Climatic Roots of Loss Aversion

Oded Galor, Viacheslav Savitskiy and Daniel C. Wainstock*

September 4, 2026

Abstract

This research explores the roots of loss aversion and the origins of the variation in its intensity across societies. It advances the hypothesis and establishes empirically that the evolution of loss aversion in the course of human history can be traced to the adaptation of humans to ancestral environments in which the subsistence consumption constraint constituted an existential cliff, exposing dynasties to the risk of extinction from adverse productivity shocks. The theory generates a seemingly paradoxical prediction. While individuals rationally favored prudent behavior and greater reliance on climatically insensitive modes of production in harsher environments, natural selection fostered loss neutrality. The spatial structure of climatic shocks further shaped loss neutrality. Greater idiosyncrasy in the spatial incidence of climatic shocks favored individuals who were sufficiently loss neutral to select into riskier but higher-expected-return modes of production. Exploiting experimentally elicited measures of loss aversion in a novel empirical design, the analysis establishes that individuals originating in regions characterized by greater exposure to adverse climatic shocks indeed have a lower predisposition toward loss aversion. Furthermore, decomposing climatic shocks into their idiosyncratic and aggregate components, the empirical analysis establishes, as predicted, that a greater degree of idiosyncrasy is associated with a lower degree of loss aversion.

Keywords: loss aversion, cultural evolution, preferences, growth, development

JEL Classification: D81, D91, Z10, O10, O40

---

*The authors are grateful to the editor and five anonymous referees for constructive comments and to Alberto Alesina, Roland Benabou, James Fenske, Stelios Michalopoulos, Louis Putterman, Felipe Valencia, and Fabrizio Zilibotti, and seminar and conference participants at the *NBER Political Economy Meeting*, *Deep Determinants of Comparative Development*, Brown, Hebrew University, and Princeton for helpful discussions. The experimental protocol received ethical approval from the University of Oxford Economics Departmental Research Ethics Committee (Economics DREC-2219681). Galor: Brown University, NBER, CEPR, IZA, and CES-Ifo. E-mail: Oded_Galor@brown.edu; Savitskiy: Culture, Cognition and Co-evolution Lab, Harvard University. E-mail: vsavitskiy@fas.harvard.edu; Wainstock: Oxford University. E-mail: daniel.crisostomowainstock@queens.ox.ac.uk.

1 Introduction

The intriguing phenomenon of loss aversion has profoundly influenced the understanding of human behavior. The predisposition to assign disproportionate weight to losses relative to comparable gains is evident across a wide range of experimental and observational studies and has far-reaching implications for economic and political outcomes. Loss aversion shapes asset prices, precautionary saving, labor-market adjustment, trade protection, voting behavior, and policy persistence.[1] Yet, despite the importance of this phenomenon for individual decision-making, aggregate economic activity, and, arguably, comparative economic development, the origins of loss aversion have remained largely obscured.

This research explores the roots of loss aversion and the origins of the variation in its intensity across societies. It advances the hypothesis and establishes empirically that the evolution of loss aversion in the course of human history can be traced to the adaptation of humans to ancestral environments in which the binding subsistence consumption constraint constituted a critical existential threshold, exposing dynasties to the risk of extinction from adverse productivity shocks. Notably, linking sensitivity to losses to the existential subsistence threshold that constrained humanity for most of its history provides a biological foundation for the reference point at the core of loss aversion.[2]

The study develops an evolutionary theory that captures the impact of this biological cliff on the evolution of loss aversion.[3] The theory generates a seemingly paradoxical prediction. While individuals rationally favored prudent behavior and greater reliance on climatically insensitive modes of production in harsher environments, natural selection fostered loss neutrality. The spatial structure of climatic shocks further shaped loss neutrality. Greater idiosyncrasy in the incidence of climatic shocks, rather than greater spatial correlation, favored individuals who were sufficiently loss neutral to select into riskier but higher expected-return modes of production. Thus, while greater frequency or severity of adverse climatic shocks made the adoption of climatically safe modes of production individually rational and increased dynastic survival in the short run, in environments characterized predominantly by idiosyncratic shocks, individuals who adopted climatically vulnerable modes of production enjoyed higher expected returns and greater expected reproductive success, and thus dominated the population in the long run. Intriguingly, these loss-neutral individuals exposed their dynasties to a greater risk of extinction, yet surviving dynasties among them expanded more rapidly and ultimately dominated the population.

Interestingly, greater frequency or severity of adverse climatic shocks reduced the degree of loss aversion regardless of whether shocks were idiosyncratic or aggregate in nature. An environment

[1] Tversky and Kahneman (1991), Kahneman et al. (1991), Benartzi and Thaler (1995), Jervis (1992), Levy (1996), Barberis and Huang (2001), Genesove and Mayer (2001), Shalev (2002), Kőszegi and Rabin (2009), Fisher and Montalto (2011), Alesina and Passarelli (2019).

[2] As humanity was liberated from the Malthusian trap and its binding subsistence constraint, sensitivity to losses persisted, yet the reference point became anchored in a wide range of contemporary economic and social benchmarks.

[3] Based on evidence from capuchin monkeys, Chen et al. (2006) argue that loss aversion is an innate and evolutionarily ancient feature of human preferences.

characterized by more severe adverse shocks raised the threshold degree of loss neutrality required for individuals to sort into climatically risky modes of production, raising the mean degree of loss neutrality within both the safe and the risky modes of production.[4]

While the severity of climatic shocks determined the fraction of individuals who were sufficiently loss neutral to sort into the risky mode of production, the degree of idiosyncrasy of these shocks determined the evolution of the relative size of this group and thus the degree of loss aversion in the long run. In environments characterized by idiosyncratic adverse climatic shocks, extinction among some loss-neutral dynasties engaged in risky production was offset by favorable realizations among others, yielding higher reproductive success than among loss-averse individuals in the safe, lower-output production mode, and allowing the loss-neutral group to dominate the population in the long run. In contrast, in environments where climatic shocks were more aggregate, adverse realizations drove the entire population engaged in risky production asymptotically toward extinction. Greater spatial correlation therefore tended to increase the degree of loss aversion in the long run, whereas greater idiosyncrasy tended to increase the degree of loss neutrality.

Viewed through the lens of the ancestral regions of contemporary populations, the theory generates two distinct empirical predictions. Individuals whose ancestral origins lie in regions characterized by greater exposure to adverse climatic shocks would be expected to exhibit lower loss aversion. Furthermore, holding adverse-shock frequency and severity constant and decomposing climatic shocks into idiosyncratic and aggregate components, individuals whose ancestral origins lie in regions characterized by greater idiosyncrasy of climatic shocks would be expected to exhibit lower loss aversion, whereas those whose ancestral origins lie in regions characterized by greater spatial correlation of these shocks would be expected to exhibit higher loss aversion.

A central challenge in assessing the impact of ancestral climatic shocks on loss aversion is the construction of reliable measures of loss aversion that permit ancestral influences to be distinguished from those of the contemporary environment. Measures of loss aversion may partly reflect current economic incentives and institutional conditions, obscuring the underlying predisposition toward loss aversion.

Addressing this challenge, this study develops a novel experimental methodology that disentangles the underlying predisposition toward loss aversion from confounding country-specific characteristics. It implements the epidemiological approach in an experimental setting, eliciting measures of loss aversion through an incentivized experiment among descendants of migrants from diverse ancestral origins, born and currently residing in the same host country. Their exposure to a com-

[4] Robson (1996) provides a biological foundation for expected and non-expected utility by deriving the fitness criterion for evaluating reproductive gambles under idiosyncratic and aggregate uncertainty, showing that when reproductive risk is idiosyncratic across individuals, the selected type maximizes expected offspring, whereas when reproductive risk is correlated across individuals, the selected type maximizes expected log offspring. In contrast, this study examines theoretically and empirically how the frequency or severity of adverse climatic shocks shapes the evolution of *loss aversion*. Notably, the central prediction that greater exposure to adverse climatic shocks selects for greater loss neutrality holds regardless of whether climatic shocks are idiosyncratic or aggregate. Holding adverse-shock frequency and severity constant, the nature of the shocks governs the relative long-run reproductive success of individuals engaged in the safe and risky modes of production and, consequently, the long-run prevalence of relatively loss-averse and loss-neutral individuals.

mon host-country environment allows variation across ancestral groups to more accurately reflect differences in the underlying predisposition toward loss aversion.[5]

Moreover, since the measured degree of loss aversion depends on the underlying degree of risk aversion, the experimental design elicits both dimensions, rather than imposing a common degree of risk aversion, permitting reliable cross-ancestry comparisons of loss aversion and allowing the analysis to establish that the proposed hypothesis pertains uniquely to loss aversion and is orthogonal to risk aversion. While imposing a common degree of risk aversion in the estimation of loss aversion is less consequential for within-country comparisons, since it preserves individuals' rank ordering, in the absence of direct experimental measures of risk aversion, variation in this dimension distorts cross-societal comparisons of loss aversion, affecting not only the magnitude of the differences but potentially also their ranking.

A further empirical challenge is the construction of a credible proxy for ancestral exposure to severe climatic shocks. The analysis draws on monthly global observations of the self-calibrating Palmer Drought Severity Index at a 0.5° spatial resolution over the twentieth century to proxy for the frequency of consequential drought conditions at each ancestral location. Moreover, motivated by the theory's secondary predictions, each shock is further decomposed into spatially correlated and idiosyncratic components according to the extent to which adjacent locations experienced the same shock. Reassuringly, the twentieth-century measure of drought exposure is highly correlated at both the grid-cell and country levels with its historical counterpart over 1–1500 CE, indicating that it provides a credible proxy for the climatic conditions that prevailed during the formative period of this evolutionary process. [6]

Using these measures, the empirical analysis examines the model's testable predictions. In accordance with the theory, it establishes that individuals originating in regions characterized by greater exposure to adverse climatic shocks have a lower predisposition toward loss aversion. Moreover, decomposing climatic shocks into their idiosyncratic and aggregate components, the analysis further establishes that greater idiosyncrasy is associated with lower loss aversion. Notably, these ancestral climatic shocks are not associated with *risk aversion*, consistent with the hypothesis that the evolutionary mechanism operated through the existential consequences of adverse shocks near the subsistence threshold rather than through exposure to uncertainty per se.

The analysis uncovers the enduring imprint of ancestral climatic shocks on loss aversion. Differences in ancestral climatic environments appear to lie at the roots of persistent differences in loss aversion across contemporary populations. Strikingly, the cross-societal dispersion of behavioral traits that shape contemporary economic and political choices bears the legacy of evolutionary adaptation to climatic shocks that, under the binding subsistence threshold, placed human survival at the edge of extinction over most of history.

[5] An additional advantage of this design is that participants are assessed under an identical experimental protocol and incentive structure, eliminating a source of heterogeneity that may arise when loss-aversion estimates are derived in different societal settings.

[6] This historical measure, constructed using the Paleo Hydrodynamics Data Assimilation reconstruction, cannot serve as the baseline measure of climatic exposure, since it relies on proxy-based reconstruction rather than direct observation and is substantially noisier, with markedly coarser spatial and temporal resolution.

Related Literature

The research relates to several strands of the literature. It provides the first account of the geographical origins of loss aversion and of the forces underlying its variation across societies. While existing explanations of loss aversion have emphasized cognitive mechanisms governing the perception and valuation of gains and losses (Stewart et al., 2006; Khaw et al., 2021), this study identifies an evolutionary mechanism through which ancestral exposure to adverse climatic shocks shaped the predisposition toward loss aversion. It further contributes to the literature on the evolution of preferences (e.g., Robson, 1996; Robson and Samuelson, 2011; Galor and Moav, 2002; Galor and Michalopoulos, 2012) and the roots of cultural evolution (e.g., Bisin and Verdier, 2001; Alesina et al., 2013; Galor and Özak, 2016; Buggle and Durante, 2021; Galor et al., 2026), while laying the conceptual and empirical groundwork for incorporating cross-societal variation in loss aversion into the study of comparative economic development (e.g., Ashraf and Galor, 2013; Spolaore and Wacziarg, 2013; Michalopoulos and Papaioannou, 2020; Mayshar et al., 2022).

2 An Evolutionary Theory of Loss Aversion

This section advances an evolutionary theory that captures the critical role of climatic shocks in the evolution of loss aversion over the course of human history. The theory suggests that contemporary variation in the intensity of loss aversion across individuals can be traced to the adaptation of their ancestral populations to adverse climatic shocks that affected agricultural productivity and reproductive success during the Malthusian era.

The model captures a fundamental asymmetry that the Malthusian environment generated in the evolution of attitudes toward gains and losses in productivity, molding the predisposition toward loss aversion. In light of existing evidence that resources per capita during most of human existence were near the subsistence consumption level, and that temporary surpluses were ultimately counterbalanced by population growth, lineages subjected to sufficiently severe adverse transitory productivity shocks faced a substantial risk of extinction, whereas lineages that experienced comparable ordinary climatic realizations enjoyed only a short-lived increase in reproductive success before population growth returned resources per capita toward subsistence and restored their proximity to the evolutionary cliff (Ashraf and Galor, 2011; Vollrath, 2011; Dalgaard and Strulik, 2015).

Hence, in view of the fundamental asymmetry in survival prospects associated with being above or below subsistence during the Malthusian epoch, relatively more loss-averse individuals favored safe agricultural practices that secured subsistence consumption and minimized the risk of catastrophic realizations. Relatively less loss-averse individuals, in contrast, favored riskier agricultural practices with higher expected returns, at the cost of a greater risk of extinction.

The central implication of the model is that greater exposure to adverse climatic shocks shifted the long-run composition of ancestral populations toward greater loss neutrality, despite the higher survival probability associated with loss aversion. Although greater loss aversion increased the survival probability of a dynasty in environments characterized by more severe adverse shocks, natural selection favored greater loss neutrality in the long run. Individuals characterized by greater loss aversion sorted into safe, but lower-return, production modes, preserving the livelihood of their dynasties while limiting their reproductive expansion. Relatively loss-neutral individuals, in contrast, sorted into riskier production modes with a higher expected return, exposing their dynasties to a greater risk of extinction but endowing the surviving dynasties with greater reproductive potential.

In environments characterized by greater exposure to adverse climatic shocks, the extent to which the reproductive advantage of relatively more loss-neutral dynasties translated into population expansion depended on the spatial structure of climatic shocks. Greater idiosyncrasy increased the likelihood that some risky dynasties escaped adverse realizations and expanded as favorable realizations generated higher income and, consequently, higher fertility, increasing the prevalence of relatively more loss-neutral types in the population.

2.1 The Basic Structure of the Model

Consider an overlapping-generations economy in a Malthusian stage of development. In every time period the economy is populated by a continuum of two-period-lived individuals who are identical in all respects except for their degree of loss aversion, which is transmitted within each dynasty without alteration. Individuals have access to two production modes: a safe mode that generates subsistence consumption and replacement fertility, and a risky mode with a higher expected return that generates higher yield and fertility in an ordinary climatic state and consumption below subsistence, leading to dynastic extinction, in an adverse climatic state. Individuals allocate their income between consumption and fertility, while facing a subsistence consumption constraint. Since richer individuals had higher reproductive success in the Malthusian environment, the degree of loss aversion, through its effect on the choice of production mode, affects reproductive success and thus the long-run composition of loss aversion in society.

2.2 Production

In every time period two production modes are feasible in the economy: a safe mode and a risky mode. The safe mode of production is associated with hunting, gathering, or the cultivation of crops that are largely insensitive to climatic conditions (e.g., roots and tubers), while the risky mode of production is associated with the cultivation of crops that generate a higher expected return but are more sensitive to climatic shocks (e.g., cereal crops).[7]

[7] Appendix B.2 establishes that the central sorting mechanism and the positive effect of adverse-shock severity on the sorting cutoff are robust to partial risk sharing, including the possibility that lower spatial correlation expands insurance opportunities. Risk sharing can induce a larger share of relatively loss-averse individuals to choose the risky mode of production. Yet, a unique interior sorting cutoff continues to exist, and greater adverse-shock frequency or severity continues to raise this cutoff, provided that insurance remains insufficient to raise the adverse-state realization to the subsistence level.

The yield generated by individual i of generation t using the safe mode of production, y_{it}^s, is independent of climatic conditions and is constant across individuals and generations.

$$y_{it}^s = \bar{y}, \tag{1}$$

where the yield under the safe mode of production, $\bar{y}$, exceeds the subsistence consumption constraint (i.e., $\bar{y} > \tilde{c}$). In particular, in line with the defining characteristics of the Malthusian epoch, this yield permits each individual to satisfy the subsistence consumption constraint while maintaining fertility at the replacement level. Hence, in the absence of technological progress, if all individuals chose the safe mode of production, the economy would remain in a Malthusian steady state in which consumption is constant at the subsistence level and population is constant as well.

In contrast, the yield generated by individual i of generation t using the risky mode of production, y_{it}^r, depends on climatic realizations. Under ordinary climatic conditions, which prevail with probability $p \in (0, 1)$, the yield is y_h, whereas when an adverse climatic shock of severity $\Delta > 0$ occurs, with probability $1 - p$, the yield is $y_l(\Delta)$, where $y_l'(\Delta) < 0$. Thus,

$$y_{it}^r = \begin{cases} y_h > \bar{y} & \text{with probability } p, \\ y_l(\Delta) < \tilde{c} & \text{with probability } 1 - p. \end{cases} \tag{2}$$

The risky mode offers a higher expected yield than the safe mode for every admissible degree of adverse-shock severity. In particular, it is assumed that $y_l(\Delta) \geq 0$ and $py^h > \bar{y}$, which imply that $\mathbb{E}[y_{it}^r] > \bar{y}$ for all Δ. Hence, an adverse climatic shock reduces the yield under the risky mode below the subsistence consumption constraint, whereas under ordinary climatic conditions the risky mode generates a yield above that obtained under the safe mode of production and therefore, in a Malthusian environment, permits greater reproductive success.[8] Moreover, consistent with the Malthusian environment, it is assumed that $y_h \leq y^*$, where y^* is the threshold level of income below which the subsistence consumption constraint binds.

Climatic realizations under the risky mode of production are spatially correlated across individuals within each generation, while being independent across generations. In particular, a fraction $\mu \in [0, 1]$ of individuals experience an identical (aggregate) realization of the climatic shock drawn from the distribution specified in (2), whereas a fraction $1 - \mu$ experience idiosyncratic realizations drawn independently from the same distribution. The nature of the shock experienced by an individual is independent of the individual's degree of loss neutrality.

2.3 Individuals

In every period t, a continuum I of two-period-lived individuals, $i \in I$, is born. Individuals are ex ante identical in all respects except for their degree of loss aversion with respect to consumption. The degree of loss aversion is heterogeneous across individuals within each generation and is

[8] The model highlights two dimensions of adverse climatic shocks that shape the evolution of loss aversion: their frequency, $1 - p$, and their severity, Δ. Accordingly, exposure to adverse climatic shocks may vary through either dimension. The empirical analysis examines the frequency of severe droughts in the baseline specification and, in the Appendix, employs an alternative measure that explicitly incorporates variation in both the frequency and severity of droughts. The results are qualitatively unchanged.

transmitted intergenerationally, from parent to child, without alteration. The initial distribution of types is uniform.

In the first period of their life, childhood, individuals are passive economic agents and consume part of their parental resources. In the second period of their life, adulthood, individuals are active economic agents. Based on their degree of loss aversion, they choose their preferred mode of production, produce, and optimally allocate the resulting yield between consumption and child rearing.

2.3.1 Preferences, Constraints, and Optimization

Preferences

The preferences of an adult i in period t are represented by a utility function that captures the spirit of prospect theory (Tversky and Kahneman, 1991). It reflects the asymmetric utility that individuals attribute to gains and losses in consumption relative to a reference point. Loss aversion in individuals' preferences induces behavior that mitigates the immediate risk of dynastic extinction. Yet the strategy that best protects a dynasty from extinction in the present need not be the strategy that maximizes its reproductive success over the long run, a Darwinian calculus that even perfectly rational individuals do not internalize.

The relevance of subsistence as the reference point reflects the inherent asymmetry associated with subsistence consumption in the Malthusian environment that characterized most of human existence. Losses in productivity due to adverse climatic shocks dropped consumption below subsistence, threatening the survival of the lineage, whereas ordinary climatic realizations raised income above subsistence and generated only transitory increases in reproductive success before population adjustment returned resources per capita toward subsistence. Hence, the biologically determined subsistence consumption constraint constituted a natural evolutionary cliff over the course of human history.[9]

The utility function of an adult i in period t, u_{it}, is defined over consumption, c_{it}, and fertility, n_{it}, with gains and losses in consumption evaluated relative to the subsistence reference point, $\tilde{c}$.[10]

$$u_{it} = u_i(c_{it}, n_{it}) = \begin{cases} [f(c_{it}) - f(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} \geq \tilde{c} \\ \lambda_i[g(c_{it}) - g(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} < \tilde{c}, \end{cases} \tag{3}$$

where f is strictly increasing and strictly concave, g is strictly increasing and strictly convex, and η is strictly increasing and strictly concave, with $\eta(0) > -\infty$.[11] The parameter $\lambda_i \in [1, \infty)$ captures

[9]Interestingly, evidence suggests that birds tend to view gambles in a similar fashion above and below the subsistence consumption level (Caraco et al., 1980).

[10]As in prospect theory, utility from consumption is normalized to zero at the reference point, so that $f(\tilde{c}) - f(\tilde{c}) = g(\tilde{c}) - g(\tilde{c}) = 0$.

[11]Appendix B.1 examines alternative specifications of preferences and constraints, including loss aversion with respect to both consumption and fertility, positive but below-replacement fertility below subsistence, a more general utility specification, the emergence of convexity with respect to losses, and adjustment of the reference point. Appendix B.3 considers conventional risk aversion.

the degree of loss aversion of individual i, where $\lambda_i = 1$ is the case of loss neutrality and larger values of λ_i correspond to greater loss aversion.

For analytical convenience, define $\theta_i \equiv 1/\lambda_i \in (0, 1]$ as the degree of loss neutrality. Thus, $\theta_i = 1$ is the case of loss neutrality, while higher values of θ_i correspond to lower degrees of loss aversion. The parameter θ_i is uniformly distributed across dynasties in period 0, i.e.,[12]

$$\theta_i \sim U(0, 1). \tag{4}$$

Budget and Subsistence Consumption Constraints

An adult i in period t is subject to two constraints: a budget constraint and a subsistence consumption constraint. Once the choice of production mode is made and uncertainty is realized, the individual allocates the resulting income, y_{it}, between consumption and fertility.

$$c_{it} + \rho y_{it} n_{it} \leq y_{it}, \tag{5}$$

where $\rho \in (0, 1)$ is the time cost of raising a child, expressed as a fraction of the adult's time endowment.

Moreover, the individual faces the subsistence consumption constraint, below which fertility is zero and the lineage therefore becomes extinct, i.e.,[13,14]

$$n_{it} = 0 \quad \text{if} \quad c_{it} < \tilde{c}. \tag{6}$$

Optimization

An adult i in period t allocates income y_{it} between consumption, c_{it}, and fertility, n_{it}, so as to maximize utility subject to the budget constraint and the subsistence consumption constraint, $n_{it} = 0$ if $c_{it} < \tilde{c}$. Given the properties of the utility function and the maintained restriction $y_{it} \leq y^h \leq y^*$, the solution to the maximization problem exists and is unique:

$$n_{it} = n_i(y_{it}) = \begin{cases} 0 & \text{if } y_{it} \in [0, \tilde{c}] \\ (y_{it} - \tilde{c})/(\rho y_{it}) > 0 & \text{if } y_{it} \in (\tilde{c}, y^h] \end{cases} \qquad c_{it} = c_i(y_{it}) = \begin{cases} y_{it} & \text{if } y_{it} \in [0, \tilde{c}] \\ \tilde{c} & \text{if } y_{it} \in (\tilde{c}, y^h] \end{cases} \tag{7}$$

In particular, if $y_{it} \in [0; \tilde{c}]$, individuals consume their entire yield, y_{it}, and have no surplus that would permit reproduction (i.e., $n_{it} = 0$). However, if $y_{it} \in (\tilde{c}, y^h]$, since the subsistence consumption constraint binds (i.e., $y^h \leq y^*$), individuals consume at the subsistence level, $\tilde{c}$, and devote the

[12] The qualitative analysis is unaffected if a non-uniform distribution is postulated.

[13] A similar structure relating subsistence consumption to fertility is used by Baudin et al. (2015). Alternatively, consumption below subsistence may result in positive but below-replacement fertility rather than zero fertility. Appendix B.1.2 demonstrates that the qualitative results are preserved under this more general formulation.

[14] Differential subsistence levels for each production mode, based on the energy requirements associated with them, would have no qualitative impact as long as the basic structure of the two production modes is maintained (i.e., the subsistence level associated with the safe mode does not exceed its output, $\tilde{c}_s \leq \bar{y}$, and the subsistence level associated with the risky mode lies between its low and high yield realizations, $y_l(\Delta) < \tilde{c}_r < y_h$).

remaining income to raising $(y_{it}-\tilde{c})/(\rho y_{it})>0$ children.[15,16] Furthermore, the yield under ordinary climatic conditions is assumed sufficiently high that, given the probability p of such a realization, the expected reproductive return associated with the risky mode exceeds replacement fertility, i.e., $p(y^h-\tilde{c})/\rho y^h>1$.

Hence, individuals who are engaged in the safe mode of production and generate an income $\bar{y}>\tilde{c}$ optimally consume at the subsistence level and, as assumed earlier, have precisely one child. Specifically, the safe yield $\bar{y}$ is assumed to be such that $\bar{y}=\tilde{c}/(1-\rho)$, permitting fertility at the replacement level (i.e., $n(\bar{y})=(\bar{y}-\tilde{c})/(\rho\bar{y})=1$).[17] Accordingly, individuals who are engaged in the risky mode of production and experience ordinary climatic conditions, generating income $y^h>\bar{y}$, have fertility above replacement (i.e., $n(y^h)=(y^h-\tilde{c})/(\rho y^h)>1$), permitting their dynasty to expand. The indirect utility associated with income realization y_{it} is therefore $v_i(y_{it})\equiv u_i(c_i(y_{it}),n_i(y_{it}))$.

2.3.2 Sorting - Choice of Production Mode

Individuals choose their preferred mode of production prior to the realization of climatic conditions so as to maximize expected utility given the distribution of climatic shocks. Their preferred mode therefore differs according to their degree of loss aversion.[18]

The expected utility u^r_{it} of adult i in period t from choosing the risky production mode, as follows from (2), (3) and (7), is therefore

$$u^r_{it}\equiv\mathbb{E}\left[v_i(y^r_{it})\right]=p\left[\eta\left(\frac{y^h-\tilde{c}}{\rho y^h}\right)\right]+(1-p)\left[\frac{1}{\theta_i}[g(y_l(\Delta))-g(\tilde{c})]+\eta(0)\right]\equiv u^r(\theta_i;\Delta,p), \tag{8}$$

where, since $g(y_l(\Delta))<g(\tilde{c})$, $\lim_{\theta_i\to 0}u^r(\theta_i;\Delta,p)=-\infty$ and $\partial u^r(\theta_i;\Delta,p)/\partial\theta_i>0\ \ \forall\,\theta_i\in(0,1)$.

Analogously, the utility, u^s_{it}, of adult i in period t from choosing the safe production mode, as follows from (1), (3) and (7) is

$$u^s_{it}=\eta(1)\equiv u^s, \tag{9}$$

which is constant across types and over time.

