.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the baseline specification re-estimated sequentially dropping the peak year of the Global Financial Crisis (2008), the peak year of the COVID-19 shock (2021), and both years simultaneously, without controls. The full-sample estimates are reported in column (1) as a reference. All macroeconomic shock variables are z-standardised and lagged by one period ($t - 1$).

Table C.2: Robustness: Excluding Crisis Episodes, Interaction Model with Unemployment

	Dependent variable: ln(Total Fertility Rate)			
	(1)	(2)	(3)	(4)
	Full Sample	Excl. GFC	Excl. COVID	Excl. Both
Unemployment	-0.0397***	-0.0384***	-0.0341**	-0.0326**
	(0.0136)	(0.0135)	(0.0128)	(0.0127)
Unemployment × Risk Taking	0.0203**	0.0200**	0.0222**	0.0219**
	(0.00872)	(0.00910)	(0.00938)	(0.00973)
Headline Consumer Price Inflation	-0.0190***	-0.0189***	-0.0172***	-0.0172***
	(0.00421)	(0.00413)	(0.00408)	(0.00400)
GDP contraction	0.0119**	0.0113**	0.0119**	0.0113**
	(0.00438)	(0.00459)	(0.00432)	(0.00452)
Urbanization	-0.104***	-0.108***	-0.106***	-0.110***
	(0.0284)	(0.0284)	(0.0307)	(0.0307)
Female Education	0.00271	0.00216	0.00590	0.00524
	(0.0112)	(0.0109)	(0.0114)	(0.0111)
Infant Mortality	0.0887***	0.0888***	0.0903***	0.0905***
	(0.0148)	(0.0149)	(0.0158)	(0.0158)
Observations	1252	1204	1189	1141
Within R^2	0.402	0.397	0.398	0.394

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). GFC: 2008. COVID: 2021.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the interaction model between unemployment and the cultural risk-taking index, re-estimated sequentially dropping the peak year of the Global Financial Crisis (2008), the peak year of the COVID-19 shock (2021), and both years simultaneously. The full-sample estimates are reported in column (1) as a reference. All macroeconomic shock variables and controls are z-standardised and lagged by one period ($t - 1$). Risk-taking is standardised but time-invariant and therefore not lagged; its main effect is absorbed by the country fixed effects and not separately identified. All specifications include country fixed effects, country-specific linear trends, and year fixed effects.

Table C.3: Robustness: Country-Specific Quadratic Trends

	Dependent variable: ln(Total Fertility Rate)		
	(1)	(2)	(3)
	Baseline No Controls	Baseline With Controls	Unemployment × Risk
Unemployment	-0.0628***	-0.0442***	-0.0395***
	(0.0157)	(0.0135)	(0.0131)
Unemployment × Risk Taking			0.0231**
			(0.00833)
Headline Consumer Price Inflation	-0.0103**	-0.0191***	-0.0183***
	(0.00411)	(0.00403)	(0.00396)
GDP contraction	0.00892	0.0112**	0.0113**
	(0.00520)	(0.00448)	(0.00440)
Urbanization		-0.100***	-0.0986***
		(0.0245)	(0.0270)
Female Education		0.00400	0.00577
		(0.0112)	(0.0115)
Infant Mortality		0.0874***	0.0884***
		(0.0136)	(0.0140)
Observations	1727	1252	1252
Within R^2	0.367	0.401	0.406

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2). Both linear and quadratic country-specific trends included.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table reports fixed-effects estimates of the baseline specification and the unemployment interaction model augmented with country-specific quadratic time trends. Both a linear trend ($\alpha_i \times t$) and a quadratic trend ($\alpha_i \times t^2$) are included for each country, allowing for non-linear fertility trajectories such as those observed in countries that experienced fertility decline followed by partial recovery. All macroeconomic shock variables and controls are z-standardised and lagged by one period ($t - 1$). Risk-taking is standardised but time-invariant and therefore not lagged; its main effect is absorbed by the country fixed effects and not separately identified.

Table C.4: Robustness: Risk Taking Purified of Patience

	Dependent variable: ln(Total Fertility Rate)	
	(1)	(2)
	Baseline	Purified
Unemployment	-0.0397***	-0.0405***
	(0.0136)	(0.0142)
Unemployment × Risk taking	0.0203**	
	(0.00872)	
Unemployment × Risk taking purified		0.0224***
		(0.00776)
Headline Consumer Price Inflation	-0.0190***	-0.0189***
	(0.00421)	(0.00422)
GDP contraction	0.0119**	0.0119**
	(0.00438)	(0.00438)
Urbanization	-0.104***	-0.103***
	(0.0284)	(0.0284)
Female Education	0.00271	0.00321
	(0.0112)	(0.0112)
Infant Mortality	0.0887***	0.0886***
	(0.0148)	(0.0150)
Observations	1252	1252
Within R^2	0.402	0.403

Standard errors in parentheses

Driscoll-Kraay standard errors, lag(2)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Note: The table compares the baseline interaction model estimated with the original risk-taking index (column 1) against a version in which the risk-taking index has been purified of its correlation with patience (column 2). The purified measure is constructed by collapsing the panel to a single cross-section of country-level means, regressing risk-taking on patience, and retaining the residuals, which are by construction orthogonal to patience. Both measures are z-standardised prior to estimation. All macroeconomic shock variables and controls are z-standardised and lagged by one period ($t - 1$). All specifications include country fixed effects, country-specific linear trends, and year fixed effects.