The Perceived Price of Redistribution: Zero-Sum Beliefs and Redistributive Demand

Martin Ljunge* Kelly Suzanne Ragan†

August 10, 2026

Abstract

Why do citizens living under the same welfare state want such different amounts of redistribution? Part of the answer, we argue, is not a difference in taste but a difference in perceived price. We place two culturally transmitted beliefs in distinct parts of a tax problem. Zero-sum beliefs set the price, through a perceived elasticity of the tax base: whether the pie is fixed or grows with effort and investment. Uncertainty avoidance sets the value of insurance against an ambiguous future. A third orientation, the weight placed on the future, multiplies both, which is why its effect cannot be signed in advance. We test the model on 10,718 second-generation respondents from 70 ancestral countries, compared within the same Western European country and survey round. Ancestral zero-sum beliefs and uncertainty avoidance both raise support for redistribution, each by between a quarter and two-fifths of the education gradient. The two readings of zero-sum thinking, a price on a pie that may not grow or appropriation from others, diverge on where in the income distribution the effect operates, so the design discriminates between them, and the evidence favors the price. The distinction is not a matter of labelling: a fixed-pie believer does not want transfers more but faces a lower perceived price for them, and unlike a taste, a price is something evidence can move. That is what technological transitions such as artificial intelligence put at stake, since they bear on exactly the beliefs we measure.

JEL codes: D72, D81, D83, H23, Z13

Keywords: preferences for redistribution; zero-sum thinking; uncertainty avoidance; cultural transmission; second-generation immigrants; perceived cost of taxation

1 Introduction

Artificial intelligence has made an old political-economy problem newly visible. Transformative technologies can raise productivity and expand the resources available for social policy, but they can also displace workers, shift bargaining power, and create risks whose incidence is difficult to know in advance. The economics of automation documents both sides: industrial robots have measurably displaced workers and reduced wages in exposed labor markets (Acemoglu and Restrepo, 2020), task displacement accounts for much of the rise in U.S. wage inequality (Acemoglu and Restrepo,

*Research Institute of Industrial Economics (IFN), Stockholm; martin.ljunge@ifn.se. Ljunge acknowledges financial support from Sandrews and the Torsten Söderberg Foundation.

†Swedish Institute for Social Research (SOFI), Stockholm University.

2022), and how the gains from artificial intelligence will be produced and distributed is the subject of active debate (Korinek and Stiglitz, 2019; Acemoglu, 2025). Policy institutions reach a similar assessment of large potential gains combined with substantial distributional uncertainty.[1] What such a transition does to the demand for redistribution therefore depends on more than who is exposed. It depends on whether citizens read an uncertain future as downside risk worth insuring against, or as growth that taxation might impede.

That is a question about beliefs, and it generalizes well beyond any one technology. People living under the same tax and welfare institutions disagree sharply about redistribution. Income position explains part of the disagreement. Another part, we argue, is not about who gains at all: it is about the perceived *price* of the policy and the value of what it insures, and both beliefs are inherited.

The standard political-economy answer begins with income position. In the canonical model of Meltzer and Richard (1981), citizens with income below the mean prefer more redistribution because proportional taxation with lump-sum transfers moves resources toward them. This logic remains central, but it is incomplete for settings in which the relevant stakes lie partly in the future. Redistribution can insure against future income risk (Varian, 1980); beliefs about mobility can reduce or increase support for redistribution (Piketty, 1995; Benabou and Ok, 2001; Alesina and La Ferrara, 2005); and beliefs about luck, effort, and fairness can sustain different redistributive regimes (Alesina and Angeletos, 2005; Benabou and Tirole, 2006). What this belief-based tradition has not delivered is an account of how beliefs about *the production of resources* and *uncertainty about the future* jointly determine redistributive demand. Citizens are not only asking whether they gain from redistribution today. They are forming beliefs about future downside risk and about whether redistribution undermines the economic process that will generate tomorrow's income.

Two persistent, culturally transmitted orientations organize that problem. The first is zero-sum beliefs: whether resources are viewed as a fixed pie to be divided or as an endogenous outcome of effort, investment, innovation, and policy. This governs the perceived price of redistribution: whether taxation is cheap reallocation of an exogenous pie or costly interference with an endogenous growth process. The second is uncertainty avoidance: the degree to which unresolved future states are experienced as threatening. This governs the value of insurance against ambiguous future shocks, and so the benefit side of the same calculation. Together they determine the two future quantities a citizen weighs: what redistribution will cost the economy she expects to live in, and what it will protect her from.

A third orientation, the weight placed on future outcomes, is not a further quantity but a multiplier on the first two. It scales both the insurance benefit that uncertainty avoidance sizes and the growth cost that zero-sum beliefs price. This is why it occupies a different position in the argument, and why its effect on redistributive demand cannot be signed in advance: it amplifies a benefit and a cost at once, and which one it amplifies more is a question about the other two orientations rather than about the future weight itself.

We formalize this in a simple model of preferred redistribution (Section 2). An agent chooses a durable tax-transfer policy that provides current redistribution and reduces exposure to an ambiguous future income shock. Taxation carries a perceived fiscal cost with a static component and a future tax-base component. Zero-sum beliefs lower perceived tax costs; in a microfounded reading, a zero-sum believer perceives a low elasticity of taxable income because there is little endogenous production to discourage. This builds on evidence that citizens hold heterogeneous beliefs about the efficiency costs of taxation (Stantcheva, 2021). Uncertainty avoidance is represented by the

[1] See IMF (2024), OECD (2023), and ILO (2023).

adverse lower bound the agent uses when evaluating the ambiguous shock. The future weight then multiplies both, entering the numerator of the preferred tax rate through insurance and its denominator through the future tax base.

Two of the model's predictions are therefore signed and one is not. Zero-sum beliefs and uncertainty avoidance both raise desired redistribution, the first by lowering the perceived price of transfers and the second by raising the value of insurance. The third prediction is a genuine trade-off, and the contrast with existing intuition is sharp: a common reading gives patience a clear political home on the anti-redistribution side, since forward-looking agents save, self-insure, and expect to climb (Benabou and Ok, 2001). In this model it has no such home. It lowers desired redistribution where perceived future tax-base costs dominate and raises it where protection against adverse future states is sufficiently valuable, or where zero-sum beliefs make taxation appear cheap. Which case a population is in is an empirical question about its other two orientations.

We bring these ideas to the data using the epidemiological approach (Fernández and Fogli, 2009; Fernández, 2011; Luttmer and Singhal, 2011), which infers the role of inherited culture from the behavior of people who share a current environment but differ in ancestral origin, and which Algan and Cahuc (2010) extend by using the inherited component of a trait measured among descendants of earlier migrants. The empirical analysis studies second-generation immigrants in the European Social Survey. Respondents live in, and were largely shaped by, the institutions and economic conditions of their European country of residence, but differ in the ancestral-country beliefs transmitted through families and communities (Bisin and Verdier, 2001). Comparing people in the same country and survey round but with different parental countries of birth holds the destination context fixed while exploiting variation in ancestral exposure. Luttmer and Singhal (2011) show with this design that preferences for redistribution have a strong culturally transmitted component; our contribution is to ask *which* inherited beliefs carry that component and why. The baseline sample retains respondents with exactly one foreign-born and one native-born parent, living in Western Europe, which assigns ancestry unambiguously from a single parent and, because parental intermarriage is itself an outcome of integration, holds destination-country minority status nearly fixed.

Three results organize the empirical section (Section 4). First, ancestral zero-sum beliefs predict more support for redistribution; the association is present when the trait enters alone and sharpens when the traits enter jointly. Second, ancestral uncertainty avoidance robustly predicts more support. Both survive wild-cluster inference, ancestral development and institutional controls, a race against eleven rival inherited attitudes, and every leave-one-out perturbation. Third, the ancestral index we use for the future weight enters negatively, but the informative result concerns the model's structure, not that coefficient. The trade-off the model places on the future weight exists only if a perceived future growth cost and an insurance value for an ambiguous future are both present. Both are estimated positive and precise, and both survive replacing the survey scale with the pre-1500 agricultural return that Galor and Özak (2016) argue selected for long-term orientation, so the trade-off is real. Where a population sits between its two arms we cannot pin down: the index is a composite that disagrees in sign with a behavioral patience measure, and the part of it that moves redistribution is not the part tracking how heavily the future is weighted (Section 7.3). We rest the paper's claims on the two belief channels the design identifies.

This paper contributes to three literatures. First, to the literature on culture and redistribution (Luttmer and Singhal, 2011; Alesina and Giuliano, 2011), which established that inherited culture matters for redistributive preferences but has been largely agnostic about which beliefs do the work; we identify specific inherited beliefs and give each a distinct theoretical role. Second, to belief-

based theories of redistribution (Piketty, 1995; Benabou and Ok, 2001; Alesina and Angeletos, 2005; Benabou and Tirole, 2006). Our model treats zero-sum thinking not as a primitive taste for transfers, nor as a belief about mobility or fairness, but as a belief about the perceived cost of taxation: a perceived elasticity of the economic pie. This connects the fairness-and-mobility tradition to the survey evidence on tax-mechanism reasoning (Stantcheva, 2021). Third, to the fast-growing literature on zero-sum thinking. The idea long predates its arrival in economics: Foster (1965, 1967) argued that peasant societies operate on an "image of limited good," the tacit premise that wealth exists in fixed supply so that one family's advance must come at another's expense. Psychology has since given the construct a validated measurement tradition (Różycka-Tran et al., 2015; Davidai and Ongis, 2019). Chinoy et al. (2026) document, in a large U.S. survey, that a zero-sum mindset strongly predicts support for redistribution and trace it to individual and ancestral experiences; Ali et al. (2025) show how zero-sum thinking can arise strategically in elections.

Relative to Chinoy et al. (2026), who establish *that* a zero-sum mindset predicts redistributive demand, we ask what kind of belief it is. "Zero-sum" bundles two distinct claims: *generativity*, whether output is fixed or is produced endogenously by effort and investment, and *appropriation*, whether one person's gains come at another's expense. The two run together in the folk concept, since Foster's limited good is a claim about a fixed pie from which he derives that gains come at others' expense, while the recent economics literature foregrounds appropriation. Embedded in a tax problem, our model isolates generativity and recasts zero-sum thinking as a perceived elasticity of the tax base: a price rather than a preference.

The generativity and appropriation readings agree on direct effects and diverge on where in the income distribution the effect operates, which is what Section 5.1 exploits; the appropriation reading is developed as the reserve variant of Appendix B. The two designs also bracket the belief's origin: their zero-sum thinking responds to own and family experience, whereas ours isolates the inherited component, measured in the ancestral country and predetermined with respect to the destination context in which preferences are expressed. Read through the model, the paper also gives Foster's image a fiscal microfoundation: a pure limited-good believer perceives a perfectly inelastic tax base, which is why redistribution appears cheap.

Beyond these literatures, the paper speaks to work on the political economy of technological change, where exposure to automation has been linked to redistributive demand and political behavior (Thewissen and Rueda, 2019; Gallego and Kurer, 2022): it offers one explanation for why exposure to the same transformation produces divergent redistributive responses.

Section 2 develops the model and states its predictions. Section 3 describes the data, sample, and estimation methods. Section 4 presents the reduced-form results and their robustness. Section 5 asks which belief zero-sum thinking is. Section 6 estimates the preferred-tax equation itself, and Section 7 asks what the ancestral measures capture. Section 8 concludes. Appendix A collects supporting tables, and Appendix B develops the extractive-belief variant of the model.

2 A Model of Redistribution with Beliefs about the Pie and the Future

This section develops a tractable model in which inherited orientations have distinct roles in the demand for redistribution. Zero-sum beliefs govern the perceived cost of taxation: an agent who views the economic pie as fixed perceives little endogenous production for taxes to distort, while an

agent who views the pie as endogenously produced perceives taxation as costly to the current and future tax base. Uncertainty avoidance governs the value of insurance against ambiguous future shocks. The future weight multiplies both: the insurance the second makes valuable and the future tax base the first prices.

The model delivers a closed-form preferred tax rate and comparative statics that map into the empirical specification. Its central theoretical payoff is that the future weight is not mechanically pro- or anti-redistribution. An agent who weights the future heavily values redistribution more as insurance against ambiguous future downside risk, but also places more weight on perceived future tax-base costs. Which force dominates depends on uncertainty avoidance and on zero-sum beliefs, which determine whether those future costs are perceived to exist at all. Each result is proved where it is stated; the extractive-belief variant is developed in Appendix B.

2.1 Environment

Agents, income, and preferences. There is a continuum of agents indexed by i, and agents live for two periods. Agent i has period-1 pre-tax income α_i. Utility is linear in consumption, which keeps the model transparent and isolates the role of ambiguity attitudes from ordinary risk aversion. The weight placed on period-2 consumption is $\beta \geq 0$, and it is convenient to work throughout with the normalized future weight

$$b \equiv \frac{\beta}{1+\beta} \in [0, 1),$$

extended to $b = 1$ as the limiting case $\beta \to \infty$, since the empirical map of Section 3.4 attains the endpoint and every expression below is continuous there.

The income process and uncertainty avoidance. Period-2 pre-tax income is $\alpha_i + \mu_i$, where μ_i is an idiosyncratic shock. The shock is ambiguous, not risky: the agent does not know its distribution, regards realizations in $[-\lambda, \lambda]$ as plausible, and evaluates policy at the adverse bound $\mu_i = -\lambda$, as in the maxmin criterion of Gilboa and Schmeidler (1989). Uncertainty avoidance is represented by that bound, $\lambda \geq 0$, which makes it a reduced-form measure of ambiguity concern: it combines the set of downside states the agent regards as plausible with the pessimism used to evaluate that set, two components that the decision-theoretic literature separates (Ghirardato et al., 2004; Chandrasekher et al., 2022) but that no country-level index does.

Remark 1 (Ambiguity, not risk)**.** *The mapping is behavioral rather than literal. Hofstede's country-level index records discomfort with unstructured or unknown situations; it is not a primitive ambiguity-aversion coefficient and it does not directly measure the support of an income shock (Hofstede, 2001). We therefore interpret ancestral uncertainty avoidance as a proxy for λ, the adverse state that receives decision weight. A smooth-ambiguity formulation (Klibanoff et al., 2005) would preserve the insurance channel; under linear utility and a known mean-zero distribution the channel would be absent, which is what separates uncertainty avoidance from ordinary risk aversion here.*

The economy and policy. Average income is $\bar{y}$. The government chooses a proportional tax rate $\tau \in [0, 1]$ on income and rebates the revenue lump-sum. Since $\tau \leq 1$, private exposure to the period-2 shock is $(1 - \tau)\mu_i$, so redistribution reduces exposure to the lower-tail state as well as transferring resources within each period.

The policy is durable: the same rate applies in both periods. This assumption does real work and we state it in the open. Durability matches the empirical object (a standing policy stance rather

than a one-year budget vote), and it is what allows the future weight to price an intertemporal trade-off at all. With period-specific rates the two periods separate and the future weight cancels from each first-order condition, leaving it nothing to trade off; we return to this once the solution is on the page and show exactly which results depend on it. The tax rate throughout is the policy preferred by an individual agent, not the implemented equilibrium rate.

2.2 Zero-Sum Beliefs and Perceived Tax Costs

Zero-sum beliefs are represented by $\theta \in [0, 1]$. Higher θ means a stronger fixed-pie view: the agent believes that resources are largely exogenous to policy and private behavior. Lower θ means a more growth-oriented view, in which taxation can distort effort, investment, or growth.

Zero-sum beliefs enter through the perceived cost of taxation. The contemporaneous efficiency cost is

$$K_s(\theta) = k_0 + (1 - \theta)k_s, \tag{1}$$

with administrative cost $k_0 > 0$ and perceived static supply-side cost $k_s \geq 0$, and the future growth or investment cost, which applies only to period-2 resources, is

$$K_g(\theta) = (1 - \theta)k_g, \qquad k_g \geq 0. \tag{2}$$

The perceived per-capita transfer in period 1 is

$$T_1(\tau, \theta) = \tau\bar{y} - \tfrac{1}{2}K_s(\theta)\tau^2\bar{y}, \tag{3}$$

and the perceived transfer in period 2 is

$$T_2(\tau, \theta) = \tau\bar{y} - \tfrac{1}{2}\left[K_s(\theta) + K_g(\theta)\right]\tau^2\bar{y}. \tag{4}$$

The quadratic form is a local approximation to what a perceived elasticity of taxable income delivers. Suppose the agent believes the aggregate tax base responds to the net-of-tax rate with elasticity $\varepsilon(\theta)$, so the perceived base is $B(\tau) = \bar{y}(1 - \tau)^{\varepsilon(\theta)}$. Then perceived revenue is $\tau B(\tau) \approx \tau\bar{y} - \varepsilon(\theta)\tau^2\bar{y}$ for small τ, which matches (3) with $K_s(\theta) = k_0 + 2\varepsilon(\theta)$ once administrative costs are added. Setting $\varepsilon(\theta) = (1 - \theta)\bar{\varepsilon}$ says that zero-sum believers perceive an inelastic tax base: there is no endogenous production to discourage. This reading connects θ directly to the elasticity and mechanism beliefs that Stantcheva (2021) measures in surveys, and it is the formal sense in which the model treats zero-sum thinking as a belief about production rather than a taste for transfers. The future component $K_g(\theta) = (1 - \theta)k_g$ is posited by analogy and not separately microfounded: it extends the same elasticity logic to a perceived period-2 investment response to taxation. The quadratic form is valid for tax rates below the perceived revenue peak $1/K_s(\theta)$, where perceived transfers increase in the rate; every fitted value in the empirical analysis satisfies this condition by a wide margin. Because the shock μ_i is idiosyncratic and averages out across agents, the perceived transfers (3)–(4) are deterministic; ambiguity enters policy evaluation only through the agent's private exposure $(1 - \tau)\mu_i$.

2.3 Individual Optimization

Period-1 consumption is $c_{i1} = (1 - \tau)\alpha_i + T_1(\tau, \theta)$, and period-2 consumption under shock μ_i is $c_{i2}(\mu_i) = (1 - \tau)(\alpha_i + \mu_i) + T_2(\tau, \theta)$. The agent chooses τ to maximize

$$V_i(\tau) = c_{i1} + \beta \min_{\mu_i \in [-\lambda, \lambda]} c_{i2}(\mu_i). \tag{5}$$

Since $\partial c_{i2}/\partial \mu_i = 1 - \tau \geq 0$, period-2 consumption is nondecreasing in the shock, so the minimand is attained at the lower bound $\mu_i = -\lambda$, giving the objective in closed form,

$$V_i(\tau) = (1-\tau)\alpha_i + T_1(\tau, \theta) + \beta\left[(1-\tau)(\alpha_i - \lambda) + T_2(\tau, \theta)\right]. \tag{6}$$

Lemma 1 (Single-peaked preferences)**.** *V_i is strictly concave in τ,*

$$V_i''(\tau) = -\left[(1+\beta)K_s(\theta) + \beta K_g(\theta)\right]\bar{y} < 0,$$

so each agent's preferences over $\tau \in [0,1]$ are single-peaked.

Proof. Differentiating V_i twice gives $V_i''(\tau) = -[(1+\beta)K_s(\theta) + \beta K_g(\theta)]\bar{y}$. Since $k_0 > 0$, $K_s(\theta) = k_0 + (1-\theta)k_s > 0$ for every $\theta \in [0,1]$ and $K_g(\theta) = (1-\theta)k_g \geq 0$, so $V_i''(\tau) < 0$; a strictly concave function on $[0,1]$ is single-peaked. □

With $A_i \equiv \bar{y} - \alpha_i$ the current redistributive motive, the first-order condition for an interior optimum solves to

$$\tau_i^\circ = \frac{A_i + b\lambda}{\left[k_0 + (1-\theta)k_s + b(1-\theta)k_g\right]\bar{y}}, \tag{7}$$

with constrained choice $\tau_i^* = \min\{1, \max\{0, \tau_i^\circ\}\}$. The propositions below refer to interior choices, with the usual qualification that corner solutions truncate the comparative statics.