Assuming that the risky production mode is sufficiently attractive that the least loss-averse type prefers it (i.e., $u^r(1;\Delta,p)>u^s$),[19] and since $u^r(\theta_i;\Delta,p)$ is continuous and strictly increasing

[15] As demonstrated empirically by Ashraf and Galor (2011), this is a plausible consumption realization in the context of the Malthusian epoch.

[16] The critical level of income, y^*, below which the subsistence consumption constraint is binding, exists and is uniquely determined by: $\eta'[(y^*-\tilde{c})/(\rho y^*)]/(\rho y^*)=f'(\tilde{c})$. A sufficient condition for the existence of $y^*>\tilde{c}$ is $\eta'(0)>\rho\tilde{c}f'(\tilde{c})$.

[17] This normalization enhances the analytical tractability of the model and has no impact on its qualitative predictions.

[18] Because enforceable intertemporal and state-contingent contracts across communities were largely absent during the Malthusian era, individuals could not generally rely on cross-community insurance against adverse climatic realizations. Moreover, within communities, climatic shocks were likely to be highly correlated, further limiting the scope for insurance.

[19] This assumption implies that $p[\eta((y^h-\tilde{c})/\rho y^h)]+(1-p)\left[[g(y_l(\Delta))-g(\tilde{c})]+\eta(0)\right]>\eta(1)$.

in θ_i, with $\lim_{\theta_i \to 0} u^r(\theta_i; \Delta, p) = -\infty$, whereas u^s is constant, it follows from the *Intermediate Value Theorem* that there exists a unique, time-invariant cutoff level of loss neutrality, $\hat{\theta}(\Delta, p) \in (0; 1)$, such that $u^r(\hat{\theta}(\Delta, p); \Delta, p) = u^s$.[20]

$$u^r(\theta_i; \Delta, p) \begin{cases} < u^s & \forall \theta_i \in (0, \hat{\theta}) \\ > u^s & \forall \theta_i \in (\hat{\theta}, 1], \end{cases} \tag{10}$$

where $\partial\hat{\theta}(\Delta, p)/\partial\Delta > 0$ and $\partial\hat{\theta}(\Delta, p)/\partial(1-p) > 0$.

Sorting The relatively more loss-averse individuals, those with $\theta_i \in (0; \hat{\theta}]$, choose the safe mode of production, whereas the relatively less loss-averse individuals, those with $\theta_i \in (\hat{\theta}; 1]$, choose the risky production mode. As depicted in Figure 1, given the uniform initial distribution of θ_i, this sorting implies an average degree of loss neutrality of $\hat{\theta}/2$ among individuals who choose the safe mode of production and $(1 + \hat{\theta})/2$ among those who choose the risky mode of production.

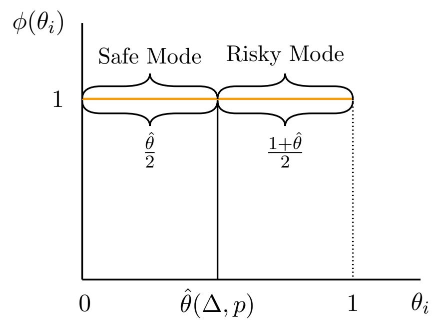

Figure 1: Sorting into the Safe and the Risky Production Modes

Notes: The figure depicts: (i) the sorting into the safe and the risky production modes, based on the degree of loss neutrality, θ_i; (ii) the average degree of loss neutrality among individuals who choose the safe mode of production, $\hat{\theta}/2$, and the average among those who choose the risky mode, $(1 + \hat{\theta})/2$. $\phi(\theta_i)$ denotes the population density of individuals characterized by degree of loss neutrality θ_i, and $\hat{\theta} = \hat{\theta}(\Delta, p)$.

The Effects of Adverse-Shock Frequency and Severity on Sorting

The cutoff level of loss neutrality, $\hat{\theta}(\Delta, p)$, above which individuals sort into the risky mode of production and below which they sort into the safe mode, increases with the frequency and severity of adverse climatic shocks (i.e., $\partial\hat{\theta}(\Delta, p)/\partial(1-p) > 0$ and $\partial\hat{\theta}(\Delta, p)/\partial\Delta > 0$).[21] Namely, since the risky mode of production becomes less attractive as adverse climatic shocks become more frequent or more severe, a higher degree of loss neutrality is required to make an individual indifferent between the two modes of production.

Hence, as depicted in Figure 2, if the severity of the adverse climatic shock increases from Δ_L to Δ_H, where $\Delta_H > \Delta_L$, holding p constant, the cutoff level increases from $\hat{\theta}(\Delta_L, p)$ to $\hat{\theta}(\Delta_H, p)$, and the average degree of loss neutrality increases in both groups: from $[\hat{\theta}(\Delta_L, p)/2]$ to $[\hat{\theta}(\Delta_H, p)/2]$

[20] $\hat{\theta}(\Delta, p) = (1-p)[g(\tilde{c}) - g(y_l(\Delta))]/[p\eta((y^h - \tilde{c})/(\rho y^h)) + (1-p)\eta(0) - \eta(1)]$.

[21] Note that the effects of adverse-shock frequency, $1 - p$, and severity, Δ, on the cutoff level of loss neutrality, $\hat{\theta}$, are orthogonal to the degree of spatial correlation in these shocks, μ.

among individuals who choose the safe mode of production, and from $[1 + \hat{\theta}(\Delta_L, p)]/2$ to $[1 + \hat{\theta}(\Delta_H, p)]/2$ among individuals who choose the risky mode of production. Likewise, holding Δ constant, a higher frequency of adverse climatic shocks (i.e., a lower p) raises $\hat{\theta}(\Delta, p)$ and increases the average degree of loss neutrality in both groups.

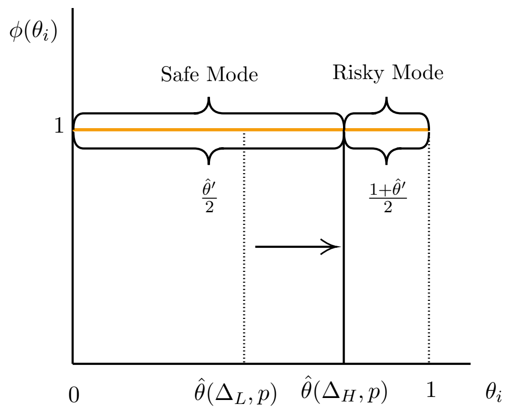

Figure 2: The Effect of Adverse-Shock Severity on Sorting

Notes: The figure depicts: (i) the effect of an increase in the severity of the adverse climatic shock from Δ_L to Δ_H, where $\Delta_H > \Delta_L$, holding p constant, on the cutoff level, raising it from $\hat{\theta}(\Delta_L, p)$ to $\hat{\theta}(\Delta_H, p)$; (ii) the corresponding increase in the average degree of loss neutrality among individuals who choose the safe and the risky production modes to $\hat{\theta}(\Delta_H, p)/2$ and $(1 + \hat{\theta}(\Delta_H, p))/2$, respectively. $\phi(\theta_i)$ denotes the population density of individuals characterized by degree of loss neutrality θ_i.

Sorting, Income, and Reproduction

The sorting of individuals across the two modes of production determines their income realizations and their reproductive success. The income of an adult i of generation t, as follows from (1) and (2), is

$$y_{it} = \begin{cases} \begin{cases} y_l(\Delta) & \text{with probability } 1-p \\ y^h & \text{with probability } p \end{cases} & \text{if} \quad \theta_i > \hat{\theta} \\ \bar{y} & \text{if} \quad \theta_i \leq \hat{\theta}. \end{cases} \tag{11}$$

Accordingly, as follows from (7) the fertility of an adult i of generation t,

$$n_{it} = \begin{cases} 0 & \text{if} \quad \theta_i > \hat{\theta} \quad \text{and} \quad y_{it} = y_l(\Delta) \\ (y^h - \tilde{c})/\rho y^h > 1 & \text{if} \quad \theta_i > \hat{\theta} \quad \text{and} \quad y_{it} = y^h \\ 1 & \text{if} \quad \theta_i \leq \hat{\theta}. \end{cases} \tag{12}$$

Moreover, climatic shocks are spatially correlated across individuals within a generation, while being independent across generations. In particular, a fraction $\mu \in [0, 1]$ of individuals who sort into the risky mode of production experience a common aggregate realization, whereas a fraction $(1 - \mu)$ experience idiosyncratic realizations. Hence, among individuals who choose the risky production mode, i.e., for all $i \in I$ such that $\theta_i > \hat{\theta}$, fertility takes the values characterized in equation (12), but these realizations are independent across individuals exposed to idiosyncratic shocks and common across individuals exposed to the aggregate shock.

2.4 The Evolution of the Composition of Loss Aversion

The evolution of the composition of loss aversion within the population is governed by the effect of the degree of loss aversion on differential reproductive success across individuals. Since loss aversion is assumed to be transmitted intergenerationally within each dynasty without alteration, if lower degrees of loss aversion are associated with greater reproductive success, types characterized by lower degrees of loss aversion will become increasingly prevalent in the population over time.

The threshold degree of loss neutrality, $\hat{\theta}$, below which individuals choose the safe production mode is time invariant. Accordingly, let $I^A \equiv i \in I \mid \theta_i \leq \hat{\theta}$ denote the set of individuals who choose the safe mode of production, and let $I^B \equiv i \in I \mid \theta_i > \hat{\theta}$ denote the set of individuals who choose the risky mode of production. As established in (12), individuals in I^A have an identical fertility rate, $n_{it}^A = 1$, whereas those in I^B face the same distribution of fertility realizations: with probability p, $n_{it}^B = (y^h - \tilde{c})/(\rho y^h) > 1$, and with probability $1-p$, $n_{it}^B = 0$. Fertility realizations within I^B, however, may be correlated according to the degree of spatial correlation in climatic shocks, μ. Thus, conditional on the choice of production mode, the degree of loss aversion has no further effect on fertility, and the distribution of θ_i within each group remains unchanged over time.

Thus, the total population in period t, N_t, is partitioned into the population of group A, N_t^A, and the population of group B, N_t^B, such that $N_t^A + N_t^B = N_t$. In particular, given the cutoff $\hat{\theta}(\Delta, p)$ and the uniform initial distribution of θ_i over $(0, 1]$, the allocation of the initial population, N_0, between the two groups is

$$N_0^A = \hat{\theta}(\Delta, p) N_0, \;\; N_0^B \;\; = [1 - \hat{\theta}(\Delta, p)] N_0. \tag{13}$$

Differential reproductive success across the two groups affects their relative sizes and therefore the evolution of the distribution of loss aversion in the population. Since all individuals in I^A have the same fertility rate, $n_{it}^A = 1$, corresponding to replacement fertility, the population of group A remains constant over time. Its share of the total population, however, changes as the population of group B evolves, affecting the distribution of loss aversion in the population. In particular,

$$N_t^A = N_0^A = \hat{\theta}(\Delta, p) N_0 \equiv N^A(\hat{\theta}(\Delta, p)), \tag{14}$$

where $\partial N_t^A / \partial \Delta > 0$ and $\partial N_t^A / \partial (1-p) > 0$ (as follows from $\partial \hat{\theta}(\Delta, p)/\partial \Delta > 0$, $\partial \hat{\theta}(\Delta, p)/\partial (1-p) > 0$, and $dN_t^A / d\hat{\theta} > 0$).

In contrast, the evolution of the population of group B depends on the sequence and spatial dependence of fertility realizations. Let $\bar{n}_t^{id}$ denote average fertility among individuals in group B under idiosyncratic climatic shocks, and let $\bar{n}_t^{ag}$ denote average fertility under aggregate climatic shocks. Under idiosyncratic climatic shocks, fertility realizations are independent across individuals. With probability p, fertility is $(y^h - \tilde{c})/(\rho y^h)$, and with probability $1-p$, it is zero. Hence, $\bar{n}_t^{id}$ is the expected number of children and is constant over time, i.e., $\bar{n}_t^{id} = p(y^h - \tilde{c})/(\rho y^h) \equiv \bar{n}^{id}$. Under aggregate climatic shocks, the fertility realization is common across all individuals. Thus, with probability p, $\bar{n}_t^{ag} = (y^h - \tilde{c})/(\rho y^h)$, and with probability $1-p$, $\bar{n}_t^{ag} = 0$. Consequently, the population of group B evolves from period t to period $t+1$ according to

$$N_{t+1}^B = \left[\mu \bar{n}_t^{ag} + (1-\mu)\bar{n}^{id}\right] N_t^B. \tag{15}$$

Iterating this law of motion from period 0 to period t yields

$$N_t^B = \prod_{s=0}^{t-1} \left[\mu \bar{n}_s^{ag} + (1-\mu)\bar{n}^{id}\right] N_0^B \equiv N^B\left(\hat{\theta}(\Delta,p), p, \mu, \{\bar{n}_s^{ag}\}_{s=0}^{t-1}\right), \tag{16}$$

where $\partial N_t^B/\partial\Delta < 0$ for all $\mu \in (0,1)$ (as follows from $\partial N_t^B/\partial\hat{\theta} < 0$ and $\partial\hat{\theta}(\Delta,p)/\partial\Delta > 0$).

The share of individuals who choose the risky mode of production in period t, β_t, is therefore

$$\beta_t \equiv \frac{N_t^B}{N_t^A + N_t^B} \equiv \beta\left(\hat{\theta}(\Delta,p), p, \mu, \{\bar{n}_s^{ag}\}_{s=0}^{t-1}\right). \tag{17}$$

Thus, β_t depends on the cutoff $\hat{\theta}(\Delta,p)$, the frequency of adverse climatic shocks, $1-p$, the degree of spatial correlation, μ, and the realized sequence of aggregate fertility outcomes.

The average degree of loss neutrality in the population in period t, $\bar{\theta}_t$, is therefore the population-weighted average of the mean degree of loss neutrality within the two groups, which, as depicted in Figure 1, are $\hat{\theta}/2$ for the safe group and $(1+\hat{\theta})/2$ for the risky group:

$$\bar{\theta}_t = \beta_t \frac{1+\hat{\theta}}{2} + (1-\beta_t)\frac{\hat{\theta}}{2} = \frac{\beta_t + \hat{\theta}}{2} \equiv \bar{\theta}\left(\hat{\theta}(\Delta,p), p, \mu, \{\bar{n}_s^{ag}\}_{s=0}^{t-1}\right), \tag{18}$$

where since an increase in either the severity, Δ, or the frequency, $1-p$, of adverse climatic shocks raises $\hat{\theta}$ but reduces β_t, its effect on $\bar{\theta}_t$ is ambiguous in the short run. Since higher values of θ_i correspond to lower degrees of loss aversion, an increase in $\bar{\theta}_t$ reflects a shift in the population toward lower loss aversion.

The long-run population dynamics underlying the evolution of loss aversion are characterized in the following lemma.

Lemma 1. *There exists a unique threshold degree of spatial correlation, $\hat{\mu}(p) \in (0,1)$, such that:*

(i) if $\mu < \hat{\mu}(p)$ (i.e., if climatic shocks are relatively idiosyncratic), the population of individuals who choose the risky mode of production grows asymptotically faster than the population of individuals who choose the safe mode and dominates the population in the long run:

$$\bar{\beta} \equiv \lim_{t\to\infty} \beta_t = 1, \qquad \bar{\theta} \equiv \lim_{t\to\infty} \bar{\theta}_t = (1+\hat{\theta})/2; \tag{19}$$

(ii) if $\mu > \hat{\mu}(p)$ (i.e., if climatic shocks are relatively aggregate), the population of individuals who choose the risky mode of production grows asymptotically more slowly than the population of individuals who choose the safe mode and vanishes in the long run:

$$\bar{\beta} = 0, \qquad \bar{\theta} = \hat{\theta}/2. \tag{20}$$

(iii) the asymptotic population growth rate of individuals who choose the risky mode of production is strictly decreasing in μ.

Proof. See Appendix A.1. □

In the baseline formulation, greater adverse-shock frequency or severity raises the sorting cutoff, $\hat{\theta}(\Delta, p)$, shifting the composition of both the safe and risky groups toward lower loss aversion. Whereas severity affects the evolution of loss aversion through this sorting margin, adverse-shock frequency also affects the relative long-run population growth of the two groups, jointly with the degree of spatial correlation in climatic shocks.

The implications of these population dynamics for the long-run evolution of loss aversion are summarized in the following proposition.

2.5 Adverse Climatic Shocks and Loss Aversion

Main Findings and Testable Implications

Proposition 1 (Adverse Climatic Shocks and Loss Aversion)**.**

Main Prediction:

- *Greater adverse-shock frequency or severity reduces long-run loss aversion:*

$$\frac{\partial\,\bar{\theta}\,[\hat{\theta}(\Delta,p),p,\mu]}{\partial(1-p)} > 0, \qquad \frac{\partial\,\bar{\theta}\,[\hat{\theta}(\Delta,p),p,\mu]}{\partial\Delta} > 0, \qquad \mu \neq \hat{\mu}(p).$$

Complementary Prediction:

- *Holding the frequency and severity of climatic shocks constant:*
  - *Greater idiosyncrasy of the shocks (i.e., greater $(1-\mu)$) weakly reduces long-run loss aversion;*
  - *Greater spatial correlation in the shocks (i.e., greater μ) weakly increases long-run loss aversion.*

In particular, for $\mu_L < \hat{\mu}(p) < \mu_H$,[22]

$$\bar{\theta}(\mu_L) > \bar{\theta}(\mu_H).$$

Proof. See Appendix A.2. □

Viewed through the lens of ancestral regions, the theory generates two distinct empirical predictions. Individuals and societies originating in regions characterized by greater adverse-shock frequency or severity would be expected to exhibit lower loss aversion.[23] Holding adverse-shock

[22] The weak monotonicity of the complementary prediction reflects the discrete two-mode structure of the baseline model. Within each region, $\mu < \hat{\mu}(p)$ or $\mu > \hat{\mu}(p)$, the same production mode dominates the population in the long run, so marginal changes in spatial correlation do not alter the asymptotic composition of loss aversion. With a richer set of production choices, different modes would become evolutionarily dominant at different degrees of spatial correlation, generating a continuous monotonic increase in long-run loss aversion as spatial correlation rises.

[23] The knife-edge case $\mu = \hat{\mu}(p)$ is excluded since the relative population size follows a zero-drift stochastic process and the limiting population composition does not exist.

frequency and severity constant, individuals and societies originating in ancestral regions characterized by greater idiosyncrasy of climatic shocks would be expected to exhibit lower loss aversion, whereas those originating in ancestral regions characterized by greater spatial correlation would be expected to exhibit higher loss aversion.[24]

Illustration of the Main Predictions

As depicted in panel (a) of Figure 3, when the degree of spatial correlation in climatic shocks is sufficiently small (i.e., $\mu < \hat{\mu}(p)$), the share of relatively loss-neutral individuals who choose the risky mode of production increases over time, ultimately approaching complete domination of the population. Since reproductive success within this group is independent of the degree of loss neutrality, the distribution of θ_i within the group remains unchanged, and its average degree of loss neutrality remains equal to $(1 + \hat{\theta})/2$. Hence, the asymptotic average degree of loss neutrality in the population is $\bar{\theta} = (1 + \hat{\theta})/2$. Since an increase in adverse-shock frequency or severity raises $\hat{\theta}$, it increases the long-run average degree of loss neutrality.

In contrast, as depicted in panel (b) of Figure 3, when the degree of spatial correlation in climatic shocks is sufficiently large (i.e., $\mu > \hat{\mu}(p)$), the share of relatively more loss-averse individuals who choose the safe mode of production increases over time, ultimately approaching complete domination of the population. Since reproductive success within this group is independent of the degree of loss neutrality, its distribution of θ_i remains unchanged, and its average degree of loss neutrality remains equal to $\hat{\theta}/2$. Hence, the asymptotic average degree of loss neutrality in the population is $\bar{\theta} = \hat{\theta}/2$. Thus, greater adverse-shock frequency or severity increases long-run loss neutrality also when the safe group dominates the population.

Since an increase in adverse-shock severity from Δ to Δ_H, where $\Delta_H > \Delta$, holding p constant, raises the cutoff from $\hat{\theta}(\Delta, p)$ to $\hat{\theta}(\Delta_H, p)$, it raises the average degree of loss neutrality within both groups. Hence, regardless of which group dominates in the long run, greater adverse-shock frequency or severity increases the long-run degree of loss neutrality: from $(1 + \hat{\theta}(\Delta, p))/2$ to $(1 + \hat{\theta}(\Delta_H, p))/2$ if $\mu < \hat{\mu}(p)$, and from $\hat{\theta}(\Delta, p)/2$ to $\hat{\theta}(\Delta_H, p)/2$ if $\mu > \hat{\mu}(p)$. Analogously, away from the threshold $\mu = \hat{\mu}(p)$, greater adverse-shock frequency increases the long-run degree of loss neutrality.

2.6 Robustness of the Main Predictions

The main predictions of the baseline model are robust to a broad range of extensions:
(a) loss aversion with respect to both fertility and consumption (Appendix B.1.1);
(b) positive fertility below the subsistence consumption constraint (Appendix B.1.2);
(c) more general forms of the utility function (Appendix B.1.3); and

[24] The discrete structure of the production model, in which individuals sort between two production modes and one group ultimately dominates the population, gives rise to the threshold characterization of the effect of spatial correlation. In a richer environment with a continuum of production choices and corresponding reproductive advantages, population composition would adjust continuously with the degree of spatial correlation. The average degree of loss neutrality would therefore decline continuously with spatial correlation. Hence, empirically, the testable prediction is unambiguous: greater spatial correlation should be associated with greater loss aversion.

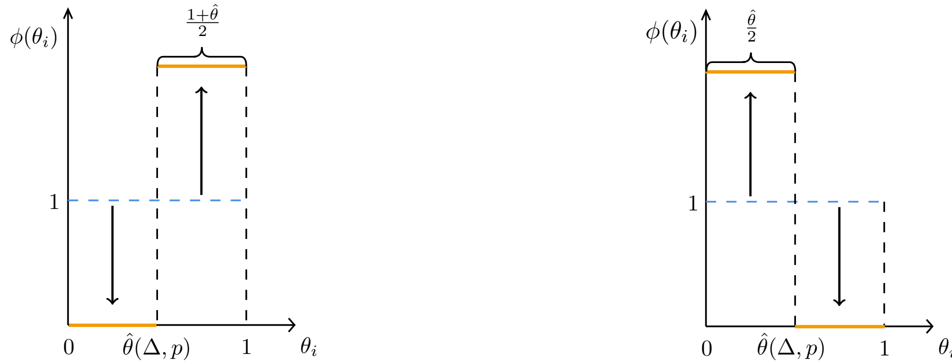

(a) Weakly correlated shocks ($\mu < \hat{\mu}(p)$) (b) Strongly correlated shocks ($\mu > \hat{\mu}(p)$)

Figure 3: Evolution of Loss Aversion under Alternative Degrees of Spatial Correlation

Notes: $\phi(\theta_i)$ denotes the population density of individuals characterized by degree of loss neutrality θ_i.

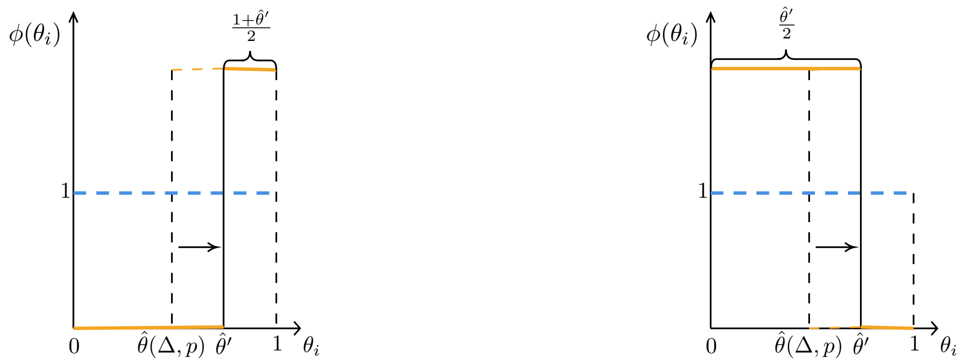

(a) Weakly correlated shocks ($\mu < \hat{\mu}(p)$) (b) Strongly correlated shocks ($\mu > \hat{\mu}(p)$)

Figure 4: Effect of Greater Adverse-Shock Frequency or Severity on the Asymptotic Distribution and Long-Run Average Degree of Loss Neutrality

Notes: The figure depicts the effect of an increase in adverse-shock frequency or severity that raises the sorting cutoff from $\hat{\theta} = \hat{\theta}(\Delta, p)$ to $\hat{\theta}' > \hat{\theta}$. For changes in adverse-shock frequency, the illustration applies to changes that do not cross the threshold $\hat{\mu}(p)$. $\phi(\theta_i)$ denotes the population density of individuals characterized by degree of loss neutrality θ_i.

(d) adjustment of the reference point (Appendix B.1.5).

The analysis further establishes that the main prediction is robust to partial risk sharing (Appendix B.2). Moreover, imperfect-intergenerational-transmission simulations in Appendix B.4 further show that, over the range considered, both predictions hold despite departures from perfect transmission. Finally, Appendix B.3 demonstrates that conventional risk aversion does not generate unambiguous predictions regarding the effects of adverse climatic shock, while Appendix B.1.4 provides an evolutionary rationale for the emergence of convexity of utility with respect to losses.

3 Data and Empirical Strategy

The theory generates two empirical predictions. First, it predicts that greater ancestral exposure to adverse climatic shocks is associated with lower loss aversion. Second, holding total exposure constant, the theory predicts that greater idiosyncrasy in the spatial distribution of those shocks is associated with lower loss aversion.

This section develops the empirical strategy used to test these predictions, describing the construction of the climatic-shock measures, the novel experimental methodology used to elicit loss aversion and risk aversion, and the epidemiological framework used to isolate ancestral influences from those of the contemporary environment.

3.1 Data and Measurement

The empirical analysis combines newly constructed measures of ancestral climatic-shock exposure with experimentally elicited measures of loss aversion. The climatic measures capture both the frequency of adverse shocks in each ancestral homeland and the extent to which these shocks were spatially idiosyncratic or aggregate. The experimental data are generated through a novel empirical design that elicits loss aversion, along with risk aversion, under an identical experimental protocol among U.S.-born individuals of diverse ancestral origins who share a common host-country environment.

3.1.1 Exposure to Climatic Shocks

Exposure to severe climatic shocks is proxied using monthly global observations of the self-calibrating Palmer Drought Severity Index (scPDSI), at a spatial resolution of $0.5° \times 0.5°$ over the twentieth century (Wells et al., 2004; van der Schrier et al., 2013). The high spatial and temporal resolution of these data permits the measurement of both the frequency of consequential drought conditions at each location and the extent to which these shocks were spatially idiosyncratic or aggregate.

Reassuringly, the twentieth-century measure of drought exposure is highly correlated at both the grid-cell and country levels with its historical counterpart over 1–1500 CE, indicating that it provides a credible proxy for the climatic conditions that prevailed during the formative period of this evolutionary process. The historical measure, constructed using the Paleo Hydrodynamics Data Assimilation reconstruction, is less suitable as the baseline measure of climatic exposure because it relies on proxy-based reconstruction rather than direct observation and is substantially noisier, with markedly coarser spatial and temporal resolution. Nevertheless, it provides an informative independent validation, reproducing the theory's central predictions.

