Dreaming rich: How inequality manufactures overoptimism

Julia Baumann[1]   Yiming Liu[2]

September 25, 2026

Abstract

The prospect of upward mobility (POUM) hypothesis posits that high inequality can be sustained with high mobility: people oppose redistribution because they expect to move up. Empirically, however, higher inequality is associated with lower mobility. We resolve this puzzle by showing that inequality generates motivated beliefs about upward mobility. In survey data across 27 countries, higher inequality is associated with greater expected mobility despite lower actual mobility, which correlates with weaker demand for redistribution. An online experiment establishes causality: random assignment to higher inequality increases overoptimism about reaching the top-earning position, holding actual mobility constant. Effort does not differ across treatments, and higher inequality reduces beliefs about reaching other high-ranking but lower-paying positions, consistent with payoff-driven motivated distortions but not performance overconfidence. This induced overoptimism persists despite accurate feedback and causally reduces demand for redistribution. We conclude by evaluating welfare implications, weighing anticipatory utility against the resulting higher inequality.

Keywords: Income inequality, Inequality acceptance, Motivated beliefs, Optimism, Overconfidence
JEL Codes: D31, D63, D83, H24

[1]Humboldt-Universität zu Berlin and WZB, julia.baumann@wzb.eu
[2]Humboldt-Universität zu Berlin and WZB, yiming.liu@hu-berlin.de
The design was approved by the institutional review board of WZB (Ethics Applications No. 208) and pre-registered at the AEA RCT Registry under AEARCTR-0014848.

1 Introduction

Ever since De Tocqueville (1835), a prominent explanation for why democratic societies tolerate high inequality is that citizens believe in upward mobility (Bénabou and Ok, 2001; Alesina et al., 2018). If people expect that they or their children will move up the income ladder, they are more willing to accept large income gaps today. This prospect of upward mobility (POUM) hypothesis implies a positive relationship between inequality and mobility: the societies that sustain the most inequality should be those where people are most likely to move up. Yet across countries the relationship runs the other way: higher inequality is associated with *lower* actual mobility, the regularity known as the Great Gatsby Curve (Corak, 2013). We replicate this pattern in Panel (a) of Figure 1 using the International Social Survey Programme (ISSP), which asks respondents in 27 countries to place both themselves and their parents on a social ladder: the share who have risen above their parents is lower where inequality is higher, the opposite of what POUM appears to require.

To resolve this tension, we start from a distinction central to POUM: the hypothesis requires only that people *believe* they will move up, not that they actually do. We propose and test a novel psychological mechanism through which high inequality itself produces this belief, leaving citizens in more unequal societies more optimistic about reaching the top even as they are less likely to get there.

Figure 1: Inequality, upward mobility, and the experience-expectation gap

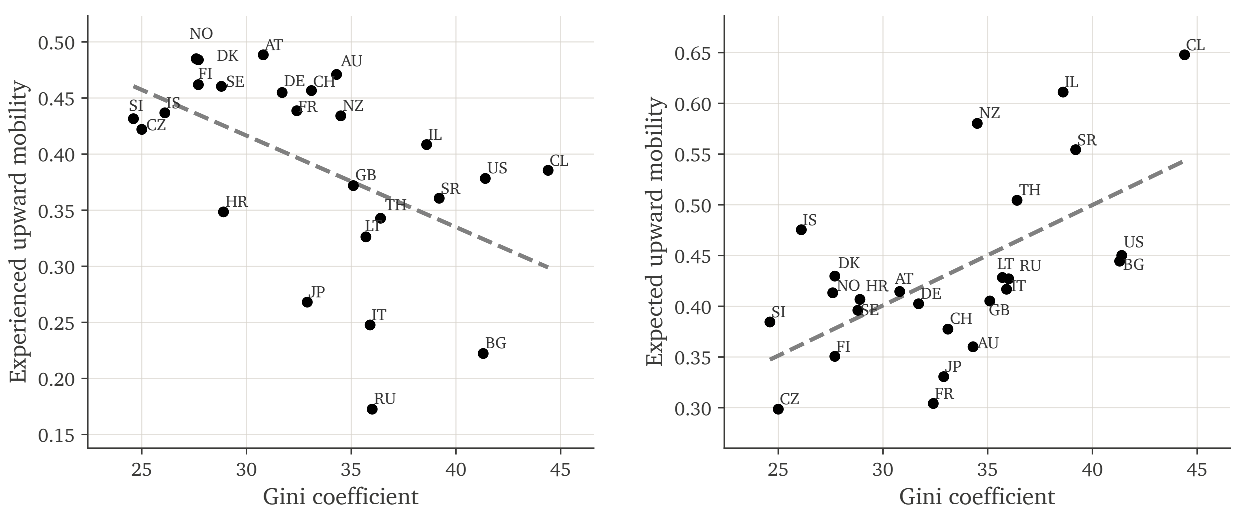

(a) Gini and actual mobility **(b) Gini and expected mobility**

Notes: This figure plots income inequality against actual and expected upward mobility across the countries in the 2019 International Social Survey Programme (ISSP), a survey that asks respondents to place their parents, themselves today, and themselves in the future on a social ladder from 1 to 10. South Africa and the Philippines, two outlier countries, are excluded from the panels. Inequality is the Gini coefficient of disposable income. Panel (a) plots actual mobility, the share of respondents who place themselves above their parents, and replicates the Great Gatsby Curve (Corak, 2013). Panel (b) plots expected mobility, the share who expect a higher status in 10 years than today. The dashed line in each panel is the linear fit.

Our proposed mechanism builds on models of anticipatory utility, in which individuals derive psychological benefit from believing in favorable future outcomes and weigh this benefit against the decision cost of acting on distorted beliefs (Brunnermeier and Parker, 2005). Consider two societies with the same total income but different distributions. In the more unequal society, income is more concentrated at the top, and the reward at the top is larger. A larger reward at the top means that individuals have more to gain psychologically from believing they will obtain it, strengthening the incentive to distort beliefs upward. Holding the decision cost fixed, this model makes a clear prediction: higher inequality induces greater (over)optimism about one's chances of reaching the top.

Testing this prediction in the same ISSP data, Panel (b) of Figure 1 shows that the share of respondents who expect to move up on the social ladder within the next 10 years is significantly *higher* in more unequal countries. Combined with the negative relationship between inequality and actual upward mobility, this produces an experience-expectation gap: citizens in more unequal societies are more optimistic about moving up despite having experienced less upward social mobility.[1]

This experience-expectation gap is consistent with our prediction, but it could reflect unobserved institutional or cultural differences across countries. To control for such differences, we exploit within-country variation in experienced inequality across birth cohorts. Following Roth and Wohlfart (2018), we measure experienced inequality as the top income concentration each respondent experienced during their formative years (ages 18 to 25) and estimate its association with expected mobility controlling for country fixed effects, subregion-by-cohort fixed effects, family-of-origin status, and six macroeconomic experiences (such as inflation and unemployment rates) during the same formative years that could independently shape expectations. A one standard deviation increase in experienced inequality is associated with an increase of 0.16 standard deviations in expected mobility on the social ladder ($p < 0.05$). Moreover, the increase in optimism concentrates on positions associated with above-mean income, consistent with our conceptual framework. Because income is unequally distributed, mean income lies above the middle rung of the ladder. Some positions therefore rank above the median in status yet still pay below the country mean. Mapping ladder rungs to household income within each country, we find that experienced inequality is associated with a higher probability of expecting a position that pays above the country mean, but not one that pays at or below it, even those ranking above the median in status.

This optimism has a direct implication for the demand for redistribution. The POUM hypothesis predicts that the more people expect to move up, the less redistribution they support. In the ISSP data we find the corresponding association: respondents with higher expected mobility report weaker demand for redistribution, holding current status and in-

[1]This experience-expectation gap is robust across multiple dimensions. It holds when inequality is measured by the top-5% income share instead of the Gini coefficient, and when mobility is measured as a continuous gap between expected and current status, as status persistence as in Corak (2013), or as a rags-to-riches transition rate on the social ladder. It also holds when South Africa and the Philippines, the two outlier countries excluded from Figure 1, are included (Appendix B). The correlation between inequality and expected mobility survives country-level controls such as average income and GDP growth.

dividual controls fixed. This pattern is consistent with a channel through which inequality raises optimism, which in turn lowers demand for redistribution.

In Section 4, we complement the survey evidence with an online experiment designed to test, in a controlled setting, whether inequality causally produces overoptimism through the motivated beliefs mechanism. Subjects are randomly assigned to societies with either high or low initial inequality. Each society consists of five individuals: a single "Rich" individual and four "Poor" individuals. In the High inequality condition, the Rich earns \$10 while each Poor earns \$0. In the Low inequality condition, the Rich earns \$2.40 while each Poor earns \$1.90. After being informed of the initial inequality, all subjects complete a quiz, with their rank determined by a combination of quiz performance (80% chance) and luck (20% chance). The highest-ranked individual, Rank 1, becomes the Rich. After taking the quiz but before learning their rank, subjects state their beliefs about becoming Rich (with incentives for accuracy), along with their beliefs about obtaining each of the remaining ranks. This design allows us to test directly whether exogenous exposure to high inequality induces greater overoptimism about obtaining the high-paying Rank 1 position.

We find that higher inequality causally raises overoptimism about reaching the high-paying position. The experiment allows us to cleanly define and measure overoptimism, because the objective probability of becoming Rich is known to be 20% for the average subject in both conditions. Subjects in both conditions display substantial overoptimism about Rank 1, believing they have over a 30% chance of becoming Rich. The effect of initial inequality is significant: subjects in the High inequality condition state a 41.3% chance of becoming Rich, compared to 34.4% in the Low inequality condition ($p < 0.01$), a 7 percentage point difference despite identical objective probabilities. This inequality-induced overoptimism cannot be explained by differences in effort or performance: we find no significant differences between conditions in quiz scores, time spent on the quiz, self-reported effort, or beliefs about others' effort.

We also find supporting evidence that the higher overoptimism observed in the High inequality condition is motivated rather than reflecting general optimism about performance. If subjects in the High inequality condition simply believed they performed better on the quiz overall, this performance optimism would systematically affect beliefs about all above-median ranks. This is in contrast to the motivated beliefs explanation, which only predicts an increase in the belief for obtaining the only high-paying position, Rank 1. To test this mechanism, we examine how beliefs change across all five possible ranks. Using simulations based on our empirical performance distribution, we find that plausible improvements in believed performance would increase both mean and median Rank-2 beliefs substantially. Yet our data show the opposite pattern: while Rank-1 beliefs rise significantly in the High inequality condition at both the mean (41.3% vs. 34.4%) and the median (35.0% vs. 24.0%), Rank-2 beliefs remain unchanged at the median (20.0% vs. 20.0%) and actually decrease at the mean (18.6% vs. 20.6%). This selective increase, appearing only for the high-paying position, aligns with the motivated beliefs explanation, but is inconsistent with performance optimism.