Remark 2 (Aggregation)**.** *By Lemma 1, preferences are single-peaked, so if the tax rate were chosen by majority rule within a polity whose members share (θ, b, λ) and differ only in income, the median-income voter's preferred rate in* (7) *would be decisive (Meltzer and Richard, 1981). The empirical design, however, uses individual stated preferences of second-generation immigrants, so the relevant object is the individual preferred tax itself, not the equilibrium policy. This is an advantage: the model's individual-level comparative statics can be tested without taking a stand on political aggregation in the destination country.*

2.4 Interpretation

Equation (7) separates the determinants of redistributive preferences into a demand side and a perceived-cost side:

$$\tau_i^\circ = \frac{\overbrace{A_i}^{\text{current redistribution}} + \overbrace{b\lambda}^{\text{future ambiguity insurance}}}{\underbrace{\left[k_0 + (1-\theta)k_s + b(1-\theta)k_g\right]\bar{y}}_{\text{perceived tax cost}}}.$$

The roles are not symmetric: two orientations set the size of a future cost and a future benefit, and the third sets how heavily both count.

(i) Zero-sum beliefs, θ, lower perceived tax costs by reducing the perceived responsiveness of the economic pie to taxation. High-zero-sum agents see redistribution as cheaper because they see less endogenous production to distort. This is the *generativity* content of zero-sum thinking, as distinct from its *appropriation* content (whether gains are taken from others), which would enter expected income rather than cost (Appendix B).

(ii) Uncertainty avoidance, λ, raises the insurance value of redistribution by making the worst plausible future income state worse.

(iii) The future weight, b, enters only in company with the other two: it multiplies λ in the numerator and the $(1-\theta)k_g$ term in the denominator. It therefore raises the value of insurance against future ambiguity and the weight on future tax-base costs at the same time, which is why its net effect is a trade-off rather than a direction. The two-sidedness has an empirical precedent: Bursztyn (2016) shows that poorer voters in Brazil choose immediate cash transfers over future-oriented public education spending, so how much the future counts, and what citizens believe converts present taxes into future resources, already govern observed redistributive choices.

The model therefore avoids the main weakness of a purely multiplicative formulation in which the cultural variables enter only as a product $\beta\theta\lambda$: each belief retains a separately identifiable role. It also improves on a static denominator formulation (the special case $k_g = 0$), in which the future weight is mechanically pro-redistribution whenever $\lambda > 0$.

2.5 Comparative Statics

Define $D(\theta, b) = k_0 + (1-\theta)k_s + b(1-\theta)k_g$ and $N_i(b) = A_i + b\lambda$, so that $\tau_i^\circ = N_i(b)/[D(\theta, b)\bar{y}]$ with $D > 0$. All statements concern interior solutions.

Proposition 1 (Zero-sum beliefs and perceived tax costs)**.** *For an interior solution,*

$$\frac{\partial \tau_i^\circ}{\partial \theta} = \frac{N_i(b)(k_s + bk_g)}{D(\theta, b)^2 \bar{y}}.$$

For agents with a positive net redistributive motive ($N_i(b) > 0$), stronger zero-sum beliefs weakly increase support for redistribution by lowering perceived tax costs, strictly so whenever $k_s + bk_g > 0$.

Proof. $\partial D/\partial\theta = -(k_s + bk_g)$; apply the quotient rule. The derivative vanishes in the degenerate case $k_s = k_g = 0$, in which the perceived cost is the administrative cost alone and zero-sum beliefs have nothing to act on. □

For agents with a negative net redistributive motive, the unconstrained solution is negative, so the constrained choice is the zero-tax corner and zero-sum beliefs leave the observed choice unchanged. The observable prediction is therefore asymmetric: zero-sum beliefs raise desired redistribution among supporters and do not move the extensive margin of support, which the extensive-margin result below states formally.

Proposition 2 (Uncertainty avoidance and ambiguity insurance)**.** *For an interior solution,*

$$\frac{\partial \tau_i^\circ}{\partial \lambda} = \frac{b}{D(\theta, b)\bar{y}} > 0 \quad \text{if } b > 0.$$

Uncertainty avoidance raises desired redistribution through the insurance motive, provided the agent places some weight on the future.

Proof. Differentiate (7) with respect to λ. □

Proposition 3 (The conditional effect of the future weight)**.** *For an interior solution,*

$$\frac{\partial \tau_i^\circ}{\partial b} = \frac{\lambda\,[k_0 + (1-\theta)k_s] - A_i(1-\theta)k_g}{D(\theta, b)^2 \bar{y}}.$$

For agents with $A_i > 0$, a higher future weight increases support for redistribution if and only if

$$\lambda > \lambda_i^*(\theta) \equiv \frac{A_i(1-\theta)k_g}{k_0 + (1-\theta)k_s}.$$

The threshold $\lambda_i^*(\theta)$ is strictly decreasing in θ whenever $A_i > 0$ and $k_g > 0$, and equals zero at $\theta = 1$.

Proof. Apply the quotient rule to (7): $\partial\tau_i^\circ/\partial b = [\lambda D - N_i(b)(1-\theta)k_g]/(D^2\bar{y})$, and the numerator simplifies to $\lambda[k_0 + (1-\theta)k_s] - A_i(1-\theta)k_g$ once the terms in $b\lambda(1-\theta)k_g$ cancel. Monotonicity of the threshold follows from $\partial\lambda_i^*/\partial\theta = -A_i k_g k_0/[k_0 + (1-\theta)k_s]^2 \leq 0$. □

Proposition 3 turns the future weight from a simple direct effect into a conditional force. A forward-looking agent can be anti-redistribution when she believes redistribution erodes the future tax base, and pro-redistribution when she sees redistribution as affordable insurance against future downside risk. The first case prevails when uncertainty avoidance is low and taxation is believed to carry future costs; the second when uncertainty avoidance is high or the economy is viewed as fixed-pie.

Remark 3 (What durability buys)**.** *Proposition 3 is the one result that requires the same rate to apply in both periods. With period-specific rates the problem separates into $\tau_1 = A_i/[K_s(\theta)\bar{y}]$ and $\tau_2 = (A_i + \lambda)/\{[K_s(\theta) + K_g(\theta)]\bar{y}\}$, and the future weight cancels from each first-order condition: it multiplies neither the insurance benefit nor the future tax-base cost, and the trade-off disappears. Propositions 1 and 2 survive instrument by instrument, since zero-sum beliefs still lower both denominators and uncertainty avoidance still raises the period-2 numerator.*

Remark 4 (Heterogeneous expected income growth)**.** *The baseline model holds relative income fixed across periods. Suppose instead that agent i expects relative income growth Δ_i, so that period-2 pre-tax income is $\alpha_i + \Delta_i$ with Δ_i mean zero in the population. The interior preferred tax becomes $\tau_i^\circ = [A_i + b(\lambda - \Delta_i)]/[D(\theta, b)\bar{y}]$: expected upward mobility enters exactly where the ambiguity bound does, with the opposite sign. It is negative insurance, in the spirit of the prospect-of-upward-mobility mechanism of Benabou and Ok (2001). Two implications follow. First, the effect of the future weight on the income gradient is no longer unambiguous: for income position p with $\alpha'(p) > 0$ and expected-growth gradient $\Delta'(p)$,*

$$\frac{\partial^2 \tau_i^\circ}{\partial b\, \partial p} = \frac{\alpha'(p)(1-\theta)k_g - \Delta'(p)\,[k_0 + (1-\theta)k_s]}{D(\theta, b)^2 \bar{y}},$$

which is positive in the baseline model ($\Delta' = 0$) but turns negative when expected relative income growth rises steeply enough with position. Second, the support threshold becomes $\alpha_i < \bar{y} + b(\lambda - \Delta_i)$, so expected mobility, unlike zero-sum beliefs, moves the extensive margin of support.

Proposition 4 (Extensive margin of support)**.** *The agent supports strictly positive redistribution, $\tau_i^* > 0$, if and only if $\alpha_i < \bar{y} + b\lambda$. The income threshold for support is weakly increasing in the future weight and in uncertainty avoidance, strictly so when the other future parameter is positive, and is independent of zero-sum beliefs and of all cost parameters.*

Proof. $\tau_i^\circ > 0$ iff $N_i(b) = A_i + b\lambda > 0$ iff $\alpha_i < \bar{y} + b\lambda$, since $D(\theta, b)\bar{y} > 0$. Neither θ nor (k_0, k_s, k_g) enters the condition. □

Remark 5 (Two meanings of zero-sum thinking)**.** *Proposition 4 separates the fixed-pie tax-cost reading of zero-sum thinking from an interpersonal-extraction reading, formalized in Appendix B,*[2] *in which θ enters expected future income itself. Under the tax-cost reading, θ scales the perceived cost of a policy the agent may or may not want, so it cannot create or destroy the desire for a positive tax; under extraction, below-mean agents expect to be extracted from, which moves the support threshold with θ and makes the zero-sum effect decline across the income distribution, turning oppositional among the well-off (Proposition 5 in Appendix B). The tax-cost reading concerns* generativity *and the extraction reading* appropriation*; the WVS item bundles both in its two anchors, the generative "wealth can grow so there is enough for everyone" and the appropriative "get rich at the expense of others," so no single item separates them, but where in the income distribution the effect operates does.*

2.6 The Trade-Off at Observed Country Values

Proposition 3 is substantive only if observed populations lie on both sides of the threshold $\lambda_i^*(\theta)$. Figure 1 checks this: it plots the model-implied preferred-tax index against the future weight, setting zero-sum beliefs and uncertainty avoidance to empirical percentiles from the 66-country WVS/EVS–Hofstede calibration panel, with the cost parameters fixed at scenario values. Those values are normalization choices, not estimates, so the figure illustrates the model's nonlinear region and makes no calibration claim.

Two features stand out. When uncertainty avoidance is at its lower quartile and zero-sum beliefs are weak, the preferred-tax index falls in the future weight: future tax-base costs dominate. When uncertainty avoidance is at its upper quartile, or when zero-sum beliefs are strong enough to make taxation appear cheap, the slope turns positive: ambiguity insurance dominates. Observed country values populate both regions, so the conditionality is not a theoretical curiosity confined to a corner of the parameter space.

[2]The extraction reading is the standard one in the zero-sum literature: Foster (1965)'s image of limited good casts one family's gain as another's loss, and Chinoy et al. (2026) define zero-sum thinking as the view that any gain by one party necessitates a loss by another, with support for redistribution following precisely because the wealth of some is seen to come at the expense of others.

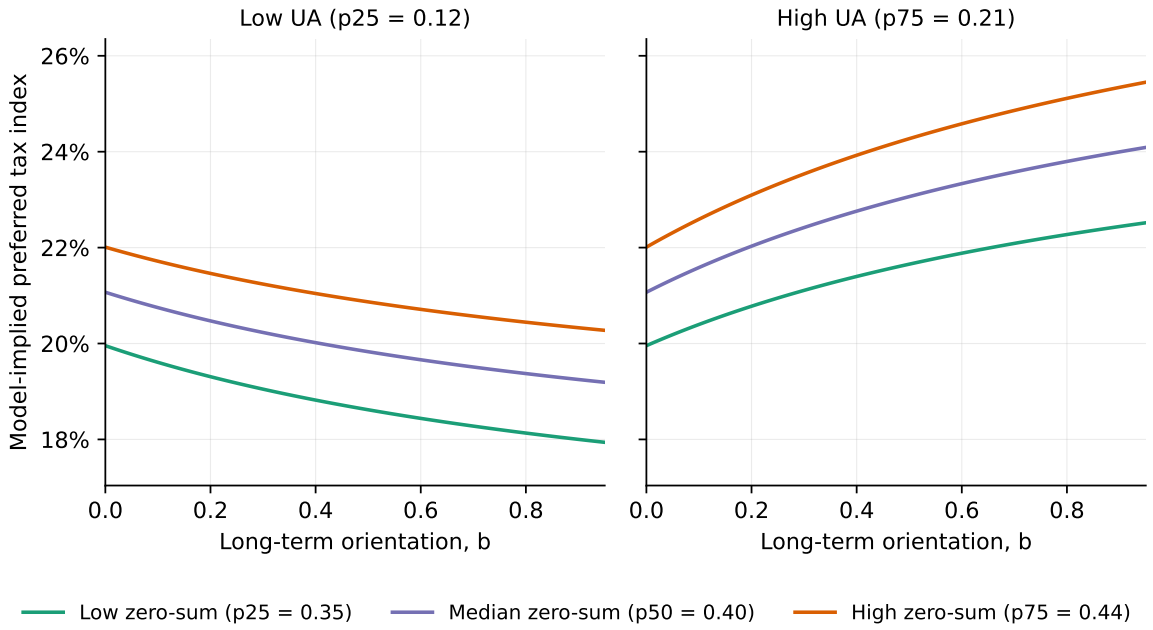


Figure 1: The conditional effect of the future weight at observed country values

Notes: Curves show the model-implied preferred-tax index (7) as a function of the future weight b. Zero-sum beliefs θ (line colors) are set to the 25th, 50th, and 75th percentiles and the uncertainty shock λ (panels) to the 25th and 75th percentiles of the 66-country core sample (WVS/EVS zero-sum beliefs, higher = stronger zero-sum views; Hofstede uncertainty avoidance mapped as $\lambda = 0.25\,\text{UAI}/100$). Cost parameters are scenario normalizations ($k_0 = 0.35$, $k_s = 1.00$, $k_g = 1.25$, $A = 0.20$), not estimates.

2.7 Predictions

The propositions map into predictions for the epidemiological design, in which ancestral-country measures proxy (θ, λ, b) and education proxies income position. Two are signed and testable here; the third is signed only conditionally and, as Section 7.3 shows, requires a measure of the future weight that this design does not have.

P1. Zero-sum beliefs (+). Ancestral zero-sum beliefs raise support for redistribution (Proposition 1).

P2. Uncertainty avoidance (+). Ancestral uncertainty avoidance raises support for redistribution (Proposition 2).

P3. The future weight (conditional). Its effect is negative where the perceived future growth cost $A_i(1-\theta)k_g$ exceeds the insurance benefit $\lambda[k_0+(1-\theta)k_s]$ and positive otherwise (Proposition 3). Testing this requires a credible ancestral measure of b, which is a measurement question taken up in Section 7.3.

Beyond direct effects, the tax-cost model and its extractive rival make opposite predictions about *where in the income distribution* the zero-sum effect operates. That contrast is the paper's

mechanism test (Section 5). The model also implies interactions among the traits themselves; those are products of ancestral country means identified off 70 ancestral clusters, and Section 7.2 shows the design can resolve only implausibly large ones.

Lemma 2 (Measurement bridge)**.** *Let the observed response be $R_i = h(\tau_i^*)$ for a strictly increasing link h, possibly nonlinear and possibly observed only through fixed ordinal cutpoints. Then, wherever τ_i° is interior:*

(i) for any argument $x \in \{\theta, \lambda, b, A_i\}$, R_i is increasing in x if and only if τ_i° is, so the directional predictions of Propositions 1–3 are preserved;

(ii) the support threshold of Proposition 4 is preserved: $R_i > h(0)$ if and only if $\alpha_i < \bar{y} + b\lambda$;

(iii) cross-partial signs are not preserved in general: for twice differentiable h,

$$\frac{\partial^2 R_i}{\partial x \, \partial z} = h'(\tau_i^\circ)\frac{\partial^2 \tau_i^\circ}{\partial x \, \partial z} + h''(\tau_i^\circ)\frac{\partial \tau_i^\circ}{\partial x}\frac{\partial \tau_i^\circ}{\partial z},$$

so a linear interaction coefficient in the observed response can differ in sign from the model's cross-partial unless h is affine.

Proof. (i) is composition with a strictly increasing function. (ii) follows because h is strictly increasing and $\tau_i^* > 0$ iff $\alpha_i < \bar{y} + b\lambda$ (Proposition 4). (iii) is the displayed chain-rule identity: the second term is nonzero whenever h is not affine and both first-order effects are nonzero. For an ordinal response generated by fixed cutpoints applied to a latent index strictly increasing in τ_i°, (i) and (ii) apply to the probabilities of upper response categories. □

Lemma 2 does more than license the linear specification. It tells us which tests survive the fact that we observe an ordinal agreement item rather than a tax rate, and which do not. Directional statements, including those within a subpopulation, are robust to any increasing link. Interaction coefficients are not. Section 5 builds its mechanism test out of the first kind for exactly this reason.

2.8 Discussion

The model connects three ideas usually treated separately. Redistribution is not only about current income: voters reason about future risks and future economic costs. Voters differ in their perceived economic mechanism: a high-zero-sum agent does not simply like redistribution more as a taste; she perceives fewer productive responses to distort, which is a statement about elasticities, not preferences. And the weight on the future is not inherently left or right.

This structure turns cultural variables into objects inside a public-finance trade-off: a future weight, an ambiguity bound, and a perceived tax-cost function. The elasticity reading of zero-sum beliefs also yields a sharp out-of-sample prediction. Randomized survey experiments show that information about inequality, taxation, and mobility moves redistributive preferences (Kuziemko et al., 2015; Alesina et al., 2018); in this model, messages about the efficiency costs of taxation should move growth-oriented respondents and leave fixed-pie believers unmoved, because the latter perceive no margin for taxes to distort. Heterogeneous treatment effects by measured zero-sum beliefs are the model's clearest experimental test. The model treats the orientations as inherited primitives; endogenizing zero-sum beliefs through motivated reasoning or experienced mobility, as the evidence in Chinoy et al. (2026) suggests, is a natural next model but would break the predetermined character of beliefs on which the empirical design relies.

3 Data and Empirical Methods

3.1 ESS Outcome and Baseline Sample

The individual-level data are European Social Survey rounds 2–11. The outcome is agreement with the statement that "the government should take measures to reduce differences in income levels" (`gincdif`), recoded so that higher values mean more support for redistribution. Respondents in the analysis sample were born in their country of residence and report parental countries of birth. Round 1 is unavailable for this design rather than excluded by choice: it records whether each parent was born in the country of residence but not which country, so no ancestral measure can be merged. Parental country of birth first appears in round 2 and is carried by a succession of variables across later rounds (`fbrncnt` in rounds 2–3, `fbrncnta` in 4–5, `fbrncntb` in 6–8, `fbrncntc` in 9–11), which the build harmonizes into a single ancestral country code.

The baseline sample imposes two restrictions. First, respondents must have exactly one foreign-born and one native-born parent, with both parental statuses observed; ancestral traits are assigned from the foreign-born parent's country of birth. This assigns ancestry unambiguously from a single parent and requires no averaging across two foreign ancestries. It also has a consequence that matters for interpretation: because parental intermarriage is itself an outcome of integration, the restriction holds destination-country minority status nearly fixed: 96.7 percent of baseline respondents do not report belonging to an ethnic minority, and 92 percent do not report belonging to a group that is discriminated against, against 80 percent on both counts among respondents with two foreign-born parents. That keeps the ancestral trait from being entangled with destination-country marginalization, and it is why we treat the restriction as a design choice rather than a cost.

It bears on external validity in particular. Every respondent in the design was born in the country of residence, and on the integration markers the ESS records this group sits close to the native majority. Estimates here therefore speak more directly to how inherited beliefs operate in European electorates at large than estimates drawn from households in which both parents are immigrants, where ancestral origin and minority status are confounded. The restriction remains a selection on intermarriage, and the estimand accordingly applies to an integrated part of the second generation; Section 4.2 sets out both what that buys and what it costs, since the zero-sum result behaves differently in precisely the subpopulation where minority status is common.

Second, respondents must reside in Western Europe: AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IS, IT, LU, MT, NL, NO, PT, or SE. This excludes post-socialist Eastern Europe and the Baltics as well as Russia, Ukraine, Turkey, and Israel as residence countries, so that destination contexts are comparable long-standing European democracies with developed welfare states. The estimation sample contains 10,718 respondents, 70 ancestral countries, and 19 residence countries (no eligible observation resides in Malta) across 154 country-by-round cells.