Since the theory pertains to adverse climatic shocks with potentially consequential effects on subsistence, the baseline measure defines a climatic shock as a drought realization of at least moderate severity. This threshold identifies drought conditions of sufficient severity to pose consequential risks to agricultural production and subsistence.[25] Accordingly, for grid cell i in month t, a climatic shock is defined as

$$D_{it} = \begin{cases} 1 & \text{if} \quad \text{scPDSI}_{it} \leq -2 \\ 0 & \text{otherwise.} \end{cases} \tag{21}$$

For each climatic-shock realization, its spatial character is determined by the incidence of contemporaneous shocks in neighboring locations. Let q_{it} denote the fraction of up to eight immediately adjacent land cells that also experience a climatic shock in month t. The spatial nature of each shock is therefore decomposed into an aggregate component and an idiosyncratic component:

$$D_{it} = \underbrace{q_{it}D_{it}}_{\text{aggregate component}} + \underbrace{(1-q_{it})D_{it}}_{\text{idiosyncratic component}}. \tag{22}$$

A shock is more aggregate when a larger fraction of neighboring locations experiences a shock in the same month, whereas it is more idiosyncratic when nearby locations are comparatively unaffected.[26]

Given the T months in the sample, overall shock exposure, E_i, aggregate shock exposure, E_i^A, and idiosyncratic shock exposure, E_i^I, for each cell i are given by

$$E_i = \frac{1}{T}\sum_{t=1}^{T} D_{it}; \qquad E_i^A = \frac{1}{T}\sum_{t=1}^{T} q_{it}D_{it}; \qquad E_i^I = \frac{1}{T}\sum_{t=1}^{T}(1-q_{it})D_{it}. \tag{23}$$

By construction,

$$E_i = E_i^A + E_i^I. \tag{24}$$

Overall exposure therefore measures the frequency of climatic shocks at location i, while aggregate and idiosyncratic exposure measure the components of this exposure associated with spatially shared and spatially localized shocks.

For country c, overall shock exposure, E_c, is the land-area-weighted average of grid-cell exposure within the country:

$$E_c = \sum_{i \in \mathcal{I}_c} \omega_{ic} E_i, \tag{25}$$

where ω_{ic} denotes the share of country c's land area represented by cell i. Aggregate and idiosyncratic shock exposure are therefore:

$$E_c^A = \sum_{i \in \mathcal{I}_c} \omega_{ic} E_i^A; \qquad E_c^I = \sum_{i \in \mathcal{I}_c} \omega_{ic} E_i^I. \tag{26}$$

[25] Under the Palmer classification, values of the scPDSI of −1, −2, and −3 correspond to the onset of mild, moderate, and severe drought, respectively (Wells et al., 2004). In the historical setting relevant to this evolutionary process, greater reliance on rain-fed agriculture and limited irrigation, storage, insurance, and market integration made droughts of at least moderate severity particularly consequential for agricultural production and subsistence. In view of these considerations, the baseline adopts the moderate-drought threshold, while Table F.2 establishes that the results are robust to the alternative mild- and severe-drought thresholds.

[26] Robustness to broader definitions of the relevant spatial neighborhood is examined in Appendix Table E.3, where the measure is reconstructed using all land grid cells within 100 and 200 kilometers of the focal cell. As suggested by the partial-risk-sharing extension in Appendix B.2, expanding the geographic reach of trade, storage, or other forms of risk sharing would be expected to attenuate the evolutionary effect of spatially idiosyncratic shocks, since adverse local realizations can increasingly be offset by transfers from unaffected locations. The effect would nevertheless persist as long as risk sharing remains incomplete.

Since the same land-area weights are applied to overall, aggregate, and idiosyncratic exposure, the cell-level decomposition is preserved at the country level: $E_c = E_c^A + E_c^I$.

Accounting for post-1500 population movements, the country-level measures are adjusted to reflect the ancestral composition of each country's present-day population, weighting the climatic exposure of each source country by its ancestral contribution to that population. The resulting global distribution of ancestry-adjusted overall climatic-shock exposure, depicted in Figure 5, reveals substantial variation across countries in the climatic conditions experienced by their ancestral populations.

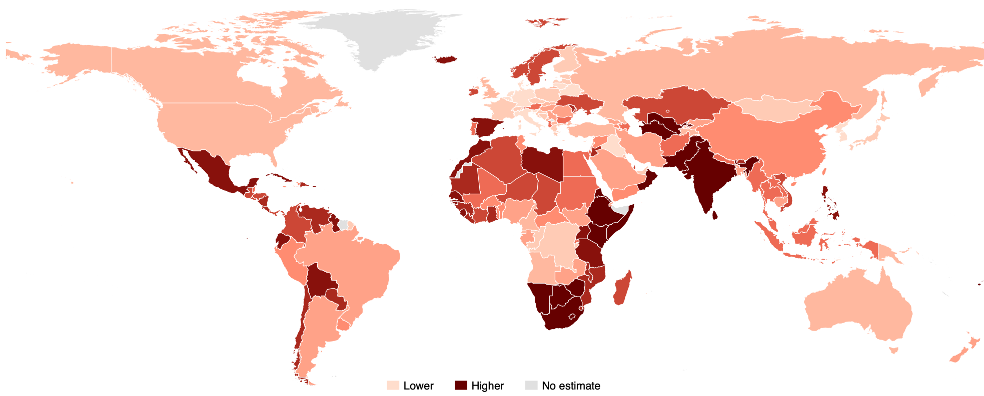

Figure 5: Global Distribution of Ancestry-Adjusted Climatic-Shock Exposure

Notes: The map displays ancestry-adjusted overall exposure to drought. Darker shades indicate greater ancestry-adjusted exposure to climatic shocks; gray denotes countries for which the measure is unavailable.

The adjustment recognizes that the climatic conditions relevant for the formation of present-day predispositions may differ from those prevailing in the territories where populations currently reside. For example, in constructing Argentina's ancestry-adjusted climatic exposure, the population of Italian descent is assigned the climatic conditions of Italy rather than those of Argentina, whereas the population of Spanish descent is assigned the climatic conditions of Spain.[27]

3.1.2 Loss Aversion

The measure of loss aversion is constructed using a novel experimental methodology that disentangles the underlying predisposition toward loss aversion from confounding country-specific characteristics. It implements the epidemiological approach in an experimental setting, eliciting measures of loss aversion through an incentivized experiment among U.S.-born descendants of migrants from more than one hundred ancestral homelands who currently reside in the United States.[28]

[27]The country-level exposure measures are ancestry-adjusted using the post-1500 population-migration matrix of Putterman and Weil (2010). Reassuringly, the baseline findings remain qualitatively unchanged when the sample is restricted to homelands that did not experience substantial post-1500 migration (Figure E.4).

[28]Respondents were asked to report their primary and, where applicable, secondary ancestry. Consistent with the pattern observed among millions of respondents in the American Community Survey (ACS), about half reported a single ancestry. The baseline assigns ancestral climatic exposure according to the respondent's primary ancestry, while

The novel experimental design addresses two distinct sources of heterogeneity that could otherwise obscure differences in the underlying predisposition toward loss aversion: country-specific characteristics and experiment-specific characteristics. The common U.S. setting holds constant the broad economic, cultural, and institutional environment across ancestral groups, while the common experimental setting holds constant the conditions under which loss aversion is elicited. In particular, all participants are assessed under an identical experimental protocol and incentive structure, eliminating heterogeneity that may arise from differences in experimental procedures, monetary stakes, or institutional settings.[29]

Moreover, since the measured degree of loss aversion depends on the underlying degree of risk aversion, the experimental design elicits both dimensions, rather than imposing a common degree of risk aversion, permitting reliable cross-ancestry comparisons of loss aversion and allowing the analysis to establish that the proposed hypothesis pertains uniquely to loss aversion and is orthogonal to risk aversion. While imposing a common degree of risk aversion in the estimation of loss aversion is less consequential for within-country comparisons, since it preserves individuals' rank ordering, variation in risk aversion across ancestral groups would distort cross-ancestry comparisons of loss aversion, affecting not only the magnitude of the differences but potentially also their ranking.

Accordingly, the experiment first elicits each respondent's degree of risk aversion and then incorporates this respondent-specific estimate into the construction of the loss-aversion measure. Risk aversion is elicited through a sequence of incentivized choices between a 50–50 lottery yielding either $x_H > 0$ or zero and a guaranteed payment varying from zero to x_H. As the guaranteed payment increases across choices, the respondent eventually switches from preferring the lottery to preferring the guaranteed payment. The guaranteed payment at which this switch occurs identifies respondent i's certainty equivalent, ce_i. If the utility function is defined over experimental payoffs as $u_i(x) = x^{r_i}$ for $x \geq 0$, where r_i captures the respondent's degree of risk aversion, with $r_i = 1$ corresponding to risk neutrality, $r_i < 1$ to risk aversion, and $r_i > 1$ to risk seeking, then at the certainty equivalent, indifference between the lottery and the guaranteed payment implies $0.5x_H^{r_i} + 0.5(0)^{r_i} = ce_i^{r_i}$, yielding

$$\widehat{r}_i = \frac{\ln(0.5)}{\ln(ce_i/x_H)}. \tag{27}$$

The resulting estimate, $\widehat{r}_i$, identifies respondent i's degree of risk aversion, with lower values of $\widehat{r}_i$ corresponding to greater risk aversion and $\widehat{r}_i = 1$ to risk neutrality.

Loss aversion is then elicited through a second sequence of incentivized choices. In each choice, respondents select between a 50–50 lottery whose two possible outcomes are a gain $y > 0$ and a loss $-c < 0$, and a guaranteed payoff of zero. The gain y is held fixed across choices, while the

retaining respondents who report a second ancestry. The findings are robust to restricting the sample to respondents reporting a single ancestry or to assigning respondents with two ancestries the average climatic exposure of their two ancestral homelands (Table F.4).

[29] The experiment was conducted on Prolific following a protocol described in detail in **??**.

magnitude of the potential loss, c, varies. Reference-dependent utility is specified as

$$u_i(x) = \begin{cases} x^{r_i} & \text{if } x \geq 0 \\ -\lambda_i(-x)^{r_i} & \text{if } x < 0. \end{cases} \tag{27}$$

This specification is the experimental counterpart of the utility function introduced in the theory section, with the reference point normalized to zero and λ_i corresponding to the same loss-aversion parameter. Thus, $\lambda_i = 1$ is the case of loss neutrality, where larger values of λ_i correspond to greater loss aversion, and $\theta_i \equiv 1/\lambda_i$ corresponds to the degree of loss neutrality used in the theoretical analysis.[30]

Let $c_i^* > 0$ denote the loss magnitude at which respondent i is indifferent between the mixed lottery and the guaranteed payoff of zero, then indifference implies $0.5y^{\widehat{r}_i} - 0.5\lambda_i(c_i^*)^{\widehat{r}_i} = 0$, and the degree of loss aversion of individual i, $\widehat{\lambda}_i$, is

$$\widehat{\lambda}_i = \left(\frac{y}{c_i^*}\right)^{\widehat{r}_i}. \tag{28}$$

Thus, reliable cross-societal comparisons of loss aversion require accounting for heterogeneity in risk aversion.

The resulting individual estimates of loss aversion are aggregated within ancestries using the median level in the population,[31] limiting the influence of the pronounced right tail of the individual-level distribution.[32] The global distribution of the resulting ancestry-level measure of loss aversion is depicted in Figure 6.

3.2 Empirical Strategy

3.2.1 Empirical Framework

A central empirical challenge in examining the long-run influence of ancestral exposure to adverse climatic shocks on loss aversion stems from systematic differences in contemporary economic, institutional, and social environments across populations exposed to different historical realizations of such shocks. Cross-country comparisons therefore confound the influence of ancestral exposure to adverse climatic shocks with that of the contemporary environment.

Addressing this challenge, the empirical analysis adopts an epidemiological approach that exploits variation in ancestral origins within a common host environment. It compares U.S.-born

[30] The experimental elicitation does not impose the simplifying theoretical restriction $\lambda_i \geq 1$; estimates below one are retained and indicate that losses receive less weight than equivalent gains.

[31] The comparative static in Proposition 1 applies equally to the median in the baseline model. If the risky group dominates, the median of θ is $(1 + \widehat{\theta})/2$; if the safe group dominates, it is $\widehat{\theta}/2$. Since $\lambda = 1/\theta$ is monotonically decreasing, greater adverse-shock frequency or severity lowers the median of λ in either case.

[32] Although the use of the median limits the influence of extreme individual-level estimates within each ancestry, the resulting ancestry-level distribution contains a single extreme upper-tail observation. Iraq's value is 20 times the sample median of 1.55 and almost 6 times the second-highest ancestry-level estimate; it is therefore excluded from the baseline analysis as an extreme experimental anomaly. Once Iraq is excluded, one-sided winsorization caps the remaining two outliers, Serbia and Angola. The lower tail exhibits no comparable discontinuity and is therefore left unchanged. Importantly, the statistical significance and qualitative conclusions are unchanged when Iraq is retained and winsorization is applied, when Serbia and Angola are not winsorized, or when the distribution is winsorized symmetrically (Figure E.2).

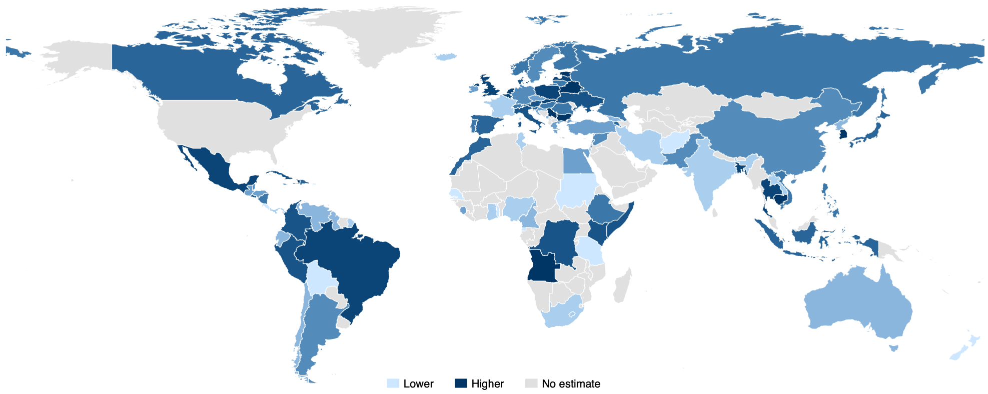

Figure 6: Global Distribution of Loss Aversion across Ancestries

Notes: The map displays experimentally elicited median loss aversion for each ancestry represented in the estimation sample. Darker shades indicate greater loss aversion; gray denotes countries for which no ancestry estimate is available. Contemporary national borders identify the homelands associated with respondents' reported ancestries.

individuals of different ancestries who currently reside in the United States. These individuals share the same national institutions, legal system, macroeconomic environment, and broad social setting, while their ancestral populations experienced markedly different historical realizations of adverse climatic shocks. The common-host framework therefore permits cross-ancestry differences in loss aversion to be examined while substantially limiting the confounding influence of the contemporary environment.

3.2.2 Experimental Design

Within this epidemiological framework, the study conducts a novel incentivized experiment that elicits individual-specific measures of loss aversion among respondents from diverse ancestral origins. The common host-country setting substantially limits heterogeneity in the contemporary economic and institutional environments across ancestral groups, while the common experimental setting ensures that all respondents face the same protocol, incentive structure, monetary stakes, and interface.

Notably, the experimental design separately elicits risk aversion and loss aversion for each respondent. This feature addresses an important limitation of existing cross-societal measures of loss aversion, which commonly impose a uniform degree of risk aversion. Since measured loss aversion depends on utility curvature, imposing a common degree of risk aversion can absorb cross-population heterogeneity in risk aversion into estimates of loss aversion, potentially distorting both the magnitude and the ranking of cross-ancestry differences. Moreover, the separate elicitation of the two preference parameters permits the analysis to establish whether the predicted relationship with ancestral exposure to adverse climatic shocks is unique to loss aversion and orthogonal to risk aversion.

3.2.3 Identification Challenges

Ancestral climatic shocks were realized long before the loss aversion of present-day respondents was observed, eliminating reverse causality as a possible explanation for the documented relationship. Yet, establishing the influence of ancestral exposure to adverse climatic shocks on loss aversion hinges on distinguishing this influence from alternative historical processes that could generate a similar cross-ancestry pattern.

The observed relationship could reflect the influence of other geographic characteristics that are correlated with both ancestral exposure to adverse climatic shocks and the evolution of loss aversion. In particular, frequent severe climatic shocks may be associated with lower agricultural suitability, while the resulting scarcity could itself have fostered greater loss aversion. Likewise, the timing of the Neolithic transition may be correlated with long-run exposure to climatic shocks while independently shaping behavioral adaptation through the duration of reliance on agriculture. The empirical analysis therefore accounts for the caloric suitability of ancestral homelands for agriculture and for the timing of the Neolithic transition, separating the influence of adverse climatic shocks from these broader features of the ancestral environment.

One alternative historical process that could generate the observed relationship is sorting across locations. At the time founder populations settled novel environments, however, they were unlikely to possess the accumulated knowledge required to infer the long-run incidence and spatial structure of adverse climatic shocks, limiting the scope for systematic sorting according to preexisting loss aversion. Moreover, settlement largely reflected collective rather than individual choices, further constraining sorting on individual differences in loss aversion. Such sorting may nevertheless have emerged subsequently, as knowledge of local climatic conditions accumulated. If individuals closer to loss neutrality then relocated disproportionately toward environments characterized by greater exposure to adverse climatic shocks, part of the observed relationship could reflect selective migration rather than evolutionary adaptation through differential reproductive success.[33,34]

This possibility is addressed more directly using the Columbian Exchange. The introduction of previously unavailable crops after 1500 altered the crops best suited to particular locations and, consequently, their agricultural growth cycles. These changes altered the timing of effective exposure to adverse climatic shocks within already settled locations. Conditional on pre-1500 exposure, the resulting variation was driven by the availability and agro-climatic suitability of newly introduced crops rather than by the loss aversion of local populations or their sorting across locations. The Columbian Exchange therefore provides a quasi-natural experiment for examining whether quasi-exogenous changes in effective climatic-shock exposure predict subsequent differences in loss aversion.

[33] Such sorting would qualify the interpretation of the underlying evolutionary mechanism, but would not overturn the broader conclusion that cross-regional differences in exposure to adverse climatic shocks constituted a historical force shaping persistent differences in loss aversion across ancestries.

[34] Selection into the experiment could not plausibly generate the documented relationship. Within ancestry groups, participation or completion of the survey would have to depend on loss aversion in a manner that varies systematically with ancestral exposure to adverse climatic shocks, disproportionately selecting more loss-neutral individuals from ancestries characterized by greater historical exposure to such shocks.

The introduction of previously unavailable crops in the post-1500 period altered the locally optimal crop and, consequently, given differences in the timing of agricultural growth cycles, effective exposure to adverse climatic shocks. For each grid cell, crop-specific caloric suitability is used to identify the locally optimal crop among those available prior to 1500 and among the expanded set available following the Columbian Exchange. FAO GAEZ crop calendars then identify the interval between successive growth cycles for each of these crops, permitting climatic-shock exposure to be measured under the agricultural calendars implied by the pre- and post-1500 crop sets. Notably, since the locally optimal crops are determined based on potential agro-climatic suitability rather than observed cultivation, and the pre- and post-1500 measures reflect the same underlying climatic history, the resulting change in climatic-shock exposure isolates the variation arising from the expansion of the feasible crop set rather than endogenous crop adoption or changes in the underlying climatic environment.

Selective migration from ancestral homelands to the United States may generate systematic differences between migrants and their populations of origin. Although respondents themselves were born in the United States, their ancestors who migrated to the United States need not have been representative of their ancestral populations. Yet, loss-averse individuals would have had stronger incentives to leave environments characterized by adverse climatic realizations, and migrants from more shock-prone homelands would be selected toward greater, rather than lower, loss aversion. Selective migration would therefore tend to attenuate, rather than generate, the negative relationship between ancestral exposure to adverse climatic shocks and loss aversion.[35]

3.2.4 Discriminating Tests of the Mechanism

The empirical strategy subjects the proposed evolutionary mechanism to two discriminating tests. In view of the unique relevance of the existential cliff associated with the subsistence consumption constraint to loss aversion, rather than risk aversion, the analysis uses experimentally elicited risk aversion as a placebo outcome for the relationship between ancestral exposure to adverse climatic shocks and loss aversion. Risk aversion and loss aversion are elicited from the same respondents under identical experimental conditions and are therefore subject to the same sample composition, participation, and experimental environment. Yet they capture distinct dimensions of preferences: risk aversion governs attitudes toward uncertainty more generally, whereas loss aversion captures the asymmetric weighting of losses relative to equivalent gains. A relationship with loss aversion accompanied by the absence of a corresponding relationship with risk aversion therefore provides a discriminating test of the proposed evolutionary mechanism.

The theory also predicts that the spatial structure of adverse climatic shocks shaped the evolution of loss aversion. Holding overall exposure to adverse shocks constant, greater idiosyncrasy is predicted to be associated with lower loss aversion, whereas greater spatial correlation is pre-

[35] Generic selection of migrants according to behavioral preferences would not suffice to generate the documented relationship; such selection would have to vary systematically with ancestral exposure to adverse climatic shocks in precisely the direction predicted by the theory.

dicted to be associated with higher loss aversion. The empirical analysis tests this implication by decomposing adverse climatic shocks into their idiosyncratic and aggregate components and, more directly, by holding total exposure constant while varying its spatial composition.

The two tests imply a distinctive empirical pattern: ancestral climatic shocks must predict the asymmetric weighting of losses rather than attitudes toward risk more generally, and greater idiosyncrasy of these shocks must be associated with lower loss aversion, whereas greater spatial correlation must be associated with higher loss aversion. Evidence consistent with each of these patterns would therefore lend distinctive credence to the proposed evolutionary mechanism.

3.3 Empirical Model

The baseline empirical specification tests the central prediction of the theory according to which ancestral homelands characterized by greater exposure to adverse climatic shocks are associated with lower loss aversion among their U.S.-born descendants. In particular, it estimates the linear regression:

$$LA_a = \alpha + \beta E_a + X_a'\Gamma + \varepsilon_a, \tag{29}$$

where LA_a is the median experimentally elicited loss aversion among individuals of ancestry a, and E_a is ancestry-adjusted overall climatic-shock exposure in ancestral homeland a. The vector X_a comprises the ancestry-adjusted caloric suitability and timing of the Neolithic transition in the ancestral homeland, accounting for differences in agricultural potential and in the historical duration of exposure to agriculture that may have independently shaped the evolution of loss aversion.

The coefficient β captures the relationship between overall climatic-shock exposure in the ancestral homeland and loss aversion, conditional on the characteristics of the ancestral environment. The central prediction of the theory is that $\beta < 0$: greater exposure to adverse climatic shocks selected for greater loss neutrality rather than greater loss aversion.

The analysis subsequently examines the secondary predictions concerning the impact of the spatial structure of climatic shocks on the evolution of loss aversion. Decomposing climatic shocks into their idiosyncratic and aggregate components, the theory predicts that a greater degree of idiosyncrasy is associated with a lower degree of loss aversion.

Let E_a^I and E_a^A denote, respectively, the ancestry-adjusted idiosyncratic and aggregate components of overall climatic-shock exposure in the ancestral homeland a, i.e., $E_a = E_a^I + E_a^A$. As an initial assessment, the two components are examined separately to establish whether both idiosyncratic and aggregate climatic-shock exposure are associated with lower loss aversion, as predicted by the theory, estimating the equations:

$$LA_a = \alpha_k + \beta_k\, E_a^k + X_a'\Gamma_k + \varepsilon_a^k, \qquad k \in \{I, A\}. \tag{30}$$

Moreover, it provides an initial assessment of the prediction that greater idiosyncrasy is associated with lower loss aversion than greater aggregate shock, i.e., that $\beta_I < \beta_A < 0$.

A conclusive formal test of the prediction that greater idiosyncrasy is associated with lower loss aversion requires increasing the idiosyncratic component of climatic-shock exposure while holding overall exposure constant. Let $s_a = E_a^I/E_a$ denote the idiosyncratic share of exposure and let $\bar{s}$ denote its sample mean. Recentered idiosyncratic exposure $\widetilde{E}_a^I = E_a^I - \bar{s}E_a = E_a(s_a - \bar{s})$.

The corresponding empirical specification is therefore

$$LA_a = \widetilde{\alpha} + \beta_T\, E_a + \delta\, \widetilde{E}_a^I + X_a'\widetilde{\Gamma} + \widetilde{\varepsilon}_a. \tag{31}$$

When $\widetilde{E}_a^I = 0$, shocks have the sample-average spatial composition; hence, β_T captures the association between overall climatic-shock exposure and loss aversion when the division of exposure between its idiosyncratic and aggregate components is maintained at the sample-average shares. Holding overall climatic-shock exposure fixed, a one-unit increase in $\widetilde{E}_a^I$ corresponds to a one-unit increase in the idiosyncratic component of climatic-shock exposure and an equal decline in its aggregate component. The coefficient δ therefore captures the effect of increasing the idiosyncratic component while holding overall exposure constant. The theory predicts $\delta < 0$, while $\beta_T < 0$ captures the predicted negative relationship between overall climatic-shock exposure and loss aversion at the sample-average spatial composition.[36]

4 Results

4.1 Main Findings

Table 1 presents the paper's central empirical finding. Consistent with the main prediction of the theory, ancestral homelands characterized by greater exposure to adverse climatic shocks are associated with lower loss aversion among their U.S.-born descendants. This relationship is economically substantial, statistically significant, and remarkably stable across alternative specifications.

[36]Without recentering, the coefficient on overall climatic-shock exposure would be interpreted at $E_a^I = 0$, corresponding to the empirically remote case in which all climatic shocks are aggregate.

Table 1: Loss Aversion and Ancestral Climatic Shocks: Baseline Estimates

	Loss aversion			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−2.39***	−2.21***	−2.59***	−2.37***
	(0.82)	(0.84)	(0.85)	(0.87)
Caloric suitability		✓		✓
Timing of Neolithic transition			✓	✓
Dep. var. mean	1.59	1.59	1.60	1.60
Observations	102	102	98	98
Adjusted R^2	0.065	0.083	0.076	0.10

Notes: The unit of observation is an ancestry. Column (2) accounts for ancestral caloric suitability; Column (3) for the timing of the Neolithic transition separately, and column (4) includes both controls. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Column (1) reports the unconditional relationship between ancestral climatic-shock exposure and loss aversion. The estimated relationship is negative and statistically significant at the 1% level. Thus, ancestries originating in homelands that experienced more frequent adverse climatic shocks exhibit systematically lower loss aversion. This finding accords with the seemingly paradoxical central prediction of the theory: although greater exposure to adverse climatic shocks increased the immediate value of prudent behavior, over the course of human history it selected for greater loss neutrality.

Columns (2)–(4) examine whether this relationship reflects broader characteristics of the ancestral environment that may have been correlated with climatic-shock exposure while independently shaping the evolution of preferences. Column (2) accounts for ancestral caloric suitability, capturing differences in the agricultural potential of ancestral homelands. The relationship remains negative, highly significant, and similar in magnitude. Column (3) instead accounts for the timing of the Neolithic transition, capturing differences in the historical duration of exposure to agriculture. Again, the estimated relationship remains negative and statistically significant at the 1% level. Column (4) includes both ancestral caloric suitability and the timing of the Neolithic transition. The relationship remains negative and statistically significant at the 1% level. The stability of the estimate across columns indicates that the negative association between ancestral climatic-shock exposure and loss aversion is not driven by these broader characteristics of the ancestral environment.