We further test whether this overoptimism is motivated by adding a Feedback treatment, in which subjects learn their quiz score and their objective probability of becoming Rich before stating their beliefs. If the overoptimism reflected missing information rather than motivated beliefs, accurate feedback should eliminate the effect of initial inequality on beliefs. Feedback does reduce overoptimism: the average Rank-1 belief falls from 37.8% in the Baseline treatment to 31.4% in the Feedback treatment ($p < 0.01$). Yet most subjects remain overoptimistic at the individual level, stating a belief above their own objective probability. The effect of initial inequality also persists. Subjects in the High inequality condition still state a higher Rank-1 belief than those in the Low inequality condition (33.7% vs. 29.3%, $p < 0.10$). Because each subject is told their objective probability, we can measure belief distortion directly as the deviation of the stated belief from that probability. This deviation is larger in the High inequality condition (16.2 vs. 11.6 percentage points, $p < 0.10$), and most pronounced among below-median performers. The persistence of inequality-induced overoptimism despite accurate feedback, especially among below-median performers, is consistent with motivated beliefs but not with missing information.

The inequality-induced overoptimism causally raises inequality acceptance. In a second period, before learning their actual rank, subjects choose their preferred income allocation between the Rich and the Poor, with one randomly selected subject's choice implemented for the entire society. Subjects in the High inequality condition allocate on average 48.0% to the Rich, compared to 40.8% in the Low inequality condition ($p < 0.01$). To establish that this effect runs through beliefs rather than through a direct shift in fairness views or anchoring, we exploit the Feedback treatment as a source of exogenous variation in beliefs. Because the feedback tells each subject their own objective probability of becoming Rich, it moves beliefs in opposite directions depending on their quiz score: it lowers the overoptimism of low performers and raises the beliefs of high performers. The effect of the feedback on the share to the Rich is flat to negative where the feedback lowers beliefs and positive where it raises them. A formal instrumental-variables analysis, using the interaction of feedback and quiz score as the instrument, confirms that beliefs causally raise inequality acceptance. Together with the ISSP evidence that higher expected mobility is associated with weaker demand for redistribution, these results identify a channel through which inequality sustains itself by inducing overoptimism that makes it more acceptable.

This paper contributes to the Prospect of Upward Mobility (POUM) literature (Piketty, 1995; Bénabou and Ok, 2001) by endogenizing mobility beliefs. Previous work documents that mobility beliefs matter for redistribution preferences but does not explain why beliefs are systematically too optimistic in high-inequality societies. Alesina et al. (2018) show that Americans are more optimistic than Europeans about social mobility and that these beliefs correlate with lower demand for redistribution, but do not address why citizens in more unequal societies are more optimistic despite facing lower actual mobility, as documented by the Great Gatsby Curve (Corak, 2013; Chetty et al., 2017). In Piketty (1995), agents rationally update beliefs based on their own mobility experience. In this model, beliefs should on average correspond to actual mobility, yet they do not. Closest to our conceptual setting is Bénabou and Tirole (2006), who endogenize beliefs about the fairness of the economic system. In their setting, individuals are motivated to form beliefs that justify how outcomes in their society are determined, through merit or luck. Our mechanism operates through anticipatory utility rather than system justification. We resolve the tension between actual and expected social mobility by showing that inequality itself shapes mobility beliefs. Because higher inequality increases the reward at the top, it raises the anticipatory utility from believing that one will get there (Brunnermeier and Parker, 2005), strengthening the incentive to distort mobility beliefs upward. Whereas the system justification mechanism emphasizes the value of viewing economic outcomes as fair and effort as rewarded, our mechanism operates through the value of optimistic expectations about one's own future position. Our mechanism thus provides a microfoundation for the POUM hypothesis that is consistent with the Great Gatsby Curve, and offers an explanation for the persistent overoptimism about upward mobility in the US (Davidai and Gilovich, 2015).

We also ask whether information can correct these belief distortions. Our findings speak to the growing literature on information provision and belief updating about economic conditions, in which a series of influential experiments provide subjects with accurate information about inequality, income rank, or tax policy, studying the effects on beliefs and preferences (Cruces et al., 2013; Kuziemko et al., 2015; Karadja et al., 2017; Stantcheva, 2021; Hvidberg et al., 2023). This literature treats misperceptions as correctable: providing information should update beliefs. Our framework offers a different perspective. Because belief distortions serve a psychological function, the anticipatory utility of expecting to reach the top, they may resist correction. Our Feedback treatment provides direct evidence for this distinction: subjects given their own objective probability of becoming Rich still display higher Rank-1 beliefs in the High inequality condition ($p < 0.10$). Missing information alone therefore cannot explain the effect of initial inequality on overoptimism. This complements Kuziemko et al. (2015), who find that information about inequality moves perceptions but barely shifts policy preferences, and Hoy and Mager (2021), who find that individuals told they are poorer than expected become *less*, not more, concerned about inequality. Together, these patterns suggest that belief distortions about economic prospects can operate through motivated channels that information alone cannot address. Moreover, using US data from the General Social Survey, we show that expected status is positively associated with current happiness, controlling for current status, especially among low-income respondents. This suggests that the beliefs that sustain inequality provide psychological benefits to those disadvantaged by it, which raises a welfare dilemma: correcting overoptimistic mobility beliefs may improve decisions in voting for policies but reduce the anticipatory utility of those who need it most.

We also contribute to the literature on motivated beliefs and overconfidence. A large body of research documents that individuals hold overconfident beliefs driven by motivational factors (Bénabou and Tirole, 2002; Brunnermeier and Parker, 2005; Eil and Rao, 2011; Schwardmann and van der Weele, 2019; Zimmermann, 2020; Huffman et al., 2022). This literature has primarily focused on how *intrapersonal* factors, such as ego concerns, risk preferences, or prior performance, shape motivated beliefs. We shift the focus to a *macroeconomic* determinant: we show that the degree of motivated overoptimism responds to the level of inequality in the economic environment. Furthermore, while anticipatory utility is central to theoretical accounts of why individuals distort beliefs (Brunnermeier and Parker, 2005), direct evidence that distorted beliefs actually deliver it has been scarce. One notable exception is Engelmann et al. (2024), who establish that individuals distort their perceptions as a result of anticipatory disutility about losses. The present study contributes a positive analog to this mechanism, where individuals distort beliefs to experience positive anticipatory utility in anticipation of high future gains. Our happiness result provides suggestive field evidence that optimistic beliefs deliver anticipatory utility.

Our individual-level survey evidence contributes to the growing literature on how macroeconomic experiences during formative years persistently shape economic beliefs and behavior. Malmendier and Nagel (2011) show that individuals who experienced low stock market returns during their youth take less financial risk throughout their lives. Subsequent work has extended this to inflation expectations (Malmendier and Nagel, 2016), consumption behavior (Malmendier and Shen, 2024), and political preferences (Alesina and Fuchs-Schündeln, 2007). Giuliano and Spilimbergo (2025) survey this literature. Most closely related to our work, Roth and Wohlfart (2018) demonstrate that experienced macroeconomic conditions, including unemployment and inequality, during one's formative years persistently affect redistribution preferences. We extend this approach to *mobility beliefs*: individuals who experienced higher inequality during their formative years hold more optimistic expectations about reaching the top of the social ladder, even decades later. Consistent with our motivated beliefs mechanism, this optimism is concentrated on positions associated with above-mean income and does not reflect a uniform upward shift.

Our paper also contributes to the literature on social preferences and the demand for redistribution.[2] This literature has established that the inequality people are willing to accept depends on its source and on the fairness views they hold. We add a determinant that operates upstream of these fairness judgments: the level of inequality itself. Higher inequality raises the anticipatory utility of believing one will reach the top and thereby raises optimism about getting there. In our experiment, this induced overoptimism causally raises inequality acceptance. In the ISSP data, higher expected mobility is in turn associated with weaker demand for redistribution. Inequality acceptance thus responds both to how inequality is determined and to a motivated belief about where one expects to end up in the future. This motivated belief is a channel through which the economic environment itself shapes social preferences.

Finally, our cross-country survey evidence contributes to the literature on the Great Gatsby Curve. As coined by Krueger (2012), this literature documented that children in more unequal societies face lower chances of moving up the income ladder (Corak, 2013; Chetty et al., 2017; Durlauf and Seshadri, 2018; Durlauf et al., 2022). We complement this by documenting an "experience-expectation gap" of the Great Gatsby Curve: while actual mobility is lower in more unequal societies, *expected* mobility is higher. Citizens in more unequal countries are more likely to expect upward mobility, and this optimism is concentrated at the top of the ladder. This experience-expectation gap has implications for how we interpret the policy consequences of the Great Gatsby Curve: if citizens in high-inequality societies are more

[2]A large literature has studied the nature of social preferences (Fehr et al., 1993; Rabin, 1993; Fehr and Schmidt, 1999; Bolton and Ockenfels, 2000; Andreoni and Miller, 2002; Charness and Rabin, 2002; Engelmann and Strobel, 2004; Bellemare et al., 2008; Bartling et al., 2015; Cappelen and Tungodden, 2019; Fehr and Vollmann, 2022; Exley and Kessler, 2024), including how the source of inequality shapes inequality acceptance (Konow, 1996, 2000; Alesina et al., 2001; Cappelen et al., 2007; Konow, 2009; Cappelen et al., 2010; Krawczyk, 2010; Balafoutas et al., 2013; Cappelen et al., 2013; Bauer et al., 2014; Durante et al., 2014; Mollerstrom et al., 2015; Akbaş et al., 2019; Cassar and Klein, 2019; Almås et al., 2020; Sugden and Wang, 2020; Müller and Renes, 2021; Cappelen et al., 2022; Barr et al., 2023; Fehr et al., 2024).

optimistic about their prospects despite facing lower actual mobility, then inequality may be more politically sustainable than the objective mobility data alone would suggest.

The remainder of the paper proceeds as follows. Section 2 presents our conceptual framework, and Section 3 the cross-country survey evidence from the International Social Survey Programme. We describe our experimental design in Section 4 and our main experimental results in Section 5. Section 6 concludes.

2 Conceptual framework

We develop a simple model of motivated belief formation based on the optimal expectations framework of Brunnermeier and Parker (2005). In this framework, individuals derive utility both from realized outcomes and from anticipation of future outcomes. Since optimistic beliefs enhance anticipatory utility, individuals have an incentive to distort their beliefs upward. However, distorted beliefs lead to worse decisions, and thus lower realized outcomes, creating a decision cost. Optimal beliefs balance these two forces. A key insight of Brunnermeier and Parker (2005) is that it is always optimal to be at least slightly optimistic: starting from rational beliefs, a small belief distortion leads to *first-order* gains in anticipatory utility but only *second-order* decision costs. We show that this distortion responds to the economic environment: higher inequality increases the incentive to hold optimistic beliefs about reaching the top.