Table 1 summarizes the sample. Mean support for redistribution is 3.73 on the five-point scale with a standard deviation of 1.06, the yardstick for the magnitudes reported below. Because the ancestral list is the identifying variation, we report it in full: Table A1 gives every ancestral country with its trait values and respondent count, and Table A2 the residence-country composition. The ancestral distribution is unbalanced, in the way second-generation samples always are, and the degree matters for what follows. Germany, Italy, the United Kingdom, France, and Poland together supply 45 percent of respondents; the median ancestry contributes 32 respondents, the lower quartile 16, and the 35 smallest ancestries together account for only 5 percent of the sample. The effective

number of ancestral clusters is well below the nominal 70. Three implications follow and each is addressed below: the clustered standard errors are supplemented by a wild-cluster bootstrap, no result may rest on a single ancestry, and the ancestral trait means themselves must be shown to be precisely estimated in their source surveys.

Table 1: Descriptive statistics, baseline estimation sample

	Mean	S.D.	Min	Max	N
Outcome					
Support for redistribution (1–5)	3.73	1.06	1.00	5.00	10,718
Ancestral traits (native scales)					
Zero-sum belief, $11 - $ `e041` (1–10)	4.64	0.43	3.27	6.05	10,718
Uncertainty avoidance, Hofstede UAI (0–100)	65.7	20.6	8.0	99.0	10,718
Long-term orientation, Hofstede (0–100)	54.5	19.7	4.0	100.0	10,718
Individual characteristics					
Education, ISCED (1–7)	3.55	2.19	0.00	7.00	10,718
Age (years)	45.6	18.7	15.0	98.0	10,718
Female	0.518	0.500	0.000	1.000	10,718
Ancestral countries					70
Residence countries					19
Country-by-round cells					154
Respondents					10,718

Notes: Baseline sample of Table 2: ESS rounds 2–11, residence in Western Europe, exactly one foreign-born and one native-born parent, complete cases on the joint specification. Ancestral traits are shown on their native scales (the regressions use the standardized versions described in Section 3.2): zero-sum belief is the ancestral mean of $11 -$ `e041`, higher = stronger zero-sum belief; the two Hofstede indices run 0–100, higher = more of the named orientation. The outcome standard deviation of 1.06 is the yardstick for the magnitudes in Section 4.

3.2 Ancestral-Country Measures

The ancestral traits are measured at the level of the foreign-born parent's birth country. Zero-sum beliefs come from the WVS/EVS wealth-accumulation item (`e041`), a ten-point scale on which respondents place themselves between two stated poles: "people can only get rich at the expense of others" at one end and "wealth can grow so there is enough for everyone" at the other. We recode it as $11 -$ `e041` and average by country, so that higher values place a country nearer the fixed-pie pole. Because the scale is bipolar, a move toward that pole simultaneously endorses an appropriative statement and rejects a generative one, so a high score is consistent with either reading of zero-sum thinking in Remark 5. The item cannot separate them, which is why Section 5.1 discriminates between them by where in the income distribution the effect operates rather than by the wording of the measure. Uncertainty avoidance is Hofstede's uncertainty-avoidance index, and the candidate measure of the future weight is Hofstede's long-term-orientation (pragmatism) dimension; for both, higher values mean more of the named orientation (Hofstede, 2001).

The provenance of the second Hofstede dimension deserves a note. Long-term orientation began as the "Confucian dynamism" factor of the Chinese Value Survey and was available for barely twenty countries. The version used here is Minkov's, computed from World Values Survey items and available for 84 countries, which is what makes an ancestral design of this scale possible at

all. The label is not a time-preference measure: the dimension loads on thrift, perseverance, and adaptability to changed circumstances against respect for tradition and social obligation, which is why Hofstede renamed it *pragmatism*. Both poles carry a temporal orientation, one forward and one backward, and neither is a discount rate. Section 7.3 takes up which of them the estimates pick up, carries the question of what the index actually measures considerably further, and reaches a negative conclusion; readers who want that answer before seeing the estimates should read it first. The three ancestral measures are empirically distinct: pairwise correlations across the 70 ancestries never exceed 0.06 in absolute value, and at the respondent level the largest is 0.18 (Table A5).

Scale of the ancestral measures. The traits enter standardized, and the normalization is a source-file one that we state precisely because it governs how the coefficients should be read. Each measure is standardized across the ancestral-country universe of its source (one observation per country, equal weight) before being merged to respondents. It is therefore *not* standardized within the estimation sample, and the two differ: in the baseline sample the standardized zero-sum measure has a standard deviation of 0.68, uncertainty avoidance 0.96, and long-term orientation 0.84. Education, by contrast, is standardized over all round-2 and later respondents in the analysis file, not within the baseline, where its standard deviation is 1.06 rather than one. A coefficient is thus the effect of a one-unit move on the source-country scale, and to compare traits with education on a like-for-like footing the trait coefficients must be multiplied by the standard deviations just given. Section 4 reports magnitudes both ways.

Education is the socioeconomic-position control: the harmonized ISCED level, standardized. We use education in place of reported household income as the measure of permanent income because the ESS income measure is missing in several rounds and changes its category definitions, whereas education is observed consistently in every round and is a strong predictor of permanent income. Education is itself potentially influenced by ancestral traits, so it is a proxy for position, not a causal control; the income-gradient specifications of Table A9 therefore report a predetermined alternative, the mean education of the respondent's parents, alongside own education.

3.3 Reduced-Form Specification

The reduced-form model is

$$y_{iacr} = \beta_\theta \, \mathrm{ZS}_a + \beta_\lambda \, \mathrm{UA}_a + \beta_b \, \mathrm{LTO}_a + \gamma_e \, \mathrm{Educ}_i + \gamma' X_i + \delta_{cr} + u_{iacr}, \tag{8}$$

where i indexes respondents, a the foreign-born parent's birth country, c the residence country, and r the ESS round. X_i contains age, age squared, and gender, u_{iacr} is the disturbance, and δ_{cr} are country-of-residence-by-round fixed effects, so all identifying comparisons are within a destination country and survey round. Standard errors are clustered by ancestral country, the level at which the traits vary. We estimate (8) with each ancestral trait entered alone and with all three entered jointly, on the identical estimation sample. Ancestral GDP per capita is deliberately excluded from the baseline (it is itself an outcome of the traits and absorbs their variation); it returns in robustness.

Identification requires two conditions, stated here once. The country-by-round fixed effects remove the institutional and macroeconomic conditions shared by respondents in the same destination and survey wave. Identification additionally requires that, conditional on those fixed effects and the pre-treatment covariates, ancestral-country traits are independent of unobserved determinants of redistributive preferences, and that the traits do not merely proxy for other transmitted origin-country attributes. The ancestral development, institutional, and regional controls of Section 4.1

support the second condition against its obvious candidates, though subtler ancestral factors cannot be excluded. Under these assumptions the coefficients are effects of ancestral exposure, and we use "effect" in that sense throughout.

One further issue is specific to epidemiological designs. Emigrants are not random draws from their origin country, and the ancestral means are measured on those who stayed. The WVS/EVS surveys run from 1981 to 2014 and the Hofstede indices were built between the 1970s and the 2000s, in most cases after the respondent's foreign-born parent had already left. If emigrants were a noisy but unbiased draw the consequence would be attenuation and the estimates below would understate the associations; the harder case is systematic selection that varies across origins in a way correlated with redistributive preferences in Europe. Three diagnostics show the estimates are not driven by measured ancestral development, by a single migration history, or by self-reported minority status: the ancestral controls of Section 4.1 move the coefficients by at most 0.004, the leave-one-ancestry-out results show no origin carries the findings, and the baseline sample is the least selected part of the second generation on the dimension where selection is most visible. They do not eliminate unobserved migrant or intermarriage selection, and readers should take the ancestral measures as proxies for a transmitted belief rather than exact readings of it.

3.4 Semi-Structural Methods: NLS and GMM

The reduced form tests directions. To connect magnitudes to the model, we also estimate the preferred-tax equation (7) directly. The measurement bridge maps observables into model objects,

$$\theta_i = \frac{\text{ZeroSum}_{a(i)} - 1}{9}, \qquad b_i = \frac{\text{LTO}_{a(i)}}{100}, \qquad \lambda_i = \lambda_{\max} \frac{\text{UAI}_{a(i)}}{100}, \tag{9}$$

and the position term uses education,

$$A_i = A_0 + A_e \left(\frac{3.5 - \text{ISCED}_i}{7} \right), \tag{10}$$

so that lower education raises the current redistributive motive under the maintained lifetime-income interpretation. Education therefore moves from the control vector into the structural numerator. The outcome is the ESS item rescaled to the unit interval, $y_i = (\text{Redistribution}_i - 1)/4 \in \{0, 0.25, 0.5, 0.75, 1\}$, treated as a cardinal proxy for the latent preferred-tax index. This cardinality assumption is stronger than the monotonicity the reduced form requires, which is why we call the estimates semi-structural. The mapping of the Hofstede dimension to b carries the qualification of Section 7.3: it is what the estimator needs to fill the b slot, not a claim that the index is an estimated discount factor.

Nonlinear least squares. With k_0 normalized to 0.35 (a scale normalization: the survey outcome cannot separately identify a free measurement scale and all numerator and denominator parameters), the parameter vector $\psi = (A_0, A_e, \lambda_{\max}, k_s, k_g)$ is estimated by minimizing the mean squared residual of $e_i(\psi) = y_i - \tau_i^*(\psi) - X_i'\gamma - \delta_{c(i)r(i)}$, where the controls X_i (age, age squared, gender, marital status, labor market status) and the country-by-round effects are profiled out at each nonlinear iteration. The fixed effects are normalized to sample-weighted mean zero, preserving the grand mean that identifies the level of the tax index. All five parameters are constrained to be at least 0.001, a numerical floor, not a substantive restriction; the two estimates that matter for the model, $\lambda_{\max}$ and k_g, sit well above it. Inference differs by table and we state it once here. The

parameter estimates of Table 7 carry ancestry-clustered sandwich standard errors with normal-reference significance levels. Only the average marginal effects of Table 8 carry cluster-bootstrap intervals, which resample the 70 ancestries and refit the model on each draw; those are the preferred summary because the marginal effects, not the individual cost parameters, are what the exercise identifies well.

Two-step GMM. The GMM estimator uses ten moments: a constant and residualized functions of θ_i, b_i, UAI, education position, and their theory-relevant products, with an ancestry-cluster-robust second-step weighting matrix. Because the moments are transformations of the model inputs rather than excluded instruments, GMM serves as a minimum-distance and specification check (via the J test), not as a solution to any remaining endogeneity of inherited culture, and agreement between NLS and GMM is not independent validation of the mechanism.

4 Reduced-Form Results

Table 2 reports the baseline estimates of (8). Columns (1)–(3) enter one ancestral trait at a time; column (4) enters all three jointly. All columns use the identical sample of 10,718 respondents.

Table 2: Ancestral main effects on support for redistribution: Western Europe, one foreign-born and one native-born parent

	(1)	(2)	(3)	(4)
Zero-sum beliefs	0.040**			0.042***
	(0.015)			(0.013)
Uncertainty avoidance		0.035***		0.038***
		(0.010)		(0.010)
Long-term orientation			-0.020	-0.025**
			(0.013)	(0.011)
Education	-0.103***	-0.102***	-0.102***	-0.101***
	(0.015)	(0.016)	(0.016)	(0.016)
Observations		10,718		
Ancestry clusters		70		
Residence countries		19		
Country-by-round cells		154		

Notes: The outcome is support for government reduction of income differences, recoded so higher values mean more support. Zero-sum beliefs are the ancestral WVS/EVS wealth item recoded as $11 - \texttt{e041}$ (higher = stronger zero-sum belief); uncertainty avoidance and long-term orientation are the Hofstede indices (higher = more of the orientation). The three ancestral measures are standardized across the ancestral-country universe of their source (one observation per country, equal weight); in the estimation sample their standard deviations are 0.68, 0.96, and 0.84 respectively. Education (ISCED) is standardized over all round-2 and later respondents in the analysis file, where its standard deviation within the baseline is 1.06. Each column includes age, age squared, gender, and country-of-residence-by-round fixed effects. The sample requires exactly one foreign-born and one native-born parent (both statuses observed) and residence in AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IS, IT, LU, MT, NL, NO, PT, or SE; all four columns use the identical complete-case sample. Ancestral traits are assigned from the foreign-born parent's birth country. Standard errors, in parentheses, are clustered by ancestral country; wild-cluster bootstrap p-values are in Table A3. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Ancestral zero-sum beliefs raise support: 0.040 (s.e. 0.015) alone and 0.042 (0.013) jointly, consistent with P1. Ancestral uncertainty avoidance likewise raises support: 0.035 (0.010) alone and 0.038 (0.010) jointly, consistent with P2. Ancestral long-term orientation enters negatively, −0.020 (0.013) alone and −0.025 (0.011) jointly; Section 7.3 takes up what that coefficient is measuring. The similarity of each trait's coefficient across stepwise and joint specifications indicates that the ancestral measures are not relabeling one another.

Magnitudes. On the source-country scale of Section 3.2, a one-unit move raises support by 0.042 points for zero-sum beliefs and 0.038 for uncertainty avoidance. Converted to the standard deviation of the variation the design actually exploits (the spread across the ancestries being compared), the effects are 0.029 and 0.037 points, or 2.7 and 3.5 percent of the outcome standard deviation of 1.06. The education gradient is −0.101 on that scale, which is −10.1 percent of the outcome standard deviation once put on the same in-sample basis. The two belief channels are therefore between a quarter and two-fifths of the education gradient, which is the natural in-sample yardstick. The education effect is itself the income-position gradient that the motive term A_i in (7) requires: higher permanent income lowers the current redistributive motive.

Per-standard-deviation effects understate what the design shows, because the ancestral distribution is wide. Combining the two belief channels into the fitted ancestral index implied by column (4), the index has a standard deviation of 4.1 percent of the outcome standard deviation across respondents; the gap between the tenth and ninetieth percentile ancestries is 9.6 percent of an outcome standard deviation, and 23 percent between the extreme ancestries, Chinese and Portuguese. Two respondents living in the same country, surveyed in the same round, with the same education, age, and gender, differ by a quarter of a standard deviation in redistributive support if their foreign-born parents came from opposite ends of the inherited-belief distribution. That is what the inherited component amounts to in this sample.

4.1 Inference and Robustness

Inference survives the small-cluster correction. Because all identifying variation is at the level of 70 ancestral countries, Table A3 reports wild-cluster bootstrap p-values (Rademacher weights, null imposed, 9,999 replications) for every coefficient of Table 2. In the joint specification they are 0.005 for zero-sum beliefs and 0.003 for uncertainty avoidance; the two belief channels are unambiguous. Long-term orientation is 0.061: negative and imprecisely estimated under wild-cluster inference.

The estimates are robust to ancestral-country controls (Table 3): adding standardized ancestral log GDP per capita (PWT, 2000), ancestral rule of law (WBGI), and ancestral-region fixed effects, alone and together, moves the three trait coefficients by at most 0.004, and the development controls themselves enter small and insignificant. They are also robust to the cardinality assumption. Re-estimating the joint specification as an ordered logit with the same fixed effects and clustering reproduces the pattern (Table A8): the average marginal effect of a one-unit stronger ancestral zero-sum belief on the probability of answering "agree strongly" is 1.2 percentage points, with 1.3 for uncertainty avoidance, against a baseline of 25 percent.

Table 3: Robustness to ancestral development, institutions, and region

	(1)	(2)	(3)	(4)
Zero-sum beliefs	0.042***	0.042***	0.040***	0.042***
	(0.013)	(0.013)	(0.013)	(0.013)
Uncertainty avoidance	0.038***	0.039***	0.042***	0.036***
	(0.010)	(0.011)	(0.011)	(0.011)
Long-term orientation	-0.025**	-0.027**	-0.028**	-0.024*
	(0.012)	(0.013)	(0.013)	(0.014)
Education	-0.101***	-0.100***	-0.100***	-0.100***
	(0.016)	(0.016)	(0.016)	(0.016)
Ancestral log GDP pc		0.013	-0.009	-0.066
		(0.019)	(0.032)	(0.060)
Ancestral rule of law			0.019	0.021
			(0.027)	(0.027)
Ancestral-region FE	No	No	No	Yes
Observations	10,606	10,606	10,606	10,606
Ancestry clusters	69	69	69	69

Notes: Baseline joint specification (Table 2, column 4) on the constant subsample with nonmissing ancestral controls. Ancestral log GDP per capita is standardized log PWT real GDP per capita in 2000; ancestral rule of law is the standardized WBGI estimate; ancestral region fixed effects use the QoG `ht_region` classification. All columns include education, age, age squared, gender, and country-of-residence-by-round fixed effects; standard errors in parentheses are clustered by ancestral country. Coding directions as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A further concern is that the traits stand in for other inherited attitudes the literature has linked to redistributive preferences. Table 4 races the joint specification against eleven rivals, each entered one at a time as an ancestral-country mean: the four remaining Hofstede dimensions, the belief that success reflects luck rather than effort (`e040`), social trust (from the WVS item and from the Global Preferences Survey, Falk et al., 2018), attitudes toward competition, and the GPS risk-taking, altruism, and reciprocity measures. The three factors keep their sign in every race and remain significant at conventional levels throughout, with zero-sum beliefs the most stable, moving within $[+0.038, +0.042]$. Most rivals add no independent signal. Two nulls are worth naming. The luck-versus-effort belief enters at essentially zero, so the zero-sum coefficient reflects the fixed-pie wealth belief specifically rather than a broad fatalism. And inherited trust, the best-established ancestral trait in this literature (Algan and Cahuc, 2010), likewise adds nothing conditional on the three factors: it carries real explanatory power for growth and institutions, but not for redistributive demand specifically. The one rival that is individually significant is GPS risk-taking (+0.024), which leaves the three factors intact; it captures ordinary risk tolerance, a channel the model keeps distinct from ambiguity-driven uncertainty avoidance. That association is partial rather than standalone. The two measures correlate -0.33 across ancestries, and removing the Hofstede index from the same specification roughly halves the risk-taking coefficient, to +0.011 (0.010). Risk taking adds something on top of ancestral uncertainty avoidance; it does not stand in for it, which is why Section 6 finds it cannot take over the uncertainty slot of the structural model.