The economic magnitude of the estimated relationship is substantial. Based on the specification in column (4), a move from the lowest to the highest ancestral climatic-shock exposure in the estimation sample is associated with a decline in loss aversion equal to 43% of the sample mean. To place this magnitude in the context of the observed cross-ancestry distribution, applying the implied decline to an ancestry initially characterized by the median degree of loss aversion would move it to approximately the $3rd$ percentile of the distribution.

Figure E.1 depicts the unconditional relationship corresponding to column (1). Each observation represents an ancestry, and the fitted line reveals a pronounced negative association between ancestral climatic-shock exposure and experimentally elicited loss aversion, providing a graphical counterpart to the regression evidence.

4.1.1 Robustness of the Main Findings

The baseline relationship is remarkably stable across alternative ways of constructing both the ancestry-level preference measure and ancestral climatic-shock exposure.

First, since loss aversion is aggregated within ancestries, the ancestry-level measure is estimated more precisely for ancestries represented by larger numbers of respondents. Figure E.3 therefore progressively raises the minimum number of respondents required for an ancestry to enter the estimation sample. The estimated coefficient remains negative and stable as increasingly sparsely represented ancestries are excluded, reinforcing the baseline finding across ancestries measured with progressively greater precision.

Second, the finding is preserved under alternative thresholds for identifying consequential climatic shocks (Table F.2), alternative geographic radii for defining neighboring locations and the spatial correlation of climatic shocks (Table F.3), measuring drought exposure on the intensive rather than the extensive margin (Table F.5), and accounting for differences in weather-station coverage (Table F.6).

Third, the finding is robust to: (i) alternative treatments of respondents reporting multiple ancestries, either restricting the sample to respondents reporting a single ancestry or assigning climatic exposure according to the average exposure of the primary and secondary ancestries (Table F.4); (ii) restricting the sample to Old World ancestries and progressively increasing the minimum percentage of the present-day population that is native to the ancestral homeland prior to 1500 (Figure E.4); and (iii) alternative treatments of extreme loss-aversion estimates (Figure E.2).

4.2 Individual-Level Analysis

The ancestry-level analysis provides the natural baseline for examining the theory's predictions. The evolutionary mechanism concerns persistent differences across ancestral populations, and aggregating the experimentally elicited estimates within ancestries permits each ancestral group to receive equal weight while limiting the influence of idiosyncratic individual observations.

The individual-level analysis provides a complementary assessment of the baseline relationship from a different empirical perspective. Rather than assigning equal weight to each ancestry, it retains each respondent as an observation and therefore implicitly provides greater weight to ancestries represented by larger numbers of respondents. Moreover, it permits the analysis to account directly for individual characteristics such as sex, age, and education. The individual-level analysis therefore examines whether the cross-ancestry relationship identified in the baseline analysis persists when the aggregation structure is relaxed and observable individual characteristics are taken into account.[37]

Table 2: Loss Aversion and Ancestral Climatic Shocks: Individual-Level Estimates

	Loss aversion			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−1.17***	−1.18***	−1.01***	−0.99**
	(0.38)	(0.38)	(0.38)	(0.39)
Sex		✓	✓	✓
Age			✓	✓
Education				✓
Dep. var. mean	1.82	1.82	1.82	1.82
Observations	1424	1424	1424	1424
Adjusted R^2	0.0046	0.0048	0.013	0.018

Notes: This table reports individual-level regressions of loss aversion on ancestral climatic shocks. Individual loss aversion is winsorized in the upper tail at the 80th percentile. Age enters linearly and quadratically. Education fixed effects distinguish between high school or less, some college or an associate degree, a bachelor's degree, and a graduate or professional degree. Standard errors, clustered at the ancestry level, are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 2 presents the results. Column (1) estimates the relationship at the individual level without respondent-level controls. As in the ancestry-level analysis, greater ancestral exposure to climatic shocks is associated with significantly lower loss aversion. Column (2) accounts for sex, leaving the estimated relationship virtually unchanged. Column (3) further accounts for age and age squared. The magnitude declines modestly, but the relationship remains negative and highly statistically significant. Finally, column (4) incorporates education fixed effects, distinguishing among major levels of educational attainment. The estimated relationship remains nearly unchanged relative to column (3) and continues to be statistically significant.

The persistence of the relationship across these specifications reinforces the ancestry-level evidence from two complementary directions. It shows that the main finding does not depend on assigning equal weight to ancestral groups and that it cannot be attributed to systematic differences in sex, age, or educational attainment across individuals of different ancestries.

4.3 Distinguishing Loss Aversion from Risk Aversion: A Placebo Test

A distinctive implication of the theory is that ancestral exposure to adverse climatic shocks shaped loss aversion specifically, rather than attitudes toward risk more generally. Risk aversion governs individuals' response to uncertainty in economic outcomes, whereas loss aversion captures

[37]The individual-level analysis cannot benefit from median aggregation and therefore applies one-sided upper-tail winsorization at the 80th percentile. This treatment addresses the substantially longer right tail of the individual estimates, some of which reach several thousand. The results are robust to alternative winsorization thresholds.

the disproportionate weight assigned to losses relative to comparable gains. The proposed evolutionary mechanism operates through the existential consequences of adverse climatic realizations near the subsistence threshold, where a sufficiently severe loss could threaten dynastic survival. It therefore provides a biological foundation for the asymmetric treatment of losses, rather than for a general aversion to risk. Accordingly, ancestral climatic shocks should predict loss aversion but not risk aversion.

The experimental design permits this distinction to be examined directly. Risk aversion and loss aversion are elicited from the same respondents under identical experimental conditions and are therefore subject to the same sample composition, participation, and experimental environment. Yet they capture distinct dimensions of preferences: risk aversion governs attitudes toward uncertainty more generally, whereas loss aversion captures the asymmetric weighting of losses relative to equivalent gains. Moreover, since the measured degree of loss aversion depends on the underlying degree of risk aversion, jointly eliciting both dimensions prevents cross-ancestry variation in risk preferences from being absorbed into estimated loss aversion. A relationship with loss aversion accompanied by the absence of a corresponding relationship with risk aversion therefore provides a discriminating test of the proposed evolutionary mechanism.

Table 3 presents the results. Columns (1) and (2) use loss aversion as the dependent variable. Column (1) reproduces the negative and statistically significant relationship between ancestral climatic shocks and loss aversion. Column (2) accounts directly for experimentally elicited risk aversion. Risk aversion is itself significantly associated with loss aversion, underscoring the importance of accounting for heterogeneity in this dimension of preferences. Nevertheless, accounting for risk aversion leaves the relationship between ancestral climatic shocks and loss aversion virtually unchanged. The climatic gradient in loss aversion therefore does not reflect cross-ancestry differences in risk aversion.

Columns (3) and (4) provide the placebo test by using risk aversion as the dependent variable. In sharp contrast to the relationship with loss aversion, ancestral climatic shocks are unrelated to risk aversion. This absence of a relationship is evident both unconditionally in column (3) and after accounting for loss aversion in column (4). Hence, ancestral climatic shocks systematically predict loss aversion while displaying no corresponding relationship with risk aversion.

This contrasting pattern is consistent with the hypothesis that the evolutionary mechanism operates through the existential consequences of adverse shocks near the subsistence threshold rather than through exposure to uncertainty per se. The placebo test therefore provides a discriminating validation of the proposed mechanism, linking ancestral climatic conditions specifically to the asymmetric treatment of losses and gains rather than to risk preferences more broadly.

Table 3: Ancestral Climatic Shocks and Risk Aversion: Placebo Test

	Loss aversion		Risk aversion	
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−2.39***	−2.31***	0.52	−0.086
	(0.82)	(0.80)	(0.86)	(0.57)
Risk aversion		−0.16**		
		(0.082)		
Loss aversion				−0.25
				(0.27)
Dep. var. mean	1.59	1.59	1.19	1.19
Observations	102	102	102	102
Adjusted R^2	0.065	0.10	0.0021	0.043

Notes: The unit of observation is an ancestry. Columns (1) and (2) use loss aversion as the dependent variable, while columns (3) and (4) use risk aversion as a placebo outcome. Column (2) controls for risk aversion, and column (4) controls for loss aversion. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

4.4 Spatial Composition of Climatic Shocks

Having established the central prediction that greater ancestral exposure to adverse climatic shocks is associated with lower loss aversion, the analysis turns to the theory's complementary prediction concerning the spatial structure of these shocks. This prediction provides a particularly informative test of the proposed evolutionary mechanism because the theory assigns distinct roles to overall exposure to adverse shocks and to their degree of spatial idiosyncrasy.

Greater adverse-shock exposure is predicted to favor loss neutrality regardless of whether the shocks are predominantly idiosyncratic or aggregate. Greater severity raises the degree of loss neutrality required for individuals to sort into climatically risky modes of production, shifting the composition of individuals in both the safe and risky modes toward greater loss neutrality. The spatial structure of the shocks, however, governs the relative long-run reproductive success of these groups. When adverse shocks are more idiosyncratic, extinction among some relatively loss-neutral dynasties engaged in risky production is accompanied by favorable realizations among others, permitting the higher reproductive returns of the risky mode to translate into greater long-run population growth. Greater spatial correlation instead exposes a larger fraction of the risky group to common adverse realizations, reducing its relative long-run reproductive success. The theory therefore predicts that, holding overall exposure to adverse climatic shocks constant, greater idiosyncrasy should be associated with lower loss aversion.

Table 4 examines these two predictions sequentially. Column (1) reproduces the baseline relationship between overall ancestral climatic-shock exposure and loss aversion. Columns (2) and (3) then decompose overall exposure into its idiosyncratic and aggregate components and examine each separately. Both components are negatively and statistically significantly associated with loss aversion. This pattern reinforces the central prediction of the theory: greater exposure to adverse climatic shocks is associated with greater loss neutrality regardless of the spatial nature of those shocks. Moreover, the negative association is substantially more pronounced for idiosyncratic exposure, pointing in the direction of the theory's complementary prediction. Column (4) provides the direct test of that complementary prediction. It includes overall climatic-shock exposure together with the recentered idiosyncratic component, permitting the spatial composition of shocks to vary while their overall frequency is held constant. The coefficient on overall exposure remains negative and statistically significant. Hence, even after distinguishing the spatial composition of climatic shocks, greater overall exposure continues to be associated with lower loss aversion, further reinforcing the paper's central finding.

Table 4: Loss Aversion and the Spatial Composition of Ancestral Climatic Shocks

	Loss aversion			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−2.39***			−3.52***
	(0.82)			(0.95)
Ancestral idiosyncratic climatic shocks		−15.2***		−18.3**
		(4.01)		(7.69)
Ancestral aggregate climatic shocks			−2.52**	
			(0.99)	
Dep. var. mean	1.59	1.59	1.59	1.59
Observations	102	102	102	102
Adjusted R^2	0.065	0.11	0.050	0.11

Notes: The unit of observation is an ancestry. Columns (1)–(3) relate loss aversion separately to total, idiosyncratic, and aggregate climatic-shock exposure. Column (4) includes total exposure and recentered idiosyncratic exposure jointly. The latter is defined as $CS_a^I - \bar{s}CS_a^T$, where $\bar{s}$ is the sample mean of the idiosyncratic share CS_a^I/CS_a^T. The coefficient on total exposure therefore captures its association with loss aversion at the sample-average spatial composition, while the coefficient on recentered idiosyncratic exposure captures the difference between idiosyncratic and aggregate exposure, holding total exposure constant. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

More importantly, holding overall climatic-shock exposure fixed, an increase in the idiosyncratic component, accompanied by an equal decline in the aggregate component, is associated with a further statistically significant decline in loss aversion. Thus, ancestries exposed to the same overall frequency of adverse climatic shocks exhibit systematically lower loss aversion when a greater share of those shocks was spatially idiosyncratic rather than aggregate. This is precisely the additional variation isolated by the theory's complementary prediction.

The results therefore reveal a distinctive correspondence between the empirical evidence and the two forces identified by the theory. Greater overall exposure to adverse climatic shocks is associated with lower loss aversion irrespective of their spatial nature, reinforcing the central prediction that adverse-shock severity favored loss neutrality. Conditional on overall exposure, greater idiosyncrasy is associated with an additional reduction in loss aversion, confirming the complementary prediction that the spatial structure of shocks shaped the long-run prevalence of relatively loss-neutral individuals. The ability of the evidence to distinguish these two forces is consistent with the distinct roles assigned to them by the theory.

4.5 Historical Climatic Shocks

The evolutionary mechanism developed in the paper operated over many centuries and generations, when recurrent adverse climatic shocks shaped differential reproductive success across populations. The relevant empirical object is therefore the persistent cross-sectional variation in climatic-shock exposure across ancestral locations. Reassuringly, the twentieth-century measure used in the baseline analysis is highly correlated, at both the grid-cell and country levels, with climatic-shock exposure reconstructed over 1–1500 CE. This persistence indicates that the high-resolution twentieth-century measure provides a credible proxy for the relative climatic conditions that prevailed across locations during the formative period of the evolutionary process.

The historical reconstruction provides an independent means of validating this interpretation. It extends directly into the period in which the proposed evolutionary mechanism operated, while the twentieth-century observational record provides substantially greater spatial and temporal precision. More importantly, climatic-shock measures constructed directly from the historical reconstruction provide an independent test of whether the theory's central predictions hold during the formative period itself.

4.5.1 Historical Climatic Reconstruction

The historical analysis draws on the Paleo Hydrodynamics Data Assimilation (PHYDA) reconstruction of Steiger et al. (2018), which provides global estimates of the Palmer Drought Severity Index over the Common Era. PHYDA therefore permits climatic-shock exposure to be measured over the historical period in which the proposed evolutionary mechanism operated. Its proxy-based nature and substantially coarser spatial and temporal resolution make the high-resolution scPDSI record preferable for the baseline analysis, while providing a valuable independent historical counterpart constructed from different underlying information and methods.[38]

The construction otherwise parallels the baseline measure. Climatic shocks are decomposed into idiosyncratic and aggregate components according to their contemporaneous incidence in neighboring locations, aggregated across space using land-area weights, and ancestry-adjusted using the same post-1500 population-migration matrix.

[38] Appendix D provides a detailed description of the PHYDA reconstruction and its measurement properties relative to the contemporary scPDSI record.

4.5.2 Persistence of Climatic-Shock Exposure

The central question is whether the high-resolution twentieth-century measure captures persistent cross-sectional differences in climatic-shock exposure across locations. Figure 7 provides direct evidence that it does, comparing climatic-shock exposure reconstructed over 1–1500 CE with exposure measured from the twentieth-century scPDSI record. Panel A examines the relationship across spatial grid cells, while Panel B aggregates the measures to the country level.

The relationship is strongly positive and highly statistically significant at both levels of aggregation. Locations characterized by greater exposure to climatic shocks during the twentieth century were also systematically more exposed during the premodern period in which the proposed evolutionary process operated. This correspondence is particularly informative given that the two measures cover non-overlapping periods, are constructed from different underlying climatic information, and differ substantially in their spatial and temporal resolution.

The persistence of climatic-shock exposure provides the empirical bridge between the historical mechanism and the baseline measure. It indicates that the twentieth-century scPDSI record captures a persistent component of ancestral climatic conditions rather than climatic realizations specific to the twentieth century. The the higher-resolution contemporary record precisely measured contemporary record can therefore serve as a credible proxy for historical ancestral exposure.

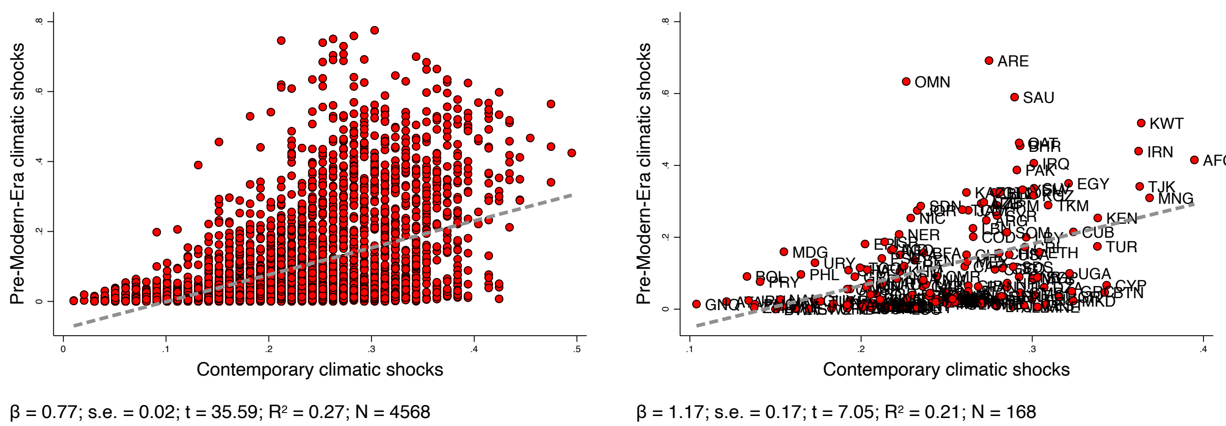

Figure 7: Persistence of Climatic-Shock Exposure Across Historical and Contemporary Periods

Notes: This figure compares historical climatic-shock exposure over $1-1500$ with contemporary exposure over $1902-2000$. Panel A presents the relationship across land grid cells, with the contemporary data aggregated to the spatial units of the coarser PHYDA grid. Panel B presents the corresponding area-weighted measures for countries covered by both data sources. The dashed line depicts the fitted linear relationship. Coefficients, robust standard errors, t-statistics, R^2, and sample sizes are reported below each panel.

4.5.3 Loss Aversion and Historical Climatic Shocks

Having established the persistence of climatic-shock exposure across the historical and contemporary periods, the analysis next examines whether climatic conditions reconstructed directly over

1–1500 CE independently reproduce the theory's predictions. This provides a particularly informative validation, since the climatic variation is measured during the period in which the proposed evolutionary process operated and is derived from information and methods entirely distinct from those underlying the baseline measure.

Table 5 reports the results. Column (1) establishes a negative and statistically significant relationship between loss aversion and overall historical climatic-shock exposure. The paper's central and seemingly paradoxical prediction therefore holds when ancestral exposure is measured from the historical reconstruction: populations originating in regions characterized by greater adverse-shock exposure exhibit greater loss neutrality.

Columns (2) and (3) decompose historical exposure into its idiosyncratic and aggregate components and examine each separately. Both are negatively and significantly associated with loss aversion, reinforcing the central prediction that greater exposure to adverse climatic shocks favored loss neutrality regardless of their spatial nature. Moreover, the relationship is substantially more pronounced for idiosyncratic exposure, consistent with the theory's complementary prediction concerning the spatial structure of shocks. Column (4) provides the direct compositional test. As in the main analysis, overall climatic-shock exposure and recentered idiosyncratic exposure enter jointly, permitting the spatial composition of shocks to vary while holding their overall frequency constant. Greater idiosyncratic exposure is associated with a further statistically significant decline in loss aversion, confirming the prediction that greater idiosyncrasy favored the long-run prevalence of relatively loss-neutral individuals. Importantly, overall historical exposure also remains negatively and significantly associated with loss aversion, further reinforcing the theory's central prediction.

Table 5: Loss Aversion and Historical Ancestral Climatic Shocks

	Loss aversion			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks (historical)	−0.98**			−3.70***
	(0.40)			(1.14)
Ancestral idiosyncratic climatic shocks (historical)		−7.96***		−14.9**
		(2.53)		(5.98)
Ancestral aggregate climatic shocks (historical)			−0.97**	
			(0.48)	
Dep. var. mean	1.59	1.59	1.59	1.59
Observations	96	96	96	96
Adjusted R^2	0.031	0.077	0.021	0.096

Notes: This table reports ancestry-level regressions of loss aversion on historical measures of ancestral climatic shocks constructed from the posterior ensemble mean of the (PHYDA) PDSI reconstruction over 1–1500. The June–August reconstruction is used for Northern Hemisphere land cells and the December–February reconstruction for Southern Hemisphere land cells. A historical shock is defined as PDSI less than or equal to −1. Column (4) includes total and recentered idiosyncratic exposure jointly. Recentered idiosyncratic exposure equals $I_a - \bar{s}T_a$, where $\bar{s}$ is the sample mean of I_a/T_a. Its coefficient captures the difference between idiosyncratic and aggregate exposure, holding total exposure constant; the coefficient on total exposure is evaluated at the sample-average spatial composition. The estimation sample contains six fewer ancestries than the corresponding baseline specification because the coarser PHYDA grid does not permit constructing historical measures for several baseline countries. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The historical reconstruction therefore independently reproduces both dimensions of the empirical evidence: greater overall exposure to adverse climatic shocks is associated with lower loss aversion, and, conditional on overall exposure, greater idiosyncrasy is associated with an additional reduction in loss aversion. This correspondence is particularly informative given that the coarser spatial and temporal resolution of the historical reconstruction makes localized shocks, and hence their idiosyncratic component, more difficult to identify. The historical evidence thus both validates the twentieth-century measure as a proxy for persistent ancestral climatic conditions and establishes that the theory's predictions hold when climatic shocks are reconstructed directly over the period in which the evolutionary process operated.

4.6 Quasi-Natural Experiment: The Columbian Exchange

Sorting of loss-averse individuals into less shock-prone locations could have generated a relationship between ancestral exposure to adverse climatic shocks and contemporary loss aversion independently of evolutionary adaptation through differential reproductive success. Yet systematic sorting along this dimension would have been prohibitively difficult when founder populations settled novel ecological niches, as they lacked the accumulated knowledge required to assess the long-run incidence and spatial structure of adverse climatic shocks. Moreover, settlement reflected collective decision-making rather than individual choice, so heterogeneity in loss aversion within migrating groups placed an additional hurdle on systematic sorting across locations.

Nevertheless, to address the unlikely possibility that some degree of historical sorting may have occurred, the Columbian Exchange provides an additional source of identification based on variation in climatic-shock exposure that arose only after populations had already settled their locations. It generates quasi-exogenous variation in effective climatic-shock exposure through the introduction of new crops.

The introduction of previously unavailable crops in the post-1500 period altered the locally optimal crop and, consequently, given differences in the timing of agricultural growth cycles, effective exposure to adverse climatic shocks. For each grid cell, crop-specific caloric suitability is used to identify the locally optimal crop among those available prior to 1500 and among the expanded set available following the Columbian Exchange. FAO GAEZ crop calendars then identify the interval between successive growth cycles for each of these crops, permitting climatic-shock exposure to be measured under the agricultural calendars implied by the pre- and post-1500 crop sets. Notably, since the locally optimal crops are determined based on potential agro-climatic suitability rather than observed cultivation, and the pre- and post-1500 measures reflect the same underlying climatic history, the resulting change in climatic-shock exposure isolates the variation arising from the expansion of the feasible crop set rather than endogenous crop adoption or changes in the underlying climatic environment.

The analysis focuses on the interval between successive growth cycles. Since the crop-specific growth cycle typically lasts approximately three to four months, this interval comprises the remaining eight to nine months before the onset of the next growth cycle. Climatic conditions during this period shape the moisture conditions prevailing at the onset of the next growth cycle, affecting crop establishment as well as subsequent performance.

The identifying strength of this natural experiment is also its limitation as a measure of the climatic shocks relevant to the evolutionary mechanism. The Columbian Exchange isolates quasi-exogenous variation in one particular channel through which drought could affect subsistence: the timing of drought exposure relative to the crop-specific agricultural cycle. The existential consequences of adverse climatic shocks, however, were considerably broader. Drought could affect livestock productivity and survival, access to drinking water, fisheries, and the availability of wild plants, game, and other non-cultivated food sources, each governed by seasonal patterns that need not coincide with the agricultural cycle. Hence, the Columbian Exchange serves as a discriminating identification test, while the baseline year-long measure of climatic vulnerability more closely captures the broad evolutionary mechanism through which adverse climatic shocks could threaten subsistence across multiple livelihood channels.

Table 6: **Crop-Specific Climatic Shocks and the Columbian Exchange**

	Loss aversion				
	(1)	(2)	(3)	(4)	(5)
Ancestral climatic shocks	−2.33***				−2.37***
	(0.82)				(0.82)
Ancestral climatic shocks (crop-specific pre-1500)		−1.93**		−1.99**	
		(0.90)		(0.91)	
Ancestral climatic shocks (crop-specific post-1500)			−1.98**		
			(0.91)		
Δ Ancestral climatic shocks (crop-specific)				−52.7**	−52.1**
				(26.3)	(25.6)
Dep. var. mean	1.60	1.60	1.60	1.60	1.60
Observations	101	101	101	101	101
Adjusted R^2	0.062	0.045	0.047	0.076	0.092

Notes: Ancestral climatic shocks denote the baseline year-long measure of drought exposure. Pre-1500 and post-1500 crop-specific climatic shocks measure exposure during the interval between successive growth cycles of the locally optimal crop under the crop sets available before and after the Columbian Exchange, respectively. Both measures use the same underlying climatic data; the pre-1500 and post-1500 designations refer only to the crop set used to determine the relevant exposure window. Δ Ancestral climatic shocks is the post-1500 measure minus its pre-1500 counterpart. Column (4) conditions on pre-1500 crop-specific exposure and column (5) on baseline climatic exposure. Crop-specific measures are unavailable for Iceland, reducing the sample from 102 to 101 ancestries. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6 reports the results. Climatic-shock exposure under both the pre-1500 and post-1500 crop regimes is negatively and significantly associated with loss aversion, providing further support for the paper's central hypothesis. More importantly, the Columbian-Exchange-induced change in climatic-shock exposure is negative and statistically significant conditional on pre-1500 crop-specific exposure. This relationship remains negative and significant when the analysis instead conditions on the baseline year-long measure of climatic vulnerability. Including this Columbian-Exchange variation leaves the estimated relationship between baseline climatic-shock exposure and loss aversion virtually unchanged. Since this quasi-exogenous variation arose within already settled locations, its association with contemporary loss aversion indicates that the baseline relationship cannot be explained solely by the sorting of populations across climatic environments.

5 Concluding Remarks

This research advances the hypothesis that systematic differences in loss aversion across societies can be traced to adaptation to ancestral environments in which the subsistence consumption constraint constituted an existential cliff. The resulting evolutionary force generates a seemingly paradoxical prediction. Greater frequency or severity of adverse climatic shocks made prudent behavior and greater reliance on climatically insensitive modes of production individually rational, yet selected over the course of human history for greater loss neutrality. Moreover, holding adverse-shock frequency and severity constant, greater idiosyncrasy favored the long-run prevalence of relatively loss-neutral individuals, whereas greater spatial correlation favored greater loss aversion. Thus, the long-run evolutionary response to a harsher climatic environment need not mirror the prudent behavioral response that the same environment induces at the individual level. Indeed, environments in which losses were potentially most consequential ultimately generated populations that were less, rather than more, loss averse.

The empirical evidence is consistent with both dimensions of this prediction. Exploiting experimentally elicited measures of loss aversion among U.S.-born individuals from diverse ancestral origins, the analysis establishes that greater ancestral exposure to adverse climatic shocks is associated with substantially lower loss aversion. Holding overall exposure constant, greater idiosyncrasy of these shocks is associated with a further reduction in loss aversion, reproducing the distinct role assigned by the theory to their spatial structure. The relationship is specific to loss aversion rather than risk aversion, and is reproduced using climatic conditions reconstructed directly over 1–1500 CE.

References

Alesina, A., Giuliano, P. and Nunn, N. (2013). On the origins of gender roles: Women and the plough, *The Quarterly Journal of Economics* **128**(2): 469–530.

Alesina, A. and Passarelli, F. (2019). Loss aversion in politics, *American Journal of Political Science* **63**(4): 936–947.

Ashraf, Q. and Galor, O. (2011). Dynamics and stagnation in the malthusian epoch, *The American Economic Review* **101**(5): 2003–2041.