Setup. Consider a society with n individuals. There are n ranked positions $j \in \{1, 2, ..., n\}$, where the top k positions ("Rich") each pay income y_R and the bottom $n-k$ positions ("Poor") each pay $y_P < y_R$. Total income is $Y = k \cdot y_R + (n-k) \cdot y_P$, and mean income is $\bar{y} = Y/n$. Income inequality is captured by π, the share of total income going to Rich positions, which implies $y_R = \pi Y/k$ and $y_P = (1-\pi)Y/(n-k)$. Higher π indicates greater inequality.

Each individual i has an objective probability p_{0ij} of ending up in position j, which may vary across individuals due to differences in ability. Before learning their position, individuals form beliefs about the probability of occupying each position.

Optimal beliefs. Following Brunnermeier and Parker (2005), each individual chooses subjective beliefs $\boldsymbol{\mu}_i = (\mu_{i1}, ..., \mu_{in})$ over positions to maximize well-being, which trades off anticipatory utility against decision cost:[3]

$$W_i(\boldsymbol{\mu}_i) = \underbrace{\sum_{j=1}^{n} \mu_{ij} \cdot u(y_j)}_{\text{anticipatory utility}} - \underbrace{\sum_{j=1}^{n} \alpha \left(\mu_{ij} - p_{0ij}\right)^2}_{\text{decision cost}} \quad \text{s.t.} \quad \sum_{j=1}^{n} \mu_{ij} = 1 \tag{1}$$

[3]The optimization over $\boldsymbol{\mu}_i$ is an "as if" characterization of the outcome of a belief formation process. As Brunnermeier and Parker (2005) noted, their "optimal expectations framework is a simple model that captures complex brain processes" regardless of whether the mapping from objective to subjective probabilities is deliberate or automatic.

The first term is anticipatory utility: it represents the utility individuals experience from contemplating their future income under subjective beliefs. Holding a higher belief about occupying a high-income position increases anticipatory utility. The second term is the decision cost. As in Brunnermeier and Parker (2005), distorted beliefs lead to suboptimal decisions and this reduces realized outcomes. The quadratic form for the decision cost (where $\alpha > 0$) presented here makes the problem more tractable, but our qualitative results hold under any strictly convex cost that is minimized at the objective probability.[4]

Assuming linear utility $u(x) = x$ for tractability and an interior solution, the Lagrangian for the constrained optimization yields the first-order condition for each position j:

$$y_j - 2\alpha(\mu_{ij} - p_{0ij}) = \lambda \tag{2}$$

where λ is the multiplier on the adding-up constraint. Summing over j and using $\sum_j \mu_{ij} = \sum_j p_{0ij} = 1$ gives $\lambda = \bar{y}$, the mean income across positions. The optimal belief distortion for position j is therefore:

$$\mu^*_{ij} - p_{0ij} = \frac{y_j - \bar{y}}{2\alpha} \tag{3}$$

This expression has a clean interpretation: the belief distortion for each position is proportional to how much that position's income deviates from the mean. Individuals become overoptimistic about positions that pay above-mean income ($y_j > \bar{y}$), and underestimate their probability of ending up in positions that pay below-mean income ($y_j < \bar{y}$).

Note the key asymmetry from Brunnermeier and Parker (2005): at the objective beliefs $\mu_{ij} = p_{0ij}$, the marginal anticipatory benefit of shifting probability toward a high-income position is $y_j - \bar{y}$ (first order), while the marginal decision cost $2\alpha(\mu_{ij} - p_{0ij}) = 0$ (second order). It therefore always pays to distort beliefs at least slightly toward optimism about high-income positions.

How inequality affects beliefs. Consider an increase in inequality π holding total income Y fixed: $y_R = \pi Y/k$ increases and $y_P = (1-\pi)Y/(n-k)$ decreases, while $\bar{y} = Y/n$ remains unchanged. From Equation 3, for Rich positions ($j \leq k$), the distortion $\frac{y_R - \bar{y}}{2\alpha}$ is increasing in π, since $\frac{\partial y_R}{\partial \pi} = Y/k > 0$. For Poor positions ($j > k$), by contrast, the distortion $\frac{y_P - \bar{y}}{2\alpha}$ is decreasing in π, since $\frac{\partial y_P}{\partial \pi} = -Y/(n-k) < 0$.[5]

This yields two testable predictions:

> *Prediction 1: Higher inequality increases overoptimism about reaching high-income positions.*

[4] Alternative specifications include the Kullback–Leibler divergence used in Caplin and Leahy (2019). In our experiment, the Binarized Scoring Rule used to incentivize belief reports provides a concrete source of decision cost: inaccurate reports reduce the expected bonus from the belief elicitation.

[5] We hold total income Y fixed as inequality rises, but our predictions are qualitatively the same if we instead hold y_R fixed and lower y_P, or hold y_P fixed and raise y_R, since in each case the gap between y_R and y_P widens, which is what the distortion responds to. We fix total income because it is the canonical Pigou–Dalton and Lorenz notion of rising inequality, and it is consistent with how we treat inequality elsewhere: the cross-country regressions control for GDP per capita, and the experiment holds total income fixed across conditions.

Prediction 2: Higher inequality does not increase beliefs about positions that pay below the mean, even if those positions are above the median rank.

Prediction 2 provides a test to distinguish motivated beliefs from general performance optimism. Under performance optimism, individuals in high-inequality environments believe they performed better on a task determining rank. The general intuition of this alternative is that higher believed performance should lead to a more across-the-board increase in beliefs about above-median positions.[6] The conceptual framework, by contrast, predicts that belief distortion targets specifically those positions with above-mean income. In our setting with $k < n/2$ (fewer Rich than Poor), some above-median-rank positions still pay below-mean income. These positions allow us to tell the two theories apart: the model predicts that beliefs about positions that rank above the median but pay below the mean income should *not* increase with inequality, and may in fact decrease, while performance optimism generally predicts an increase.

Robustness of the framework. The conceptual framework assumes that subjects form beliefs without anticipating the Period 2 allocation choice, which matches our Baseline treatment, where this choice is introduced only after beliefs are stated. In Appendix F, we show that the model's main results, strict overoptimism about reaching the top and a positive effect of inequality on this overoptimism, continue to hold when subjects fully anticipate the Period 2 allocation choice, with beliefs determined in a fixed point together with the allocation they expect to implement. Through this belief, higher Period 1 inequality also raises inequality acceptance in Period 2.

3 Evidence from ISSP survey data

Our conceptual framework predicts a positive relationship between inequality and beliefs about reaching the top of the income distribution. We first examine this prediction in large-scale cross-country survey data, in which individuals have long-term exposure to inequality and face actual economic outcomes. We then ask whether the same pattern holds within countries, exploiting variation in the inequality that different birth cohorts experienced during their formative years.

We use data from the International Social Survey Programme (ISSP).[7] The ISSP 2019 "Social Inequality V" module asks respondents to place themselves on a social ladder running from 1 (bottom) to 10 (top) at three points in time: where their family stood when they were growing up (v42), where they stand now (v41), and where they expect to be in 10 years

[6] The exact prediction of performance optimism for each individual rank depends on the performance distribution and the magnitude of believed improvement. We provide a detailed analysis in Appendix E.1.

[7] Data were collected between November 2018 and May 2022, depending on the country. The ISSP provides individual weights to improve representativeness. The results in this section do not use these weights but are qualitatively and quantitatively robust to including them (available on request). Summary statistics for countries with a Gini coefficient above and below the median are in Appendix Table C.1.

(v43).[8] These three questions let us construct, for each respondent, a measure of *actual* upward mobility (current status minus family-of-origin status, $v_{41} - v_{42}$) and a measure of *expected* upward mobility (expected status minus current status, $v_{43} - v_{41}$). The two measures allow us to compare, within the same population, how inequality relates to mobility that has already happened and to mobility that people anticipate.

3.1 The Great Gatsby Curve and the experience-expectation gap

Our framework predicts that in more unequal countries, where the reward at the top is larger, people hold more optimistic beliefs about moving up. Panel (b) of Figure 1 in the introduction shows this pattern. Across the 25 countries in the figure[9], average expected mobility is higher where inequality is higher. The cross-country correlation between the Gini coefficient and expected mobility is 0.60 ($p < 0.01$), and 0.51 with all 27 countries ($p < 0.01$).

This relationship has the opposite sign of the Great Gatsby Curve. Panel (a) of Figure 1 is a conceptual replication of that curve (Corak, 2013): in more unequal countries, the average respondent reports less actual mobility, a cross-country correlation of -0.50 ($p < 0.05$), and -0.43 with all 27 countries ($p < 0.05$). The two relationships have opposite signs even though both describe mobility on the same 1 to 10 ladder for the same respondents. Expected mobility therefore has the positive relationship with inequality that the POUM hypothesis requires (Bénabou and Ok, 2001), whereas actual mobility does not.

Consistent with our framework, this optimism is concentrated at the top of the ladder, where the rewards are largest. Respondents in more unequal countries (Gini above the median) are more likely than those in more equal countries (Gini below the median) to expect one of the top two rungs (rung 9 or 10) and less likely to expect the upper-middle rungs 7 and 8. The share of respondents expecting a top-two position is about 11% in more unequal countries, compared to about 7% in more equal ones (Appendix Table C.4 and Appendix Figure C.8). The difference is thus mainly a shift of probability mass from the upper-middle rungs to the top two rungs rather than a uniform upward shift of the whole distribution. It does not appear in how respondents place themselves *currently*: in more unequal countries, respondents are if anything more likely to place themselves in the bottom half today (Appendix Figure C.9).

We define the share of respondents who report having moved up at least one rung compared to their status growing up as the binary measure of actual mobility and the share of respondents who report expecting to move up by at least one rung as the binary measure of expected social mobility. With these binary measures of mobility, across all 27 countries, average expected

[8]The exact wording for current status is "In our society there are groups which tend to be towards the top and groups which tend to be towards the bottom. Below is a scale which runs from top to bottom. Where would you put yourself on this scale?" For family-of-origin status, it is "Where did the family that you grew up in fit in then?" For expected status, it is "Thinking ahead 10 years from now, where do you think you will be on this scale?"

[9]Figure 1 excludes South Africa and the Philippines as outliers in terms of their Gini coefficient and expected upward mobility, respectively. Figure B.7 shows that the correlations are robust to including them.

mobility exceeds average actual mobility by 6.3 percentage points. This expected-minus-actual mobility gap is larger in more unequal countries. The correlation between the gap and the Gini coefficient is 0.61 on the binary measure ($p < 0.01$), and the discrete measure of number of rungs moved or expected to move gives the same pattern. On the binary mobility measure, the expected-minus-actual gap is largest in the Philippines, followed by Chile and Russia, and smallest in France, followed by Czechia and Finland. Appendix Table B.2 reports the full ranking and the underlying correlations.