Table 4: Horse race: the three ancestral factors versus rival inherited attitudes

Rival added (one at a time)	Rival	Zero-sum	Unc. avoid.	Long-term	N	Anc.
(baseline, no rival)	—	+0.042***	+0.038***	-0.025**	10,718	70
		(0.013)	(0.010)	(0.011)		
Luck vs effort (e040)	+0.002	+0.042***	+0.037***	-0.025**	10,718	70
	(0.010)	(0.013)	(0.010)	(0.012)		
Individualism (Hofstede)	+0.008	+0.042***	+0.040***	-0.027**	10,718	70
	(0.012)	(0.013)	(0.011)	(0.012)		
Power distance (Hofstede)	+0.006	+0.040***	+0.035**	-0.023*	10,718	70
	(0.014)	(0.015)	(0.014)	(0.012)		
Masculinity (Hofstede)	-0.002	+0.040***	+0.038***	-0.024**	10,718	70
	(0.011)	(0.014)	(0.010)	(0.012)		
Indulgence (Hofstede)	+0.019	+0.042***	+0.047***	-0.021*	10,685	69
	(0.012)	(0.012)	(0.011)	(0.012)		
Social trust (WVS)	+0.011	+0.042***	+0.046***	-0.027**	10,718	70
	(0.011)	(0.013)	(0.012)	(0.012)		
Competition good/harm (WVS)	+0.010	+0.041***	+0.035***	-0.025**	10,718	70
	(0.008)	(0.013)	(0.010)	(0.011)		
Social trust (GPS)	-0.017	+0.038***	+0.023*	-0.037***	9,356	47
	(0.012)	(0.014)	(0.013)	(0.013)		
Risk taking (GPS)	+0.024**	+0.041***	+0.044***	-0.029**	9,356	47
	(0.010)	(0.013)	(0.012)	(0.012)		
Altruism (GPS)	-0.002	+0.040**	+0.031**	-0.031**	9,356	47
	(0.010)	(0.015)	(0.013)	(0.013)		
Neg. reciprocity (GPS)	-0.003	+0.041***	+0.033**	-0.031**	9,356	47
	(0.009)	(0.013)	(0.013)	(0.013)		

Notes: Each row adds one rival ancestral-country attitude to the joint specification of Table 2, column 4 (baseline design: Western European residence, one foreign-born and one native-born parent; country-by-round fixed effects; standard errors clustered by ancestral country). The three factors are on the paper's standardized scale; each rival is standardized within its estimation sample. "Zero-sum," "Unc. avoid.," and "Long-term" are the coefficients on the three ancestral factors in that race, and "Rival" the added trait's own coefficient. Higher values denote more zero-sum belief, more uncertainty avoidance, more long-term orientation, and more of the named rival (the WVS trust row reverses `a165` so that higher = more trust). N is respondents and "Anc." the number of ancestral clusters; the Global Preferences Survey measures cover far fewer ancestries (47) than the WVS and Hofstede measures (up to 70), so those rows rest on a smaller sample. The GPS patience measure is not raced here; it is a candidate measure of the paper's own third construct rather than a rival, and is treated in Table A10. Adding any rival moves the three-factor coefficients by at most 0.012 relative to the same sample estimated without the rival. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

One rival deserves separate treatment because it is the construct most often proposed as an alternative to uncertainty avoidance. Societal *tightness*, the strength of social norms and intolerance of deviance, is the live competitor to Hofstede's index, and if ancestral UA were really measuring tightness the second result would be a relabelling. Table A14 adds it, using Gelfand's six-item scale as administered in the International Study of Metanorms across 57 societies in 2019–20 (Eriksson et al., 2021). The two constructs are distinct in the direction that matters: across the covered ancestries tightness correlates -0.35 with the uncertainty-avoidance index, so this is a genuine rival. Adding it changes nothing: the uncertainty-avoidance coefficient moves from 0.033 to 0.033 on the covered subsample of 8,384 respondents, the other factors are unchanged to three decimals, and tightness itself enters at 0.000 with a standard error of 0.014, an interval tight enough to

exclude effects above roughly two-thirds of the zero-sum coefficient. A second discriminant check, reported with the semi-structural estimates in Section 6, puts the behavioral risk-taking measure of Falk et al. (2018) into the model's ambiguity slot in place of the Hofstede index; it carries none of the effect, which is what Remark 1 predicts of a channel driven by concern about unresolved states rather than by tolerance of known risk. The channel also has a direct individual-level counterpart. Two round-3 ESS items ask how worried respondents are that their income in old age will be inadequate and that they will not be able to retire when they would like: worry about adverse future states, which is what λ is. Ancestral uncertainty avoidance predicts both, at +0.064 (s.e. 0.030) and +0.081 (0.026) with the sample restrictions of Section 3.1 relaxed to gain power on a single-round module, and it does not predict saving for old age (−0.027, 0.028), which the future-weight measures do. The two channels separate on the outcomes the model says should separate them.

Splitting by the foreign-born parent and by respondent gender (Table A6) leaves zero-sum beliefs and uncertainty avoidance stable: all eight coefficients are positive, of similar magnitude, and significant at conventional levels. The long-term-orientation point estimate is larger through foreign-born fathers (−0.039, s.e. 0.015) than mothers (−0.020, s.e. 0.017), and larger still for father–son pairs (−0.054, s.e. 0.027), but the cross-cell differences are imprecise and we do not read them as establishing same-sex transmission. Parent-specific transmission of economic preferences is nonetheless well documented (Dohmen et al., 2012), and which parent carries a trait appears to depend on the trait: Ljunge (2014), studying inherited trust in exactly this design, finds transmission significantly stronger on the mother's side.

Finally, no single ancestry or destination country drives the results (Table A7). Across 70 leave-one-ancestry-out and 19 leave-one-residence-country-out re-estimations, every coefficient keeps its sign in every run; zero-sum beliefs and uncertainty avoidance remain significant at five percent in all 89 runs. Five of the six trait-by-dimension coefficient ranges span at most 0.015. The exception is uncertainty avoidance across residence countries, ranging from 0.024 to 0.043: excluding Germany, which supplies 14 percent of the sample, lowers it to 0.024 (s.e. 0.009), a third below baseline but still significant at five percent, while excluding Switzerland raises it to 0.043. The two largest German-speaking destinations trade influence between the two positively signed channels, but no exclusion overturns a sign or a qualitative conclusion.

4.2 Parental Composition and the Boundary of the Zero-Sum Result

The results above establish that the two belief channels are real. A separate question is who they are about. The baseline restricts attention to respondents with exactly one foreign-born parent, and relaxing that restriction changes exactly one of the three results. Table A4 estimates the joint specification separately by parental composition in the same residence countries. Uncertainty avoidance and long-term orientation are stable: 0.055 and −0.042 among respondents with two foreign-born parents against 0.038 and −0.025 in the baseline, with the differences far from significant. Zero-sum beliefs are not. The coefficient is −0.039 (s.e. 0.020) among the two-immigrant-parent respondents against +0.042 (0.013) in the baseline, and in a pooled model over both compositions ($N = 14{,}939$) interacting all three traits with a two-parent indicator the zero-sum differential is −0.071 (0.031), $p = 0.025$, while the other two differentials are insignificant ($p = 0.69$ and 0.29).

Two explanations can be ruled out. The reversal is not an artifact of averaging two ancestries: it is present, slightly stronger, among two-immigrant-parent respondents whose parents come from the *same* country, where no averaging occurs (−0.041, s.e. 0.021, $N = 3{,}718$), while the mixed-origin group in which averaging actually happens is a noisy zero. Nor is it the work of particular cells:

the coefficient is negative in all 66 leave-one-ancestry-out and all 19 leave-one-residence-country-out re-estimations, and dropping the large non-European origins that dominate the two-parent sample (Turkey, Morocco, Pakistan and the South Asian origins) leaves it between −0.037 and −0.039 throughout.

Roughly half of the gap is that these are different populations. Having one native-born parent is itself the product of intermarriage, and intermarriage rates differ enormously across origin groups: Turkey supplies 17.9 percent of the two-parent sample and 2.3 percent of the baseline; Morocco 7.9 against 2.0 percent; Germany 5.1 against 14.2 percent. Restricting both samples to the 66 ancestries they share changes nothing on its own, leaving the baseline at +0.041; reweighting that sample to the two-parent ancestry distribution moves it to +0.021, about half the gap. This is a matter of weights rather than of support. Part of the remainder plausibly tracks minority status in the destination. Respondents with two foreign-born parents are six times as likely to report belonging to an ethnic minority (19.9 against 3.3 percent), and within the two-parent sample the negative coefficient is larger among minority respondents (−0.071, s.e. 0.055) than among the rest (−0.017, s.e. 0.022), against +0.059 (0.015) for non-minority baseline respondents, a difference in point estimates that the minority subsample of 666 respondents is too small to resolve. Controlling for the same indicators on the pooled sample points the same way. The gap between the two samples, measured as the interaction of zero-sum belief with an indicator for two foreign-born parents, attenuates from −0.071 (s.e. 0.031) to −0.067 (0.036) once minority and discrimination indicators are added, and to −0.062 (0.039) once both are also interacted with zero-sum belief. Part of that movement is a sample-size effect rather than mediation, since neither item is fielded in every round and the pooled sample falls from 14,939 to 11,965.

We draw a bounded conclusion. For respondents from origin groups where intermarriage is rare and minority status is salient, the ancestral zero-sum mean plausibly marks membership of a group facing exclusion in the destination as much as it proxies an inherited belief about the elasticity of the pie, and there is no reason the model's logic should govern the first of those. About a third of the gap is left unexplained by composition and minority status together. Two implications follow, pulling in opposite directions. The result strengthens the case for the baseline design, which holds minority status nearly fixed. It also bounds the zero-sum finding to populations where minority status is not salient. The uncertainty-avoidance result carries no such boundary.

5 Which Belief Is Zero-Sum Thinking?

The model's mechanism claim is that zero-sum thinking operates on the perceived price of redistribution rather than on expected income. Its rival, formalized in Appendix B, is that zero-sum believers expect to be extracted from. The two agree on direct effects, so the direct-effects evidence of Section 4 cannot separate them. They disagree about *where in the income distribution* the zero-sum effect operates, and that is what this section uses.

5.1 The Cutpoint Profile and the Two Meanings

Table 5 sets out what each reading implies and what the data show. The two blocks do different work. On the margins in the first block the readings agree, which is why the headline estimates of Section 4 cannot separate them however precisely they are measured: both models make stronger zero-sum belief raise desired redistribution, and both make the income gradient steeper. The adjudication has to come from the second block, where the readings make opposite predictions.

Under the tax-cost reading θ scales the perceived cost of a policy the agent may or may not want, so its effect is positive at every income position among supporters and it cannot move the threshold at which desired redistribution becomes positive (Proposition 4). Under extraction $\partial\tau^{\text{ext}}/\partial\theta = bA_i/(k_0\bar{y})$ declines across the income distribution and turns negative above the mean (Proposition 5 in Appendix B).

Table 5: The two readings of zero-sum thinking: predictions and evidence

Margin	Generative (tax cost)	Appropriative (extraction)	Observed
Both readings make the same prediction			
Direct effect of θ	+	+	0.042^{***} (0.013)
Income-position gradient	steeper	steeper	0.005 (0.010)
The readings diverge			
θ among high-position respondents	+ at every position	− above the mean	0.006^{**} to 0.020^{**}, positive at all four cutpoints
Support threshold in θ	independent of θ	decreasing in θ	0.002 (0.003), on a 97 percent base rate

Notes: Predictions are those of the tax-cost model of Section 2 and of the extractive variant of Appendix B, from Propositions 1, 4 and 5. Observed estimates are the zero-sum coefficients from Table 2 column 4, from the own-education column of Table A9, and from the two panels of Table 6; standard errors clustered by ancestral country. Neither model's income-gradient prediction is borne out, and by Lemma 2(iii) an interaction coefficient in an ordinal response need not carry the sign of either model's cross-partial, which is why the argument rests on the directional comparisons of the second block. The support-threshold row is reported for completeness: with 97 percent of respondents already above the lowest cutpoint the design cannot resolve movement there for any regressor. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 6 estimates linear probability models of crossing each cutpoint of the five-category item, in the full sample and among above-median-education respondents.

Table 6: Cutpoint profile of the ancestral effects

	$P(y \geq 2)$	$P(y \geq 3)$	$P(y \geq 4)$	$P(y \geq 5)$
Panel A: full baseline sample				
Zero-sum beliefs	0.002	0.011**	0.012**	0.016**
	(0.003)	(0.005)	(0.006)	(0.006)
Uncertainty avoidance	0.002	0.005	0.017***	0.014***
	(0.002)	(0.004)	(0.005)	(0.005)
Long-term orientation	-0.003*	-0.006	-0.011**	-0.004
	(0.002)	(0.004)	(0.005)	(0.006)
Education	-0.006***	-0.036***	-0.038***	-0.021***
	(0.002)	(0.005)	(0.007)	(0.005)
Mean of outcome	0.97	0.84	0.67	0.25
Panel B: above-median education				
Zero-sum beliefs	0.006**	0.016**	0.010	0.020**
	(0.003)	(0.007)	(0.009)	(0.007)
Uncertainty avoidance	0.003	-0.000	0.009	0.007
	(0.003)	(0.005)	(0.007)	(0.007)
Long-term orientation	-0.003	-0.014**	-0.018**	-0.008
	(0.002)	(0.006)	(0.007)	(0.008)
Mean of outcome	0.97	0.82	0.65	0.23
Observations (A / B)		10,718 / 5,163		
Ancestry clusters (A / B)		70 / 70		

Notes: Linear probability models of crossing each cutpoint of the five-category item, with the covariates and fixed effects of Table 2, column 4, estimated on the full baseline sample (Panel A) and among above-median-education respondents (Panel B). Panel B contains no position interactions, so its coefficients are effects within that subsample. Under the fixed-pie tax-cost reading of zero-sum thinking the zero-sum effect should be positive throughout and concentrated at intensity margins; under the interpersonal-extraction reading it should decline across position and be negative among high-position respondents. Standard errors in parentheses are clustered by ancestral country. Coding directions as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Panel A gives the tax-cost profile. At the lowest cutpoint, $P(y \geq 2)$, the zero-sum effect is 0.002 (s.e. 0.003), essentially zero, and it then grows monotonically along the intensity margins: 0.011 at $P(y \geq 3)$, 0.012 at $P(y \geq 4)$, and 0.016 at $P(y \geq 5)$, all significant. Zero-sum beliefs operate on how much redistribution supporters want, not on whether one supports it.

That pattern is consistent with the tax-cost mechanism, but it cannot carry the argument on its own, for two reasons. The first is power. With 97 percent of respondents already above the lowest cutpoint, the design has very little room to detect movement in *any* regressor there. The cleanest benchmark is education, comfortably the strongest predictor in the paper: at this cutpoint the education gradient is itself only -0.006, and the smallest ancestral effect detectable at eighty percent power is 0.007, larger than the entire education gradient. At the next cutpoint the same calculation gives 0.014 against an education gradient of -0.036. The lowest cutpoint is not a demanding test that a variable failed; it is a test this sample cannot run, which is also why we do not claim support for Proposition 4's positive predictions for uncertainty avoidance and the future weight, both of which are indistinguishable from zero there.

The second reason is that the bottom category of an agreement item is not an observed choice

of zero redistribution. A respondent who prefers a positive but modest rate may well disagree that government should reduce income differences. Lemma 2(ii) preserves the structural support threshold only if we know which observed category corresponds to $h(0)$, and we do not. The semi-structural estimates of Section 6 make the point concrete: with $A_0 = 0.263$ and $A_e = 0.029$, the fitted redistributive motive A_i lies between 0.249 and 0.278 across the entire education range, so the fitted model places *every* respondent strictly inside the region where positive redistribution is preferred. The structural corner is not in the sample.

The discriminating evidence is therefore the position profile, and it is the kind of evidence Lemma 2 protects. Directional statements, including directional statements within a subpopulation, survive any strictly increasing link between the preferred rate and the observed response; interaction coefficients do not, which is why we rest the comparison on the former and report the trait-by-position interactions in the appendix (Table A9) without building on them. Extraction predicts that the zero-sum effect declines with position and becomes negative among the well-off. Among above-median-education respondents it does neither: the zero-sum effect is positive at all four cutpoints and significant at three of them, running 0.006 at the lowest cutpoint and 0.020 at "agree strongly" (the intermediate cutpoints are 0.016 and 0.010, so the profile rises overall without being monotone). It is if anything stronger in this group than in the full sample, not weaker. The sign that would favour extraction does not appear at any cutpoint, in either panel, or under any placement of the corner we have tried.[3] One caveat on reading magnitudes: because education is standardized in the full sample, the coefficient in a specification that also contains position interactions is evaluated at the full-sample mean, which lies at the lower edge of the above-median group. The cutpoint estimates in Panel B of Table 6 contain no such interactions and are free of this ambiguity, which is why the comparison rests on them.

On this evidence the fixed-pie tax-cost reading is the better description of what the WVS item transmits. What the design shows is a monotonicity comparison rather than a knife-edge test: the effect does not decline across position and does not turn negative where extraction says it must.

A cross-country check settles a narrower question about the measure itself. One worry is that the wealth item records not a belief but an *experience*, of living under predatory or corrupt institutions, in which case its country-level values should track institutional quality. They do not. Across the WVS/EVS countries the zero-sum mean is essentially uncorrelated with control of corruption ($r = 0.07$), rule of law (0.05), government effectiveness (0.03), and the Corruption Perceptions Index (0.03). What it tracks is worldview: it is higher where citizens report little freedom of choice and control over their lives (WVS `a173`, $r = -0.39$) and where others are seen as taking advantage rather than dealing fairly (-0.32; Table A22). The two correlations point in different directions. Low perceived control is close to what the generativity reading requires: if effort is not seen to change outcomes, output is not seen to respond to effort. Low interpersonal trust is closer to the appropriation reading. The exercise establishes the narrow point, that the measure behaves like a belief about the game rather than a report on extractive institutions, and not the broad one, which is why the mechanism argument rests on the position evidence above.

[3]Suppose "disagree" also signals a preferred rate of zero, so the support margin becomes $P(y \geq 3)$ with a base rate of 84 rather than 97 percent. The zero-sum effect there is 0.011 (s.e. 0.005), and among above-median-education respondents 0.016 (0.007). Neither is the exact zero the tax-cost reading predicts at a true corner (under this mapping they are better read as intensity effects), but Proposition 5 in Appendix B requires a *negative* estimate at the support margin, and no candidate placement delivers one.

5.2 Placebo Outcomes

The model gives zero-sum beliefs a specific job, lowering the perceived price of redistribution, so they should predict redistributive demand and not attitudes toward government in general. Table A13 estimates the joint specification on two outcomes the mechanism has no claim on, standardized for comparability: support for further European unification and trust in the police. The zero-sum effect on redistribution (0.039 in standardized units) has no counterpart on either (-0.006 and -0.001). Uncertainty avoidance does not load on unification and enters police trust *negatively*, the opposite of what a diffuse pro-state disposition would produce. These are the two outcomes on which a generalized pro-government disposition would show up most clearly, and it does not, though two placebos cannot establish that the traits predict nothing beyond redistributive demand.

6 Semi-Structural Estimates of the Preferred Tax Rate

Table 7 reports NLS and two-step GMM estimates of the preferred-tax equation on the same baseline sample. The estimation sample is slightly smaller than the reduced-form sample (10,010 respondents, 70 ancestry clusters, 148 country-by-round cells) because the structural specification requires nonmissing values for the additional controls.

Table 7: Semi-structural preferred-tax estimates: Western Europe, one foreign-born and one native-born parent, education-based position

Parameter	NLS	Two-step GMM
Midpoint-education redistribution motive, A_0	0.263^{***}	0.268^{***}
	(0.012)	(0.011)
Education-gap loading, A_e	0.029^{***}	0.031^{***}
	(0.007)	(0.007)
UA shock scale, $\lambda_{\max}$	0.030^{***}	0.031^{***}
	(0.010)	(0.008)
Static tax-cost loading, k_s	0.047	0.062^{**}
	(0.033)	(0.029)
Future growth-cost loading, k_g	0.071^{***}	0.072^{***}
	(0.023)	(0.019)
Observations	10,010	10,010
Ancestry clusters	70	70
Country-by-round fixed effects	148	148
Mean fitted tax index	0.682	0.681
Within partial correlation	0.077	0.077
GMM J test p-value	–	0.928

Notes: The outcome is the ESS redistribution item rescaled to $[0, 1]$ (higher = more support). Model inputs follow equations (9)–(10): ancestral zero-sum beliefs (higher = stronger zero-sum belief) map to θ, Hofstede LTO to b, Hofstede UAI to λ (all on unit scales), and education enters the structural numerator, not the control vector. Both estimators profile out age, age squared, gender, marital status, labor-market status, and country-by-round fixed effects; k_0 is normalized to 0.35. Standard errors, in parentheses, are clustered by ancestral country. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The estimates give the model's mechanisms empirical content, within the limits the exercise

imposes. The education-gap loading A_e is positive and precise under both estimators (0.029 and 0.031): lower education raises the current redistributive motive. The uncertainty-avoidance scale $\lambda_{\max}$ is positive and significant (0.030 and 0.031), so ancestral uncertainty avoidance operates as an insurance motive inside the numerator. On the cost side, the static loading k_s is small and imprecise under NLS (0.047) and significant under GMM (0.062), while the future growth-cost loading k_g is positive and precise under both (0.071 and 0.072). No fitted observation is censored at either tax corner, and the J test does not reject the overidentifying restrictions ($p = 0.928$).