Ashraf, Q. and Galor, O. (2013). The "out of africa" hypothesis, human genetic diversity, and comparative economic development, *The American Economic Review* **103**(1): 1–46.

Barberis, N. and Huang, M. (2001). Mental accounting, loss aversion, and individual stock returns, *The Journal of Finance* **56**(4): 1247–1292.

Baudin, T., de la Croix, D. and Gobbi, P. E. (2015). Fertility and childlessness in the united states, *American Economic Review* **105**(6): 1852–82.
URL: *http://www.aeaweb.org/articles?id=10.1257/aer.20120926*

Benartzi, S. and Thaler, R. H. (1995). Myopic loss aversion and the equity premium puzzle, *The quarterly journal of Economics* **110**(1): 73–92.

Bisin, A. and Verdier, T. (2001). The economics of cultural transmission and the dynamics of preferences, *Journal of Economic theory* **97**(2): 298–319.

Buggle, J. C. and Durante, R. (2021). Climate risk, cooperation and the co-evolution of culture and institutions, *The Economic Journal* **131**(637): 1947–1987.

Caraco, T., Martindale, S. and Whittam, T. S. (1980). An empirical demonstration of risk-sensitive foraging preferences, *Animal Behaviour* **28**(3): 820–830.

Chen, M. K., Lakshminarayanan, V. and Santos, L. R. (2006). How basic are behavioral biases? evidence from capuchin monkey trading behavior, *Journal of Political Economy* **114**(3): 517–537.

Dalgaard, C.-J. and Strulik, H. (2015). The physiological foundations of the wealth of nations, *Journal of Economic Growth* **20**(1): 37–73.

Fisher, P. J. and Montalto, C. P. (2011). Loss aversion and saving behavior: Evidence from the 2007 us survey of consumer finances, *Journal of Family and Economic Issues* **32**(1): 4–14.

Galor, O., Klemp, M. C. and Wainstock, D. C. (2026). Cultural evolution. Brown University.

Galor, O. and Michalopoulos, S. (2012). Evolution and the growth process: Natural selection of entrepreneurial traits, *Journal of Economic Theory* **147**(2): 759–780.

Galor, O. and Moav, O. (2002). Natural selection and the origin of economic growth, *The Quarterly Journal of Economics* **117**(4): 1133–1191.

Galor, O. and Özak, Ö. (2016). The agricultural origins of time preference, *The American Economic Review* **106**(10): 3064–3103.

Genesove, D. and Mayer, C. (2001). Loss aversion and seller behavior: Evidence from the housing market, *The Quarterly Journal of Economics* **116**(4): 1233–1260.

Jervis, R. (1992). Political implications of loss aversion, *Political psychology* pp. 187–204.

Kahneman, D., Knetsch, J. L. and Thaler, R. H. (1991). Anomalies: The endowment effect, loss aversion, and status quo bias, *The journal of economic perspectives* **5**(1): 193–206.

Khaw, M. W., Li, Z. and Woodford, M. (2021). Cognitive imprecision and small-stakes risk aversion, *The Review of Economic Studies* **88**(4): 1979–2013.

Kőszegi, B. and Rabin, M. (2009). Reference-dependent consumption plans, *American Economic Review* **99**(3): 909–936.

Levy, J. S. (1996). Loss aversion, framing, and bargaining: The implications of prospect theory for international conflict, *International Political Science Review* **17**(2): 179–195.

Mayshar, J., Moav, O. and Pascali, L. (2022). The origin of the state: Productivity or appropriability?, *Journal of Political Economy* .

Michalopoulos, S. and Papaioannou, E. (2020). Historical legacies and african development, *Journal of Economic Literature* **58**(1): 53–128.

Putterman, L. (2008). Agriculture, diffusion and development: Ripple effects of the neolithic revolution, *Economica* **75**(300): 729–748.

Putterman, L. and Weil, D. N. (2010). Post-1500 population flows and the long-run determinants of economic growth and inequality, *The Quarterly journal of economics* **125**(4): 1627–1682.

Robson, A. J. (1996). A biological basis for expected and non-expected utility, *Journal of economic theory* **68**(2): 397–424.

Robson, A. J. and Samuelson, L. (2011). The evolutionary foundations of preferences, *Handbook of social economics*, Vol. 1, Elsevier, pp. 221–310.

Shalev, J. (2002). Loss aversion and bargaining, *Theory and decision* **52**(3): 201–232.

Spolaore, E. and Wacziarg, R. (2013). How deep are the roots of economic development?, *Journal of Economic Literature* **51**(2): 325–369.

Steiger, N. J., Smerdon, J. E., Cook, E. R. and Cook, B. I. (2018). A reconstruction of global hydroclimate and dynamical variables over the common era, *Scientific Data* **5**: 180086.
URL: *https://doi.org/10.1038/sdata.2018.86*

Stewart, N., Chater, N. and Brown, G. D. A. (2006). Decision by sampling, *Cognitive Psychology* **53**(1): 1–26.

Tversky, A. and Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model, *The quarterly journal of economics* **106**(4): 1039–1061.

van der Schrier, G., Barichivich, J., Briffa, K. R. and Jones, P. D. (2013). A scPDSI-based global data set of dry and wet spells for 1901–2009, *Journal of Geophysical Research: Atmospheres* **118**(9): 4025–4048.

Vollrath, D. (2011). The agricultural basis of comparative development, *Journal of Economic Growth* **16**(4): 343–370.

Wells, N., Goddard, S. and Hayes, M. J. (2004). A self-calibrating palmer drought severity index, *Journal of Climate* **17**(12): 2335–2351.

Appendix (Supplementary Online Material)

Contents

A Proof of Lemma 1 and Proposition 1

A.1 Proof of Lemma 1

Proof. The asymptotic composition of the population is determined by the relative long-run population growth rates of groups A and B. Define the asymptotic logarithmic population growth rate of group $j \in \{A, B\}$ as

$$\bar{n}^j \equiv \lim_{t\to\infty} \frac{1}{t} \ln N_t^j, \qquad j = A, B.$$

As follows from (14), $N_t^A = N_0^A$ for all t, and therefore $\bar{n}^A = 0$, reflecting the fact that the population of group A remains constant over time. For group B, equation (16) implies

$$\frac{1}{t} \ln N_t^B = \frac{1}{t} \sum_{s=0}^{t-1} \ln \left[\mu \bar{n}_s^{ag} + (1-\mu)\bar{n}^{id}\right] + \frac{1}{t} \ln N_0^B.$$

Since aggregate climatic realizations are independent across generations, the Strong Law of Large Numbers implies that, almost surely, for $\mu < 1$,

$$\begin{aligned}
\bar{n}^B(\mu, p) &= \mathbb{E} \ln \left[\mu \bar{n}_s^{ag} + (1-\mu)\bar{n}^{id}\right] \\
&= p \ln \left[\mu \frac{y^h - \tilde{c}}{\rho y^h} + (1-\mu) p \frac{y^h - \tilde{c}}{\rho y^h}\right] + (1-p) \ln \left[(1-\mu) p \frac{y^h - \tilde{c}}{\rho y^h}\right] \\
&= p \ln [\mu + (1-\mu)p] + (1-p) \ln [(1-\mu)p] + \ln \left(\frac{y^h - \tilde{c}}{\rho y^h}\right).
\end{aligned}$$

As $\mu \to 1$, the idiosyncratic component vanishes and an adverse aggregate climatic realization eliminates group B. Accordingly, $\lim_{\mu\to 1} \bar{n}^B(\mu, p) = -\infty$.

At $\mu = 1$, the argument is direct. Since an adverse aggregate realization occurs with probability $1 - p > 0$ independently in every generation, such a realization occurs in finite time almost surely. Once it occurs, $N_t^B = 0$ thereafter, and hence $\beta_t \to 0$ almost surely. Moreover, for $\mu \in (0, 1)$,

$$\frac{\partial \bar{n}^B(\mu, p)}{\partial \mu} = (1-p) \left[\frac{p}{\mu + (1-\mu)p} - \frac{1}{1-\mu}\right] = -\frac{\mu(1-p)}{[\mu + (1-\mu)p](1-\mu)} < 0.$$

Thus, the asymptotic population growth rate of group B is strictly decreasing in μ, establishing part (iii).

At $\mu = 0$,

$$\bar{n}^B(0, p) = \ln \left[p \frac{y^h - \tilde{c}}{\rho y^h}\right] > 0,$$

where the inequality follows from the assumption that the expected reproductive return associated with the risky mode exceeds replacement fertility, i.e., $p(y^h - \tilde{c})/(\rho y^h) > 1$.

Since $\bar{n}^A = 0$, $\bar{n}^B(0, p) > 0$, $\lim_{\mu\to 1} \bar{n}^B(\mu, p) = -\infty$, and $\bar{n}^B(\mu, p)$ is continuous and strictly decreasing for $\mu \in [0, 1)$, there exists a unique $\hat{\mu}(p) \in (0, 1)$ such that $\bar{n}^B(\hat{\mu}(p), p) = \bar{n}^A = 0$. Consequently,

$$\begin{aligned}
\bar{n}^B(\mu, p) > \bar{n}^A = 0 \quad \text{for } \mu < \hat{\mu}(p), \\
\bar{n}^B(\mu, p) < \bar{n}^A = 0 \quad \text{for } \mu > \hat{\mu}(p).
\end{aligned}$$

If $\mu < \hat{\mu}(p)$, group B grows asymptotically faster than group A, and therefore

$$\frac{N_t^B}{N_t^A} \longrightarrow \infty.$$

Hence, as follows from (17),

$$\bar{\beta} \equiv \lim_{t\to\infty} \beta_t = 1.$$

Moreover, as follows from (18), $\bar{\theta} \equiv \lim_{t\to\infty} \bar{\theta}_t = (1+\hat{\theta})/2$. Thus, group B dominates the population in the long run, establishing part (i).

If $\mu > \hat{\mu}(p)$, group B grows asymptotically more slowly than group A, and therefore $(N_t^B/N_t^A) \longrightarrow 0$. Hence, $\bar{\beta} = 0$, and as follows from (18), $\bar{\theta} = \hat{\theta}/2$. Thus, group B vanishes in the long run, establishing part (ii). □

A.2 Proof of Proposition 1

Proof. As established in Section 2.3.2,

$$\frac{\partial \hat{\theta}(\Delta, p)}{\partial \Delta} > 0.$$

Moreover, as follows from the expression for $\hat{\theta}(\Delta, p)$,

$$\frac{\partial \hat{\theta}(\Delta, p)}{\partial (1-p)} = \frac{[g(\tilde{c}) - g(y_l(\Delta))]\left[\eta\left(\frac{y^h - \tilde{c}}{\rho y^h}\right) - \eta(1)\right]}{\left[p\eta\left(\frac{y^h - \tilde{c}}{\rho y^h}\right) + (1-p)\eta(0) - \eta(1)\right]^2} > 0,$$

where the inequality follows since $g(\tilde{c}) > g(y_l(\Delta))$, $(y^h - \tilde{c})/(\rho y^h) > 1$, and η is strictly increasing.

As follows from Lemma 1, if $\mu < \hat{\mu}(p)$, $\bar{\theta} = (1 + \hat{\theta}(\Delta, p))/2$, whereas if $\mu > \hat{\mu}(p)$, $\bar{\theta} = \hat{\theta}(\Delta, p)/2$. Hence, for $\mu \neq \hat{\mu}(p)$,

$$\frac{\partial \bar{\theta}}{\partial \Delta} = \frac{1}{2}\frac{\partial \hat{\theta}(\Delta, p)}{\partial \Delta} > 0, \qquad \frac{\partial \bar{\theta}}{\partial (1-p)} = \frac{1}{2}\frac{\partial \hat{\theta}(\Delta, p)}{\partial (1-p)} > 0.$$

Since a higher degree of loss neutrality corresponds to a lower degree of loss aversion, an increase in adverse-shock frequency or severity reduces long-run loss aversion both when climatic shocks are relatively idiosyncratic, $\mu < \hat{\mu}(p)$, and when they are relatively aggregate, $\mu > \hat{\mu}(p)$, establishing the main prediction.

Moreover, Lemma 1 implies that

$$\bar{\theta} = \begin{cases} \dfrac{1+\hat{\theta}}{2} & \mu < \hat{\mu}(p), \\[2ex] \dfrac{\hat{\theta}}{2} & \mu > \hat{\mu}(p). \end{cases}$$

Thus, holding adverse-shock frequency and severity constant, a lower degree of spatial correlation, μ, and therefore greater idiosyncrasy of climatic shocks, $(1-\mu)$, is associated with a weakly higher long-run degree of loss neutrality and hence a weakly lower degree of loss aversion. Conversely, greater spatial correlation is associated with a weakly lower long-run degree of loss neutrality and hence a weakly higher degree of loss aversion, establishing the complementary prediction.

□

B Robustness to Alternative Preferences and Constraints

B.1 Evolution of Loss Aversion: Alternative Utility Functions & Constraints

B.1.1 Loss Aversion with Respect to Fertility and Consumption

This subsection demonstrates that the main predictions of the baseline model regarding the effects of the frequency or severity of an adverse shock and the degree of its spatial correlation on the evolution of loss aversion remain unaltered if loss aversion is defined over both consumption and fertility.

Suppose that individuals are loss averse with respect to both consumption and fertility. In general, this could require the introduction of separate degrees of loss aversion and reference points with respect to consumption and fertility. To preserve the one-dimensional heterogeneity of the baseline model, suppose that the same degree of loss aversion, $1/\theta_i$, applies to losses in both consumption and fertility, and let replacement fertility, $\tilde{n} = 1$, be the reference point with respect to fertility.

The utility of an adult i in period t, u_{it}, is:

$$u_{it} = \begin{cases} [f(c_{it}) - f(\tilde{c})] + [\eta(n_{it}) - \eta(1)] & \text{if } c_{it} \geq \tilde{c} \text{ and } n_{it} \geq 1, \\ \frac{1}{\theta_i}[g(c_{it}) - g(\tilde{c})] + [\eta(n_{it}) - \eta(1)] & \text{if } c_{it} < \tilde{c} \text{ and } n_{it} \geq 1, \\ [f(c_{it}) - f(\tilde{c})] + \frac{1}{\theta_i}[z(n_{it}) - z(1)] & \text{if } c_{it} \geq \tilde{c} \text{ and } n_{it} < 1, \\ \frac{1}{\theta_i}\{[g(c_{it}) - g(\tilde{c})] + [z(n_{it}) - z(1)]\} & \text{if } c_{it} < \tilde{c} \text{ and } n_{it} < 1, \end{cases}$$

where, as in the baseline model, $f(\cdot)$ and $\eta(\cdot)$ are concave and strictly monotonically increasing, while $g(\cdot)$ and $z(\cdot)$ are convex and strictly monotonically increasing.

Proposition A1. *The main prediction of the baseline model—that greater frequency or severity of an adverse shock favors the evolution of greater loss neutrality—remains unaltered when loss aversion is defined over both consumption and fertility. In addition, the secondary prediction that a greater degree of spatial correlation in adverse shocks favors greater loss aversion is preserved.*

Proof. Given the production structure and constraints of the baseline model, the safe production mode generates $(c_{it}, n_{it}) = (\tilde{c}, 1)$, the ordinary climatic realization of the risky mode generates $(c_{it}, n_{it}) = (\tilde{c}, n^h)$, where $n^h \equiv (y^h - \tilde{c})/(\rho y^h) > 1$, and the adverse realization generates $(c_{it}, n_{it}) = (y^l(\Delta), 0)$, where $y^l(\Delta) < \tilde{c}$. Thus, expected utility from the risky production mode is

$$u^r_{it} = p\left[\eta(n^h) - \eta(1)\right] + \frac{1-p}{\theta_i}\{[g(y_l(\Delta)) - g(\tilde{c})] + [z(0) - z(1)]\} \equiv u^r(\theta_i; \Delta, p).$$

Since

$$g(y_l(\Delta)) - g(\tilde{c}) < 0 \qquad \text{and} \qquad z(0) - z(1) < 0,$$

it follows that

$$\lim_{\theta_i \to 0} u^r(\theta_i; \Delta, p) = -\infty \qquad \text{and} \qquad \frac{\partial u^r(\theta_i; \Delta, p)}{\partial \theta_i} > 0 \quad \forall\, \theta_i \in (0, 1].$$

The utility generated by the safe production mode is

$$u^s_{it} = 0 \equiv u^s.$$

Assuming that the risky production mode is sufficiently attractive that the least loss-averse individual prefers it,

$$u^r(1;\Delta,p) > u^s,$$

there exists a unique and time-invariant cutoff $\hat{\theta}(\Delta,p) \in (0,1)$ such that

$$u^r\Big(\hat{\theta}(\Delta,p);\Delta,p\Big) = u^s.$$

In particular,

$$\hat{\theta}(\Delta,p) = \frac{(1-p)\left[g(\tilde{c}) - g(y_l(\Delta)) + z(1) - z(0)\right]}{p\left[\eta(n^h) - \eta(1)\right]}.$$

Since $y_l'(\Delta) < 0$,

$$\frac{\partial\hat{\theta}(\Delta,p)}{\partial\Delta} = -\frac{(1-p)g'(y_l(\Delta))y_l'(\Delta)}{p\left[\eta(n^h)-\eta(1)\right]} > 0.$$

Moreover,

$$\frac{\partial\hat{\theta}(\Delta,p)}{\partial(1-p)} = \frac{g(\tilde{c}) - g(y_l(\Delta)) + z(1) - z(0)}{p^2\left[\eta(n^h)-\eta(1)\right]} > 0.$$

Thus, allowing for loss aversion with respect to fertility changes the level of the cutoff but leaves its response to the frequency and severity of an adverse shock unchanged. Moreover, conditional on the choice of production mode, fertility under each realization is identical to that in the baseline model. The population dynamics associated with the frequency and spatial correlation of adverse shocks are therefore unchanged. Hence, for $\mu \neq \hat{\mu}(p)$, greater adverse-shock frequency or severity increases long-run loss neutrality. Since individuals with $\theta_i \leq \hat{\theta}(\Delta,p)$ choose the safe mode whereas those with $\theta_i > \hat{\theta}(\Delta,p)$ choose the risky mode, the secondary prediction that a greater degree of spatial correlation favors greater loss aversion is preserved.

B.1.2 Positive Fertility Below the Subsistence Consumption Constraint

This subsection demonstrates that the main predictions of the baseline model regarding the effects of the frequency or severity of an adverse shock and the degree of its spatial correlation on the evolution of loss aversion remain unaltered if individuals whose resources fall below subsistence retain positive, but below-replacement, fertility.

Proposition A2. *Suppose that consumption below the subsistence level results in a fixed positive but below-replacement fertility rate, $n^l \in (0,1)$, that is independent of adverse-shock severity, and that the geometric reproductive return to risky production under fully aggregate climatic shocks is below replacement, i.e., $p\ln n^h + (1-p)\ln n^l < 0$. The main prediction of the baseline model—that greater frequency or severity of an adverse shock favors the evolution of greater loss neutrality—remains unaltered. In addition, the secondary prediction that a greater degree of spatial correlation in adverse shocks favors greater loss aversion is preserved.*

Proof. The baseline model assumes that consumption below the subsistence level, $\tilde{c}$, results in zero fertility and therefore in the termination of the lineage. This assumption provides a transparent representation of the evolutionary cliff associated with subsistence, but the qualitative results do not require adverse climatic realizations to generate complete infertility.

Suppose that, following an adverse realization under the risky mode of production, fertility is a fixed $n^l \in (0,1)$, independent of Δ, rather than zero. Thus, an adverse realization causes the affected dynasty to contract rather than disappear immediately. Given the budget constraint, consumption in the adverse state is $c^l(\Delta) \equiv \left(1 - \rho n^l\right) y_l(\Delta) < \tilde{c}$. Since $y_l'(\Delta) < 0$,

$$\frac{dc^l(\Delta)}{d\Delta} < 0.$$

As in the baseline model, let fertility under the favorable realization of the risky production mode be:

$$n^h \equiv \frac{y^h - \tilde{c}}{\rho y^h} > 1$$

Expected utility from the risky mode is now

$$u^r(\theta_i; \Delta, p) = p\eta(n^h) + (1-p)\left[\frac{1}{\theta_i}\Big(g(c^l(\Delta)) - g(\tilde{c})\Big) + \eta(n^l)\right],$$

whereas utility from the safe mode remains $u^s = \eta(1)$.

Maintaining the assumption that the risky production mode is sufficiently attractive that the least loss-averse type prefers it, $u^r(1; \Delta, p) > u^s$, there exists a unique cutoff degree of loss neutrality, $\hat{\theta}(\Delta, p) \in (0,1)$, such that $u^r(\hat{\theta}(\Delta, p); \Delta, p) = u^s$. In particular,

$$\hat{\theta}(\Delta, p) = \frac{(1-p)\left[g(\tilde{c}) - g(c^l(\Delta))\right]}{p\eta(n^h) + (1-p)\eta(n^l) - \eta(1)}.$$

The denominator is positive under $u^r(1; \Delta, p) > u^s$. Hence,

$$\frac{\partial \hat{\theta}(\Delta, p)}{\partial \Delta} = -\frac{(1-p)g'(c^l(\Delta))c^{l\prime}(\Delta)}{p\eta(n^h) + (1-p)\eta(n^l) - \eta(1)} > 0.$$

Moreover,

$$\frac{\partial \hat{\theta}(\Delta, p)}{\partial (1-p)} = \frac{\left[g(\tilde{c}) - g(c^l(\Delta))\right]\left[\eta(n^h) - \eta(1)\right]}{\left[p\eta(n^h) + (1-p)\eta(n^l) - \eta(1)\right]^2} > 0.$$

Thus, as in the baseline model, greater adverse-shock frequency or severity raises the cutoff degree of loss neutrality and increases the average degree of loss neutrality within both the safe and risky groups.

It remains to establish that allowing positive fertility in the adverse state preserves the long-run population dynamics. Under idiosyncratic climatic shocks, average fertility among individuals who choose the risky mode is

$$\bar{n}^{id} = pn^h + (1-p)n^l.$$

Since the baseline assumption $pn^h > 1$ continues to hold, $\bar{n}^{id} > 1$.

Hence, when climatic realizations are entirely idiosyncratic, the population choosing the risky mode grows faster than the population choosing the safe mode, whose fertility remains at replacement.

Under aggregate climatic shocks, fertility among individuals choosing the risky mode is n^h with probability p and n^l with probability $1-p$. Consequently, for a degree of spatial correlation $\mu \in [0,1]$, the asymptotic logarithmic population growth rate of the risky group is

$$\bar{n}^B(\mu, p) = p \ln \left[\mu n^h + (1-\mu)\bar{n}^{id}\right] + (1-p) \ln \left[\mu n^l + (1-\mu)\bar{n}^{id}\right].$$

Differentiating with respect to $\mu > 0$,

$$\frac{\partial \bar{n}^B(\mu, p)}{\partial \mu} = -\frac{\mu p(1-p)(n^h - n^l)^2}{[\mu n^h + (1-\mu)\bar{n}^{id}]\,[\mu n^l + (1-\mu)\bar{n}^{id}]} < 0$$

Thus, as in the baseline model, greater spatial correlation reduces the asymptotic population growth rate of individuals who choose the risky mode.

At $\mu = 0$,

$$\bar{n}^B(0, p) = \ln \left[p n^h + (1-p) n^l\right] > 0.$$

Suppose further that the geometric reproductive return to risky production under fully aggregate climatic shocks is below replacement:

$$p \ln n^h + (1-p) \ln n^l < 0,$$

or equivalently, $(n^h)^p (n^l)^{1-p} < 1$. It follows that

$$\bar{n}^B(1, p) = p \ln n^h + (1-p) \ln n^l < 0.$$

Since $\bar{n}^B(\mu, p)$ is continuous and strictly decreasing in μ, there exists a unique threshold $\hat{\mu}(p) \in (0,1)$ such that $\bar{n}^B(\hat{\mu}(p), p) = 0$.

Hence, if $\mu < \hat{\mu}(p)$, individuals who choose the risky mode grow asymptotically faster than those who choose the safe mode and dominate the population in the long run, whereas if $\mu > \hat{\mu}(p)$, they grow asymptotically more slowly and vanish in the long run.

Moreover, since n^l is independent of adverse-shock severity, $\hat{\mu}(p)$ is independent of Δ. Thus, exactly as in the baseline model,

$$\bar{\theta} = \begin{cases} \dfrac{1 + \hat{\theta}(\Delta, p)}{2}, & \mu < \hat{\mu}(p), \\[2ex] \dfrac{\hat{\theta}(\Delta, p)}{2}, & \mu > \hat{\mu}(p). \end{cases}$$

Therefore, for $\mu \neq \hat{\mu}(p)$,

$$\frac{\partial \bar{\theta}}{\partial \Delta} = \frac{1}{2} \frac{\partial \hat{\theta}(\Delta, p)}{\partial \Delta} > 0, \qquad \frac{\partial \bar{\theta}}{\partial (1-p)} = \frac{1}{2} \frac{\partial \hat{\theta}(\Delta, p)}{\partial (1-p)} > 0.$$

The result follows from the positive effects of adverse-shock frequency and severity on $\hat{\theta}(\Delta, p)$, the strict decline of the asymptotic population growth rate of the risky group with μ, and the existence of a unique $\hat{\mu}(p) \in (0,1)$ established above.

B.1.3 Robustness to More General Forms of the Utility Function

This subsection demonstrates that the main predictions of the baseline model regarding the effects of the frequency or severity of an adverse shock and the degree of its spatial correlation on the evolution of loss aversion remain unaltered if a more general specification of the utility function is adopted.

Proposition A3. *The main prediction of the baseline model, that greater frequency or severity of an adverse shock favors the evolution of greater loss neutrality, remains unaltered under a more general specification of the utility function. In addition, the secondary prediction that a greater degree of spatial correlation in adverse shocks favors greater loss aversion is preserved.*

Proof. Consider the utility function

$$u_{it} = u_i(c_{it}, n_{it}) = \begin{cases} [f(c_{it}) - f(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} \geq \tilde{c}, \\ \frac{1}{\theta_i}[g(c_{it}) - g(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} < \tilde{c}, \end{cases}$$

where $f(\cdot)$, $g(\cdot)$, and $\eta(\cdot)$ are continuously differentiable and strictly monotonically increasing functions, and $\eta(0) > -\infty$. No restrictions on the curvature of these functions are required. In particular, $g(\cdot)$ may be concave rather than convex. As in the baseline model, utility from consumption is normalized to zero at the reference point, $\tilde{c}$.

To preserve the Malthusian allocation structure of the baseline model, suppose that when income exceeds $\tilde{c}$, an individual consumes at the subsistence level and allocates the remaining income to child rearing, whereas when income falls below $\tilde{c}$, it is fully consumed. Following the production structure of the baseline model, expected utility from the risky production mode is

$$u^r_{it} = p\eta\left(\frac{y^h - \tilde{c}}{\rho y^h}\right) + (1-p)\left[\frac{1}{\theta_i}\Big(g(y^l(\Delta)) - g(\tilde{c})\Big) + \eta(0)\right] \equiv u^r(\theta_i; \Delta, p),$$

whereas utility from the safe production mode is

$$u^s_{it} = \eta(1) \equiv u^s.$$

Since $g(y^l(\Delta)) - g(\tilde{c}) < 0$, it follows that $\lim_{\theta_i \to 0} u^r(\theta_i; \Delta, p) = -\infty$ and $\partial u^r(\theta_i; \Delta, p)/\partial \theta_i > 0$ for all $\theta_i \in (0, 1]$.