Robustness. The experience-expectation gap holds under alternative measures of inequality and of mobility, and in the full sample of 27 countries (Appendix B). First, replacing the Gini coefficient with the top-5% or top-1% income share leaves both signs unchanged (Appendix Figure B.2). Second, actual mobility falls and expected mobility rises in inequality whether mobility is the continuous measure, status persistence, or the "rags-to-riches" transition rate (Appendix Figures B.3, B.4, and B.5). Third, both signs are unchanged when the two outlier countries, South Africa and the Philippines, are included (Appendix Figure B.1). Fourth, the association between inequality and expected mobility is robust to country-level controls for log GDP per capita, 10-year average GDP growth, unemployment, an index for democracy (V-Dem polyarchy), and country-mean years of education, and its coefficient is larger with the full control set (Appendix Figure B.6 and Appendix Table B.1).

Result 1: *Across countries, higher inequality is associated with higher expected mobility but lower actual mobility. This optimism about future mobility is concentrated at the top of the status ladder.*

3.2 Experienced inequality during formative years

The cross-country relationship could reflect unobserved institutional or cultural differences across countries. To control for such differences, we exploit variation *within* countries, across birth cohorts, in the inequality that people were exposed to when they were young. Following Roth and Wohlfart (2018), we measure experienced inequality as the average top-5% income share in a respondent's country over the ages 18 to 25, the "impressionable years" during which economic experiences have a lasting effect on beliefs (Malmendier and Nagel, 2011; Giuliano and Spilimbergo, 2025). Because different cohorts in the same country came of age under different top-5% income shares, this measure varies within a country even after we absorb everything fixed about the country itself. We standardize it within the sample (one standard deviation is 9.5 percentage points of the top-5% income share, roughly the gap between Northern Europe and the US). We restrict the sample to respondents aged 29 or older so that the formative years have ended for everyone we study.[10]

Figure 2 reports how experienced inequality relates to expected mobility, $v_{43} - v_{41}$, across four specifications, with the underlying coefficients in Appendix Table C.2. We use expected mobility as the dependent variable rather than the level of expected status because the level mechanically confounds optimism with current status: a cohort that already places itself

[10]For respondents younger than 29, the formative years overlap with or end just before the present, and the measure therefore does not capture exposure that is firmly in the past. Restricting to age 29 and older ensures that the formative years ended at least three years before the survey for the entire sample.

higher will report a higher expected status for that reason alone, and taking the difference removes this mechanical relationship. The estimating equation is

$$(v_{43,i} - v_{41,i}) \;=\; \beta\,\overline{\text{top5}}^{\,\text{std}}_{c(i),b(i)} \;+\; X_i'\gamma \;+\; \alpha_{c(i)} \;+\; \delta_{r(c(i))\times b(i)} \;+\; \varepsilon_i,$$

where $\overline{\text{top5}}^{\,\text{std}}_{c(i),b(i)}$ is standardized experienced inequality for cohort $b(i)$ in country $c(i)$, one standard deviation of which is 9.5 percentage points of the top-5% income share. The vector X_i collects three individual controls: age, an indicator for female, and years of education. Birth cohorts differ systematically in all three characteristics, and each can independently affect expected mobility. We control for them so that the coefficient on experienced inequality does not simply reflect cohort composition. The term $\alpha_{c(i)}$ is a country fixed effect, which absorbs everything fixed about a country: its culture, its institutions, its average level of inequality, and its average optimism. The term $\delta_{r(c(i))\times b(i)}$ is a subregion-by-cohort fixed effect, which lets the age and cohort profile of expectations differ across subregions. With both sets of fixed effects, identification comes only from within-country, across-cohort variation in experienced inequality net of any subregion-wide cohort trend. Standard errors are clustered at the country level.

The most demanding specification adds a further set of controls, each a country-level average over the same formative years: the unemployment rate, GDP-per-capita growth, CPI inflation, the level of electoral democracy, exposure to a banking or currency crisis, and exposure to a rare disaster, together with the respondent's family-of-origin status (v_{42}). A large body of work shows that macroeconomic experiences during the formative years have a lasting effect on economic beliefs and risk-taking (Malmendier and Nagel, 2011; Giuliano and Spilimbergo, 2025): inflation experienced at a young age shapes later inflation beliefs (Malmendier and Nagel, 2016), recessions and aggregate shocks shape preferences and beliefs about luck and redistribution (Giuliano and Spilimbergo, 2025), adverse financial shocks leave "scars" affecting long-term consumption and saving behavior (Malmendier and Shen, 2024), and coming of age under a different political regime shapes redistribution preferences (Alesina and Fuchs-Schündeln, 2007). Because any of these experiences could have been correlated with the top-5% income share when a cohort was young, we hold them fixed and ask whether experienced inequality is still associated with expected mobility on its own.

The coefficient decreases as we add successively more demanding controls, but remains positive and significant. Specification (1) of Figure 2 regresses expected mobility on standardized experienced inequality alone, with no controls: a one standard deviation increase in experienced inequality is associated with 0.466 ladder rungs higher expected mobility (SE 0.135, $p < 0.01$, $N = 22{,}558$). Specification (2) adds the three individual controls (age, an indicator for female, and years of education) so that the estimate no longer reflects differences in age structure, schooling, or gender composition across cohorts. The coefficient falls modestly to 0.335 (SE 0.113, $p < 0.01$, $N = 21{,}879$). Specification (3) further adds country and subregion-by-cohort fixed effects and is the identifying comparison of interest: it contrasts adjacent cohorts within the same country, net of any subregion-wide cohort trend, rather than comparing across countries that differ in many other ways. The coefficient is 0.288 (SE 0.104, $p < 0.01$, $N = 21{,}879$), which translates to about 0.18 standard deviations of expected mobility. Specification (4) adds the six macroeconomic experiences measured over the same formative years (the unemployment rate, GDP-per-capita growth, CPI inflation, electoral democracy, banking-or-currency-crisis exposure, and rare-disaster exposure), together with family-of-origin status. This specification tests whether experienced inequality is standing in for one of these macroeconomic experiences. The coefficient is 0.252 (SE 0.111, $p < 0.05$, N = 19,152) and remains positive and statistically significant at the 5% level. Coming of age under higher inequality is associated with higher expected mobility decades later, even when comparing adjacent cohorts in the same country.

Figure 2: The effect of experienced inequality on expected mobility across specifications

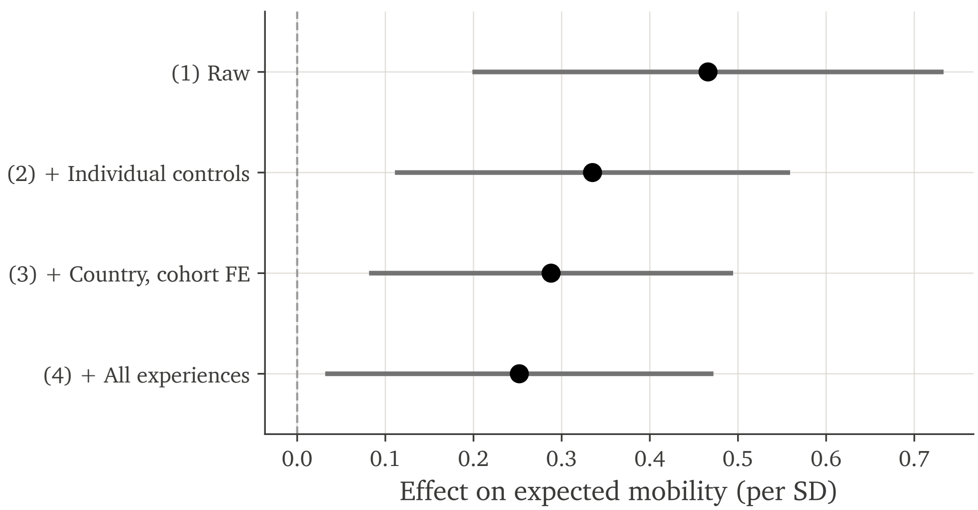

Notes: This figure shows the coefficient on standardized experienced inequality, the average top-5% income share in the respondent's country over ages 18 to 25 (one standard deviation = 9.5 percentage points), across four specifications. The outcome is expected mobility, expected status in 10 years minus current status ($v_{43} - v_{41}$), on the 1 to 10 social ladder. Specification (1) has no controls, and each further specification adds to the previous one: (2) individual controls; (3) country and subregion-by-cohort fixed effects; (4) the six macroeconomic experiences and family-of-origin status. The sample includes ISSP 2019 respondents aged 29 or older. Error bars show 95% confidence intervals based on standard errors clustered at the country level. The underlying regression estimates are reported in Appendix Table C.2.

The framework also makes a more specific prediction than "more optimism on average." Belief distortion should be positive for positions that pay above-mean income and negative for those that pay below-mean income, even when the below-mean positions rank above the median. We test this in income terms by mapping each ladder rung to the mean household income of respondents at that rung within their country (Appendix Figure C.3). The mapping places the country-mean income between rungs 7 and 8. Expecting rung 8 or higher therefore means expecting a position that pays above the country mean. We estimate Specification (3) with current status and the experienced unemployment rate as additional controls (N = 19,600). A one standard deviation increase in experienced inequality is associated with a 6.0 percentage point higher probability of expecting a position that pays above the country mean, and by construction with an equal decrease in the probability of expecting a position at or below it ($p < 0.01$, Appendix Figure C.4). A continuous version, which relates experienced inequality to the income of the expected position relative to the country mean, gives the same result: a one standard deviation increase in experienced inequality is associated with the expected position's income being 5.5% higher relative to the country mean ($p < 0.01$). Thus, the shift in probability mass associated with higher inequality is not evenly distributed across the ladder. Instead it concentrates on the positions that are associated with above-mean income in the respective country, which is consistent with motivated beliefs and in contrast to a uniform upward shift.

Robustness. Appendix C.1 reports four robustness checks. First, the relationship between experienced inequality and expected mobility is not driven by a single birth cohort whose formative years coincided with a particular episode: estimating it separately for seven age groups (Appendix Figure C.5) gives a positive coefficient in every age group, significant across the younger and middle-aged groups and present, if noisier, among the oldest. Second, the top-1% and top-10% income shares give the same pattern as the top-5% income share (Appendix Figure C.6). Third, in a binned scatter plot of the residualized variables, the binned means lie close to the linear fit (Appendix Figure C.1). Fourth, a placebo that uses inequality experienced at ages 0 to 7, *before* the formative years, gives a small and statistically insignificant coefficient, as does a post-formative window (ages 34 to 41). Only the 18 to 25 window has a positive, significant coefficient (Appendix Figure C.7). All coefficients in this section are replicated in both Python and Stata.