A positive k_g establishes less than it may appear to. Within the model's measurement mapping, k_g is the parameter that permits a negative association between the future-weight input and redistributive support; a positive fitted value therefore attributes that association to a perceived future tax-base cost rather than imposing it. It is not independent evidence that respondents perceive such a cost, since the exercise uses the same outcome and the same ancestral variables as the reduced form and the GMM moments are transformations of the model inputs rather than excluded instruments. And as Section 7.3 reports, k_g depends on the index used for b: it falls to 0.021 with a standard error of 0.019 when the behavioral patience measure is substituted, though it holds at 0.052 (0.018) when the survey scale is replaced by the Galor and Özak (2016) agricultural return.

Table 8: Average marginal effects on the model-implied preferred-tax index

Input	NLS	Delta-method s.e.	Bootstrap s.e.	Bootstrap 95% interval	Two-step GMM
Zero-sum belief (θ)	0.145	(0.046)	(0.052)	[0.050, 0.243]	0.166
Uncertainty avoidance (u)	0.041	(0.013)	(0.016)	[0.001, 0.063]	0.041
Long-term orientation (b)	-0.023	(0.013)	(0.018)	[-0.058, 0.010]	-0.021
Bootstrap replications			1,000		
Ancestry clusters			70		

Notes: Entries are sample averages of numerical derivatives of the fitted constrained preferred-tax function τ^* with respect to the indicated input on its zero-to-one scale. Delta-method standard errors propagate the ancestry-clustered parameter covariance of Table 7 through the marginal-effect function; bootstrap standard errors and percentile intervals resample the 70 ancestral countries with replacement and refit the NLS on each draw. The bootstrap intervals are the preferred summary here, since the parameter constraint and the nonlinearity both make a symmetric normal interval a poor approximation; Table 7 reports sandwich standard errors only. The final column repeats the two-step GMM point estimates. Coding directions as in Table 7.

Table 8 translates the parameters into average marginal effects. The ordering matches the reduced form and the two belief channels are estimated with enough precision to stand on their own. Zero-sum beliefs have the largest positive average marginal effect (0.145 NLS, 0.166 GMM), operating through the perceived-cost denominator, with a bootstrap interval of $[0.050, 0.243]$. Uncertainty avoidance is positive and also excludes zero (0.041 and 0.041; interval $[0.001, 0.063]$). The future-weight input is negative on average (-0.023 and -0.021) with a bootstrap interval of $[-0.058, 0.010]$ that contains zero.

Proposition 3 does not sign that average, which is why it is the wrong quantity to judge the model by. The proposition says the effect of the future weight is positive where ambiguity insurance dominates and negative where the perceived growth cost does, so a population straddling the threshold $\lambda_i^*(\theta)$ delivers a small average of mixed sign by construction. What the proposition requires is that both arms of the trade-off exist separately, and the estimates deliver exactly that: $\lambda_{\max}$ and k_g are both positive and precisely estimated, so neither the insurance term nor the

growth-cost term is a degenerate zero. The quantity that carries information alongside the average is therefore the distribution of $\partial\tau_i^{\circ}/\partial b$ across the sample.

Table A11 makes that comparison across the three candidate measures of the future weight, and it sharpens Section 7.3 in one direction and qualifies it in another. In the direction of support, it delivers the estimates Section 7.3 anticipated: the growth-cost loading k_g is positive and precise whether the future weight is measured by the Hofstede index or by the Galor–Özak agricultural return, and weakens only under GPS patience. The zero-sum and uncertainty-avoidance marginal effects are unchanged across all three. In the direction of qualification, the three measures place the sample in three different positions relative to the threshold: 8 percent of respondents lie on the insurance side under the Hofstede mapping, 22 percent under the agricultural return, and 98 percent under GPS patience. The model's conditionality is therefore consistent with all three, but which regime a given population occupies is not identified independently of the measure used for b, which is the measurement problem of Section 7.3 restated in structural terms.

Varying the measure in the uncertainty slot turns the same exercise into a test the model could have failed. Substituting the behavioral risk-taking measure of Falk et al. (2018) for the Hofstede uncertainty index removes the insurance channel: $\lambda_{\max}$ falls from 0.030 (s.e. 0.010) to 0.006 (0.025) and its marginal effect from 0.039 to 0.008 (0.031). That is the outcome Remark 1 requires. The two measures correlate -0.33 across ancestries, so risk taking is a different construct rather than a noisier reading of the same one, and a motive built on concern about unresolved states should not be reproducible by tolerance of known risk. Had the risk measure reproduced it, the ambiguity interpretation would have been the casualty.

Figure 2 shows the fitted preferred-tax surface against the future-weight input at low and high percentiles of zero-sum beliefs and uncertainty avoidance. It is the estimated counterpart of Figure 1: where the scenario illustration showed that both signs of the slope are attainable at observed country values, the fitted surface shows where this sample sits under the estimated parameters.

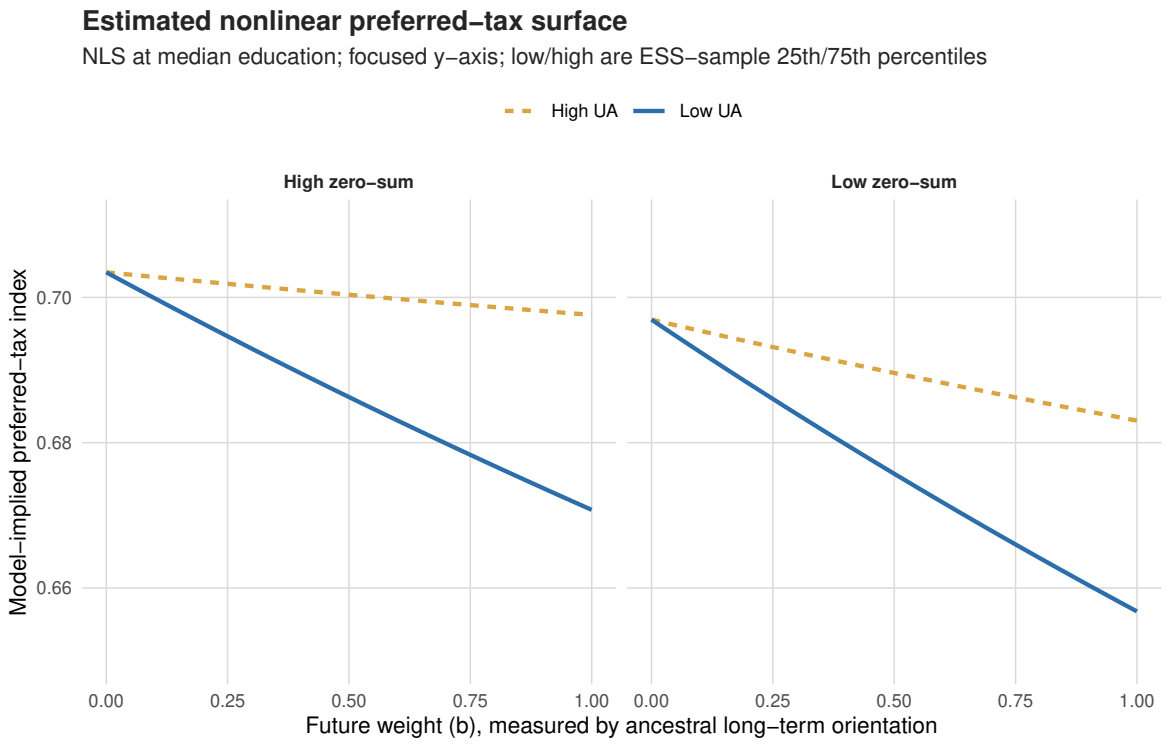


Figure 2: Estimated response surface of the model-implied preferred-tax index

Two limitations bound the exercise and both are visible in the diagnostics. The within partial correlation is 0.077, so the structural specification explains a small share of the residual variation; and the Jacobian condition numbers of 2,200 to 3,800 confirm that the decomposition of k_s, k_g, and $\lambda_{\max}$ is weakly conditioned, so the combined nonlinear index and its marginal effects are far better identified than the individual cost parameters. Readers should interpret the marginal effects, not the raw parameters.

Table A20 probes the maintained choices. The k_0 normalization matters for parameter levels and not at all for economic content, and the invariance is exact: the preferred-tax function is homogeneous of degree zero in $(A_0, A_e, \lambda_{\max}, k_s, k_g, k_0)$ jointly, so re-normalizing k_0 rescales the estimated vector proportionally and leaves fitted values and average marginal effects unchanged (0.145, 0.041, -0.023 throughout). Those columns are an exact computational check on the estimator rather than substantive robustness. Replacing the equally spaced $(R-1)/4$ scoring with midrank (ridit) scores, which uses only the ordinal content of the item plus the sample distribution, leaves the pattern intact: no parameter changes sign or significance status and the marginal effects move modestly (0.155, 0.048, -0.021), so the cardinality assumption is not driving the conclusions. Across 70 leave-one-ancestry-out re-estimations (Table A21), every average marginal effect keeps its sign in every run, with tight ranges (zero-sum 0.117 to 0.178; uncertainty avoidance 0.033 to 0.047).

7 What the Ancestral Measures Capture

Three questions about the ancestral measures have been deferred. The first is whether the traits reach the respondents at all. The second is whether sampling error in the country means is large

enough to matter. The third is what the index standing in for the future weight actually captures. All three are questions about measurement rather than about the model, which is why they come after the estimates. The first two change nothing. The third changes what we are willing to say about the third channel, and it is why the paper rests on the other two.

7.1 Do the Ancestral Traits Reach the Respondent?

The design assumes that an ancestral-country mean proxies a belief transmitted to the respondent. That assumption has a testable implication: the ancestral trait should predict the respondent's own attitudes and behavior in the domain it names, and not only their view of redistribution. Table A15 collects what the ESS allows on this, one outcome at a time, using the joint specification of (8) throughout.

Uncertainty avoidance has the strongest case, and it is the one trait the ESS proxies in every round. Ancestral uncertainty avoidance predicts the narrow security-values index at +0.023 (s.e. 0.014), with the broader security-and-order index positive but weaker at +0.014 (0.014). Two round-3 items about anticipated old-age hardship, estimated across all second-generation respondents, are stronger: worrying about income in old age at +0.064 (0.030) and about being unable to afford retirement at +0.081 (0.026). An ancestral scale built decades before these respondents were surveyed therefore predicts, in the second generation, the specific anxiety about future provision that the model's insurance motive runs on.

The long-term-orientation index reaches respondents through behavior rather than through stated style. It predicts whether someone saves to live comfortably in old age, at +0.089 (s.e. 0.045) in the baseline design and +0.093 (0.030) with the sample restrictions relaxed, and the agricultural return of Galor and Özak (2016) predicts the same saving outcome at +0.097 (0.032), a first stage that does not run through a survey scale at all. On the ESS planning items it does not reach them, at −0.017 (0.030) for planning for the future and −0.012 (0.034) for liking to plan; those nulls survive every sample enlargement the ESS allows, and no measure we have tried moves them. We read the pair as showing that the index is not inert, and that a stated planning style is the wrong place to look for it.[4] Section 7.3 returns to the saving result, which divides into components in the same way the redistribution result does.

Zero-sum beliefs have the thinnest first stage of the three. The ESS fields a perceived-welfare-cost battery in two rounds and a wealth-gap-fairness item in one, and neither is estimated with useful precision, at −0.020 (0.027) and −0.061 (0.056). That channel rests on its reduced-form behavior rather than on a first stage.

7.2 Ancestral Measurement Quality

The zero-sum measure is a single WVS/EVS item averaged over country samples of different sizes and waves, so sampling error in the ancestral means could attenuate the estimates. It does not. Recomputing the means from the 1981–2014 WVS/EVS microdata reproduces the build's ancestral values exactly for all 70 matched ancestries, including the Serbia-Montenegro-era codes `CS` (1996) and `CSS` (2006), which are retained as separate codes instead of being silently dropped. The median

[4]The pattern extends to accumulation. Across the 84 countries with a Hofstede score, long-term orientation correlates 0.47 with the World Bank Human Capital Index, 0.069 (0.015) per standard deviation and 0.039 (0.008) conditional on log GDP per capita, and predicts gross capital formation conditional on income, 2.58 (0.83) percentage points of GDP, though the bivariate association is weaker. Gross *national* savings is a null, unsurprisingly for a series that combines public and private saving. These are descriptive correlations rather than identified estimates.

ancestral mean rests on 2,597 respondents (minimum 296), drawn from a median of two waves, and the implied first-stage reliability (the share of cross-country variance in the estimated means that is signal rather than sampling noise) averages 0.991 and never falls below 0.950 (Table A16). Re-estimating with wave-adjusted means, empirical-Bayes shrunken means, both corrections, or reliability weights moves no coefficient by more than 0.002 (Tables A17–A18). This establishes that sampling error in the source means is small; it does not establish cross-cultural comparability, temporal stability, or migrant representativeness, which are separate questions.

An independent country-level measure can be extracted from the Chinoy et al. (2026) replication package, though it turns out to carry little information. Their public file records origin countries only for first-generation immigrants, so the extractable index is the mean of their four-item zero-sum index among 1,421 first-generation respondents spread over 126 birth countries: a median of ten respondents per country, for an implied reliability of 0.28 against 0.991 for the WVS/EVS means. The two indices are essentially uncorrelated across our ancestries, which reliabilities of this size plus migrant selection produce mechanically. On the 38 covered ancestries the WVS/EVS zero-sum effect is 0.034 (s.e. 0.018) and is unchanged by adding the Chinoy index, whose own coefficient is indistinguishable from zero (Table A19). The WVS/EVS item is substantially more reliable and remains our primary measure; the convergence with Chinoy et al. (2026) is conceptual rather than a country-index swap.

The same arithmetic settles what a set of nulls does and does not mean. The pairwise and triple interactions among the three traits that the model implies are individually and jointly insignificant (joint $p = 0.92$), with point estimates between 0.002 and 0.018 in absolute value. Two readings of that are available and the reliability results choose between them. It is not attenuation: with first-stage reliability near one, products of ancestral means inherit almost no additional noise, so the nulls are not an artifact of noisy zero-sum means. It is power, and the shortfall is severe. The smallest interaction this design could detect at eighty percent power runs from 0.041 to 0.076 across the four terms, against main effects of 0.038 and 0.042: the design cannot resolve an interaction even as large as the effects it is built to estimate. These nulls are therefore not evidence against the model's interaction predictions. They are evidence that this design cannot test them, which is why the paper rests on the direct effects. The one channel this program cannot quantify is measurement error in the Hofstede indices, which carry no sampling-based reliability estimate. The next subsection addresses that directly for the index where it matters most.

7.3 The Third Channel: What the Ancestral Long-Term-Orientation Index Measures

Three measures can fill the future-weight slot, and they do not agree. Two deliver a negative estimate: the Hofstede long-term-orientation index of Table 2 and, independently of it, the agricultural return of Galor and Özak (2016). The third, the behavioral patience measure of the Global Preferences Survey, delivers a positive one. This section sets out what that pattern establishes (more than one might expect) and what it leaves open. Table A11 collects the estimates.

Take the agricultural return first, because it is the one measure here that is not a survey scale at all. Galor and Özak (2016) argue that a high return to agricultural investment selected for long-term orientation, and they measure that return as the yield of the crops available before 1500 relative to the length of their growth cycle, the return per unit of waiting. We construct it in their ancestry-adjusted form from their replication data.[5] It is predetermined by millennia, carries no

[5]Built as pre-1500 crop yield divided by pre-1500 crop growth cycle, both ancestry-adjusted, from Galor and Özak

survey response style, and covers 68 of our 70 ancestries.

It behaves as their argument predicts. Across our ancestries it correlates +0.52 with the Hofstede index, reproducing their cross-validation, and it carries the effort content the label promises: +0.44 with the thrift-and-perseverance component of Minkov's index against −0.45 with its pride component. It is also close to orthogonal to the paper's other two channels, at +0.13 with ancestral uncertainty avoidance and −0.09 with zero-sum belief. Substituted for the Hofstede index it gives the same negative sign in the reduced form, −0.016 (s.e. 0.012) against −0.020 (0.010), imprecise in both cases. It also preserves the mechanism. Anticipating the semi-structural estimates of Section 6, which fits the preferred-tax equation itself, the future growth-cost loading k_g comes out at 0.052 (0.018) under the agricultural return against 0.071 (0.023) under the Hofstede index (Table A11). That is the substantive result of this section: the growth-cost channel that makes Proposition 3's trade-off operative is not an artifact of one survey scale. A Hofstede dimension built in the 1990s and an agroclimatic quantity fixed millennia ago both deliver it.

The disagreement is with the third measure. The Global Preferences Survey provides a behavioral patience measure, elicited from intertemporal choice, for 47 of our 70 ancestries (87 percent of baseline respondents), and on the identical sample it reverses the sign: Hofstede long-term orientation gives −0.031 (s.e. 0.013) and GPS patience +0.024 (0.011), both significant at five percent, on measures that correlate +0.39 with one another. Entering both sharpens each (−0.048 and +0.042). This rules out treating them as two indicators of one latent future weight, and does so without any estimation: if LTO $= \ell_1 b^* + e_1$ and GPS $= \ell_2 b^* + e_2$ with loadings sharing a sign, which a +0.39 correlation requires, then both reduced-form coefficients must carry the sign of the structural parameter. They do not. A composite index, a factor score, and an instrumental-variables correction are therefore all unavailable. The reversal propagates: k_g is 0.021 (0.019), indistinguishable from zero, when b is measured by patience.

The negative coefficient can be traced to a particular part of the index. Minkov's version is built from World Values Survey items for thrift (`a038`), perseverance (`a039`), national pride (`g006`), and making one's parents proud (`d054`), with the pride items at the short-term pole. Grouping the first two into an EFFORT component and the last two into a PRIDE component reconstructs the published index well: across 67 countries, EFFORT − PRIDE correlates +0.85 with it. Its single strongest correlate, though, is reverse-coded national pride (−0.80), against +0.54 for thrift and +0.40 for perseverance. Entered separately the two components divide the work unevenly: alone, EFFORT gives −0.010 (0.009) and PRIDE +0.019 (0.010), and with both entered EFFORT is −0.002 (0.010). Because the published index is close to EFFORT minus PRIDE, this locates its negative coefficient in the pride pole: ancestries with stronger national and parental pride support redistribution more.[6] The decomposition therefore delivers a specific conclusion rather than a general doubt: whatever the third coefficient reflects, it is not the thrift-and-perseverance content that would make it a reading of the model's future weight.

The pride pole is not free of temporal content, however, and it is worth being exact about what kind. Its high end is an orientation toward inherited norms and national history, which is why the original Chinese Value Survey placed respect for tradition at the *short*-term pole. The index therefore contrasts two temporal orientations rather than separating a temporal component from a non-temporal one. Time-perspective research likewise treats past and future orientation as distinct

(2016)'s country file. Their own specification enters yield and cycle jointly and estimates yield positive and cycle negative, which is what the ratio imposes.