Assuming that the risky production mode is sufficiently attractive that a loss-neutral individual prefers it, $u^r(1; \Delta, p) > u^s$, there exists a unique cutoff $\hat{\theta}(\Delta, p) \in (0, 1)$ such that $u^r(\hat{\theta}(\Delta, p); \Delta, p) = u^s$. In particular,

$$\hat{\theta}(\Delta, p) = \frac{(1-p)\left[g(\tilde{c}) - g(y^l(\Delta))\right]}{p\eta\left(\frac{y^h - \tilde{c}}{\rho y^h}\right) + (1-p)\eta(0) - \eta(1)}.$$

The denominator is positive under $u^r(1; \Delta, p) > u^s$. Since $g'(\cdot) > 0$ and $y^{l\prime}(\Delta) < 0$,

$$\frac{\partial \hat{\theta}(\Delta, p)}{\partial \Delta} = -\frac{(1-p)g'(y^l(\Delta))y^{l\prime}(\Delta)}{p\eta\left(\frac{y^h - \tilde{c}}{\rho y^h}\right) + (1-p)\eta(0) - \eta(1)} > 0.$$

Moreover,

$$\frac{\partial \hat{\theta}(\Delta, p)}{\partial (1-p)} = \frac{\left[g(\tilde{c}) - g(y^l(\Delta))\right]\left[\eta\left(\frac{y^h - \tilde{c}}{\rho y^h}\right) - \eta(1)\right]}{\left[p\eta\left(\frac{y^h-\tilde{c}}{\rho y^h}\right) + (1-p)\eta(0) - \eta(1)\right]^2} > 0.$$

Thus, greater frequency or severity of an adverse shock raises the cutoff degree of loss neutrality, preserving the main prediction of the baseline model. Moreover, conditional on the choice of production mode, fertility under each climatic realization is identical to that in the baseline model. The population dynamics associated with the frequency and spatial correlation of adverse shocks are therefore unchanged. Hence, for $\mu \neq \hat{\mu}(p)$, greater adverse-shock frequency or severity increases long-run loss neutrality, while the secondary prediction that a greater degree of spatial correlation in adverse shocks favors greater loss aversion is preserved.

B.1.4 The Emergence of Convexity with Respect to Losses

This subsection demonstrates that, under some conditions, convexity of utility with respect to losses can emerge as an evolutionarily dominant feature of preferences. In particular, when greater risk-taking below the reference point creates a possibility of crossing the subsistence threshold and reproducing, individuals whose utility is sufficiently convex with respect to losses may choose a more dispersed production mode, whereas individuals whose utility is concave with respect to losses choose a less dispersed mode.

Proposition A4. *Suppose that individuals can choose between two production modes with the same expected yield, where the more dispersed mode permits consumption to exceed the subsistence level with positive probability, whereas the less dispersed mode does not. There exist parameter configurations under which individuals whose utility is concave with respect to losses choose the less dispersed production mode, while individuals whose utility is sufficiently convex with respect to losses choose the more dispersed production mode. If the reproductive return associated with the more dispersed mode is sufficient for persistence, convexity with respect to losses is evolutionarily favored and can emerge as an evolutionarily dominant feature of preferences.*

Proof. Suppose that individuals choose between two risky production modes with identical expected yields, y^0, but different degrees of dispersion. The less dispersed production mode generates

$$y^1_{it} = \begin{cases} y^0 + \varepsilon_1 & \text{with probability } 1/2, \\ y^0 - \varepsilon_1 & \text{with probability } 1/2, \end{cases}$$

whereas the more dispersed production mode generates

$$y^2_{it} = \begin{cases} y^0 + \varepsilon_2 & \text{with probability } 1/2 \\ y^0 - \varepsilon_2 & \text{with probability } 1/2, \end{cases}$$

where $0 < \varepsilon_1 < \varepsilon_2$ and $y^0 - \varepsilon_2 > 0, \quad y^0 + \varepsilon_1 < \tilde{c} < y^0 + \varepsilon_2 \leq y^*$.

Thus, both realizations of the less dispersed mode lie below subsistence, whereas the favorable realization of the more dispersed mode lies above subsistence. The two modes nevertheless generate the same expected yield.

For any given degree of loss aversion, $1/\theta_i$, consider two types of individuals. Individuals who are risk averse with respect to losses have utility

$$u_{it}^{RA} = \begin{cases} [f(c_{it}) - f(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} \geq \tilde{c}, \\ \frac{1}{\theta_i}[g^{RA}(c_{it}) - g^{RA}(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} < \tilde{c}, \end{cases}$$

where $g^{RA}(\cdot)$ is strictly monotonically increasing and concave. Individuals who are risk-loving with respect to losses have utility

$$u_{it}^{RL} = \begin{cases} [f(c_{it}) - f(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} \geq \tilde{c}, \\ \frac{1}{\theta_i}[g^{RL}(c_{it}) - g^{RL}(\tilde{c})] + \eta(n_{it}) & \text{if } c_{it} < \tilde{c}, \end{cases}$$

where $g^{RL}(\cdot)$ is strictly monotonically increasing and convex.

Under the less dispersed production mode, both realizations lie below subsistence and therefore fertility is zero. Under the more dispersed mode, the favorable realization exceeds subsistence and generates fertility

$$n_h^{(2)} \equiv \frac{y^0 + \varepsilon_2 - \tilde{c}}{\rho(y^0 + \varepsilon_2)} > 0,$$

whereas its adverse realization generates zero fertility.

For either type $j \in \{RA, RL\}$, the difference between expected utility from the more dispersed and the less dispersed production modes is therefore

$$E[u_{it}^j(y_{it}^2)] - E[u_{it}^j(y_{it}^1)] = \frac{1}{2}\left[\eta(n_h^{(2)}) - \eta(0)\right] + \frac{1}{2\theta_i}\Phi^j,$$

where $\Phi^j \equiv g^j(\tilde{c}) - g^j(y^0+\varepsilon_1) - g^j(y^0-\varepsilon_1) + g^j(y^0-\varepsilon_2)$. Equivalently, $\Phi^j = \left[g^j(\tilde{c}) - g^j(y^0+\varepsilon_1)\right] - \left[g^j(y^0-\varepsilon_1) - g^j(y^0-\varepsilon_2)\right]$. Since $\tilde{c} - (y^0+\varepsilon_1) < (y^0-\varepsilon_1) - (y^0-\varepsilon_2) = \varepsilon_2 - \varepsilon_1$, and the first interval lies entirely to the right of the second, concavity of $g^{RA}(\cdot)$ implies

$$\Phi^{RA} < 0.$$

Thus, for parameter configurations in which the fertility gain from the more dispersed mode is not sufficiently large to compensate for this utility loss,

$$E[u_{it}^{RA}(y_{it}^2)] - E[u_{it}^{RA}(y_{it}^1)] < 0,$$

and individuals whose utility is concave with respect to losses choose the less dispersed production mode.

By contrast, because the first increment in Φ^{RL} occurs at a higher consumption level than the second, sufficiently strong convexity of $g^{RL}(\cdot)$ can imply $\Phi^{RL} > 0$.
For example, if $g^{RL}(c) = e^{kc}$, then $\Phi^{RL} > 0$ for all sufficiently large k. Hence, for sufficiently convex utility with respect to losses,

$$E[u_{it}^{RL}(y_{it}^2)] - E[u_{it}^{RL}(y_{it}^1)] > 0,$$

and these individuals choose the more dispersed production mode.

The evolutionary implication follows from the subsistence constraint. Individuals choosing the less dispersed mode remain below subsistence under either realization and therefore do not reproduce. In contrast, individuals choosing the more dispersed mode cross the subsistence threshold with probability 1/2 and reproduce in the favorable realization. If these realizations are idiosyncratic across dynasties and $(1/2)n_h^{(2)} > 1$, the population choosing the more dispersed mode grows, whereas the population choosing the less dispersed mode vanishes. Thus, individuals whose utility is sufficiently convex with respect to losses increase their representation in the population, and convexity with respect to losses can emerge as an evolutionarily dominant feature of preferences.

B.1.5 Evolution of the Reference Point

This subsection examines the robustness of the baseline predictions to a permanent change in the subsistence consumption level when the reference point adjusts to the contemporary level of subsistence. It also clarifies the condition under which such reference-point adjustment is favored by evolutionary selection.

Suppose that the subsistence consumption level changes permanently from $\tilde{c}_0$ to $\tilde{c}_1$. To isolate the implications of reference-point adjustment from changes in the Malthusian benchmark, maintain the normalization that the safe mode generates replacement fertility at the contemporary subsistence level, so that $\bar{y}(\tilde{c}_1) = \tilde{c}_1/(1-\rho)$. Assume also that

$$y^l(\Delta) < \tilde{c}_1 < \bar{y}(\tilde{c}_1) < y^h \leq y^*(\tilde{c}_1),$$

where $y^*(\tilde{c}_1)$ denotes the critical income level associated with the contemporary subsistence level, and that

$$p\frac{y^h - \tilde{c}_1}{\rho y^h} > 1.$$

Proposition A5. *Suppose that, following a permanent change in the subsistence consumption level, the reference point adjusts to the contemporary subsistence level. Then the main prediction of the baseline model—that greater adverse-shock frequency or severity favors the evolution of greater loss neutrality—remains unaltered. In addition, the complementary prediction that a greater degree of spatial correlation in adverse shocks favors greater loss aversion is preserved.*

Moreover, suppose that the rule governing reference-point adjustment is transmitted intergenerationally. If individuals whose reference point adjusts choose the evolutionarily dominant production mode following the change in subsistence, whereas otherwise identical individuals with a fixed reference point choose the evolutionarily dominated production mode, reference-point adjustment is evolutionarily favored.

Proof. Consider first individuals whose reference point adjusts to the contemporary subsistence level, $\tilde{c}_1$. Fertility under the favorable realization of the risky mode is

$$n^h(\tilde{c}_1) \equiv \frac{y^h - \tilde{c}_1}{\rho y^h} > 1,$$

while fertility following an adverse realization is zero. Expected utility from the risky mode is therefore

$$u^r(\theta_i; \Delta, p, \tilde{c}_1) = p\eta\left(n^h(\tilde{c}_1)\right) + (1-p)\left[\frac{1}{\theta_i}\left(g(y^l(\Delta)) - g(\tilde{c}_1)\right) + \eta(0)\right],$$

whereas utility from the safe mode is $u^s = \eta(1)$.

Assuming, as in the baseline model, that the least loss-averse individual prefers the risky mode, there exists a unique interior cutoff $\hat{\theta}(\Delta, p; \tilde{c}_1)$ satisfying $u^r(\hat{\theta}; \Delta, p, \tilde{c}_1) = u^s$. In particular,

$$\hat{\theta}(\Delta, p; \tilde{c}_1) = \frac{(1-p)\left[g(\tilde{c}_1) - g(y^l(\Delta))\right]}{p\eta\left(n^h(\tilde{c}_1)\right) + (1-p)\eta(0) - \eta(1)}.$$

Since $y^{l\prime}(\Delta) < 0$,

$$\frac{\partial\hat{\theta}(\Delta, p; \tilde{c}_1)}{\partial\Delta} = -\frac{(1-p)g'(y^l(\Delta))y^{l\prime}(\Delta)}{p\eta\left(n^h(\tilde{c}_1)\right) + (1-p)\eta(0) - \eta(1)} > 0.$$

Moreover,

$$\frac{\partial\hat{\theta}(\Delta, p; \tilde{c}_1)}{\partial(1-p)} = \frac{\left[g(\tilde{c}_1) - g(y^l(\Delta))\right]\left[\eta\left(n^h(\tilde{c}_1)\right) - \eta(1)\right]}{\left[p\eta\left(n^h(\tilde{c}_1)\right) + (1-p)\eta(0) - \eta(1)\right]^2} > 0.$$

Thus, following adjustment of the reference point to the contemporary subsistence level, greater adverse-shock frequency or severity continues to raise the sorting cutoff.

The population dynamics are also analogous to those in the baseline model. The asymptotic logarithmic population growth rate of individuals choosing the risky mode is

$$\bar{n}^B(\mu, p; \tilde{c}_1) = p\ln[\mu + (1-\mu)p] + (1-p)\ln[(1-\mu)p] + \ln n^h(\tilde{c}_1).$$

As in the proof of Lemma 1,

$$\frac{\partial\bar{n}^B(\mu, p; \tilde{c}_1)}{\partial\mu} < 0,$$

while $\bar{n}^B(0, p; \tilde{c}_1) = \ln\left[pn^h(\tilde{c}_1)\right] > 0$ and $\lim_{\mu\to 1}\bar{n}^B(\mu, p; \tilde{c}_1) = -\infty$.

Hence there exists a unique threshold $\hat{\mu}(p; \tilde{c}_1) \in (0, 1)$ such that the risky group dominates in the long run when $\mu < \hat{\mu}(p; \tilde{c}_1)$ and the safe group dominates when $\mu > \hat{\mu}(p; \tilde{c}_1)$. Consequently,

$$\bar{\theta} = \begin{cases} \dfrac{1 + \hat{\theta}(\Delta, p; \tilde{c}_1)}{2}, & \mu < \hat{\mu}(p; \tilde{c}_1), \\ \dfrac{\hat{\theta}(\Delta, p; \tilde{c}_1)}{2}, & \mu > \hat{\mu}(p; \tilde{c}_1). \end{cases}$$

Thus, for $\mu \neq \hat{\mu}(p; \tilde{c}_1)$, greater adverse-shock frequency or severity increases long-run loss neutrality regardless of which production mode is evolutionarily dominant, whereas greater spatial correlation shifts evolutionary dominance toward the relatively more loss-averse group.

Finally, consider individuals who differ only in whether their reference point adjusts following the permanent change in subsistence. If adjustment induces one type to choose a production mode with asymptotic logarithmic population growth rate $\bar{n}^*$, while individuals retaining the fixed reference point choose the alternative mode with growth rate $\bar{n}^o < \bar{n}^*$, then the population ratio of the two types satisfies

$$\frac{N_t^*}{N_t^o} \longrightarrow \infty.$$

Reference-point adjustment is therefore evolutionarily favored whenever it induces individuals to choose the evolutionarily dominant production mode. □

B.2 Robustness to Partial Risk Sharing

This subsection examines the robustness of the central sorting mechanism and the comparative statics with respect to adverse-shock frequency and severity to partial risk sharing. It allows individuals who choose the risky mode of production to partially insure climatic risk through transfers from individuals experiencing different climatic realizations, and permits the scope for such risk sharing to vary with the degree of spatial correlation in climatic shocks. Risk sharing can induce a larger share of relatively loss-averse individuals to choose the risky mode of production. Nevertheless, as long as insurance remains partial, a unique interior sorting cutoff continues to exist and greater adverse-shock frequency or severity raises this cutoff. Moreover, under sufficiently limited community-level risk sharing, the long-run population dynamics preserve the main prediction of the baseline model: Greater adverse-shock frequency or severity increases the long-run degree of loss neutrality. Thus, neither the central sorting mechanism nor its main evolutionary implication relies on the complete absence of risk sharing.

Let the expected yield generated by the risky mode of production be

$$y^e(\Delta, p) \equiv py^h + (1-p)y_l(\Delta).$$

As assumed in the baseline model, $y^e(\Delta, p) > \bar{y} > \tilde{c}$.

Let $\tau \in [0, 1]$ denote the degree of risk sharing among individuals who choose the risky mode of production. Risk sharing generates a mean-preserving contraction of the distribution of risky yields. Thus, the post-transfer yields associated with the favorable and adverse climatic realizations are, respectively,

$$y^{h,\tau}(\Delta, p) = (1-\tau)y^h + \tau y^e(\Delta, p),$$

and

$$y^{l,\tau}(\Delta, p) = (1-\tau)y_l(\Delta) + \tau y^e(\Delta, p).$$

Hence,

$$py^{h,\tau}(\Delta, p) + (1-p)y^{l,\tau}(\Delta, p) = y^e(\Delta, p),$$

such that risk sharing does not alter the expected resources generated by the risky mode. When $\tau = 0$, there is no risk sharing and the baseline model is obtained. When $\tau = 1$, risk is completely pooled and both realizations generate the expected risky yield, $y^e(\Delta, p)$.

Since $y_l(\Delta) < \tilde{c} < y^e(\Delta, p)$, there exists a unique degree of risk sharing,

$$\tau^*(\Delta, p) \equiv \frac{\tilde{c} - y_l(\Delta)}{y^e(\Delta, p) - y_l(\Delta)} = \frac{\tilde{c} - y_l(\Delta)}{p[y^h - y_l(\Delta)]} \in (0, 1),$$

such that

$$y^{l,\tau}(\Delta, p) \begin{cases} < \tilde{c}, & \tau < \tau^*(\Delta, p), \\ = \tilde{c}, & \tau = \tau^*(\Delta, p), \\ > \tilde{c}, & \tau > \tau^*(\Delta, p). \end{cases}$$

Thus, $\tau^*(\Delta, p)$ is the critical degree of risk sharing at which the adverse realization ceases to constitute a loss relative to the subsistence reference point. Moreover,

$$\frac{\partial \tau^*(\Delta, p)}{\partial \Delta} = \frac{[-y_l'(\Delta)]p(y^h - \tilde{c})}{\{p[y^h - y_l(\Delta)]\}^2} > 0,$$

and

$$\frac{\partial \tau^*(\Delta, p)}{\partial(1-p)} = \frac{\tilde{c} - y_l(\Delta)}{p^2[y^h - y_l(\Delta)]} > 0.$$

Consequently, greater adverse-shock frequency or severity raises the degree of risk sharing required to eliminate the adverse-state loss.

For every $\tau < \tau^*(\Delta, p)$, the adverse realization remains below subsistence. Moreover, since both y^h and $y^e(\Delta, p)$ exceed $\bar{y}$,

$$y^{h,\tau}(\Delta, p) > \bar{y} > \tilde{c}.$$

Hence, conditional on a favorable realization, fertility is

$$n^{h,\tau}(\Delta, p) \equiv \frac{y^{h,\tau}(\Delta, p) - \tilde{c}}{\rho y^{h,\tau}(\Delta, p)} > 1,$$

whereas conditional on an adverse realization fertility is zero.

The expected utility associated with the risky mode, for a given degree of risk sharing $\tau < \tau^*(\Delta, p)$, is therefore

$$u^r(\theta_i; \Delta, p, \tau) = p\eta\left(n^{h,\tau}(\Delta, p)\right) + (1-p)\left[\frac{1}{\theta_i}\left(g(y^{l,\tau}(\Delta, p)) - g(\tilde{c})\right) + \eta(0)\right].$$

The utility associated with the safe mode remains $u^s = \eta(1)$,
where, since $y^{l,\tau}(\Delta, p) < \tilde{c}$,

$$\frac{\partial u^r(\theta_i; \Delta, p, \tau)}{\partial \theta_i} > 0, \qquad \lim_{\theta_i \to 0} u^r(\theta_i; \Delta, p, \tau) = -\infty.$$

Moreover, holding Δ and τ constant,

$$\begin{aligned}\frac{\partial u^r(\theta_i; \Delta, p, \tau)}{\partial(1-p)} = & \left[\frac{1}{\theta_i}\left(g(y^{l,\tau}(\Delta, p)) - g(\tilde{c})\right) + \eta(0)\right] - \eta\left(n^{h,\tau}(\Delta, p)\right) \\ & - \tau[y^h - y_l(\Delta)]\left[p\eta'\left(n^{h,\tau}(\Delta, p)\right)\frac{\tilde{c}}{\rho[y^{h,\tau}(\Delta, p)]^2} + \frac{1-p}{\theta_i}g'\left(y^{l,\tau}(\Delta, p)\right)\right] < 0.\end{aligned}$$

The inequality follows since $y^{l,\tau}(\Delta, p) < \tilde{c}$, $n^{h,\tau}(\Delta, p) > 1$, and $g(\cdot)$ and $\eta(\cdot)$ are strictly increasing. Thus, a higher frequency of adverse shocks lowers expected utility from risky production at every fixed degree of partial risk sharing.

Risk sharing opportunities need not be identical across environments. Let $\bar{\tau}(\mu) \in [0, 1]$ denote the maximum feasible degree of risk sharing, where

$$\bar{\tau}'(\mu) \leq 0,$$

reflecting the possibility that lower spatial correlation in climatic shocks facilitates transfers across individuals experiencing different climatic realizations. Individuals choosing the risky mode may select any degree of risk sharing $\tau \in [0, \bar{\tau}(\mu)]$. Their expected utility from risky production is therefore

$$U^r(\theta_i; \Delta, p, \bar{\tau}(\mu)) \equiv \max_{\tau \in [0, \bar{\tau}(\mu)]} u^r(\theta_i; \Delta, p, \tau).$$

Provided that $\bar{\tau}(\mu) < \tau^*(\Delta, p)$, every feasible degree of risk sharing leaves the adverse realization below subsistence. Hence, $U^r(\theta_i; \Delta, p, \bar{\tau}(\mu))$ is continuous and strictly increasing in θ_i, while $\lim_{\theta_i \to 0} U^r(\theta_i; \Delta, p, \bar{\tau}(\mu)) = -\infty$. Since $\tau = 0$ remains feasible and the least loss-averse individual prefers the risky mode in the baseline model, there exists a unique interior cutoff $\hat{\theta}(\Delta, p, \bar{\tau}(\mu))$. Moreover, if $\Delta_H > \Delta_L$, greater adverse-shock frequency or severity lowers utility from risky production at every feasible degree of risk sharing, implying $U^r(\theta_i; \Delta_H, p, \bar{\tau}(\mu)) < U^r(\theta_i; \Delta_L, p, \bar{\tau}(\mu))$ for every θ_i. Likewise, greater adverse-shock frequency lowers utility from risky production at every feasible degree of risk sharing. It follows that the sorting cutoff is strictly increasing in adverse-shock frequency and severity; wherever the cutoff is differentiable, $\partial\hat{\theta}(\Delta, p, \bar{\tau}(\mu))/\partial\Delta > 0$ and $\partial\hat{\theta}(\Delta, p, \bar{\tau}(\mu))/\partial(1-p) > 0$. These observations yield the following proposition:

Proposition A6 (Risk Sharing and Sorting)**.** *Suppose that the maximum feasible degree of risk sharing satisfies*

$$\bar{\tau}(\mu) < \tau^*(\Delta, p).$$

Then:

1. *There exists a unique interior cutoff* $\hat{\theta}(\Delta, p, \bar{\tau}(\mu)) \in (0, 1)$ *such that individuals with* $\theta_i \leq \hat{\theta}(\Delta, p, \bar{\tau}(\mu))$ *choose the safe mode of production, whereas individuals with* $\theta_i > \hat{\theta}(\Delta, p, \bar{\tau}(\mu))$ *choose the risky mode.*

2. *Greater adverse-shock frequency or severity raises the cutoff:*
$$\frac{\partial\hat{\theta}(\Delta, p, \bar{\tau}(\mu))}{\partial\Delta} > 0, \qquad \frac{\partial\hat{\theta}(\Delta, p, \bar{\tau}(\mu))}{\partial(1-p)} > 0.$$

3. *An expansion in feasible risk sharing weakly reduces the cutoff.*

The proposition establishes that the central sorting mechanism and the adverse-shock-frequency and severity comparative statics of the sorting cutoff are preserved under partial risk sharing. In particular, as long as insurance remains partial such that an adverse realization continues to generate a loss relative to subsistence, there exists a unique interior sorting cutoff, and greater adverse-shock frequency or severity raises this cutoff. Thus, the baseline result does not require the absence of risk sharing, but only that risk sharing was insufficient to fully insure individuals against subsistence-threatening losses, a condition that plausibly characterized most of the Malthusian era and was especially relevant during the formative stages of human evolution.

Long-Run Evolutionary Implications. The preceding analysis permits individuals to choose among feasible risk-sharing arrangements and therefore establishes the robustness of the sorting mechanism under a general form of partial risk sharing. To characterize the corresponding long-run evolutionary dynamics, consider, as a simplifying assumption, the case in which risk sharing is organized at the community level, consistent with the historically prevalent organization of risk sharing through community-based arrangements rather than individualized state-contingent contracts. Thus, all individuals who choose the risky mode of production are subject to a common degree of risk sharing, $\tau^c(\mu) \in [0, \bar{\tau}(\mu)]$. Suppose that

$$\tau^c(\mu) < \tau^*(\Delta, p)$$

for every admissible degree of adverse-shock frequency and severity. Thus, the adverse realization remains below subsistence and generates zero fertility.

For a given μ, expected utility from risky production under the common risk-sharing arrangement is

$$u^r(\theta_i; \Delta, p, \tau^c(\mu)).$$

As established above, this utility is strictly increasing in θ_i and strictly decreasing in both Δ and $1-p$. Hence, maintaining the assumption that the least loss-averse type prefers risky production, there exists a unique interior cutoff,

$$\hat{\theta}^c(\Delta, p, \mu) \in (0, 1),$$

such that individuals with $\theta_i \leq \hat{\theta}^c(\Delta, p, \mu)$ choose the safe mode, whereas individuals with $\theta_i > \hat{\theta}^c(\Delta, p, \mu)$ choose the risky mode. Moreover,

$$\frac{\partial \hat{\theta}^c(\Delta, p, \mu)}{\partial \Delta} > 0, \qquad \frac{\partial \hat{\theta}^c(\Delta, p, \mu)}{\partial (1-p)} > 0.$$

Since all individuals who choose the risky mode are subject to the same risk-sharing arrangement, conditional on the choice of production mode the degree of loss neutrality has no further effect on fertility. Consequently, exactly as in the baseline model, the distribution of θ_i within each production group remains unchanged over time. The average degree of loss neutrality is therefore $\hat{\theta}^c(\Delta, p, \mu)/2$ within the safe group and $[1 + \hat{\theta}^c(\Delta, p, \mu)]/2$ within the risky group.

Let

$$n^{h,c}(\Delta, p, \mu) \equiv \frac{y^{h,\tau^c(\mu)}(\Delta, p) - \tilde{c}}{\rho y^{h,\tau^c(\mu)}(\Delta, p)} > 1$$

denote fertility under the ordinary climatic realization among individuals who choose the risky mode. Under idiosyncratic climatic shocks, average fertility in this group is therefore

$$\bar{n}^{id,c} = p n^{h,c}(\Delta, p, \mu),$$

whereas under aggregate climatic shocks,

$$\bar{n}_t^{ag,c} = \begin{cases} n^{h,c}(\Delta, p, \mu) & \text{with probability } p \\ 0 & \text{with probability } 1-p. \end{cases}$$

Hence, following the proof of Lemma 1, the asymptotic logarithmic population growth rate of the risky group is

$$\bar{n}^{B,c}(\Delta, p, \mu) = p \ln[\mu + (1-\mu)p] + (1-p) \ln[(1-\mu)p] + \ln n^{h,c}(\Delta, p, \mu),$$

whereas the asymptotic population growth rate of the safe group remains equal to zero.

Consider first $\mu > \hat{\mu}(p)$, where $\hat{\mu}(p)$ is the threshold degree of spatial correlation in the baseline model. Since risk sharing lowers the yield under the ordinary realization,

$$y^{h,\tau^c(\mu)}(\Delta, p) \leq y^h,$$

and therefore

$$n^{h,c}(\Delta, p, \mu) \leq n^h.$$

It follows that

$$\bar{n}^{B,c}(\Delta, p, \mu) \leq \bar{n}^B(\mu, p) < 0.$$

Thus, whenever the safe group dominates in the baseline model, it continues to dominate under partial risk sharing.