Result 2: *Within countries, respondents who experienced higher inequality during their formative years expect more upward mobility decades later. This optimism about future mobility is concentrated on status positions associated with above-mean income.*

3.3 Expected mobility and demand for redistribution

For optimism to sustain inequality, it must be associated with weaker demand for redistribution. At the individual level, holding current status fixed, respondents who expect to move further up the ladder support less redistribution, the association the POUM hypothesis requires (Bénabou and Ok, 2001). Figure 3 plots the coefficient on expected mobility across specifications that successively add individual controls, country fixed effects, household income, and political orientation. The relationship stays negative and significant throughout ($p < 0.01$).[11]

Result 3: *Higher expected mobility is associated with weaker demand for redistribution.*

To test the mechanism causally, we next turn to an experiment that randomly assigns inequality and measures beliefs at the individual level, in a setting in which inequality and actual social mobility are not determined jointly with beliefs and institutions.

[11] Demand for redistribution is measured by agreement that government should reduce income differences between the rich and the poor, rescaled so that higher values indicate stronger support for redistribution.

Figure 3: Expected mobility and demand for redistribution

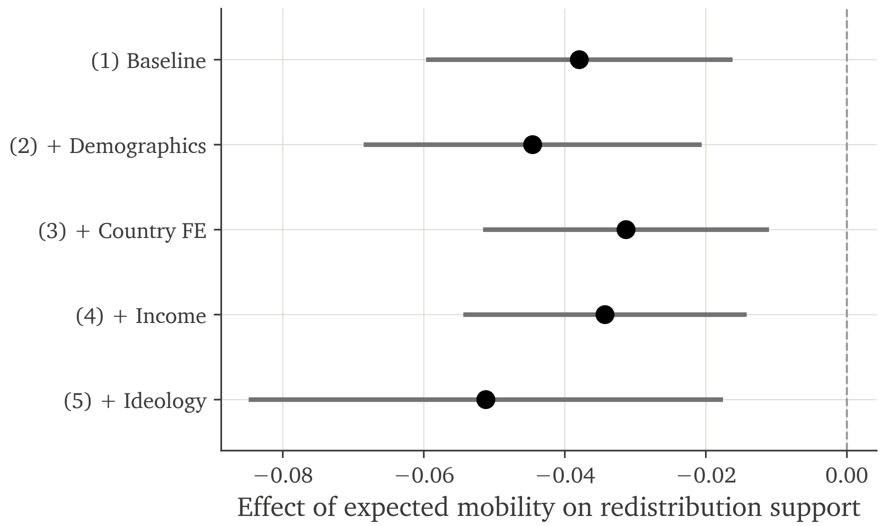

Notes: This figure shows the coefficient on individual expected upward mobility (expected status in 10 years minus current status, $v_{43} - v_{41}$) from regressions of demand for redistribution on expected mobility, holding current status fixed. Demand for redistribution is measured by agreement that government should reduce income differences between the rich and the poor, reverse-coded so higher values indicate stronger support. Specification (1) controls only for current status, and each further specification adds to the previous one: (2) individual controls (age, an indicator for female, and years of education), (3) country fixed effects, (4) household-income quintile, and (5) left-right political orientation. The sample includes ISSP 2019 respondents. Error bars show 95% confidence intervals based on standard errors clustered at the country level.

4 Experimental design

This study employs a between-subjects experimental design.[12] We first present the Baseline treatment, which consists of a High inequality condition and a Low inequality condition. We then introduce the Feedback treatment, which differs from the Baseline treatment by providing subjects with accurate feedback after completing the quiz.

Overview. The experiment consists of two periods. In Period 1, subjects face a fixed income structure that creates inequality between one "Rich" and four "Poor" members in their five-person group. Who becomes Rich or Poor depends on their relative rank, determined by performance on a quiz with some random noise. Under this given inequality, subjects work on the quiz. After they finish the quiz but before they know their ranks, we elicit their beliefs about their chances of becoming Rich. In Period 2, subjects choose their preferred level of inequality between the Rich and the Poor. They choose their preferred allocation knowing that the ranking from Period 1 will also determine their position in Period 2, and still without knowing their actual rank. Thus, subjects make their allocation choice under beliefs formed in Period 1. Only after they make this choice do we determine their ranks and income in both periods. One period is randomly selected for payment.

[12]The design was approved by the institutional review board of WZB on July 26 2023 (Ethics Application No. 208) and pre-registered at the AEA RCT Registry under AEARCTR-0014848.

Figure 4: Experimental design and timeline

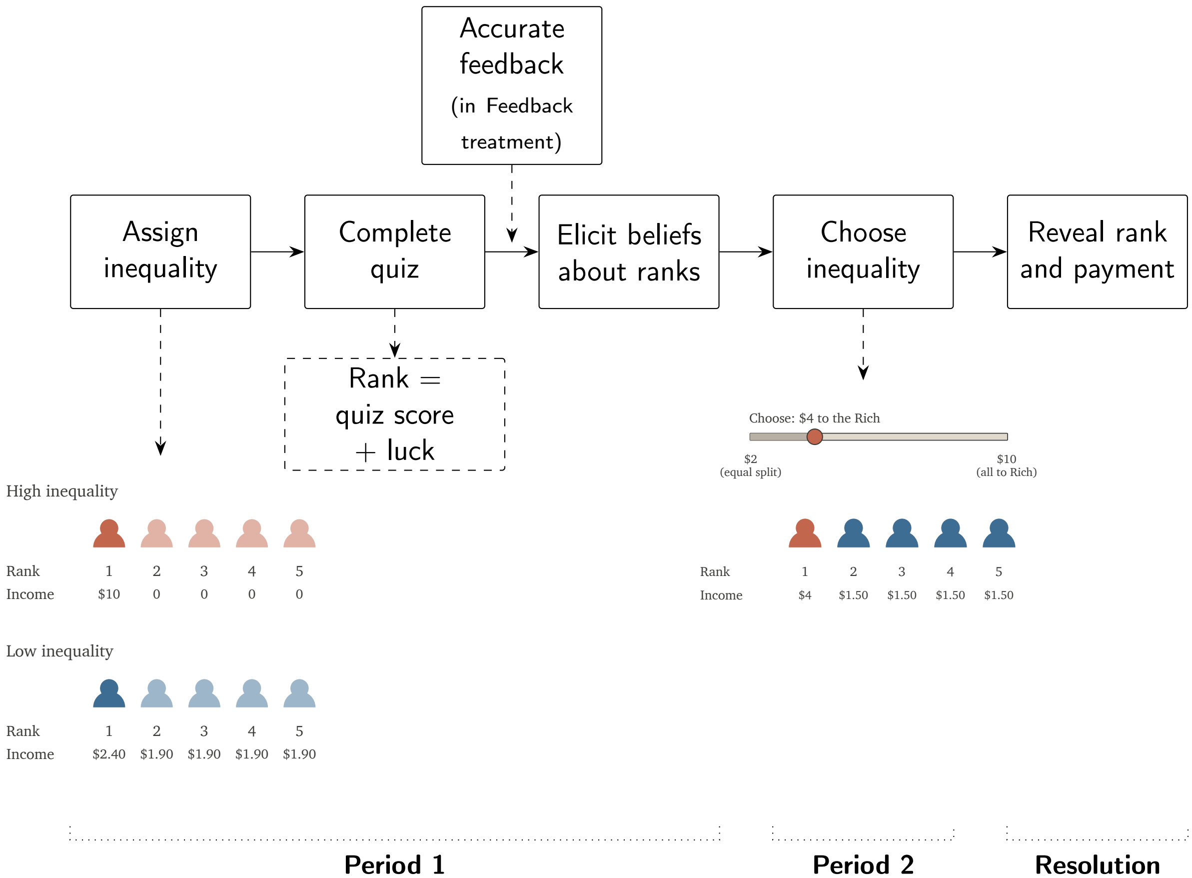

Notes: This figure shows the timeline of the experiment. The panels below *Assign inequality* show the Period 1 income at each of the five ranks under the High and Low inequality conditions. The panel below *Choose inequality* illustrates the Period 2 allocation choice: subjects choose how much of the $10 goes to the Rich (here $4), with the remainder split equally among the four Poor (here $1.50 each). Dashed elements mark the rank, which is determined by quiz performance and luck once the quiz ends but stays hidden until the end: subjects state beliefs about each rank without knowing it, the same rank applies in Period 2, and it is revealed only at payment. The Rank-1 belief, the belief about becoming Rich, is incentivized. One of the two periods is paid at random. Subjects are aware of the existence of Period 2 in Period 1, but they do not know about the choice they need to make in Period 2. Feedback is provided only in the Feedback treatment.

The main manipulation is to assign either a high or low income inequality to a group in Period 1. Figure 4 provides an overview of the main experimental stages.[13]

We deliberately design an abstract experiment rather than mimicking a real economy with its institutions, such as labor markets and taxes, and its life-cycle dynamics. This allows us to isolate the psychological mechanism from our conceptual framework: how the level of inequality a person faces shapes the belief they form about reaching the top. Two design choices follow. First, the experimental society contains no mobility process to learn from

[13]In the experimental implementation, there are two variations of the Baseline treatment, in which we varied when exactly subjects completed the main part of the quiz. As pre-registered, we pool these two variations since there were no significant differences between them. Details are described in Appendix D.

and no repeated play. Real economies give people decades of feedback about their own and others' mobility, and this accumulated experience makes field beliefs hard to interpret. We abstract it away so that the belief a subject reports is formed under a known inequality and not influenced by an idiosyncratic mobility history. Our goal is to identify the mechanism cleanly, not to mimic the characteristics of the field setting that motivates it. The predicted mechanism exists without past mobility history, and thus introducing mobility into our design would reintroduce confounders that the experiment exists to remove.

Second, the experiment allows us to measure overoptimism directly, which is not possible with the survey.[14] We set the objective probability of becoming Rich and elicit the belief. Overoptimism is the belief minus this probability. We measure it at the group level in the Baseline treatment, where the average objective probability is 20%, and at the individual level in the Feedback treatment, where each subject is told their own objective probability.

Period 1 income inequality. In the experimental "society", there are always one Rich subject (referred to as the "high income earner" in the instructions) and four Poor subjects (referred to as the "low income earner" in the instructions). Who becomes the Rich is determined by their rank in the society. Subjects' incomes by rank in each inequality condition are shown in the income panels of Figure 4.

In the "High" inequality condition, the Rich (Rank 1) earns \$10, while the Poor (Ranks 2-5) earn \$0. In the "Low" inequality condition, the Rich earns \$2.40, while the Poor (Ranks 2-5) earn \$1.90. Total group income is \$10 in both conditions, so mean income is \$2 in both: the manipulation changes only the concentration of income at the top. This matches the constant-total-income comparative static in our conceptual framework and ensures the manipulation varies inequality rather than the scale of stakes. The income structure in Period 1 is announced at the beginning of the experiment, and subjects are not aware of the other inequality condition.

Ranking. A subject's rank (Rank 1 vs. Ranks 2-5) is determined by quiz performance or luck. With 80% probability, rankings are assigned based on performance in a quiz. The quiz consists of solving 10 Raven matrices within a 10-minute time limit. Performance is defined as the number of correct answers. Ties in performance are broken randomly in ranking. Subjects are informed about how their income depends on their relative rank and there is no additional payment based on their absolute performance. Before working on the 10 quiz questions, we also provide two example Raven matrices to familiarize them with this type of task, but they do not count towards their performance.