[6] That association is interpretable in its own right, as pride standing for in-group identification and identification as a basis for transfers to co-nationals, but it belongs to a literature on welfare nationalism, not to the model tested here, and we do not pursue it.

dimensions rather than opposite ends of one scale (Zimbardo and Boyd, 1999), and our measures behave that way: pride correlates −0.59 with GPS patience across ancestries, so the two do track a common backward-versus-forward axis. That axis is not what predicts redistribution. Entered together the two carry the same sign and both strengthen: pride +0.054 (s.e. 0.014) and patience +0.064 (0.016), against +0.017 and +0.024 entered alone. What moves the outcome is therefore what remains in each once their shared temporal variation is removed.[7]

The same split appears on a second outcome, which matters because it shows the result is not particular to the redistribution question. The saving result of Section 7.1 is the natural place to look, since it is the one outcome besides redistribution on which the index reaches respondents at all. Decomposed, saving divides exactly as redistribution did: PRIDE −0.100 (0.045), so lower ancestral pride means more private saving, against EFFORT −0.009 (0.037). One component of the index carries both results, and it is the same one in each case. The long-term-orientation-by-position interaction of Table A9 is specific to this index in the same way: with GPS patience it is a third the size and insignificant. Its sign runs opposite to the closest prior estimate, since Wang (2018) finds patience flattening rather than steepening the class cleavage in redistributive demand, and it is the sign that the heterogeneous-growth extension of Remark 4 predicts.[8]

Two conclusions follow about the third channel, and they point in different directions. The model's structure is supported by measures that share no common source: both arms of Proposition 3's trade-off are estimated positive and precise, and the growth-cost arm survives replacing a survey scale with pre-1500 agroclimate. What the evidence does not support is attributing the third coefficient to the future weight specifically. The index is a composite, and the component that moves redistribution is its pride pole rather than its thrift-and-perseverance pole, so the association is real without being a reading of how heavily ancestors weighted the future. We keep Hofstede long-term orientation as the estimated measure on grounds that do not depend on the answer it gives: it covers all 70 ancestries where GPS patience covers 47, so the baseline sample is defined by it, and its sign is the one the agricultural return corroborates. We report its coefficient accordingly, as a negative and imprecisely estimated ancestral association that we do not attribute to the future weight, and leave open where a population sits between the two arms of the trade-off, a question on which the three measures disagree.

8 Conclusion

This paper asks which inherited beliefs organize the demand for redistribution, and answers with two that occupy distinct places in a public-finance problem. Uncertainty avoidance sets the value of insurance against an ambiguous future. Zero-sum beliefs set the perceived price of the policy, through a perceived elasticity of the tax base. Among second-generation Europeans compared within the same country and survey round, a standard deviation of ancestral zero-sum belief raises

[7] Suppression of this size on 45 ancestries should not be pushed further than the sign: the robust statement is that two measures correlating −0.59 predict redistribution in the same direction rather than in opposite directions, which a common backward-versus-forward axis would not produce.

[8] One test of that extension points toward a future-weight reading, and we report it rather than rest on it. Estimating within each destination country the slope of income position on education and splitting at the median slope, the interaction is −0.057 (s.e. 0.016) where education buys the most income and 0.000 (0.012) elsewhere; entered continuously, its interaction with the standardized country gradient is −0.024 (0.010). That is what heterogeneous expected growth predicts. It is not decisive. The moderator varies across only the sixteen destination countries for which both measures exist, while the standard errors cluster on seventy ancestries, so the apparent precision overstates what a destination-level split can establish. A second implication of the same reading fails: the interaction is no sharper among the young, at 0.001 (0.017) below 36 against −0.045 (0.023) at 50 and above.

support for redistribution by 2.7 percent of an outcome standard deviation and ancestral uncertainty avoidance by 3.5 percent (between a quarter and two-fifths of the education gradient), and the gap between the extreme ancestries is 23 percent of a standard deviation. Both survive wild-cluster bootstrap inference, eleven rival inherited attitudes, ancestral development and institutional controls, an ordered-response specification, and every leave-one-out perturbation we can construct.

The paper's other contribution is to say what kind of belief zero-sum thinking is. The literature that established its political importance treats it as a claim about appropriation: one person's gain is another's loss. Embedded in a tax problem, our model isolates a different content, generativity: whether output responds to effort and investment at all. It recasts the belief as a price rather than a preference. The two readings agree on direct effects and diverge on where in the income distribution the effect operates. The evidence favors the price: among above-median-education respondents, where extraction predicts opposition, ancestral zero-sum belief predicts more support at every cutpoint of the outcome, and the belief tracks a low-agency worldview rather than any measure of institutional quality. Read this way, the paper gives Foster's image of limited good a fiscal microfoundation: a pure limited-good believer perceives a perfectly inelastic tax base, which is exactly why redistribution looks cheap.

Three limitations bound these conclusions. First, the design holds the destination context fixed but cannot rule out that ancestral traits proxy for other ancestral factors; the development, institutional, and regional controls address the obvious candidates and not the subtle ones, and neither they nor the leave-one-out results eliminate unobserved migrant or intermarriage selection. Second, the zero-sum result is bounded in population. It is estimated among respondents with one native-born parent and reverses sign among those with two foreign-born parents, a group drawn disproportionately from origins where intermarriage is rare and minority status in the destination is salient, and for whom the ancestral mean plausibly marks group position as much as inherited belief. About half that gap is ancestry composition and part of the rest tracks minority status, but not all of it, and whether the perceived-cost mechanism operates differently for Europe's most visibly distinct immigrant-origin populations is a question this design cannot settle. The uncertainty-avoidance result carries no such qualification.

Third, the model's most distinctive theoretical claim is only partly testable here. Its structure is confirmed: both arms of the trade-off in Proposition 3 are separately identified, with a positive insurance scale and a positive growth-cost loading that survives replacing the Hofstede index with the predetermined agricultural return of Galor and Özak (2016). What cannot be established is where a population sits between them. Proposition 3 says the weight a citizen places on the future has no fixed political sign: it argues against redistribution when the future reads as a growth cost and for it when the future reads as insurable risk. That result stands as theory. What we cannot supply is an ancestral measure that isolates the future weight from what else it carries. The available index disagrees in sign with the behavioral patience measure on the same respondents, and decomposing it shows its redistribution-relevant content is national and parental pride rather than thrift or perseverance (Section 7.3). Testing the conditionality requires either a behavioral ancestral measure with wider country coverage or more ancestral clusters than 70 (pooling with ISSP second-generation data is the obvious route), and we regard it as the most valuable open question the model raises.

Two further extensions follow from what the paper does establish. Inherited beliefs are, by construction, not manipulated, so the paper cannot say what would happen if a citizen's beliefs about the efficiency costs of taxation were changed. That is a different causal question, and the model's sharpest untested prediction is a signed interaction: information about efficiency costs

should move growth-oriented respondents and leave fixed-pie believers unmoved, because the latter perceive no margin for taxes to distort. Such an experiment and this design are complements, since the predetermined variation studied here is something no experiment can manufacture.

What This Implies for Technological Transitions

We opened with artificial intelligence and should close on it. What these results supply is the prior structure any account of AI and redistributive demand has to reckon with, and it has a specific shape. They do not measure attitudes toward AI or identify whom its arrival displaces.

Exposure is not destiny. The debate is usually conducted as though measuring who is displaced would tell us who demands compensation, but the same disruption reaches a population already sorted by whether it reads an uncertain future as insurable downside risk or as growth that taxation might impede, and our estimates say both readings are present, inherited, and measurable before the technology arrives. Which one dominates is what Proposition 3 signs: the insurance channel prevails where uncertainty avoidance is high or where zero-sum beliefs make taxation appear cheap, and the growth-cost channel prevails otherwise. Figure 3 makes that concrete for the technology now arriving. Across OECD countries the model-implied response to a weightier future is negatively related to the IMF's AI Preparedness Index and to how long a country has had majority internet use: the populations meeting this transition best prepared sit disproportionately on the growth-cost side of the threshold, where a more consequential future lowers the demand for redistribution rather than raising it. Roughly half the countries shown fall on each side, so the conditionality is not degenerate at the national level. The alignment is not itself a discovery: preparedness loads on the same growth-oriented, uncertainty-tolerant disposition as the model's cost side, and the figure is an illustration at scenario parameters rather than an estimate. What it makes concrete is that the countries meeting this technology first are, on average, those whose prior beliefs point away from redistribution as the response.

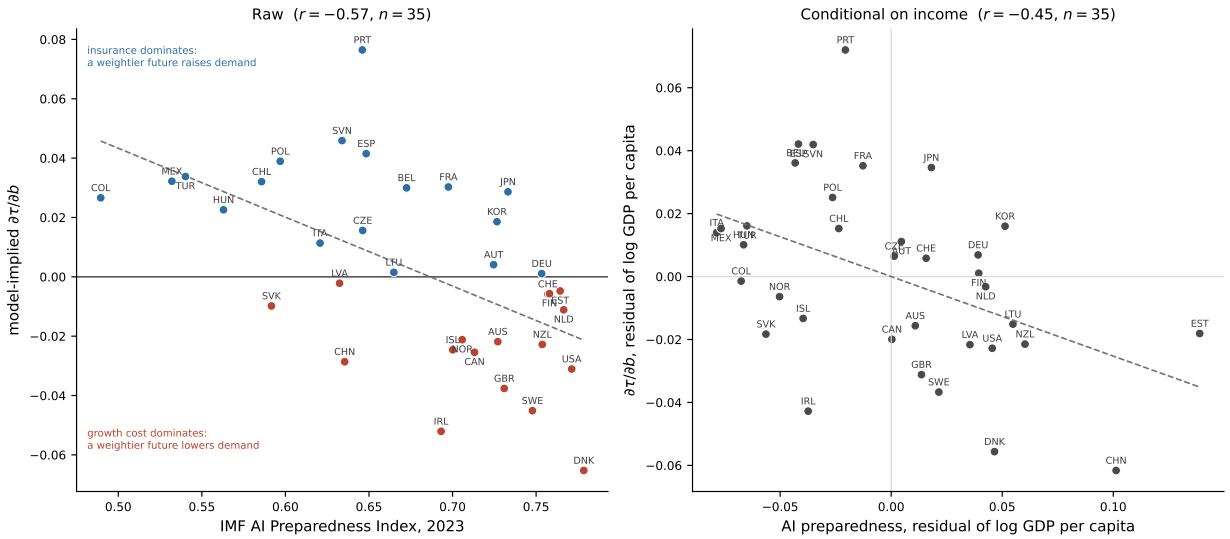


Figure 3: Technology adoption and the model-implied response to a weightier future

Notes: Each point is an OECD member or China. The vertical axis is $\partial\tau_i^{\circ}/\partial b$ from Proposition 3, evaluated at each country's WVS/EVS zero-sum mean and Hofstede uncertainty-avoidance score: it is positive where ambiguity insurance dominates, so that a weightier future raises desired redistribution, and negative where the perceived future growth cost dominates. Note that the quantity plotted uses the zero-sum and uncertainty-avoidance measures and the scenario cost parameters; it does not use the ancestral long-term-orientation index discussed in Section 7.3. The horizontal axis is the IMF AI Preparedness Index for 2023. The right panel plots both variables as residuals from log GDP per capita. Eighteen of the 35 countries shown lie above the threshold and 17 below. The same negative alignment appears using years of majority internet use, measured as 2026 minus the year each country's internet use first passed half the population ($r = -0.52$ raw, -0.38 conditional on income), so both technology measures point the same way. *This figure is illustrative and is not an estimate.* The cost parameters are the scenario normalizations of Section 2.6 rather than fitted values, the sample is 35 countries, and AI preparedness correlates 0.85 with log GDP per capita across the full sample of economies. The figure makes the model's implication concrete for a familiar policy question; it is not an independent test of it.

Finally, the perceived price of redistribution is itself in play. If zero-sum thinking is a belief about whether output responds to effort and investment, then a technology that visibly raises productivity is an argument against redistribution for those who see the pie as growing, and a visible redistribution of a fixed pie for those who do not. That is where the paper's framing and its mechanism meet. Inherited beliefs about the pie and about uncertainty are not a fixed backdrop against which technological change plays out; they are the terms on which citizens will evaluate it. And because one of them is a price rather than a preference, a belief about economic mechanics, not a taste, it is the one input that evidence can move.

A Appendix Tables

Table A1: Ancestral countries in the baseline sample

Anc.	N	ZS	UAI	LTO	$\bar{y}$	Anc.	N	ZS	UAI	LTO	$\bar{y}$
DE	1,517	4.58	65	83	3.63	TH	30	3.75	64	32	3.50
IT	991	4.18	75	61	3.72	VN	30	4.41	30	57	3.90
GB	888	4.95	35	51	3.78	IR	30	6.05	59	14	3.53
FR	750	4.94	86	63	3.79	CL	29	4.22	86	31	3.93
PL	696	4.17	93	38	3.83	EE	29	4.71	60	82	3.38
US	389	4.60	46	26	3.75	VE	29	4.27	76	16	3.93
AT	331	3.83	70	60	3.63	ZA	28	4.20	49	34	3.64
IE	318	4.88	35	24	3.64	CN	28	3.52	30	87	3.82
ES	315	5.37	86	48	3.92	LV	25	4.12	63	69	3.52
CZ	299	5.10	74	70	3.82	BG	22	5.22	85	69	4.32
NL	290	4.80	53	67	3.72	JP	21	4.96	92	88	3.71
FI	263	4.73	59	38	3.72	PE	20	3.78	87	25	3.90
DK	252	4.75	23	35	3.63	IQ	20	4.85	85	25	3.75
SE	251	4.96	29	53	3.59	MX	20	4.07	82	24	3.80
TR	251	5.09	85	46	4.12	NG	18	5.25	55	13	3.61
ID	247	3.86	48	62	3.32	CO	18	3.71	80	13	3.61
BE	238	4.61	94	82	3.74	LT	17	4.54	65	82	3.24
HU	214	4.52	82	58	3.65	AL	16	5.29	70	61	4.00
MA	211	5.08	68	14	3.93	IS	15	4.24	50	28	3.00
NO	188	4.62	50	35	3.57	GH	13	5.33	65	4	3.92
RU	176	5.09	95	81	4.00	LB	12	5.38	50	14	3.33
PT	151	5.29	99	28	3.90	KR	11	4.76	85	100	3.45
IN	116	5.53	40	51	3.56	MT	10	3.27	96	47	3.30
RO	112	4.73	90	52	3.91	DO	10	3.46	45	13	4.00
CH	109	4.49	58	74	3.72	MY	8	4.34	36	41	3.38
BR	109	3.59	76	44	3.92	NZ	8	4.77	49	33	2.88
CA	80	4.15	48	36	3.80	SG	8	5.19	8	72	4.00
HR	68	5.02	80	58	3.69	LY	7	3.58	68	23	3.57
PH	64	4.61	44	27	3.59	BD	7	4.05	60	47	3.71
SI	62	5.56	88	49	3.95	UY	6	4.41	99	26	4.33
AR	52	4.31	86	20	3.71	TT	4	3.80	55	13	3.25
EG	48	3.77	80	7	3.65	JO	2	5.56	65	16	4.50
AU	47	4.87	51	21	3.74	TW	2	3.55	69	93	3.00
SK	36	5.40	51	77	3.86	ZM	2	4.86	50	30	3.50
PK	33	5.06	70	50	3.45	BF	1	5.98	55	27	5.00

Notes: All 70 ancestral countries, ordered by respondent count and read down the left block then the right. ZS is the ancestral WVS/EVS zero-sum mean ($11 - $ `e041`, 1–10); UAI and LTO are the Hofstede indices (0–100); $\bar{y}$ is mean support for redistribution among that ancestry's respondents (1–5). Ancestry is the foreign-born parent's country of birth and is the level at which standard errors are clustered.

Table A2: Residence-country composition of the baseline sample

Residence country	Respondents	Ancestries	Mean outcome
DE	1,499	54	3.76
CH	1,339	55	3.59
FR	963	45	3.90
BE	889	43	3.82
AT	800	45	3.91
GB	793	47	3.60
NL	774	43	3.49
IE	734	30	3.84
SE	693	49	3.67
NO	541	41	3.58
DK	386	45	3.16
ES	271	28	4.00
FI	224	26	3.87
LU	178	14	3.51
GR	162	15	4.29
PT	143	20	4.12
IS	135	20	4.01
CY	98	26	4.07
IT	96	24	4.12
Total	10,718	70	3.73

Notes: "Ancestries" counts distinct ancestral countries represented among that residence country's respondents. All identifying comparisons are within a residence-country-by-round cell.

Table A3: Wild-cluster bootstrap p-values for the main-effects table

	(1)	(2)	(3)	(4)
Zero-sum beliefs	0.025			0.005
Uncertainty avoidance		0.004		0.003
Long-term orientation			0.142	0.061
Education	<0.001	<0.001	<0.001	<0.001
Observations		10,718		
Ancestry clusters		70		

Notes: Wild-cluster bootstrap p-values (Rademacher weights, null imposed, 9,999 replications, clusters = ancestral countries) for the coefficients of Table 2, column by column. They are reported here instead of in the main table so that coefficients and standard errors remain legible there.

Table A4: Parental composition and the zero-sum coefficient

Sample	Zero-sum coefficient	N	Anc.
Panel A: the contrast			
One foreign-born parent (baseline)	+0.042***	10,718	70
	(0.013)		
Both parents foreign-born	-0.039**	4,221	66
	(0.020)		
Panel B: is it ancestry averaging?			
Both, same origin (no averaging)	-0.041*	3,718	62
	(0.021)		
Both, mixed origin (averaging)	-0.018	503	58
	(0.084)		
Panel C: is it the origin mix?			
Both, excl. TR	-0.037*	3,467	65
	(0.020)		
Both, excl. TR+MA+PK	-0.039*	3,019	63
	(0.020)		
Both, excl. TR+MA+PK+IN+BD+LK	-0.038*	2,894	61
	(0.022)		
Panel D: is it ancestry composition?			
One, common ancestries (66)	+0.041***	10,703	66
	(0.013)		
Both, common ancestries (66)	-0.039**	4,221	66
	(0.020)		
One, reweighted to both-parent ancestry mix	+0.021	10,703	66
Panel E: minority status			
One-parent, non-minority	+0.059***	8,367	70
	(0.015)		
Both-parent, non-minority	-0.017	2,688	62
	(0.022)		
Both-parent, minority	-0.071	666	45
	(0.055)		

Notes: Each row reports the zero-sum coefficient from the joint specification of Table 2, column 4, estimated on the indicated sample, with the same controls, country-of-residence-by-round fixed effects, and ancestry-clustered standard errors throughout; all samples are restricted to the same residence countries. Panel B splits respondents with two foreign-born parents by whether the parents share a country of birth: ancestral traits are averaged across parents only in the mixed-origin group, so the same-origin row is free of averaging. Panel C drops the largest non-European origins in the two-parent sample. Panel D restricts both samples to the 66 ancestries they share, and reweights the baseline sample to the two-parent ancestry distribution (a weighted least-squares fit, so no standard error is reported). Panel E splits by whether the respondent reports belonging to an ethnic minority (`blgetmg`), an item not fielded in every round; the one-parent minority cell has 287 respondents and is not estimated. Pooled interaction tests of the differentials are reported in the text. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A5: Pairwise correlations of the ancestral traits, baseline sample

	Ancestry level	Respondent level
Zero-sum beliefs – Uncertainty avoidance	-0.06	-0.11
Zero-sum beliefs – Long-term orientation	0.01	-0.06
Uncertainty avoidance – Long-term orientation	0.06	0.18
Units	70 ancestries	10,718 respondents

Notes: Standardized ancestral measures; coding directions as in Table 2. Ancestry-level correlations use one observation per ancestral country in the estimation sample; respondent-level correlations weight by sample composition.