Consider next $\mu < \hat{\mu}(p)$. In the absence of risk sharing, $\bar{n}^B(\mu, p) > 0$. Moreover,

$$y^{h,\tau}(\Delta, p) = y^h - \tau(1-p)\big[y^h - y^l(\Delta)\big].$$

Since $y^l(\Delta) \geq 0$,

$$y^{h,\tau}(\Delta, p) \geq y^h[1 - \tau(1-p)].$$

Hence, uniformly over all admissible values of Δ,

$$n^{h,\tau}(\Delta, p) \geq n\Big(y^h[1-\tau(1-p)]\Big),$$

and

$$\lim_{\tau \to 0} n\Big(y^h[1-\tau(1-p)]\Big) = n^h.$$

Since $\bar{n}^B(\mu, p) > 0$, continuity implies that there exists $\tau_0(\mu, p) > 0$ such that, for every $\tau^c(\mu) < \tau_0(\mu, p)$,

$$\bar{n}^{B,c}(\Delta, p, \mu) > 0$$

for every admissible Δ. Thus, if climatic shocks are sufficiently idiosyncratic for the risky group to dominate in the baseline model, that group continues to dominate when risk sharing is sufficiently limited.

Consequently, for every $\mu \neq \hat{\mu}(p)$, sufficiently limited partial risk sharing leaves the identity of the evolutionarily dominant group unchanged. The long-run average degree of loss neutrality is therefore

$$\bar{\theta}^c(\Delta, p, \mu) = \begin{cases} \dfrac{1 + \hat{\theta}^c(\Delta, p, \mu)}{2}, & \mu < \hat{\mu}(p), \\[2ex] \dfrac{\hat{\theta}^c(\Delta, p, \mu)}{2}, & \mu > \hat{\mu}(p). \end{cases}$$

It follows that

$$\frac{\partial \bar{\theta}^c(\Delta, p, \mu)}{\partial \Delta} = \frac{1}{2}\frac{\partial \hat{\theta}^c(\Delta, p, \mu)}{\partial \Delta} > 0, \qquad \frac{\partial \bar{\theta}^c(\Delta, p, \mu)}{\partial (1-p)} = \frac{1}{2}\frac{\partial \hat{\theta}^c(\Delta, p, \mu)}{\partial (1-p)} > 0, \qquad \mu \neq \hat{\mu}(p).$$

Thus, the main prediction of the baseline model is preserved under sufficiently limited partial risk sharing: greater adverse-shock frequency or severity increases the long-run degree of loss neutrality regardless of which production group is evolutionarily dominant.

B.3 Climate and the Evolution of Risk Aversion rather than Loss Aversion

This subsection establishes that the climatic predictions of the baseline model do not follow systematically from conventional risk aversion alone. Whereas loss aversion generates sign-definite predictions regarding adverse-shock frequency, severity, and spatial correlation, conventional risk aversion does not. In particular, within a standard smooth expected-utility framework with constant relative risk aversion, feasible parameter configurations can generate opposite patterns of sorting across production modes and, consequently, opposite relationships between climatic characteristics and the long-run degree of risk aversion.

The purpose of the analysis is therefore not to establish that conventional risk aversion can never generate the same directional relationships as loss aversion. Rather, it establishes that these relationships are not implied by conventional risk aversion itself: depending on the preference configuration, greater conventional risk aversion may either increase or decrease the relative attractiveness of the risky production mode. This contrasts with the loss aversion mechanism of the baseline model, which delivers the climatic comparative statics unambiguously.

Suppose that the subsistence consumption constraint is present and binding, but individuals are loss neutral. Thus, subsistence continues to govern fertility and reproductive success, but it is not a reference point in preferences and there is no kink in utility at $\tilde{c}$.

Suppose that individuals differ instead in their degree of constant relative risk aversion with respect to consumption. In particular, let

$$v(c;\sigma) = \begin{cases} \dfrac{c^{1-\sigma}-1}{1-\sigma}, & \sigma \neq 1, \\ \ln c, & \sigma = 1, \end{cases}$$

and suppose that the utility of an adult i in period t is

$$u_{it} = (1-\gamma)v(c_{it};\sigma_i) + \gamma \ln(n_{it}+\epsilon),$$

where $\sigma_i \in [0,\infty)$ is the coefficient of relative risk aversion, $\gamma \in (0,1)$, and $\epsilon > 0$.

Proposition A7. *Conventional risk aversion does not generate sign-definite predictions regarding the effects of adverse-shock frequency, severity, and spatial correlation on the evolution of attitudes toward risk. There exist feasible parameter configurations with an interior sorting cutoff under which greater adverse-shock frequency or severity reduces the long-run degree of conventional risk aversion and a shift toward greater spatial correlation favors greater risk aversion. There also exist feasible parameter configurations under which these relationships are reversed. Hence, unlike loss aversion in the baseline model, conventional risk aversion does not systematically imply any of these climatic comparative statics.*

Proof. Individuals choose their production mode prior to the realization of climatic conditions so as to maximize expected utility. To isolate the implications of conventional risk aversion from changes in the consumption–fertility allocation, retain the production, budget, and subsistence structure of the baseline model and restrict attention to parameter configurations for which the allocation remains optimal.

For completeness, the maintained allocation can be supported by primitive parameter restrictions. Since $\tilde{c} \geq 1$ in both examples below, a sufficient condition for $c = \tilde{c}$ to be optimal whenever income exceeds subsistence is

$$\frac{\gamma}{1-\gamma} \geq y^h - \tilde{c} + \rho y^h \epsilon.$$

Indeed, under this condition,

$$(1-\gamma)\tilde{c}^{-\sigma_i} - \frac{\gamma}{y - \tilde{c} + \rho y \epsilon} \leq 0$$

for every $\sigma_i \geq 0$ and every $y \leq y^h$. Thus the baseline consumption–fertility allocation is optimal for the parameter configurations considered below.

As in the baseline model, the ordinary climatic realization of the risky mode generates consumption $\tilde{c}$ and fertility

$$n^h \equiv \frac{y^h - \tilde{c}}{\rho y^h} > 1,$$

whereas the adverse realization generates consumption $y^l(\Delta) < \tilde{c}$ and zero fertility. The safe production mode generates $(\tilde{c}, 1)$.

Expected utility from the risky production mode is therefore

$$u^r(\sigma_i, \Delta, p) = p\left[(1-\gamma)v(\tilde{c};\sigma_i) + \gamma\ln(n^h+\epsilon)\right] + (1-p)\left[(1-\gamma)v(y^l(\Delta);\sigma_i) + \gamma\ln\epsilon\right],$$

whereas utility from the safe mode is

$$u^s(\sigma_i) = (1-\gamma)v(\tilde{c};\sigma_i) + \gamma\ln(1+\epsilon).$$

Let $D(\sigma_i, \Delta, p) \equiv u^r(\sigma_i, \Delta, p) - u^s(\sigma_i)$. Then

$$D(\sigma_i, \Delta, p) = (1-p)(1-\gamma)\left[v(y^l(\Delta);\sigma_i) - v(\tilde{c};\sigma_i)\right] + A(p),$$

where $A(p) \equiv \gamma\left[p\ln(n^h+\epsilon) + (1-p)\ln\epsilon - \ln(1+\epsilon)\right]$ is independent of σ_i and Δ. An individual chooses the risky production mode if $D(\sigma_i, \Delta, p) > 0$ and the safe mode otherwise.

The effect of conventional risk aversion on the relative attractiveness of the risky mode need not have a unique sign. This can be established by two feasible parameter configurations.

Example 1. Suppose that $\tilde{c} = 1$ and $0 < y^l(\Delta) < 1$. For $\sigma_i \neq 1$,

$$\frac{\partial D(\sigma_i, \Delta, p)}{\partial \sigma_i} = (1-p)(1-\gamma)\frac{x - x\ln x - 1}{(1-\sigma_i)^2}, \qquad x \equiv [y^l(\Delta)]^{1-\sigma_i}.$$

Since $x - x\ln x - 1 < 0$ for $x \neq 1$, and the corresponding limit at $\sigma_i = 1$ is also strictly negative,

$$\frac{\partial D(\sigma_i, \Delta, p)}{\partial \sigma_i} < 0.$$

Thus, greater conventional risk aversion reduces the relative attractiveness of the risky production mode. If an interior cutoff $\hat{\sigma}$ exists, individuals with $\sigma_i < \hat{\sigma}$ choose the risky mode, whereas individuals with $\sigma_i > \hat{\sigma}$ choose the safe mode.

Example 2. Suppose instead that $\tilde{c} > 1$ and, at the climatic environment under consideration, $y^l(\Delta) = 1$. For $\sigma_i \neq 1$,

$$\frac{\partial D(\sigma_i, \Delta, p)}{\partial \sigma_i} = (1-p)(1-\gamma)\frac{-x + x\ln x + 1}{(1-\sigma_i)^2}, \qquad x \equiv \tilde{c}^{\,1-\sigma_i}.$$

Since $-x + x\ln x + 1 > 0$ for $x \neq 1$, and the corresponding limit at $\sigma_i = 1$ is also strictly positive,

$$\frac{\partial D(\sigma_i, \Delta, p)}{\partial \sigma_i} > 0.$$

Thus, at this feasible configuration, greater conventional risk aversion increases the relative attractiveness of the risky production mode. By continuity, the same ordering holds for climatic environments in a neighborhood of this configuration. If an interior cutoff $\hat{\sigma}$ exists, individuals with $\sigma_i < \hat{\sigma}$ choose the safe mode, whereas individuals with $\sigma_i > \hat{\sigma}$ choose the risky mode.

The two examples therefore establish that conventional risk aversion does not impose a unique pattern of sorting across production modes.

It remains to verify that the two configurations can be associated with an interior sorting cutoff while maintaining the consumption–fertility allocation imposed above. Let

$$B(\epsilon) \equiv p \ln(n^h + \epsilon) + (1-p)\ln \epsilon - \ln(1+\epsilon),$$

so that $A(p) = \gamma B(\epsilon)$. Since $pn^h > 1$,

$$\lim_{\epsilon \to \infty} \epsilon B(\epsilon) = pn^h - 1 > 0,$$

and therefore $B(\epsilon) > 0$ for all sufficiently large ϵ. Let $r \equiv \frac{\gamma}{1-\gamma}$. The sufficient condition supporting the maintained allocation requires

$$r \geq K(\epsilon) \equiv y^h - \tilde{c} + \rho y^h \epsilon.$$

For any sufficiently large ϵ, however, r remains otherwise unbounded above, and

$$\frac{A(p)}{(1-p)(1-\gamma)} = \frac{rB(\epsilon)}{1-p}.$$

In Example 1, an interior cutoff requires

$$\frac{rB(\epsilon)}{1-p} > 1 - y_l(\Delta).$$

Since r is unbounded above, this inequality can always be satisfied while maintaining $r \geq K(\epsilon)$. In Example 2, an interior cutoff requires

$$0 < \frac{rB(\epsilon)}{1-p} < \tilde{c} - 1.$$

The lower bound implied by the allocation condition satisfies

$$\lim_{\epsilon \to \infty} \frac{K(\epsilon)B(\epsilon)}{1-p} = \frac{\rho y^h (pn^h - 1)}{1-p}.$$

Because $pn^h > 1$ can be chosen arbitrarily close to one, this limiting lower bound can be made smaller than $\tilde{c} - 1$. Hence, for feasible primitive parameters and sufficiently large ϵ, one can choose $r \geq K(\epsilon)$ such that

$$0 < \frac{rB(\epsilon)}{1-p} < \tilde{c} - 1.$$

Thus both examples admit feasible interior sorting cutoffs while preserving the maintained consumption–fertility allocation.

Consider next the effects of adverse-shock frequency and severity. At an interior cutoff, $D(\hat{\sigma}, \Delta, p) = 0$. Since

$$\frac{\partial D}{\partial \Delta} = (1-p)(1-\gamma)[y^l(\Delta)]^{-\hat{\sigma}} y^{l\prime}(\Delta) < 0,$$

and

$$\frac{\partial D}{\partial (1-p)} = (1-\gamma)\left[v(y^l(\Delta); \hat{\sigma}) - v(\tilde{c}; \hat{\sigma})\right] + \gamma \left[\ln \epsilon - \ln(n^h + \epsilon)\right] < 0,$$

the Implicit Function Theorem implies

$$\frac{\partial \hat{\sigma}}{\partial \Delta} = -\frac{\partial D/\partial \Delta}{\partial D/\partial \sigma}, \qquad \frac{\partial \hat{\sigma}}{\partial (1-p)} = -\frac{\partial D/\partial (1-p)}{\partial D/\partial \sigma}.$$

Hence, in Example 1, where $\partial D/\partial \sigma < 0$,

$$\frac{\partial\hat{\sigma}}{\partial\Delta} < 0, \qquad \frac{\partial\hat{\sigma}}{\partial(1-p)} < 0,$$

whereas in Example 2, where $\partial D/\partial\sigma > 0$,

$$\frac{\partial\hat{\sigma}}{\partial\Delta} > 0, \qquad \frac{\partial\hat{\sigma}}{\partial(1-p)} > 0$$

locally around the parameter configuration considered. Thus, adverse-shock frequency and severity can shift the sorting cutoff toward either lower or higher degrees of conventional risk aversion.

To characterize the corresponding long-run composition, suppose for simplicity that σ_i is initially uniformly distributed on $[0, \sigma^u]$. As in the baseline model, the population choosing the risky mode dominates in the long run when $\mu < \hat{\mu}(p)$, whereas the population choosing the safe mode dominates when $\mu > \hat{\mu}(p)$.

In Example 1, relatively less risk-averse individuals choose the risky mode. Hence, for $\mu \neq \hat{\mu}(p)$,

$$\bar{\sigma} = \begin{cases} \dfrac{\hat{\sigma}}{2}, & \mu < \hat{\mu}(p), \\ \dfrac{\hat{\sigma} + \sigma^u}{2}, & \mu > \hat{\mu}(p). \end{cases}$$

Since $\partial\hat{\sigma}/\partial\Delta < 0$ and $\partial\hat{\sigma}/\partial(1-p) < 0$, greater adverse-shock frequency or severity reduces long-run conventional risk aversion locally, away from $\mu = \hat{\mu}(p)$. Moreover, a shift from relatively idiosyncratic shocks, $\mu < \hat{\mu}(p)$, to relatively aggregate shocks, $\mu > \hat{\mu}(p)$, increases long-run conventional risk aversion.

In Example 2, relatively more risk-averse individuals choose the risky mode. Hence, for $\mu \neq \hat{\mu}(p)$,

$$\bar{\sigma} = \begin{cases} \dfrac{\hat{\sigma} + \sigma^u}{2}, & \mu < \hat{\mu}(p), \\ \dfrac{\hat{\sigma}}{2}, & \mu > \hat{\mu}(p). \end{cases}$$

Since $\partial\hat{\sigma}/\partial\Delta > 0$ and $\partial\hat{\sigma}/\partial(1-p) > 0$ locally around the parameter configuration considered, greater adverse-shock frequency or severity increases long-run conventional risk aversion locally, away from $\mu = \hat{\mu}(p)$. Moreover, a shift from relatively idiosyncratic shocks to relatively aggregate shocks reduces long-run conventional risk aversion.

Thus, feasible specifications of conventional risk aversion generate opposite relationships between climatic characteristics and the evolution of attitudes toward risk. The climatic comparative statics of the baseline model are therefore not generic consequences of risk aversion. Rather, their sign-definite nature is a distinctive implication of the reference-dependent loss aversion mechanism.

B.4 Robustness to Imperfect Intergenerational Transmission

This subsection examines the robustness of the evolutionary mechanism to imperfect intergenerational transmission of predispositions toward loss neutrality. Unlike the baseline model, imperfect transmission permits the degree of loss neutrality of a dynasty to change across generations, so descendants may cross the sorting cutoff and the permanent partition of dynasties into the safe and risky groups can no longer be used to characterize the long-run dynamics. The analysis therefore proceeds in two steps. First, an analytical result establishes that imperfect transmission attenuates, but does not reverse, the direction of evolutionary selection in any given generation. Second, numerical simulations trace the resulting distributional dynamics and examine whether the severity and spatial-correlation comparative statics of the baseline model are preserved.

Suppose that the degree of loss neutrality transmitted to a child of individual i in period $t+1$ is

$$\theta_{i,t+1} = \kappa\theta_{i,t} + (1-\kappa)\bar{\theta}_t,$$

where $\bar{\theta}_t$ is the average degree of loss neutrality in the adult population in period t, and $\kappa \in [0,1]$ measures the degree of intergenerational persistence. The baseline model corresponds to perfect transmission, $\kappa = 1$, whereas $\kappa < 1$ represents imperfect transmission.

Proposition 2. *For any $\kappa \in [0,1]$ and any generation t such that $E_t[n_{i,t}] > 0$, the change in the average degree of loss neutrality from period t to period $t+1$ is*

$$\bar{\theta}_{t+1} - \bar{\theta}_t = \kappa \frac{\mathrm{Cov}_t(\theta_{i,t}, n_{i,t})}{E_t[n_{i,t}]}.$$

Hence, for every $\kappa > 0$, imperfect intergenerational transmission does not reverse the direction of evolutionary selection in any generation; conditional on the population distribution and fertility schedule in that generation, it scales the change in the population mean by κ. Moreover, when $\kappa < 1$, the intergenerational transmission mapping contracts pairwise differences in inherited predispositions—and therefore the width of the support among represented parental types—by the factor κ.

Proof. Let $E_t[\cdot]$ denote an average over the adult population in period t. Since an adult of type $\theta_{i,t}$ has $n_{i,t}$ children, the average degree of loss neutrality among members of generation $t+1$ is

$$\bar{\theta}_{t+1} = \frac{E_t\left[n_{i,t}\left(\kappa\theta_{i,t} + (1-\kappa)\bar{\theta}_t\right)\right]}{E_t[n_{i,t}]}.$$

Therefore,

$$\bar{\theta}_{t+1} = (1-\kappa)\bar{\theta}_t + \kappa\frac{E_t[n_{i,t}\theta_{i,t}]}{E_t[n_{i,t}]},$$

and thus

$$\begin{aligned}\bar{\theta}_{t+1} - \bar{\theta}_t &= \kappa\frac{E_t[n_{i,t}\theta_{i,t}] - E_t[n_{i,t}]E_t[\theta_{i,t}]}{E_t[n_{i,t}]}\\ &= \kappa\frac{\mathrm{Cov}_t(\theta_{i,t}, n_{i,t})}{E_t[n_{i,t}]}.\end{aligned}$$

The connection with the sorting mechanism of the baseline model can be seen directly. Let β_t denote the share of individuals who choose the risky mode, let $\bar{\theta}_t^A$ and $\bar{\theta}_t^B$ denote the average degree of loss neutrality among individuals who choose the safe and risky modes, respectively, and let $\bar{n}_t^B$ denote average fertility among individuals who choose the risky mode. Since fertility in the safe group is one,

$$\mathrm{Cov}_t(\theta_{i,t}, n_{i,t}) = \beta_t(1-\beta_t)(\bar{n}_t^B - 1)(\bar{\theta}_t^B - \bar{\theta}_t^A).$$

Since $\bar{\theta}_t^B > \bar{\theta}_t^A$, the sign of the change in average loss neutrality is governed by the relative reproductive success of the risky group, $\bar{n}_t^B - 1$, exactly as under perfect transmission. The parameter κ affects the magnitude, but not the sign, of this selection effect.

Finally, for any two reproducing parental types i and j represented among the offspring,

$$|\theta_{i,t+1} - \theta_{j,t+1}| = \kappa|\theta_{i,t} - \theta_{j,t}|.$$

Hence, the diameter of the offspring support is equal to κ times the diameter of the support of reproducing parental types and is therefore no greater than κ times the diameter of the adult support in period t. In the limiting case $\kappa = 0$, all children inherit $\bar{\theta}_t$, so evolutionary selection on the degree of loss neutrality is eliminated altogether. □

The proposition establishes that imperfect transmission cannot reverse the direction of selection conditional on a given population distribution and fertility realization. It does not, by itself, establish the global climatic comparative statics, because when $\kappa < 1$ descendants may cross the sorting cutoff over time and the population distribution itself evolves differently from that in the baseline model. The full distributional dynamics are therefore examined numerically.

Simulation design. The simulations begin from the uniform distribution of loss neutrality on $[0, 1]$ used in the baseline model. In every generation, individuals with $\theta_{i,t} \leq \hat{\theta}(\Delta, p)$ choose the safe mode and have fertility equal to one, whereas individuals with $\theta_{i,t} > \hat{\theta}(\Delta, p)$ choose the risky mode. As in the baseline model, risky individuals have fertility n^h under ordinary climatic conditions and zero fertility following an adverse realization. A fraction μ of risky individuals is exposed to a common aggregate realization, while the remaining fraction $1 - \mu$ experiences idiosyncratic realizations. Thus, in each generation the average fertility of the risky group is

$$\bar{n}_t^B = \mu \bar{n}_t^{ag} + (1 - \mu)\bar{n}^{id},$$

where

$$\bar{n}^{id} = pn^h > 1,$$

and

$$\bar{n}_t^{ag} = \begin{cases} n^h & \text{with probability } p, \\ 0 & \text{with probability } 1 - p. \end{cases}$$

Because the severity of the adverse shock affects the population dynamics through the sorting cutoff, $\hat{\theta}(\Delta, p)$, the simulations parameterize the two levels of adverse-shock severity directly by their associated cutoffs, holding p constant. The baseline calibration is

$$p = 0.60, \qquad n^h = 2, \qquad \hat{\theta}(\Delta, p) = 0.35, \qquad \mu = 0.25.$$

A higher degree of adverse-shock severity, $\Delta_H > \Delta$, is represented by

$$\hat{\theta}(\Delta_H, p) = 0.55,$$

whereas the higher degree of spatial correlation is

$$\mu_H = 0.75 > \mu.$$

Under perfect transmission, the implied threshold degree of spatial correlation is approximately $\hat{\mu}(p) = 0.624$, so the baseline value μ and the higher value μ_H lie on opposite sides of the threshold.

The simulations consider four degrees of intergenerational persistence,

$$\kappa \in \left\{1,\ 0.75,\ \frac{2}{3},\ 0.5\right\},$$

and three climatic environments: (Δ, p, μ), (Δ, p, μ_H), (Δ_H, p, μ). For each configuration, the distribution is simulated for 50 generations over 2,000 independent histories of the aggregate climatic shock. The idiosyncratic component is integrated using its population average, as implied by the continuum formulation of the model. Figure B.1 reports the median path of $\bar{\theta}_t$ across these histories.

Simulation results. Panel (a) of Figure B.1, corresponding to perfect intergenerational transmission, reproduces the severity and spatial-correlation predictions of the baseline model. By the end of the 50-generation simulation horizon, average loss neutrality is highest under the greater severity of the adverse shock, (Δ_H, p, μ), lowest under the greater degree of spatial correlation, (Δ, p, μ_H), and intermediate in the baseline environment, (Δ, p, μ).

Panels (b)–(d) show that the same qualitative ranking is preserved under substantial departures from perfect intergenerational transmission:

$$\text{Med}\big[\bar{\theta}_{50}(\Delta_H, p, \mu)\big] > \text{Med}\big[\bar{\theta}_{50}(\Delta, p, \mu)\big] > \text{Med}\big[\bar{\theta}_{50}(\Delta, p, \mu_H)\big]\,.$$

Thus, for each of the imperfect-transmission specifications considered, greater adverse-shock frequency or severity produces a higher degree of loss neutrality over the 50-generation simulation horizon, whereas greater spatial correlation produces a lower degree of loss neutrality and hence greater loss aversion.

Consistent with the analytical contraction result, the quantitative differences across climatic environments become smaller as κ declines. Imperfect transmission therefore weakens the persistence of evolutionary selection but, over the range considered, does not overturn the qualitative severity and spatial-correlation comparative statics of the baseline model.

Notes: Each panel reports the median simulated path of the average degree of loss neutrality, $\bar{\theta}_t$, across 2,000 independent histories of the aggregate climatic shock. The initial distribution of θ is uniform on $[0, 1]$. The baseline calibration is $p = 0.60$, $n^h = 2$, $\hat{\theta}(\Delta, p) = 0.35$, and $\mu = 0.25$. Greater adverse-shock frequency or severity is represented by $\hat{\theta}(\Delta_H, p) = 0.55$, and greater spatial correlation by $\mu_H = 0.75$.

C Data

This appendix provides additional information on the sources, construction, and interpretation of the variables used in the empirical analysis. It distinguishes the individual experimental estimates from the ancestry-level preference measures used in the baseline analysis and documents the alternative definitions of historical climatic exposure used in the sensitivity exercises.

C.1 Climatic Data

Self-calibrating Palmer Drought Severity Index. The climatic variables are constructed from the global self-calibrating Palmer Drought Severity Index (scPDSI) of van der Schrier et al. (2013). Although the underlying source is updated periodically, we define the exposure measure over the fixed 1902–2000 window, using monthly observations on a $0.5^\circ \times 0.5^\circ$ grid. The source

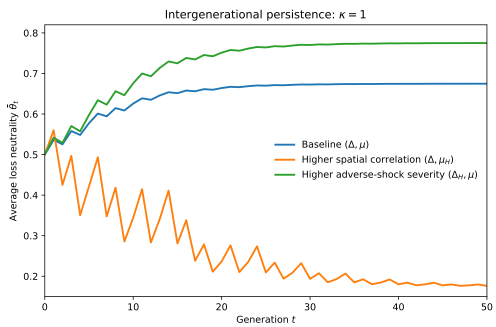

(a) $\kappa = 1$

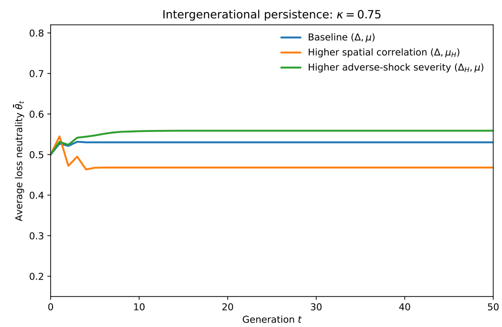

(b) $\kappa = 0.75$

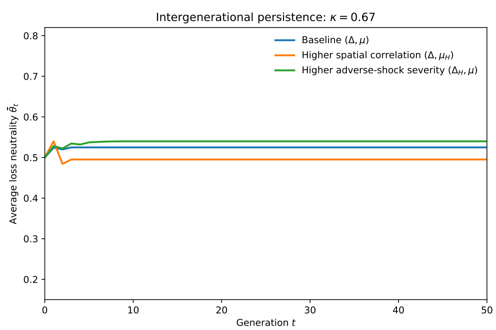

(c) $\kappa = 0.67$

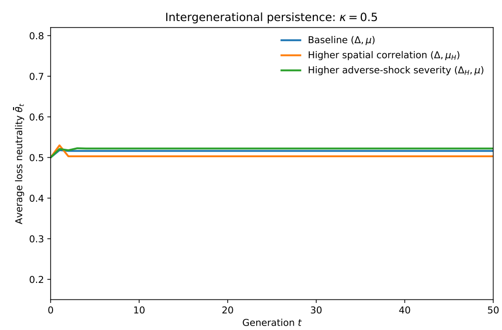

(d) $\kappa = 0.5$

Figure B.1: Evolution of Average Loss Neutrality under Imperfect Intergenerational Transmission record begins in 1901; we omit its first year so that every retained observation is preceded by a full year of climatic data, reducing sensitivity to the initialization of the Palmer index's recursive water-balance state. The index combines precipitation and temperature with fixed information on local soil and surface water-holding capacity in a two-layer water-balance model. It tracks potential evapotranspiration, soil-moisture recharge and depletion, runoff, and the accumulated duration of wet and dry spells, while calibrating its parameters to the climatic conditions of each grid cell (Wells et al., 2004). Negative values indicate conditions that are drier than those ordinarily experienced at that location. This local calibration makes the measure more suitable for comparing climatic anomalies across heterogeneous climatic regimes than the original Palmer index.