With 20% probability, rankings are assigned randomly, giving each subject an equal chance of obtaining any rank.

[14] Because expected mobility in the survey is a belief about a future that has not yet happened, it cannot be labeled overoptimistic on its own terms. The most one can do is benchmark it against actual mobility. That benchmark is suggestive: in all 27 ISSP countries the average respondent expects to climb the ladder over the coming decade, with mean expected mobility of 0.58 rungs compared to 0.27 rungs of actual mobility. On a relative ladder, where not everyone can move up, such universally positive expectations are difficult to square with accuracy. The comparison is nonetheless imperfect, because the future need not mirror the past.

Belief elicitation. We elicit subjects' beliefs about their rank within their five-person group after the inequality announcement in Period 1 and completing the Raven matrix quiz questions.

Specifically, we ask subjects: "For each possible rank, what do you think is the probability that you will achieve it?" Subjects must assign a probability to each rank, with the probabilities summing to 100.

Beliefs about being in Rank 1 are incentivized with the Binarized Scoring Rule (BSR) (Hossain and Okui, 2013). The BSR computes the distance between the subject's stated probability for being in Rank 1 and whether they actually are in Rank 1. The bonus is paid if the distance is smaller than a randomly drawn cut-off threshold, which makes truthful reporting the best strategy regardless of risk preferences. In the instructions, we follow the protocol recommended by Danz et al. (2022), i.e. we tell subjects that stating their true belief would maximize their chances of getting the bonus payment, while putting technical details of the payment rule behind an optional button.

Period 2 income allocation choice. In Period 2, subjects choose how to allocate income within their five-person group. Specifically, they decide how much of the total \$10 to allocate to the Rich, with the remainder divided equally among the four Poor. The set of options lies between \$2, i.e. fully equalizing income, and \$10, i.e. giving everything to the Rich and nothing to the Poor. One randomly selected group member's allocation choice is implemented for the entire group.

When making this choice, subjects do not know their Period 2 rank but are informed that it will be identical to their Period 1 rank. This implies that they should base their choice on their previously formed Rank-1 belief. We remind them of their previously stated Rank-1 belief on the allocation decision page.[15]

After Period 2, we collect survey data to understand the factors influencing subjects' choices. Subjects indicate how much various factors influenced their allocation choice and assess the relative importance of effort, luck, and intelligence in quiz performance.[16] We also gather standard socio-demographic and economic variables (gender, age, education, income, nationality, ethnicity, political orientation), as well as general redistribution and risk preferences. Additionally, we elicit subjects' beliefs about their own economic mobility by asking about their current relative income status and their projected status in 10 years.

After the session is complete, we determine the ranking that applies to both periods, calculate Period 1 income based on ranks and Period 2 income based on ranks and the chosen allocation, and randomly select one period for payment.

[15] As we do not elicit beliefs a second time, we cannot rule out that subjects adjust their Rank-1 belief consciously or unconsciously upon learning about the possibility to redistribute. However, the fact that their Rank-1 belief is now more consequential to their income should rather mitigate motivated beliefs (Zimmermann, 2020), which would reduce the effect of initial inequality on inequality acceptance.

[16] As a robustness check, we also ask whether subjects have used AI during the quiz and whether they think others might have used AI. Only a minority of subjects admitted to having used AI (1%) or thought others might have used AI during the quiz (7%), and our results remain qualitatively unchanged if we drop them.

Feedback treatment. The Feedback treatment serves two purposes. First, it tests whether the inequality-induced overoptimism reflects missing information rather than motivated beliefs: if subjects are simply uninformed, giving them their objective probability of becoming Rich should eliminate the effect of initial inequality on beliefs. Second, the feedback generates *exogenous* variation in beliefs that we use to instrument them. The feedback is individualized: it tells each subject their own objective probability, which is low for low performers and high for high performers. It therefore moves beliefs in opposite directions depending on quiz score, lowering the beliefs of overoptimistic subjects and confirming or raising the beliefs of high performers. We exploit this score-dependent shift by examining the treatment effect of feedback conditional on quiz score, which forms the first stage of the instrumental-variables analysis in Section 5.5.

The Feedback treatment is identical to the Baseline treatment except that subjects receive accurate feedback after completing the 10 Raven matrices but before stating their beliefs, the step marked "Accurate feedback" in Figure 4. We inform them of their score and their probability of achieving Rank 1 in a five-person group, based on their score and historical performance data.

Specifically, we calculate this probability by comparing their score to a distribution of scores from Baseline treatment subjects. For each score, we bootstrap 10,000 five-person groups and calculate how often that score is the best in the five-person group, with random tie-breaking. The probability for achieving the best score in the group is the proportion of times the score ranks the best across these simulations. The Rank-1 probability is the probability for having the best score multiplied by the probability that the rank will be determined by performance (80%), in addition to the probability for being in Rank 1 by luck (20% * 20% = 4%). Subjects can access detailed information about this calculation through an optional button.

The remainder of Period 1 proceeds as in the Baseline treatment. After receiving the feedback, subjects state their beliefs about achieving each rank in the group. Like the Baseline treatment, there is a High inequality condition and a Low inequality condition. By providing accurate feedback, we expect to observe a smaller effect of initial inequality on the Rank-1 belief and thus on the Period 2 allocation choice.

Implementation and sample. Our sample consists of US nationals recruited via the crowd-sourcing web service Prolific, with data collected in November 2024. On average, subjects earned $5 and spent 22 minutes on the study. The sample includes $N = 773$ subjects who completed the Baseline treatment and $N = 426$ who completed the Feedback treatment. The experiment was programmed using oTree (Chen et al., 2016).

5 Experimental results

Descriptive statistics and balance tests are in Appendix Table A.2. The Baseline treatment has $N = 377$ subjects in the High inequality condition and $N = 396$ in the Low inequality condition, and the balance tests reveal no significant differences in socio-demographic characteristics.

5.1 Initial inequality and beliefs

Figure 5 plots average beliefs for each rank by inequality condition and shows two results. First, subjects are overoptimistic about the high-paying position. Regardless of initial inequality, their average Rank-1 belief is well above the 20% average objective probability, while their average Rank-2 belief sits close to 20%. That subjects inflate their Rank-1 belief but not their Rank-2 belief is itself consistent with motivated beliefs: Rank 1 pays more than every other rank, whereas Rank 2 pays the same as Ranks 3-5.

Second, and central to our prediction, higher initial inequality strongly raises the Rank-1 belief. Subjects in the High inequality condition state an average Rank-1 belief of 41.3%, compared to 34.4% in the Low inequality condition, a 6.9 percentage point difference that is significant at the 1% level (Column 1 of Appendix Table A.4). Because the objective probability of becoming Rich is 20% in both conditions, it cannot explain the difference. It is only for Rank 1 that beliefs rise with inequality. For every other rank, average beliefs are lower in the High inequality condition.

Figure 5: Average beliefs for each rank by inequality condition

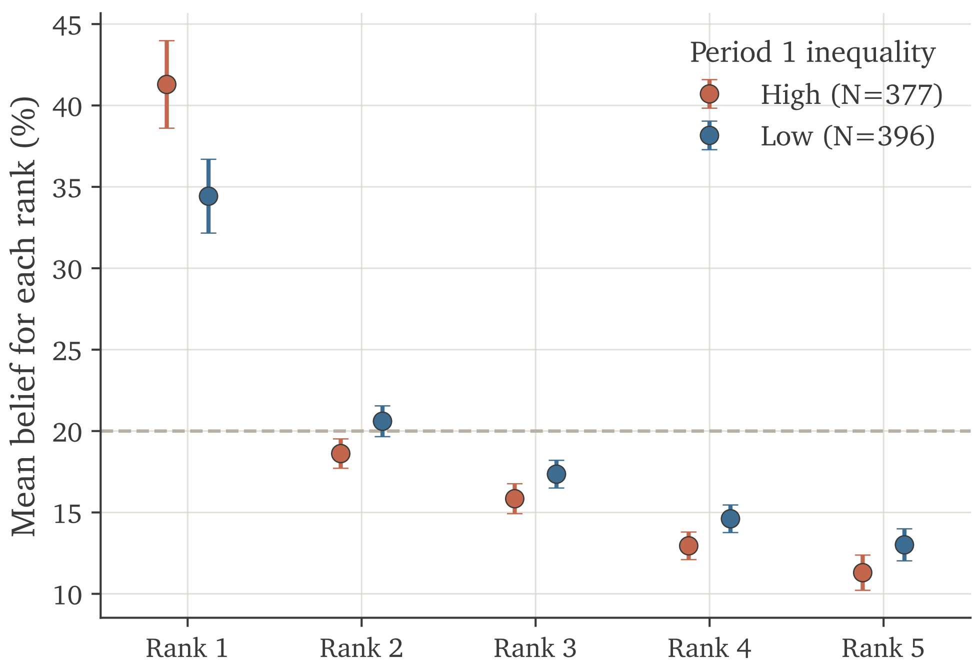

Notes: This figure shows subjects' average belief for each rank (1 to 5) by inequality condition in the Baseline treatment. The dashed line marks the 20% average objective probability. Error bars show 95% confidence intervals. The full distribution of Rank-1 and Rank-2 beliefs is in Appendix Figure A.3.

Robustness. Initial inequality induces an overall shift in the belief distribution: a Kolmogorov–Smirnov test rejects equality of the Rank-1 belief distributions ($p < 0.01$), and the share stating a Rank-1 belief strictly above 20%, i.e. an above-average chance, rises from 54.3% in the Low inequality condition to 64.7% in the High inequality condition (Appendix Figure A.3). Because rank is determined by performance with 80% probability and by luck with

20%, the objective Rank-1 probability can never exceed 84%. Only 38 of 773 subjects state a probability above this ceiling, almost all for Rank 1 and disproportionately more in the High inequality condition (26 subjects in the High inequality condition vs. 12 in the Low inequality condition, $p < 0.05$), which points to motivated beliefs rather than confusion. Excluding these subjects makes the effect of initial inequality smaller but leaves it significant, and rescaling their Rank-1 belief to account for luck leaves it nearly unchanged (Appendix Table A.5).

Result 4: *Higher initial inequality raises overoptimism about reaching the high-paying position.*

5.2 Is the overoptimism motivated?

A natural alternative to motivated beliefs is performance optimism: subjects in the High inequality condition might simply believe they performed better on the quiz. We rule this out in two steps.