Table A6: Main effects by foreign-born parent and respondent gender

	Mother FB	Father FB	Women	Men
Zero-sum beliefs	0.037**	0.040**	0.039**	0.046*
	(0.019)	(0.017)	(0.018)	(0.027)
Uncertainty avoidance	0.038***	0.038***	0.042***	0.033*
	(0.013)	(0.014)	(0.013)	(0.019)
Long-term orientation	-0.020	-0.039**	-0.020	-0.030
	(0.017)	(0.015)	(0.015)	(0.020)
Education	-0.094***	-0.106***	-0.066***	-0.141***
	(0.020)	(0.018)	(0.021)	(0.017)
Observations	5,301	5,417	5,551	5,167
Ancestry clusters	66	69	69	68

Notes: Columns 1–2 split the baseline sample by whether the mother or the father is the foreign-born parent; columns 3–4 split by respondent gender (gender is dropped as a control in the gender splits). All columns include education, age, age squared, and country-of-residence-by-round fixed effects; standard errors in parentheses are clustered by ancestral country. Coding directions as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A7: Leave-one-out stability of the three main effects

	Baseline	Min	Max	Sign stable	Sig. 5%	Runs
Leave one ancestry out (70 runs)						
Zero-sum beliefs	0.042	0.034	0.048	70/70	70/70	70
Uncertainty avoidance	0.038	0.030	0.041	70/70	70/70	70
Long-term orientation	-0.025	-0.029	-0.016	70/70	64/70	70
Leave one residence country out (19 runs)						
Zero-sum beliefs	0.042	0.037	0.049	19/19	19/19	19
Uncertainty avoidance	0.038	0.024	0.043	19/19	19/19	19
Long-term orientation	-0.025	-0.029	-0.015	19/19	14/19	19

Notes: Joint specification of Table 2, column 4, re-estimated dropping one ancestral country (top panel) or one residence country (bottom panel) at a time. "Sign stable" counts runs with the baseline sign; "Sig. 5%" counts runs significant at five percent with ancestry-clustered standard errors.

Table A8: Ordered-logit estimates of the three main effects

	Index coefficient	AME on $P(y = 1)$	AME on $P(y \geq 4)$	AME on $P(y = 5)$
Zero-sum beliefs	0.069***	-0.0019	0.0143	0.0122
	(0.024)	(0.0007)	(0.0049)	(0.0041)
Uncertainty avoidance	0.075***	-0.0020	0.0155	0.0132
	(0.019)	(0.0005)	(0.0040)	(0.0035)
Long-term orientation	-0.042*	0.0011	-0.0087	-0.0074
	(0.022)	(0.0006)	(0.0046)	(0.0039)
Education	-0.174***	0.0047	-0.0360	-0.0307
	(0.026)	(0.0009)	(0.0053)	(0.0046)
Observations		10,718		
Ancestry clusters		70		

Notes: Ordered logit on the five-category redistribution item (higher = more support) with the covariates of Table 2, column 4, and country-of-residence-by-round fixed effects entered as indicators. Standard errors in parentheses are clustered by ancestral country. Average marginal effects are sample averages of the analytic derivatives of the indicated category probabilities with respect to each standardized input. Note that in a single-index ordered model every category's marginal effect is the same scalar $\overline{\Lambda} = \mathbb{E}[\Lambda'(\kappa_1 - \eta_i)]$ times the index coefficient, so this specification cannot represent a variable that moves one margin but not another; the cutpoint-by-cutpoint estimates of Table 6 are required for that. The incidental-parameters bias from fixed-effect indicators is limited by the average cell size (≈ 70). Coding directions as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A9: Income-gradient tests: trait × position interactions

	Own education	Parental education	Above-median educ.	Below-median educ.
Zero-sum beliefs	0.043***	0.026*	0.084**	0.042
	(0.013)	(0.014)	(0.036)	(0.037)
Uncertainty avoidance	0.037***	0.028**	0.002	0.038
	(0.010)	(0.012)	(0.028)	(0.032)
Long-term orientation	-0.027**	-0.028*	0.024	0.025
	(0.012)	(0.016)	(0.032)	(0.037)
Position	-0.089***	-0.045***	-0.043	-0.133**
	(0.016)	(0.015)	(0.036)	(0.062)
Zero-sum × position	0.005	-0.004	-0.036	0.015
	(0.010)	(0.016)	(0.032)	(0.035)
UA × position	-0.012	0.011	0.021	-0.018
	(0.009)	(0.015)	(0.033)	(0.027)
LTO × position	-0.020*	-0.035***	-0.078**	0.037
	(0.011)	(0.012)	(0.032)	(0.035)
Observations	10,718	7,436	5,163	5,555
Ancestry clusters	70	70	70	66

Notes: Position is standardized own education (column 1; also used in the median splits of columns 3–4) or standardized mean parental education, a predetermined proxy (column 2). Each column includes the three trait main effects, position, the three trait × position interactions, age, age squared, gender, and country-of-residence-by-round fixed effects; own education is not separately controlled where it is the position variable. Model predictions with position coded high = high position: zero-sum × position < 0, UA × position ≈ 0, LTO × position > 0. By Lemma 2(iii) these interaction coefficients need not share the sign of the model's cross-partials once the response is ordinal, which is why Section 5.1 rests its argument on directional comparisons instead; they are reported here for completeness. Because position is standardized in the full sample, the trait main effects in columns 3–4 are evaluated at the full-sample mean, which lies at the boundary of each median split. The zero-sum × position and UA × position estimates are unchanged to within 0.001 under wave-adjusted and empirical-Bayes-shrunken zero-sum means (Table A18). Standard errors in parentheses are clustered by ancestral country. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A10: The ancestral long-term-orientation index: convergent validity and content

	Zero-sum	Unc. avoid.	Future-weight measure
Panel A: alternative measures of the future weight b, GPS-covered subsample			
Hofstede LTO (paper's measure)	+0.041***	+0.032***	-0.031**
	(0.014)	(0.012)	(0.013)
GPS patience in its place	+0.037***	+0.048***	+0.024**
	(0.013)	(0.013)	(0.011)
Both jointly: Hofstede LTO	+0.036***	+0.058***	-0.048***
	(0.012)	(0.011)	(0.015)
GPS patience			+0.042***
			(0.013)
Panel B: semi-structural future growth-cost loading k_g			
b measured by Hofstede LTO	0.068, significant at one percent		
b measured by GPS patience	0.020, not significant		
Panel C: correlation of the published Hofstede index with its WVS ingredients			
Reconstruction: EFFORT − PRIDE	+0.85		
National pride (g006), reversed	-0.80		
PRIDE: national and parental pride	-0.75		
Parental pride (d054), reversed	-0.59		
EFFORT: thrift and perseverance	+0.55		
Thrift (a038)	+0.54		
Perseverance (a039)	+0.40		
Panel D: which ingredient carries the redistribution effect			
Published Hofstede LTO	-0.030** (0.012)		
Reconstruction (EFFORT − PRIDE)	-0.017* (0.009)		
EFFORT only (thrift, perseverance)	-0.010 (0.009)		
PRIDE only (national, parental)	+0.019* (0.010)		
Both components: EFFORT	-0.002 (0.010)		
PRIDE	+0.018 (0.012)		
Panel A respondents / ancestries	9,356 / 47		
Panel C countries	67		
Panel D respondents / ancestries	9,403 / 63		

Notes: Evidence behind Section 7.3. Panel A re-estimates the joint specification of Table 2, column 4, on the subsample of ancestries with a Global Preferences Survey patience score (Falk et al., 2018), replacing the Hofstede long-term-orientation index with GPS patience and then entering both; the two measures correlate +0.39 across ancestries. Panel B reports the future growth-cost loading k_g from the semi-structural specification of Table 7 with b measured each way, on the GPS-covered sample. Panel C decomposes Minkov's index into an EFFORT component (standardized WVS thrift a038 and perseverance a039) and a PRIDE component (national pride g006 and parental pride d054), both country means over countries with at least 200 valid responses per item, and reports correlations with the published index; the pride items sit at the short-term pole, so they enter the published index reversed. Panel D re-estimates the joint specification substituting each component for the published index. Standard errors in parentheses are clustered by ancestral country throughout. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A11: Three candidate measures of the ancestral future weight

	Hofstede LTO (paper's measure)	GPS patience (behavioral)	Galor–Özak return (pre-1500 crop)
Panel A: reduced form			
Coefficient on the future-weight measure	-0.020*	+0.028**	-0.016
	(0.010)	(0.012)	(0.012)
Panel B: semi-structural parameters			
Uncertainty scale, $\lambda_{\max}$	+0.031***	+0.026**	+0.025***
	(0.010)	(0.012)	(0.007)
Static tax cost, k_s	+0.048	+0.075**	+0.044
	(0.033)	(0.036)	(0.038)
Future growth cost, k_g	+0.071***	+0.021	+0.052***
	(0.023)	(0.019)	(0.018)
Panel C: average marginal effects			
Zero-sum belief, θ	+0.145***	+0.146***	+0.136***
	(0.046)	(0.051)	(0.051)
Uncertainty avoidance	+0.040***	+0.033**	+0.042***
	(0.013)	(0.015)	(0.012)
Future weight, b	-0.022*	+0.023*	-0.013
	(0.013)	(0.012)	(0.016)
Panel D: where the sample sits relative to the Proposition 3 threshold			
Share with $\partial\tau/\partial b > 0$	8%	98%	22%
same, by ancestry	7%	94%	19%
Respondents	10,010	8,709	9,909
Ancestral countries	70	47	68
Within correlation	0.077	0.074	0.077

Notes: Each column replaces the measure filling the future-weight slot b and is estimated on its own maximal sample; zero-sum beliefs, uncertainty avoidance, the controls, the country-by-round fixed effects, and the ancestry-clustered standard errors are as in Table 2, column 4, and Table 7 throughout. Panel A standardizes each measure within its own estimation sample. In Panels B to D every candidate, including the Hofstede index, is min–max rescaled to the unit interval within its estimation sample so that the three columns are comparable with one another. That differs from the native LTO/100 mapping of Table 7, so the Hofstede column here is close to but not identical with the estimates reported there. Panel D evaluates $\partial\tau_i^\circ/\partial b$ from Proposition 3 at the fitted parameters of that column and reports the share of the sample on the insurance-dominant side of the threshold; because the parameters are fitted, this describes where each mapping places the sample rather than testing it independently. The Galor–Özak column uses their agricultural return, ancestry-adjusted pre-1500 crop yield divided by ancestry-adjusted pre-1500 crop growth cycle, constructed from the country file of their replication package; it covers 68 of the 70 ancestries. Coverage is not driving the contrast: re-estimated on the 47 ancestries the GPS measure covers, the Hofstede column gives $k_g = 0.067$ (s.e. 0.024). $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A12: The conditional LTO effect: split by ancestral uncertainty avoidance

	High-UA ancestries	Low-UA ancestries
Zero-sum beliefs	0.031**	0.042
	(0.012)	(0.039)
Uncertainty avoidance	0.083***	0.011
	(0.024)	(0.038)
Long-term orientation	-0.019	-0.030
	(0.019)	(0.019)
Education	-0.100***	-0.099***
	(0.018)	(0.026)
Observations	5,338	5,380
Ancestry clusters	33	37

Notes: Joint specification of Table 2, column 4, estimated separately for respondents whose ancestral uncertainty avoidance is above versus below the ancestry-level median. Proposition 3 would predict a less negative long-term-orientation effect in the high-UA half, and the point estimates have that ordering (-0.019 against -0.030), but with 33 and 37 clusters the contrast is far from significant and does not establish conditionality; given Section 7.3 we do not read it as evidence on the model's threshold. Standard errors in parentheses are clustered by ancestral country. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A13: Placebo outcomes

	Redistribution	EU unification further	Trust in police
Zero-sum beliefs	0.039***	-0.006	-0.001
	(0.012)	(0.018)	(0.016)
Uncertainty avoidance	0.036***	0.003	-0.032**
	(0.010)	(0.013)	(0.012)
Long-term orientation	-0.023**	0.013	0.024*
	(0.011)	(0.011)	(0.014)
Education	-0.096***	0.217***	0.095***
	(0.015)	(0.014)	(0.014)
Observations	10,718	9,276	10,687
Ancestry clusters	70	70	70

Notes: Joint specification of Table 2, column 4, on standardized outcomes. EU unification: "European unification should go further" (0–10, higher = further); trust in police (0–10, higher = more trust); the redistribution column repeats the baseline outcome in standardized units for comparison. Standard errors in parentheses are clustered by ancestral country. Ancestral coding directions as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A14: Societal tightness as a rival to ancestral uncertainty avoidance

Specification	Zero-sum	Unc. avoid.	Long-term	Tightness	N	Anc.
Baseline joint specification	+0.042***	+0.038***	-0.025**	—	10,718	70
	(0.013)	(0.010)	(0.011)			
covered subsample, no rival	+0.044***	+0.033**	-0.031**	—	8,384	41
	(0.014)	(0.012)	(0.013)			
+ ancestral tightness	+0.044***	+0.033***	-0.031**	+0.000	8,384	41
	(0.014)	(0.012)	(0.014)	(0.014)		

Notes: Joint specification of Table 2, column 4. Ancestral tightness is Gelfand's six-item tightness scale as fielded in the International Study of Metanorms in 57 societies between April 2019 and January 2020, adjusted for response style following the original procedure (country-level $\alpha = 0.80$); higher values mean stronger norms and less tolerance of deviance. It is standardized within its estimation sample. The middle row re-estimates the baseline on the covered subsample without the rival, so that the coverage change is never confounded with the rival's effect. Across the covered ancestries tightness correlates -0.35 with the Hofstede uncertainty-avoidance index. Standard errors in parentheses are clustered by ancestral country. Coding directions otherwise as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A15: Cultural-transmission first stage: do the ancestral traits reach the respondent?

Second-generation outcome	Sample	ESS rounds	Ancestral trait	Predicted sign	Estimate	N
Saves for old age	Baseline	3	LTO	+	+0.089*	1,068
					(0.045)	
Plans for the future	Baseline	3, 9	LTO	+	-0.017	2,140
					(0.030)	
Likes planning ahead	Baseline	3	LTO	+	-0.012	1,103
					(0.034)	
Worried old-age income	All 2nd gen.	3	UA	+	+0.064**	1,989
					(0.030)	
Worried cannot retire	All 2nd gen.	3	UA	+	+0.081***	1,520
					(0.026)	
Security values (narrow)	Baseline	2–11	UA	+	+0.023*	10,137
					(0.014)	
Security/order values (broad)	Baseline	2–11	UA	+	+0.014	10,144
					(0.014)	
Perceived welfare cost	Baseline	4, 8	ZS	−	-0.020	2,078
					(0.027)	
Wealth gap seen as unfair	Baseline	9	ZS	+	-0.061	984
					(0.056)	

Notes: Each row is a separate regression with a standardized ESS attitude as the outcome, all three ancestral traits entered jointly, education, age, age squared, gender, country-of-residence-by-round fixed effects, and ancestry-clustered standard errors. *Baseline* rows use the design of Table 2; the saving row alone is restricted to the common support of the Hofstede and Galor–Özak future-weight measures, while the two planning rows use their maximal outcome-specific baseline sample, so their N is larger and magnitudes are not comparable across those rows. *All 2nd gen.* relaxes the one-parent and destination restrictions for the round-3 worry items. The reported coefficient is the one on the matching ancestral trait, named in the fourth column. Outcomes are oriented so that higher values mean more of the named attitude: saving for old age from `svclvo`; planning from `plnftr` and `plprftr`; old-age worry from `wrinco` and `ablrtr`; security values from `impsafe` and `ipstrgv`, with the broad index adding `ipfrule` and `imptrad` (Schwartz 1–6 scales reversed; round-11 variants with the `a` suffix appended); perceived welfare cost from `sbstrec`, `sbbsntx`, and `sblazy`; and wealth-gap unfairness from the signed round-9 item `wltdffr`, where positive values mean differences are seen as large and unfair. The predicted sign for perceived welfare cost is negative because the model has zero-sum belief *lower* the perceived cost of redistribution. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A16: Measurement quality of the ancestral zero-sum means

Statistic (across 70 matched ancestries)	Value
WVS/EVS respondents per ancestry: min / p25 / median	296 / 1,597 / 2,597
p75 / max	4,339 / 9,158
Ancestries with fewer than 500 respondents	2
Ancestries with fewer than 1,000 respondents	6
Contributing waves per ancestry: median (range)	2 (1–4)
Survey years spanned	1989–2014
First-stage reliability w_c: mean (min)	0.991 (0.950)
Correlation of wave-adjusted with raw means	0.987

Notes: Computed from the integrated WVS/EVS 1981–2014 microdata (valid responses to `e041`, recoded $11 - $ `e041`, higher = stronger zero-sum belief); recomputed country means reproduce the build's ancestral values exactly for all 70 matched ancestries. Reliability is $w_c = \hat{\sigma}^2_{\text{between}}/(\hat{\sigma}^2_{\text{between}} + \widehat{\text{var}}_c/n_c)$ with $\hat{\sigma}^2_{\text{between}}$ by method of moments across the 98-country zero-sum universe. The Serbia-Montenegro-era codes `CS` (1996) and `CSS` (2006) are retained explicitly.

Table A17: Main effects under measurement corrections of the zero-sum means

	Baseline	Wave-adjusted	EB-shrunken	Shrunk + wave-adj.	Reliability WLS
Zero-sum beliefs	0.042***	0.041***	0.042***	0.041***	0.042***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.013)
Uncertainty avoidance	0.038***	0.036***	0.038***	0.036***	0.038***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
Long-term orientation	-0.025**	-0.026**	-0.025**	-0.026**	-0.025**
	(0.011)	(0.011)	(0.011)	(0.011)	(0.011)
Education	-0.101***	-0.101***	-0.101***	-0.101***	-0.101***
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Observations	10,718	10,718	10,718	10,718	10,718
Ancestry clusters	70	70	70	70	70

Notes: Joint specification of Table 2, column 4. Columns replace the baseline ancestral zero-sum means with wave-adjusted means (two-way country-by-wave composition adjustment), empirical-Bayes shrunken means, or both; the final column weights observations by first-stage reliability. All zero-sum variants are standardized with the baseline affine map, so coefficients are comparable across columns. Uncertainty avoidance and long-term orientation (Hofstede) are unchanged throughout. Standard errors in parentheses are clustered by ancestral country. Coding directions as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A18: Income-gradient tests under measurement corrections

	Baseline	Wave-adjusted	EB-shrunken
Zero-sum beliefs	0.043***	0.042***	0.043***
	(0.013)	(0.013)	(0.013)
Uncertainty avoidance	0.037***	0.035***	0.037***
	(0.010)	(0.010)	(0.010)
Long-term orientation	-0.027**	-0.028**	-0.027**
	(0.012)	(0.012)	(0.012)
Position (own education)	-0.089***	-0.090***	-0.089***
	(0.016)	(0.016)	(0.016)
Zero-sum × position	0.005	0.004	0.005
	(0.010)	(0.010)	(0.010)
UA × position	-0.012	-0.013	-0.012
	(0.009)	(0.009)	(0.009)
LTO × position	-0.020*	-0.020*	-0.020*
	(0.011)	(0.011)	(0.011)
Observations	10,718	10,718	10,718
Ancestry clusters	70	70	70

Notes: The own-education gradient specification of Table A9, column 1, re-estimated with the indicated zero-sum measurement variant (interactions rebuilt from the variant). Standard errors in parentheses are clustered by ancestral country. Coding directions as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A19: Robustness to the Chinoy et al. first-generation zero-sum index

	Baseline	Covered, WVS	Covered, Chinoy	Covered, both
Zero-sum beliefs (WVS/EVS)	0.042***	0.034*		0.034*
	(0.013)	(0.018)		(0.018)
Zero-sum beliefs (Chinoy et al.)			-0.008	-0.005
			(0.010)	(0.009)
Uncertainty avoidance	0.038***	0.040***	0.039***	0.040***
	(0.010)	(0.010)	(0.010)	(0.010)
Long-term orientation	-0.025**	-0.014	-0.015	-0.017
	(0.011)	(0.013)	(0.013)	(0.013)
Education	-0.101***	-0.104***	-0.103***	-0.104***
	(0.016)	(0.019)	(0.019)	(0.019)
Observations	10,718	7,109	7,109	7,109
Ancestry clusters	70	38	38	38

Notes: The Chinoy et al. index is the mean of their four-item zero-sum index (first principal component, reconstructed from the replication package following their `00_make_analysis_data.R`) among first-generation immigrants by country of birth, retained where at least five respondents underlie the mean (median cell: ten respondents; implied first-stage reliability 0.28); higher = more zero-sum, and the index is standardized within the covered estimation sample. "Covered" columns restrict the baseline sample to the 38 ancestries with a Chinoy index. All columns use the joint specification of Table 2, column 4. Standard errors in parentheses are clustered by ancestral country. Coding directions otherwise as in Table 2. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A20: Semi-structural extensions: normalization, measurement link, and parental transmission

	Baseline ($k_0 = 0.35$)	$k_0 = 0.20$	$k_0 = 0.50$	Midrank outcome	Mother foreign-born	Father foreign-born
Midpoint-education motive, A_0	0.263***	0.150***	0.375***	0.201***	0.259***	0.263***
	(0.012)	(0.007)	(0.017)	(0.013)	(0.019)	(0.014)
Education-gap loading, A_e	0.029***	0.017***	0.042***	0.029***	0.030***	0.028***
	(0.007)	(0.004)	(0.010)	(0.008)	(0.009)	(0.008)
UA shock scale, $\lambda_{\max}$	0.030***	0.017***	0.043***	0.039***	0.032***	0.025*
	(0.010)	(0.006)	(0.014)	(0.011)	(0.010)	(0.014)
Static tax-cost loading, k_s	0.047	0.027	0.067	0.072	0.044	0.035
	(0.033)	(0.019)	(0.047)	(0.046)	(0.053)	(0.037)
Future growth-cost loading, k_g	0.071***	0.041***	0.102***	0.115***	0.068**	0.081***
	(0.023)	(0.013)	(0.033)	(0.034)	(0.029)	(0.030)
AME: zero-sum beliefs	0.145	0.145	0.145	0.155	0.139	0.135
AME: uncertainty avoidance	0.041	0.041	0.041	0.048	0.044	0.033
AME: long-term orientation	-0.023	-0.023	-0.023	-0.021	-0.017	-0.042
Observations	10,010	10,010	10,010	10,010	4,977	5,033
Ancestry clusters	70	70	70	70	66	69
Mean fitted tax index	0.682	0.682	0.682	0.500	0.679	0.685
Within partial correlation	0.077	0.077	0.077	0.074	0.077	0.078
Jacobian condition number	2,275	2,275	2,275	3,826	2,394	2,191

Notes: NLS estimates of the preferred-tax equation (education-based position; Section 3.4). Columns 2–3 change the k_0 normalization; column 4 replaces equally spaced outcome scores with midrank (ridit) scores; columns 5–6 estimate separately by which parent is foreign-born. Average marginal effects are two-sided numerical derivatives of the constrained preferred-tax function on unit input scales. Standard errors in parentheses are clustered by ancestral country. Coding directions as in Table 7. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A21: Semi-structural leave-one-ancestry-out stability

	Baseline	Min	Max	Sign stable
AME: zero-sum beliefs	0.145	0.117	0.178	70/70
AME: uncertainty avoidance	0.041	0.033	0.047	70/70
AME: long-term orientation	-0.023	-0.029	-0.012	70/70

Notes: Average marginal effects from 70 NLS re-estimations, each dropping one ancestral country.