The baseline classifies a cell-month as a climatic shock when $\text{scPDSI}_{it} \leq -2$, the conventional boundary for moderate drought. Milder drought can already impair agricultural production; the purpose of the baseline cutoff is to exclude comparatively mild moisture deficits while retaining consequential events under historically limited irrigation, storage, insurance, and market integration. In the global grid, approximately 17 percent of cell-months satisfy this definition. For sensitivity, we reconstruct all climatic variables using thresholds of -1 and -3, which identify approximately 33 and 7 percent of cell-months, respectively.

Spatial scope of climatic shocks. The baseline neighborhood of cell i, denoted $\mathcal{N}_i$, consists of the up to eight other land cells that share a side or vertex with the focal cell. Formally,

$$q_{it} = \frac{1}{|\mathcal{N}_i|} \sum_{j \in \mathcal{N}_i} D_{jt}.$$

Ocean cells are excluded from the neighborhood. The variable q_{it} is therefore the fraction of adjacent land cells that experience a climatic shock in the same month as the focal cell. The monthly aggregate component is $D_{it}q_{it}$, while the monthly idiosyncratic component is $D_{it}(1 - q_{it})$. Averaging over months and then across the land area of each country produces total, aggregate, and idiosyncratic exposure in common units. Accordingly,

$$\text{E}_c = \text{E}^{\text{A}}{}_c + \text{E}^{\text{I}}{}_c.$$

Total exposure is the fraction of months in which a representative unit of national land area experiences a climatic shock. Aggregate exposure weights those shock months more heavily when nearby locations are simultaneously affected, while idiosyncratic exposure weights them more heavily when surrounding locations are unaffected. These are components of total exposure rather than separately standardized indices, allowing their magnitudes to be compared directly.

We assess sensitivity to the spatial definition by reconstructing the decomposition using all land cells within radii of 100 and 200 km rather than the immediately adjacent cells (Table F.3).

C.2 Ancestry Adjustment

The experimental design associates respondents with modern national ancestries, but the populations currently residing in a country need not descend entirely from the population residing there before the large population movements that followed 1500. We therefore ancestry-adjust the historical variables using the post-1500 migration matrix of Putterman and Weil (2010). Let m_{as} denote the share of the current population of ancestry a descended from the population residing in source country s in 1500. For a country-level historical variable Z_s, its ancestry-adjusted counterpart is

$$Z_a^{AA} = \sum_s m_{as} Z_s.$$

This transformation is applied separately to total, aggregate, and idiosyncratic climatic exposure, and to the ancestry-level controls. Thus, for example, ancestral climatic shocks measure the historical shock exposure of the source populations contributing to a present-day national ancestry rather than climatic exposure within that country's contemporary borders alone.

The ancestry-adjusted data preserve the support of the underlying source variable: a country absent from the original country-level dataset is not introduced through the adjustment. For countries covered by the migration matrix, we apply its population shares. For countries outside the matrix, the ancestry-adjusted value equals the unadjusted value in Africa, Asia, and Europe, where post-1500 population replacement was comparatively limited; it is coded as missing in the Americas and Oceania, where equating present location with ancestral origin would be especially inappropriate.

C.3 Control Variables

Caloric suitability. Caloric suitability is the index developed by Galor and Özak (2016), measuring potential agricultural output in calories per hectare per year based on the crop generating the highest caloric yield in each location. The estimates assume rain-fed cultivation and low input intensity, making the index appropriate for historical agricultural conditions. We average the grid-cell measure within countries and ancestry-adjust it using the same population-migration matrix applied to the climatic variables. This control distinguishes adverse climatic shocks from average agricultural potential.

Timing of the Neolithic transition. The timing of the Neolithic transition comes from the Agricultural Transition Data Set of Putterman (2008). It records the number of years before 2000 in which the first substantial region within a present-day country transitioned from hunting and gathering to agriculture, so larger values indicate an earlier transition. We ancestry-adjust the measure using the post-1500 migration matrix. This variable accounts for cross-ancestry differences in the duration of exposure to agriculture, the economic environment in which the proposed evolutionary mechanism operates.

Individual characteristics. The individual-level specifications progressively include a male indicator, age and age squared, and education fixed effects. Education distinguishes four categories: high school or less; some college or an associate degree; a bachelor's degree; and a graduate or professional degree. These controls account for differences in the demographic and educational composition of ancestry groups.

C.4 Variable Definitions

Table C.1: Variable Definitions

Variable	Definition
Loss aversion	Median experimentally elicited $\widehat{\lambda}_i$ within ancestry, winsorized in the upper tail at the 97.5th percentile after excluding Iraq's degenerate estimate. Higher values indicate greater loss aversion; one denotes loss neutrality.
Risk aversion	Median experimentally elicited utility-curvature parameter $\widehat{r}_i$ within ancestry. Used as a placebo outcome.
Ancestral climatic shocks	Ancestry-adjusted total historical shock exposure, TotExp_a^{AA}, using the baseline threshold scPDSI ≤ -2. Measured as a frequency in original units.
Ancestral idiosyncratic climatic shocks	Ancestry-adjusted idiosyncratic component, IdioExp_a^{AA}, based on the fraction of immediately adjacent land cells that are not simultaneously affected.
Ancestral aggregate climatic shocks	Ancestry-adjusted aggregate component, AggExp_a^{AA}, based on the fraction of immediately adjacent land cells that are simultaneously affected.
Caloric suitability	Ancestry-adjusted potential caloric output under rain-fed, low-input cultivation.
Timing of Neolithic transition	Ancestry-adjusted number of years before 2000 at which agriculture was first adopted in a substantial portion of the ancestral homeland.
Observations per ancestry	Number of valid experimental observations contributing to the ancestry-level preference measure.
Male	Indicator equal to one for male respondents.
Age	Respondent age in years; included linearly and quadratically.
Education	Highest educational attainment, grouped into four categories and included through fixed effects.

D Historical Climatic Reconstruction: Measurement

This appendix provides additional details on the Paleo Hydrodynamics Data Assimilation (PHYDA) reconstruction used to construct historical climatic-shock exposure over 1–1500 CE, its measurement properties relative to the contemporary scPDSI record employed in the baseline analysis, and the construction of the historical measures of overall, idiosyncratic, and aggregate climatic-shock exposure.

D.1 PHYDA Data and Measurement Properties

The historical analysis uses the Paleo Hydrodynamics Data Assimilation (PHYDA) product developed by Steiger et al. (2018). PHYDA provides global hydroclimatic reconstructions covering the Common Era. It combines 2,978 annually resolved paleoclimate proxy series—primarily tree rings, but also corals, ice cores, speleothems, and sediment records—with the physical structure and spatial covariance of an atmosphere–ocean climate model through data assimilation. The product provides approximately two-degree global fields of the standard Palmer Drought Severity Index

(PDSI) for annual conditions, defined from April through March of the following year, as well as for the June–August (JJA) and December–February (DJF) seasons.

PHYDA provides the closest available global historical counterpart to the contemporary climatic record and permits climatic-shock exposure to be reconstructed over the period in which the proposed evolutionary mechanism operated. Its measurement properties nevertheless differ importantly from those of the observational scPDSI record used in the baseline analysis.

First, the contemporary scPDSI record is constructed from station-based climatic observations, whereas PHYDA infers historical climatic conditions from paleoclimate proxies in conjunction with climate-model relationships. Proxy coverage is geographically uneven and becomes progressively sparser further back in time, increasing the relative influence of the model prior. Because data assimilation propagates information through model-based spatial covariance, a reconstructed grid cell need not be informed by a geographically proximate proxy. Moreover, declining proxy information reduces the variance of the posterior ensemble mean in earlier periods; this attenuation should not be interpreted as a corresponding reduction in the variability of the underlying climate (Steiger et al., 2018).

Second, the approximately two-degree spatial resolution of PHYDA is substantially coarser than the 0.5° resolution of the contemporary record. Each historical grid cell therefore averages climatic conditions over a considerably larger area. This spatial averaging can smooth localized climatic shocks and reduces the number of spatial units available for measuring whether adverse realizations occur simultaneously across neighboring locations. These features are particularly relevant for the decomposition of climatic shocks into their idiosyncratic and aggregate components. The coarser grid also prevents the construction of historical country-level measures for several small countries and islands represented in the baseline sample.

Third, PHYDA is resolved annually or seasonally, whereas the contemporary record is monthly. Seasonal aggregation can attenuate shorter-lived drought episodes and reduces the temporal precision with which the coincidence of climatic shocks across neighboring locations can be identified. Finally, PHYDA reconstructs the standard PDSI, whereas the baseline analysis employs the self-calibrated PDSI. Numerical thresholds are consequently not directly comparable across the two datasets.

These differences motivate the respective roles assigned to the two climatic datasets in the empirical analysis. The directly observed contemporary scPDSI record provides substantially finer spatial and temporal resolution and greater precision for measuring persistent cross-sectional differences in climatic-shock exposure. PHYDA, in contrast, provides an independent reconstruction based on different underlying information and methods and extending directly into the historical period in which the proposed evolutionary process operated. Its value therefore lies both in assessing the persistence of cross-regional climatic-shock exposure and in permitting an independent historical test of the theory's predictions.

D.2 Construction of Historical Climatic-Shock Exposure

The historical climatic-shock measures are constructed from the posterior ensemble mean of the PHYDA PDSI reconstruction over 1–1500 CE. Tree rings constitute the majority of the underlying proxy network and primarily record climatic conditions during the local growing season. The analysis therefore uses the JJA reconstruction for Northern Hemisphere land cells and the DJF reconstruction for Southern Hemisphere land cells. This choice strengthens the underlying proxy signal and reduces reconstruction error; it does not impose the assumption that these particular calendar months were universally the most agriculturally consequential.

Averaging across posterior ensemble members attenuates extreme PDSI realizations toward zero. A historical cell-year is therefore classified as experiencing an adverse climatic shock when the hemisphere-specific PDSI is less than or equal to -1, rather than mechanically applying the baseline scPDSI threshold of -2 to the posterior ensemble mean. The more inclusive historical threshold accommodates the attenuation induced by ensemble averaging and should not be interpreted as equating the two numerical thresholds or the two versions of the PDSI.

Historical exposure is measured over 1–1500 CE, capturing the premodern period during which the Malthusian forces and subsistence constraint at the center of the theoretical mechanism were operative. The subsequent construction parallels that of the baseline climatic variables.

The three measures are averaged over 1–1500 CE and aggregated across land cells within each country using cell-area weights. The resulting country-level measures are then ancestry-adjusted using the same post-1500 population-migration matrix employed in the baseline analysis. The historical measures therefore preserve the baseline distinction between overall climatic-shock exposure and its idiosyncratic and aggregate components while replacing the contemporary climatic record with an independent reconstruction of conditions prevailing during the historical period itself.

E Additional Figures

E.1 Baseline Scatterplot

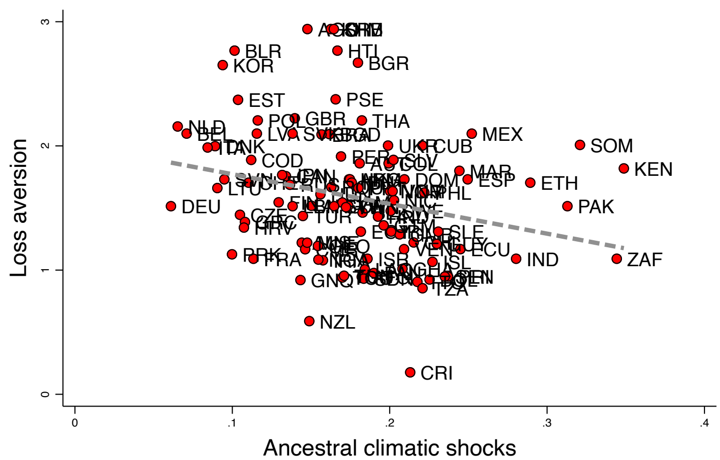

β = -2.39; s.e. = 0.82; t = -2.91; R^2 = 0.07; N = 102

Figure E.1: Loss Aversion and Ancestral Climatic Shocks

Notes: Each dot represents an ancestry. The dashed line depicts the fitted linear relationship. The coefficient, robust standard error, t-statistic, R^2, and sample size are reported below the figure.

E.2 Alternative Treatments of Extreme Loss-Aversion Values

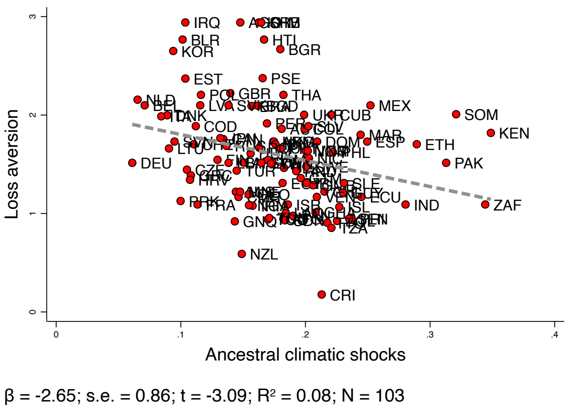

(a) Iraq retained and capped

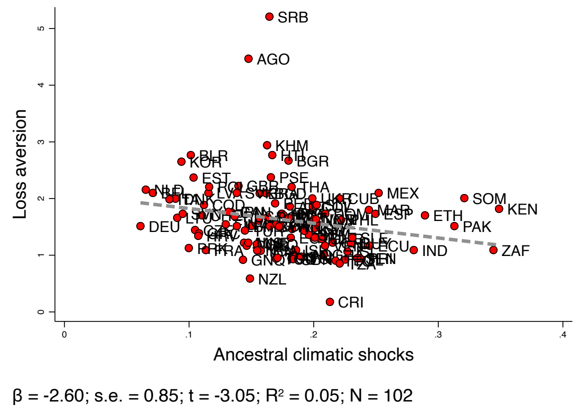

(b) Angola and Serbia unwinsorized

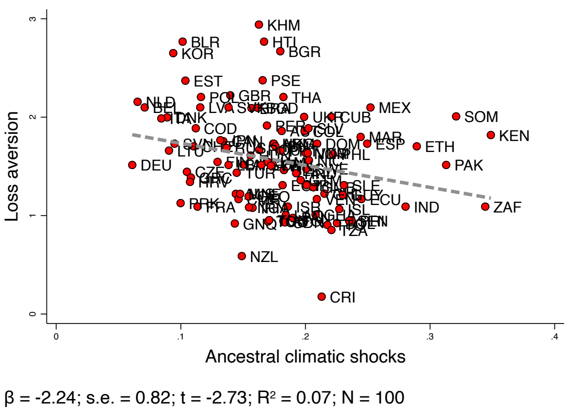

(c) Angola and Serbia excluded

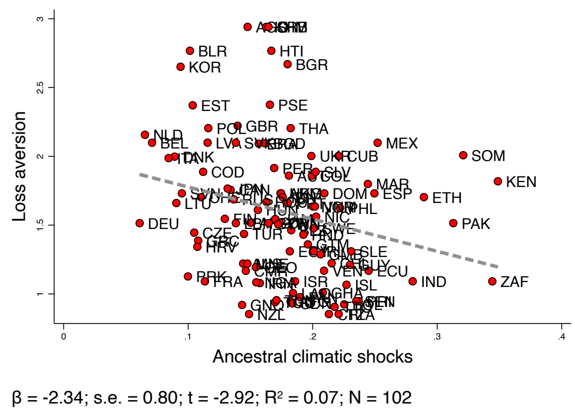

(d) Symmetric winsorization

Figure E.2: Robustness to Alternative Treatments of Extreme Loss-Aversion Values

Notes: This figure examines the robustness of the baseline relationship to alternative treatments of extreme ancestry-level loss-aversion values. The baseline excludes Iraq and caps Angola and Serbia at the 97.5th percentile. Panel (a) retains Iraq and applies the same upper-tail cap to all observations exceeding it. Panel (b) excludes Iraq but restores the original, unwinsorized values of Angola and Serbia. Panel (c) excludes Iraq, Angola, and Serbia. Panel (d) excludes Iraq and symmetrically winsorizes loss aversion at the 2.5th and 97.5th percentiles. The dependent variable is the ancestry-level median of loss aversion, and ancestral climatic-shock exposure is measured as in the baseline analysis. Dashed lines show ordinary least-squares fits. Estimated coefficients, heteroskedasticity-robust standard errors, t-statistics, R^2, and sample sizes are reported below each panel.

E.3 Minimum Ancestry Sample-Size Thresholds

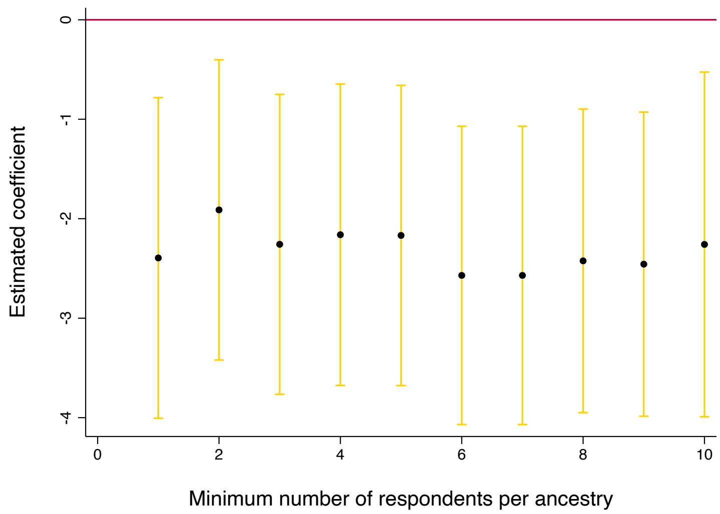

Figure E.3: Robustness to Minimum Ancestry Sample-Size Thresholds

Notes: The figure plots the estimated coefficient on ancestral climatic shocks from separate regressions of loss aversion. The horizontal axis indicates the minimum number of respondents required for an ancestry to be included in the estimation sample. At the most restrictive threshold of 10 respondents, only 70 ancestries remain in the estimation sample. Vertical bars report 95% confidence intervals based on robust standard errors.

E.4 Migration Flows and Population Composition

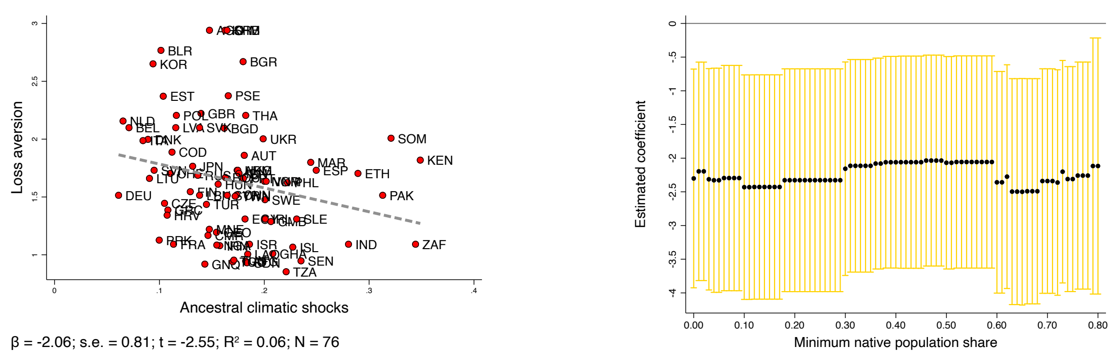

Figure E.4: Loss Aversion and Ancestral Climatic Shocks: Migration Flows and Population Composition

Notes: This figure examines whether post-1500 migration and population replacement drive the baseline relationship. Panel (a) restricts the sample to ancestries originating in the Old World. Panel (b) plots the estimated coefficient from the baseline specification as the minimum required share of the present-day population descended from the population residing in the same territory in 1500 rises from 0 to 80%. At the 80% threshold, the sample contains 60 ancestries.

F Additional Tables

F.1 Summary Statistics

Table F.1: Summary Statistics

	Mean	SD	Median	Min	Max	N
A. Ancestry-level						
Loss aversion	1.59	0.53	1.54	0.18	2.94	102
Risk aversion	1.19	0.64	1.08	0.30	6.58	102
Ancestral climatic shocks	0.18	0.06	0.17	0.06	0.35	102
Ancestral idiosyncratic climatic shocks	0.04	0.01	0.04	0.02	0.07	102
Ancestral aggregate climatic shocks	0.14	0.05	0.13	0.04	0.28	102
Observations per ancestry	15	10	14	1	34	102
B. Individual-level						
Loss aversion	1.82	0.95	1.66	0.02	3.08	1,424
Risk aversion	1.74	1.71	1.08	0.30	6.58	1,424
% of males	0.51	0.50	1.00	0.00	1.00	1,424
Age	38.50	14.29	36.00	18.00	85.00	1,424
% of college completed	0.62	0.49	1.00	0.00	1.00	1,424

Notes: The table reports the mean, standard deviation (SD), median, minimum (MIN), maximum (MAX), and number of observations (N) for the principal variables used in the analysis.

F.2 Alternative Climatic-Shock Thresholds

Table F.2: Robustness to Alternative Climatic-Shock Thresholds

	Loss aversion			
	A. scPDSI ≤ -1			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−1.90***	−1.77***	−1.94***	−1.78***
	(0.55)	(0.55)	(0.56)	(0.56)
Caloric suitability		✓		✓
Timing of Neolithic transition			✓	✓
Dep. var. mean	1.59	1.59	1.60	1.60
Observations	102	102	98	98
Adjusted R^2	0.091	0.10	0.096	0.12
	B. scPDSI ≤ -3			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−4.75**	−4.39**	−5.55***	−5.06**
	(1.94)	(2.01)	(2.06)	(2.11)
Caloric suitability		✓		✓
Timing of Neolithic transition			✓	✓
Dep. var. mean	1.59	1.59	1.60	1.60
Observations	102	102	98	98
Adjusted R^2	0.045	0.066	0.059	0.087

Notes: The unit of observation is an ancestry. This table examines robustness to alternative scPDSI thresholds for defining a climatic shock. The baseline defines a shock as scPDSI less than or equal to −2. Panel A uses the less restrictive threshold of −1, while Panel B uses the more restrictive threshold of −3. Ancestral climatic-shock exposure is the area-weighted share of cell-month observations satisfying the corresponding threshold and is ancestry-adjusted using the post-1500 population-migration matrix. Columns (2) and (3) add ancestral caloric suitability and the timing of the Neolithic transition separately; column (4) includes both controls. Because the exposure measures are shares constructed using different thresholds and enter the regressions in levels, coefficient magnitudes are not directly comparable across panels. The dependent variable is the ancestry-level median of loss aversion used in the baseline analysis. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

F.3 Alternative Definitions of Neighboring Locations

Table F.3: Robustness to Alternative Definitions of Neighboring Locations

	Loss aversion			
	A. 100km radius			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−2.40***			−3.75***
	(0.82)			(1.01)
Ancestral idiosyncratic climatic shocks		−15.3***		−18.1**
		(4.04)		(7.91)
Ancestral aggregate climatic shocks			−2.49**	
			(0.99)	
Dep. var. mean	1.59	1.59	1.59	1.59
Observations	102	102	102	102
Adjusted R^2	0.065	0.11	0.050	0.12
	B. 200km radius			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−2.40***			−3.32***
	(0.82)			(0.93)
Ancestral idiosyncratic climatic shocks		−9.67***		−10.4*
		(2.67)		(5.78)
Ancestral aggregate climatic shocks			−2.67**	
			(1.11)	
Dep. var. mean	1.59	1.59	1.59	1.59
Observations	102	102	102	102
Adjusted R^2	0.066	0.094	0.047	0.094

Notes: The unit of observation is an ancestry. This table examines robustness to alternative spatial definitions of neighboring locations. Whereas the baseline defines the relevant neighbors as the immediately adjacent land grid cells, Panels A and B define neighbors as all land grid cells located within 100 and 200 kilometers of the focal cell, respectively. Column (4) includes total exposure and recentered idiosyncratic exposure jointly. The coefficient on idiosyncratic exposure therefore captures the difference between idiosyncratic and aggregate exposure, holding total exposure constant, while the coefficient on total exposure is evaluated at the sample-average spatial composition. Because total climatic-shock exposure does not depend on the definition of neighboring locations, its estimate in column (1) is identical across panels. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

F.4 Alternative Treatments of Mixed Ancestry

Table F.4: Robustness to Alternative Treatments of Mixed Ancestry

	Loss aversion		
	Baseline	Average exposure	Single ancestry
	(1)	(2)	(3)
Ancestral climatic shocks	−2.39***		−2.27**
	(0.82)		(1.11)
Ancestral climatic shocks (avg.)		−1.48**	
		(0.71)	
Dep. var. mean	1.59	1.79	1.74
Observations	102	554	94
Adjusted R^2	0.065	0.0067	0.032

Notes: This table examines the robustness of the relationship between ancestral climatic shocks and loss aversion to alternative treatments of respondents reporting two ancestries. Column (1) presents the baseline specification, which retains respondents reporting either one or two ancestries, groups them by their primary ancestry, and assigns climatic exposure according to that ancestry. Column (2) retains the same respondent sample, groups respondents by their unique combination of reported ancestries, and assigns those reporting two ancestries the average climatic exposure of their primary and secondary ancestries. Column (3) restricts the sample to respondents reporting a single ancestry. An observation is an ancestry or ancestry-combination bin. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

F.5 Climatic-Shock Severity

Table F.5: Robustness to Measuring Climatic Shocks on the Intensive Margin

	Loss aversion			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks (severity)	−2.21**	−2.04**	−2.56***	−2.33**
	(0.85)	(0.87)	(0.90)	(0.91)
Caloric suitability		✓		✓
Timing of Neolithic transition			✓	✓
Dep. var. mean	1.59	1.59	1.60	1.60
Observations	102	102	98	98
Adjusted R^2	0.052	0.071	0.067	0.093

Notes: The unit of observation is an ancestry. This table replaces the baseline extensive-margin measure—the frequency with which scPDSI crosses the −2 threshold—with an ancestry-adjusted intensive-margin measure of the severity of climatic shocks. The dependent variable is the ancestry-level median of loss aversion used in the baseline analysis. Columns (2) and (3) add ancestral caloric suitability and the timing of the Neolithic transition separately; column (4) includes both controls. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

F.6 Controlling for Weather-Station Coverage

Table F.6: Controlling for Weather-Station Coverage

	Loss aversion			
	(1)	(2)	(3)	(4)
Ancestral climatic shocks	−2.38***	−2.27***	−2.68***	−2.53***
	(0.80)	(0.82)	(0.84)	(0.85)
Weather station density	✓	✓	✓	✓
Caloric suitability		✓		✓
Timing of Neolithic transition			✓	✓
Dep. var. mean	1.59	1.59	1.60	1.60
Observations	102	102	98	98
Adjusted R^2	0.13	0.14	0.15	0.16

Notes: This table reports ancestry-level estimates of the relationship between ancestral climatic shocks and loss aversion controlling for ancestry-adjusted weather-station coverage, measured as the average number of active precipitation stations per 100,000 square kilometers in the ancestral country during 1902–2000. Columns (2) and (4) control for ancestral caloric suitability, while columns (3) and (4) control for the timing of the Neolithic transition. Robust standard errors are reported in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.