First, there is no difference in performance or effort across conditions. Figure 6 shows that quiz scores, time on the quiz, self-reported effort, and beliefs about others' effort are all statistically indistinguishable between conditions. The absence of a difference in effort also rules out equilibrium neglect, the possibility that subjects respond to the stronger incentive in the High inequality condition by working harder while failing to anticipate that others do the same.[17]

Figure 6: Performance and effort do not differ across conditions

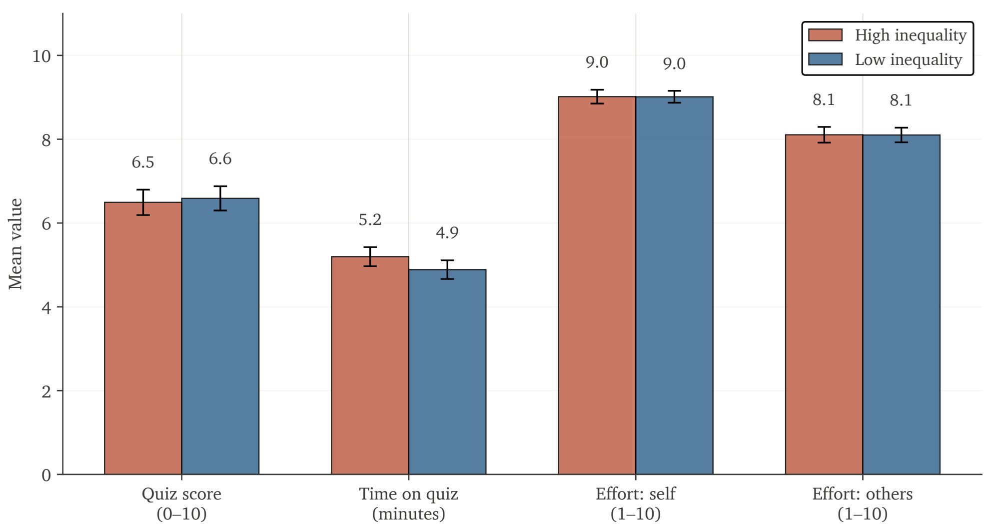

Notes: This figure shows the mean quiz score (out of 10), time on the quiz (maximum time allowed: 10 minutes), self-reported effort from 1 to 10, and beliefs about others' effort from 1 to 10, by inequality condition. Error bars show 95% confidence intervals.

[17]There is also no difference between inequality conditions in how much effort versus luck subjects attributed to their success in the quiz (see Appendix Figure A.6). This suggests that the system justification mechanism described by Bénabou and Tirole (2006) is not a driver of our results.

Second, Rank-2 beliefs distinguish the two accounts, because performance optimism and motivated beliefs make opposite predictions about Rank 2. Since Rank 1 and Rank 2 are both above-median ranks in a five-person group, a subject who believes they performed better should increase the probability they assign to *both* ranks. Motivated beliefs, by contrast, target only the high-paying position: because probabilities must sum to 100, inflating Rank 1 should leave Rank-2 beliefs unchanged or decrease them. Using our empirical performance distribution, we show that a plausible belief in better performance increases both the mean and the median Rank-2 belief (Appendix E.1). The data show the opposite (Figure 5): as Rank-1 beliefs rise in the High inequality condition, Rank-2 beliefs remain unchanged at the median (20.0% vs. 20.0%) and decrease at the mean (18.6% vs. 20.6%, $p < 0.01$). The increase appears only for the high-paying position, which is inconsistent with performance optimism but consistent with motivated beliefs.

Robustness. The Rank-2 prediction of performance optimism is robust to the choice of performance distribution. Appendix E.1 proves that, if subjects believe their own score increased while the scores of others did not, the median Rank-2 belief must increase for any performance distribution, as long as the believed improvement does not move the median performer above the 75th percentile of that distribution.

Result 5: *The inequality-induced overoptimism is motivated rather than driven by performance optimism: it appears only for the high-paying position and not for the other above-median rank, and performance and effort do not differ across conditions.*

5.3 Persistence despite accurate feedback

If the overoptimism in the Baseline treatment reflected missing information rather than motivated beliefs, telling subjects their objective probability of becoming Rich should eliminate the effect of initial inequality on beliefs. In the Feedback treatment, subjects learn their quiz score and their simulated objective probability of becoming Rich before stating beliefs (Section 4). Feedback does move beliefs: average Rank-1 beliefs fall from 37.8% in the Baseline treatment to 31.4% in the Feedback treatment ($p < 0.01$), and 30% of subjects state exactly their objective probability. However, the effect of initial inequality on the Rank-1 belief persists: subjects in the High inequality condition still state higher Rank-1 beliefs than those in the Low inequality condition (33.7% vs. 29.3%, $p < 0.10$) even after learning their objective probability. This effect is smaller than in the Baseline treatment but remains positive (Appendix Figure A.2).

Because subjects in the Feedback treatment are told their objective probability, we can measure belief distortion directly as the deviation of the stated belief from the objective probability. Panel (a) of Figure 7 shows that deviations are asymmetric: beyond the 30% who state exactly their objective probability, subjects are far more likely to state a belief above it than below it. Panel (b) shows that the upward deviation is larger in the High inequality condition. Even though every subject was told their own objective probability, those in the High inequality condition state a belief 16.2 percentage points above it on average, compared to 11.6 percentage points in the Low inequality condition ($p < 0.10$). That the deviation persists despite accurate feedback, and is larger where the incentive to distort beliefs is stronger, is consistent with motivated beliefs but not with missing information.

Figure 7: Deviations of Rank-1 beliefs from the objective probability

(a) Distribution of deviations **(b) Deviations by inequality condition**

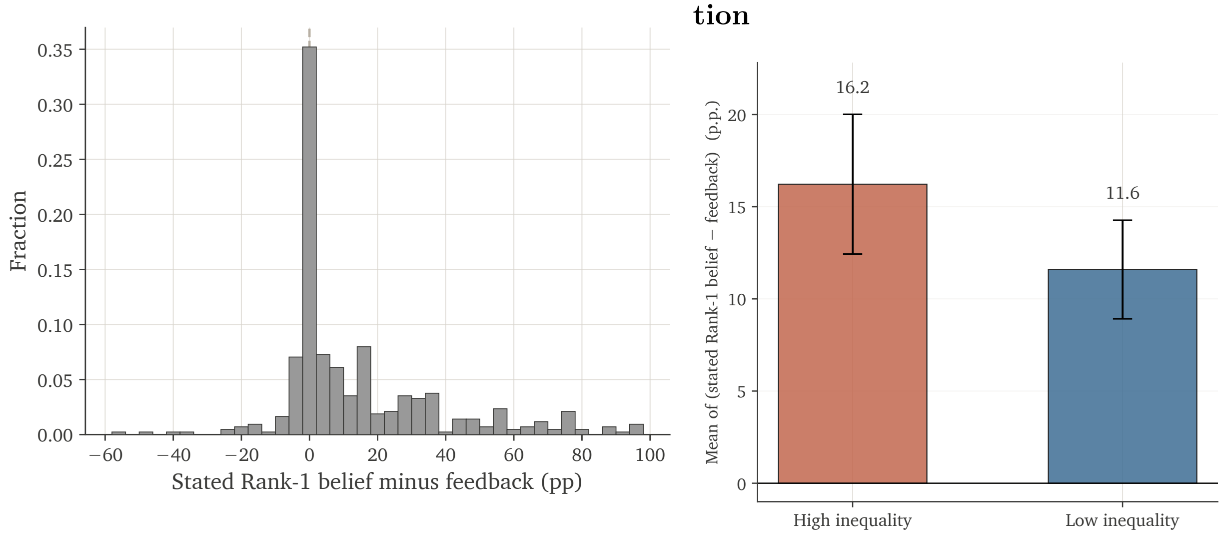

Notes: This figure shows the deviation of the stated Rank-1 belief from the objective probability provided in the feedback, in the Feedback treatment. Positive values indicate a belief above the objective probability. Panel (a) shows the full distribution of deviations. Panel (b) shows the mean deviation by inequality condition (16.2 percentage points in the High inequality condition and 11.6 in the Low inequality condition). Error bars show 95% confidence intervals.

Robustness. The upward deviation is concentrated among below-median performers, where the incentive to distort beliefs is strongest, and the difference between inequality conditions there is significant at the 5% level (Appendix Figure A.4).

Result 6: *The inequality-induced overoptimism persists despite accurate, individualized feedback, and the residual belief distortion is largest among those with the strongest motive to hold optimistic beliefs.*

5.4 Initial inequality and inequality acceptance

We now show that the shift in beliefs induced by inequality translates into behavior. In Period 2, before learning their rank, subjects choose how much of \$10 to allocate to the Rich, with the rest split equally among the four Poor and one subject's choice implemented for the group. Figure 8 shows that the share to the Rich rises monotonically with the Rank-1 belief, from about 35% in the lowest belief bin to about 60% in the highest. Consistent with the POUM hypothesis, subjects with a higher Rank-1 belief prefer more inequality.

Because higher initial inequality raises the Rank-1 belief, and the share to the Rich rises with the Rank-1 belief, higher initial inequality should raise inequality acceptance. Subjects in the High inequality condition indeed allocate 48.0% to the Rich on average, compared to

Figure 8: Share to the Rich by Rank-1 belief

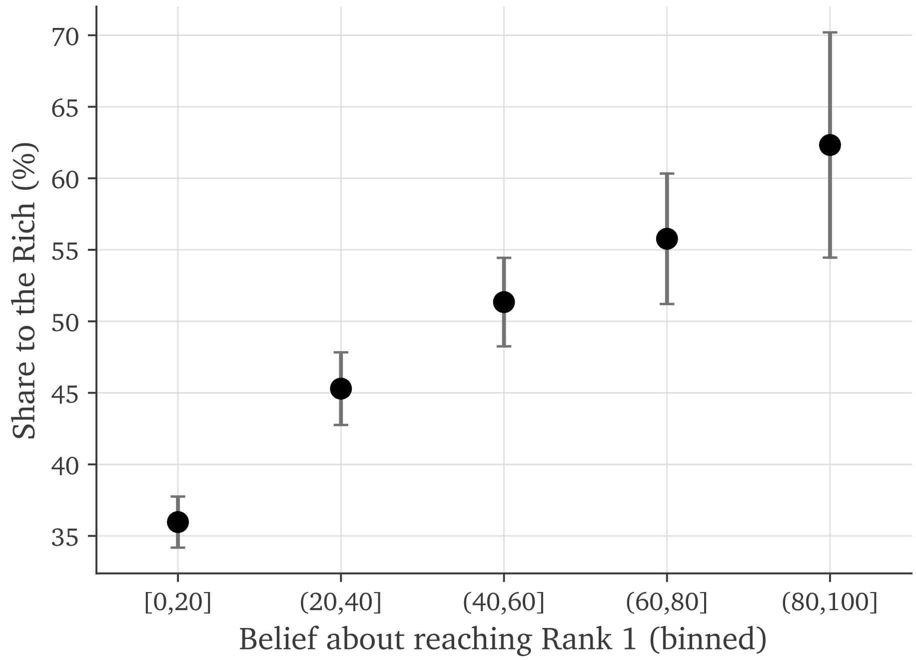

Notes: This figure shows the average share to the Rich by Rank-1 belief, pooling both inequality conditions. Rank-1 beliefs are grouped into five bins. Error bars show 95% confidence intervals.