Table A22: Cross-country correlates of ancestral zero-sum belief: institutions versus worldview

	Pearson r	Spearman ρ	n
Panel A: institutions and extraction (higher = cleaner/stronger)			
Control of corruption (WBGI)	+0.07	+0.09	93
Rule of law (WBGI)	+0.05	+0.07	93
Government effectiveness (WBGI)	+0.03	+0.04	93
Regulatory quality (WBGI)	+0.02	+0.04	93
Corruption Perceptions Index	+0.03	+0.04	92
Panel B: beliefs and worldview			
Freedom of choice and control over life (a173)	-0.39	-0.39	95
Others try to be fair vs. take advantage (a168)	-0.32	-0.30	73
Life satisfaction (a170)	-0.33	-0.35	95
Success is hard work vs. luck (e040)	+0.08	+0.06	95

Notes: Country-level Pearson and Spearman correlations of the ancestral zero-sum mean (11−`e041`, higher = more zero-sum) with governance and worldview measures, across the WVS/EVS countries for which the item is available (pairwise n shown). Governance indices are from QOG (WBGI and Transparency International); `a173`, `a168`, `a170`, and `e040` are WVS/EVS country means. Panel A shows that zero-sum belief is essentially unrelated to institutional quality; Panel B shows it tracks a low-agency, low-trust worldview. Descriptive, without controls.

B An Extractive-Belief Variant of the Model

This appendix formalizes the interpersonal-extraction reading of zero-sum thinking contrasted with the fixed-pie tax-cost model in Remark 5, and derives the margins on which the two readings make opposite predictions. It is the reserve variant referred to in the two-meanings discussion; the main text keeps the tax-cost model as primary and uses this variant to define the discriminating tests.

Motivation. The extraction reading is closest to how zero-sum thinking is defined in the literature this paper builds on, which is why it is worth formalizing rather than dismissing. Foster (1965)'s image of limited good already contains both components. Its premise is generative: land, wealth, and status are believed to exist in fixed and short supply. From that fixed pie Foster derives the appropriative conclusion that one family's advance must come at the expense of others, which is why accumulation draws envy and sanction (Foster, 1967). The two are linked in the folk concept, generativity supplying the premise and appropriation the conclusion, yet they are separable economic primitives: a fixed pie can be divided by luck rather than taking, and appropriation can occur even as the pie grows. The modern economics statement is the same. Chinoy et al. (2026) define zero-sum thinking as the view that any gain by one party necessitates a corresponding loss by another, and the redistributive implication they draw runs through extraction: if the wealth and income of some come at the expense of others, government should correct it through the income tax. The psychology tradition frames the construct the same way (Różycka-Tran et al., 2015; Davidai and Ongis, 2019). The WVS item inherits the same duality in its two anchors: one pole is generative ("wealth can grow so there is enough for everyone"), the other appropriative ("people can only get rich at the expense of others"), so moving along the scale shifts both readings at once and a single item cannot separate them. In the paper's terms, the main text isolates the *generativity* content (whether output responds to effort and investment, entering as the elasticity of the tax base), while the variant developed here formalizes its *appropriation* content, entering

expected income. The former is the novel reading; the latter, foregrounded in the recent literature, is the more literal one.

Belief. In the baseline model zero-sum thinking θ operates on the *perceived cost* of taxation, lowering it by (1)–(2), and leaves expected income unchanged. The extraction variant instead places θ in the *future income process*: period-2 pre-tax income is

$$y_{i2} = \alpha_i - \theta(\bar{y} - \alpha_i) + \mu_i = \alpha_i - \theta A_i + \mu_i, \tag{B.1}$$

with $A_i \equiv \bar{y} - \alpha_i$ as in the text. A zero-sum believer expects the below-mean ($A_i > 0$) to be extracted from and the above-mean ($A_i < 0$) to extract from others, in proportion to their distance from the mean and to the strength of the belief. The construction is literally zero-sum for a population sharing θ, and more generally whenever $\mathbb{E}[\theta_i A_i] = 0$: averaging (B.1) across agents then gives mean period-2 income $\bar{y}$, so extraction reallocates the pie without changing its size, matching the survey wording of items that ask whether one person's gain is another's loss (Chinoy et al., 2026). To isolate the extraction channel, taxation here carries only the administrative cost k_0 ($k_s = k_g = 0$), so θ operates through beliefs about *who ends up with income*, not through beliefs about *how much income taxation destroys*. This makes the two models non-nested by construction: the encompassing specification would place θ in both the numerator and the denominator, and the tests below discriminate between the polar cases rather than estimating a mixture.

Preferred tax. With linear utility, the maxmin evaluation at $\mu_i = -\lambda$, and durable policy as in the text, agent i maximizes

$$V_i(\tau) = (1-\tau)\alpha_i + T(\tau) + \beta\left[(1-\tau)(\alpha_i - \theta A_i - \lambda) + T(\tau)\right], \qquad T(\tau) = \tau\bar{y} - \tfrac{1}{2}k_0\tau^2\bar{y}.$$

The first-order condition gives the interior preferred rate

$$\tau_i^{\circ,\text{ext}} = \frac{A_i(1+b\theta) + b\lambda}{k_0\,\bar{y}}, \qquad b = \frac{\beta}{1+\beta}, \tag{B.2}$$

to be compared with the baseline $\tau_i^{\circ} = (A_i + b\lambda)/\{[k_0 + (1-\theta)k_s + b(1-\theta)k_g]\bar{y}\}$ of (7). The contrast is exact: in the baseline θ sits in the *denominator* and multiplies no income term; in (B.2) it sits in the *numerator* and multiplies the redistributive motive A_i, producing a direct future-weight $\times$ zero-sum $\times$ income term $b\theta A_i$.

The discriminating margins. Equations (7) and (B.2) agree on the signs of the direct effects but split on exactly the margins examined in Section 5.1.

Proposition 5 (Extraction fingerprints)**.** *Under the extraction variant* (B.2)*, away from corners:*

(i) Extensive margin. *The agent supports positive redistribution, $\tau_i^{\circ,\text{ext}} > 0$, if and only if $\alpha_i < \bar{y} + b\lambda/(1+b\theta)$. The support threshold is weakly decreasing in θ, strictly so when $b\lambda > 0$: stronger zero-sum beliefs then withdraw support from above-mean agents. In the tax-cost model the same threshold is $\alpha_i < \bar{y} + b\lambda$, independent of θ (Proposition 4).*

(ii) Position of the zero-sum effect. *$\partial\tau_i^{\circ,\text{ext}}/\partial\theta = bA_i/(k_0\bar{y})$, which is positive for below-mean agents ($A_i > 0$) and* negative for above-mean agents *($A_i < 0$): the zero-sum effect declines as income rises and turns oppositional among the well-off, with cross-partial $\partial^2\tau_i^{\circ,\text{ext}}/\partial A_i\,\partial\theta =$*

$b/(k_0\bar{y}) > 0$ *(equivalently,* $-b/(k_0\bar{y}) < 0$ *with respect to income* α_i*). In the tax-cost model, for any supporter the zero-sum effect* $\partial\tau_i^{\circ}/\partial\theta = (A_i + b\lambda)(k_s + bk_g)/(D(\theta, b)^2\bar{y})$ *has the sign of the numerator and is therefore* positive at every income position*, including above the mean.*

(iii) Future weight. $\partial\tau_i^{\circ,\text{ext}}/\partial b = (A_i\theta + \lambda)/(k_0\bar{y}) \geq 0$ *for every supporter: with the cost channel switched off, the extraction variant cannot generate a negative future-weight effect at all.*

Proof. Part (i): $\tau_i^{\circ,\text{ext}} > 0$ iff $A_i(1 + b\theta) + b\lambda > 0$ iff $\bar{y} - \alpha_i > -b\lambda/(1 + b\theta)$; differentiating the threshold $\bar{y} + b\lambda/(1 + b\theta)$ in θ gives $-b^2\lambda/(1 + b\theta)^2 \leq 0$, with strict inequality iff $b\lambda > 0$. Part (ii): differentiate (B.2) in θ, and again in A_i; the tax-cost expression differentiates (7) in θ with $D(\theta, b) = k_0 + (1 - \theta)(k_s + bk_g)$. Part (iii): differentiate (B.2) in b; both $A_i\theta$ and λ are nonnegative for a supporter with $A_i > 0$. □

Why the paper keeps the tax-cost model primary. Both models predict a positive direct zero-sum effect, so the direct-effects evidence does not separate them. Both also predict a steeper income gradient in zero-sum belief, and it is worth noting that this shared prediction is not borne out: the zero-sum × position interaction is 0.005 (s.e. 0.010) in Table A9, a precise zero. By Lemma 2(iii) that interaction need not carry the sign of either model's cross-partial once the response is ordinal, which is why it is not treated as decisive against either.

They part company on Proposition 5, and Section 5.1 takes the discriminating tests to the data. The extensive-margin test in part (i) cannot be run here: with 97 percent of respondents above the lowest cutpoint the margin is weakly powered for every regressor, and the fitted structural model places every respondent strictly inside the interior, so no observed category corresponds to the model's corner. The decisive test is part (ii). Among above-median-education respondents the zero-sum effect is positive at every cutpoint and, if anything, larger than in the full sample, rather than negative and declining. Part (iii) points the same way for a different reason: the negative future-weight coefficient reported in Table 2 is one the pure-extraction variant cannot produce, though this is a weaker argument than it appears, since the sign restriction is an artifact of switching off the cost channel to isolate extraction, and Section 7.3 in any case declines to interpret that coefficient as a future-weight effect. The cross-country correlations of Table A22 establish the narrower point that the measure is a belief rather than a report on extractive institutions; they do not themselves separate the two readings of that belief.

The extraction variant is retained because it is closer to the interpersonal wording of some zero-sum survey items and because it defines the tests. The pre-committed decision rule is that a second zero-sum item exhibiting the fingerprints of Proposition 5 would reopen the question.

References

Acemoglu, Daron. 2025. "The Simple Macroeconomics of AI." *Economic Policy* 40(121): 13–58.

Acemoglu, Daron, and Pascual Restrepo. 2020. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128(6): 2188–2244.

Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica* 90(5): 1973–2016.

Alesina, Alberto, and George-Marios Angeletos. 2005. "Fairness and Redistribution." *American Economic Review* 95(4): 960–980.

Alesina, Alberto, and Paola Giuliano. 2011. "Preferences for Redistribution." In *Handbook of Social Economics*, Vol. 1A, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 93–131. Amsterdam: North-Holland.

Alesina, Alberto, and Eliana La Ferrara. 2005. "Preferences for Redistribution in the Land of Opportunities." *Journal of Public Economics* 89(5–6): 897–931.

Algan, Yann, and Pierre Cahuc. 2010. "Inherited Trust and Growth." *American Economic Review* 100(5): 2060–2092.

Alesina, Alberto, Stefanie Stantcheva, and Edoardo Teso. 2018. "Intergenerational Mobility and Preferences for Redistribution." *American Economic Review* 108(2): 521–554.

Ali, S. Nageeb, Maximilian Mihm, and Lucas Siga. 2025. "The Political Economy of Zero-Sum Thinking." *Econometrica* 93(1): 41–70.

Benabou, Roland, and Efe A. Ok. 2001. "Social Mobility and the Demand for Redistribution: The POUM Hypothesis." *Quarterly Journal of Economics* 116(2): 447–487.

Benabou, Roland, and Jean Tirole. 2006. "Belief in a Just World and Redistributive Politics." *Quarterly Journal of Economics* 121(2): 699–746.

Bisin, Alberto, and Thierry Verdier. 2001. "The Economics of Cultural Transmission and the Dynamics of Preferences." *Journal of Economic Theory* 97(2): 298–319.

Bursztyn, Leonardo. 2016. "Poverty and the Political Economy of Public Education Spending: Evidence from Brazil." *Journal of the European Economic Association* 14(5): 1101–1128.

Chandrasekher, Madhav, Mira Frick, Ryota Iijima, and Yves Le Yaouanq. 2022. "Dual-Self Representations of Ambiguity Preferences." *Econometrica* 90(3): 1029–1061.

Chinoy, Sahil, Nathan Nunn, Sandra Sequeira, and Stefanie Stantcheva. 2026. "Zero-Sum Thinking and the Roots of US Political Differences." *American Economic Review* 116(3): 1052–1096.

Davidai, Shai, and Martino Ongis. 2019. "The Politics of Zero-Sum Thinking: The Relationship between Political Ideology and the Belief That Life Is a Zero-Sum Game." *Science Advances* 5(12): eaay3761.

Dohmen, Thomas, Armin Falk, David Huffman, and Uwe Sunde. 2012. "The Intergenerational Transmission of Risk and Trust Attitudes." *Review of Economic Studies* 79(2): 645–677.

Eriksson, Kimmo, Pontus Strimling, Michele Gelfand, et al. 2021. "Perceptions of the Appropriate Response to Norm Violation in 57 Societies." *Nature Communications* 12: 1481.

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

Fernández, Raquel. 2011. "Does Culture Matter?" In *Handbook of Social Economics*, Vol. 1A, edited by Jess Benhabib, Alberto Bisin, and Matthew O. Jackson, 481–510. Amsterdam: North-Holland.

Fernández, Raquel, and Alessandra Fogli. 2009. "Culture: An Empirical Investigation of Beliefs, Work, and Fertility." *American Economic Journal: Macroeconomics* 1(1): 146–177.

Foster, George M. 1965. "Peasant Society and the Image of Limited Good." *American Anthropologist* 67(2): 293–315.

Foster, George M. 1967. *Tzintzuntzan: Mexican Peasants in a Changing World*. Boston: Little, Brown and Company.

Galor, Oded, and Ömer Özak. 2016. "The Agricultural Origins of Time Preference." *American Economic Review* 106(10): 3064–3103.

Gallego, Aina, and Thomas Kurer. 2022. "Automation, Digitalization, and Artificial Intelligence in the Workplace: Implications for Political Behavior." *Annual Review of Political Science* 25: 463–484.

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci. 2004. "Differentiating Ambiguity and Ambiguity Attitude." *Journal of Economic Theory* 118(2): 133–173.

Gilboa, Itzhak, and David Schmeidler. 1989. "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics* 18(2): 141–153.

Hofstede, Geert. 2001. *Culture's Consequences: Comparing Values, Behaviors, Institutions, and Organizations Across Nations*. 2nd ed. Thousand Oaks, CA: Sage.

Gmyrek, Paweł, Janine Berg, and David Bescond. 2023. "Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality." ILO Working Paper 96. Geneva: International Labour Organization.

International Monetary Fund. 2024. "AI Will Transform the Global Economy. Let's Make Sure It Benefits Humanity." IMF Blog, January 14. Washington, DC: International Monetary Fund.

Klibanoff, Peter, Massimo Marinacci, and Sujoy Mukerji. 2005. "A Smooth Model of Decision Making under Ambiguity." *Econometrica* 73(6): 1849–1892.

Korinek, Anton, and Joseph E. Stiglitz. 2019. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, 349–390. Chicago: University of Chicago Press.

Kuziemko, Ilyana, Michael I. Norton, Emmanuel Saez, and Stefanie Stantcheva. 2015. "How Elastic Are Preferences for Redistribution? Evidence from Randomized Survey Experiments." *American Economic Review* 105(4): 1478–1508.

Ljunge, Martin. 2014. "Trust Issues: Evidence on the Intergenerational Trust Transmission among Children of Immigrants." *Journal of Economic Behavior & Organization* 106: 175–196.

Luttmer, Erzo F. P., and Monica Singhal. 2011. "Culture, Context, and the Taste for Redistribution." *American Economic Journal: Economic Policy* 3(1): 157–179.

Meltzer, Allan H., and Scott F. Richard. 1981. "A Rational Theory of the Size of Government." *Journal of Political Economy* 89(5): 914–927.

OECD. 2023. *OECD Employment Outlook 2023: Artificial Intelligence and the Labour Market*. Paris: OECD Publishing.

Piketty, Thomas. 1995. "Social Mobility and Redistributive Politics." *Quarterly Journal of Economics* 110(3): 551–584.

Różycka-Tran, Joanna, Paweł Boski, and Bogdan Wojciszke. 2015. "Belief in a Zero-Sum Game as a Social Axiom: A 37-Nation Study." *Journal of Cross-Cultural Psychology* 46(4): 525–548.

Stantcheva, Stefanie. 2021. "Understanding Tax Policy: How Do People Reason?" *Quarterly Journal of Economics* 136(4): 2309–2369.

Thewissen, Stefan, and David Rueda. 2019. "Automation and the Welfare State: Technological Change as a Determinant of Redistribution Preferences." *Comparative Political Studies* 52(2): 171–208.

Varian, Hal R. 1980. "Redistributive Taxation as Social Insurance." *Journal of Public Economics* 14(1): 49–68.

Wang, Austin Horng-En. 2018. "Patience Moderates the Class Cleavage in Demand for Redistribution." *Social Science Research* 70: 18–27.

Zimbardo, Philip G., and John N. Boyd. 1999. "Putting Time in Perspective: A Valid, Reliable Individual-Differences Metric." *Journal of Personality and Social Psychology* 77(6): 1271–1288.