40.8% in the Low inequality condition, a 7.2 percentage point difference significant at the 1% level (Column 3 of Appendix Table A.4).

Robustness. The difference also appears in the full allocation distributions (Appendix Figure A.1): a Kolmogorov–Smirnov test rejects equality of the distributions ($p < 0.01$), and the share of subjects who fully equalize income is higher in the Low inequality condition (27.0% vs. 21.0%, $p < 0.05$). The effect of initial inequality is also robust to demographic controls (available on request).

Result 7: *Higher initial inequality raises inequality acceptance.*

5.5 The belief channel is causal

The effect of initial inequality on inequality acceptance is consistent with the belief channel, but inequality might also raise the share to the Rich directly, for instance by shifting fairness views (Cappelen et al., 2007; Almås et al., 2020) or anchoring. To isolate the belief channel, we use the Feedback treatment as a source of *exogenous* variation in beliefs. Because the feedback is individualized, it moves beliefs in different directions depending on quiz score. This variation in beliefs is driven by the feedback rather than by subjects' fairness views or other preferences that might independently affect the allocation.

Figure 9 visualizes this logic. Panel (a) is the first stage: it plots how feedback shifts the Rank-1 belief, by quiz score. For quiz scores up to 8, feedback lowers the belief by 7 to 19 percentage points. For the top scorers, feedback *raises* it (by 4.8 percentage points), because they learn their objective probability is higher than they had assumed. Panel (b) shows the reduced form: the effect of the same feedback on the share to the Rich, by the same quiz-score bins. The two panels show the same pattern: where feedback lowers beliefs, the effect on the share to the Rich is flat to negative, and where feedback raises beliefs (among the top scorers), the share to the Rich rises by 7.9 percentage points.

Figure 9: Feedback effects by quiz score: belief and behavior move in the same direction

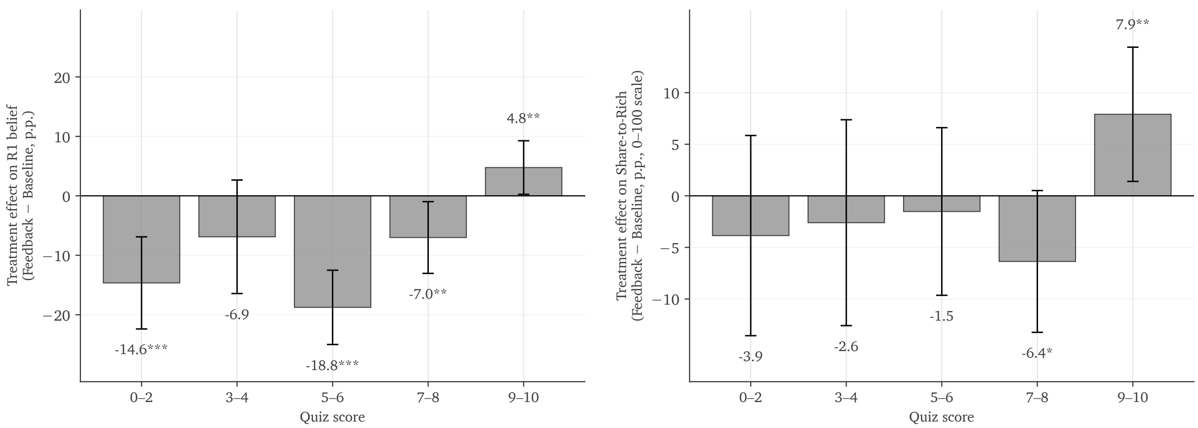

(a) Feedback effect on Rank-1 belief **(b) Feedback effect on share to the Rich**

Notes: This figure shows the treatment effect of Feedback (Feedback minus Baseline) by quiz-score bin. Panel (a) shows the effect on the Rank-1 belief (first stage). Panel (b) shows the effect on the share to the Rich (reduced form). Error bars show 95% confidence intervals. Stars denote significance at the 10/5/1% level.

We formalize this with two-stage least squares (Figure 10, with full estimates in Appendix Table A.1), instrumenting the Rank-1 belief with the interaction of feedback and quiz score and controlling for feedback, quiz score and its square, and the inequality condition. We do not use feedback on its own as the instrument, because receiving feedback could also change the allocation directly. The interaction of feedback and quiz score captures the score-dependent shift in beliefs and is strong, with a first-stage F-statistic of 23.1. The instrumented effect of the Rank-1 belief on the dollar amount allocated to the Rich is 0.043 (SE 0.020, $p < 0.05$): a 10 percentage point increase in the Rank-1 belief raises this amount by about \$0.43 out of \$10, equivalently 4.3 percentage points of the share. This IV estimate is close to, and slightly larger than, the OLS estimate of 0.032.[18]

Robustness. The identifying assumption is that feedback affects the allocation only through beliefs. The natural threat is that learning one's score also makes performance, and hence desert, more salient when choosing. Because the estimate is identified off the *interaction* of feedback and quiz score rather than the level of feedback, any uniform effect of receiving feedback, including a general rise in the salience of performance or desert, is absorbed by the

[18] An over-identified specification that adds feedback as a second instrument does not reject the over-identifying restriction (Hansen–Sargan test).

Figure 10: The belief channel: OLS and instrumental-variables estimates

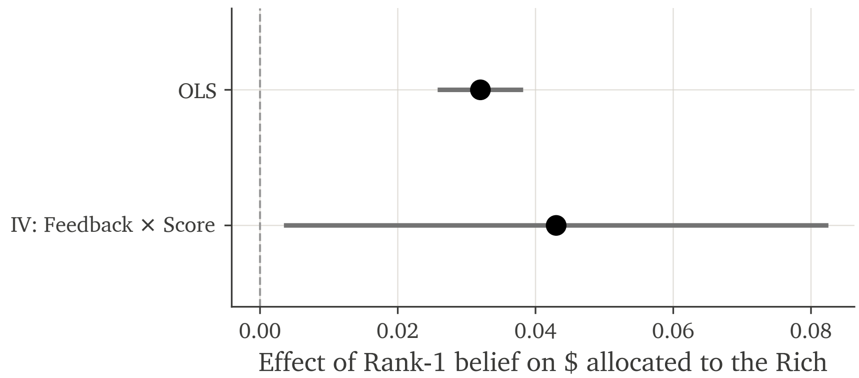

Notes: This figure shows the effect of the Rank-1 belief (0 to 100) on the dollar amount allocated to the Rich (2 to 10), estimated by OLS and by two-stage least squares instrumenting the Rank-1 belief with the interaction of feedback and quiz score (first-stage $F = 23.1$), controlling for feedback, the inequality condition, quiz score, and its square. Error bars show 95% confidence intervals. Full estimates are in Appendix Table A.1.

feedback main effect. Identification requires only that feedback not shift the performance-to-allocation mapping differentially by quiz score through some channel other than beliefs. The estimate is a local average treatment effect for the subjects whose beliefs respond to feedback, which is the relevant population for the belief channel. Beliefs are *a* causal channel but not the only one: the effect of initial inequality on the share to the Rich remains positive and significant after controlling for the Rank-1 belief (Column 4 of Appendix Table A.4). Inequality therefore also operates through other channels, such as fairness views anchored by the Period 1 income structure. The instrument shows that the belief channel is causal rather than driven by preferences that affect both beliefs and the allocation.

Result 8: *Beliefs are a causal channel through which inequality raises inequality acceptance: instrumenting the Rank-1 belief with the interaction of feedback and quiz score yields a positive, significant effect of beliefs on the share to the Rich.*

5.6 Subgroup analysis

We examine the effect of initial inequality in subgroups defined by gender, income, age, or political orientation (Appendix Figure A.8 and Appendix Tables A.7 and A.8). In the Baseline treatment, the Rank-1 belief and inequality acceptance are higher among men than women and among Republicans than Democrats, but no interaction of any subgroup with the High inequality condition is significant at the 5% level. The estimated effect of initial inequality on both outcomes is positive in every subgroup. Some of these estimates, such as the effect on the Rank-1 belief among women, are not significant at the 5% level. While the power to detect interactions is limited, the consistent direction of the effect suggests that inequality-induced overoptimism is not driven by any one subgroup.

6 Conclusion

The POUM hypothesis can reconcile high inequality with weak demand for redistribution only if citizens of unequal societies are more optimistic about moving up, yet the Great Gatsby Curve shows that these are the societies where actual mobility is lowest. This paper proposes and tests a mechanism that resolves this tension. Because a larger reward at the top raises the anticipatory utility of believing one will get there, higher inequality itself induces more optimistic beliefs about upward mobility.

The survey evidence shows the pattern in the field. Respondents in more unequal countries expect more upward mobility even as they report less actual mobility, an experience-expectation gap. Within countries, respondents who experienced higher inequality during their formative years also expect more upward mobility decades later, and this optimism is concentrated on positions associated with above-mean income rather than taking the form of a uniform upward shift. Respondents with higher expected mobility report weaker demand for redistribution, the association the POUM hypothesis requires.

In the experiment, we causally pin down the mechanism. Subjects randomly assigned to a more unequal society become more overoptimistic about reaching the high-paying position. This overoptimism is motivated rather than driven by performance optimism: it appears only for the high-paying position and not for the other above-median rank, it is not accompanied by differences in performance or effort, and it persists even when subjects are told their own objective probability of becoming Rich. The induced overoptimism in turn raises inequality acceptance when subjects are asked to choose their preferred level of inequality. Using the interaction of feedback and quiz score as an instrument for the Rank-1 belief, we show that beliefs are a causal channel through which inequality raises inequality acceptance, though not the only one.

The two halves of the paper point to the same mechanism. In the experiment, inequality raises overoptimism and overoptimism raises inequality acceptance. In the survey, inequality is associated with higher expected mobility, and higher expected mobility with weaker demand for redistribution. Together, these findings offer a behavioral account of why higher inequality need not generate the demand for redistribution that the median-voter logic predicts. One way inequality may sustain itself is by weakening its own political correction when it distorts beliefs about who will benefit from it.

The same beliefs that are associated with weaker demand for redistribution may also provide anticipatory utility to those who hold them. Using US data from the General Social Survey, we find that higher expected status is associated with higher current happiness, controlling for current status, and that this association is concentrated among low-income respondents. We read this evidence as suggestive rather than causal, but it points to a welfare dilemma: the optimism that helps sustain inequality may also be a source of anticipatory utility for those whom redistribution would help most. Correcting these beliefs could improve decisions in voting for policies while lowering the anticipatory utility of the disadvantaged, a trade-off that informational policies would need to weigh.

Finally, our mechanism has implications for the American Dream. Declining mobility is often interpreted as evidence that the dream is broken. Our findings suggest that belief in the dream can persist when actual mobility falls, because part of the value of this belief lies in anticipatory utility, which increases with inequality. This may help explain why optimism about mobility is most pronounced in the most unequal societies, where it can reflect motivated beliefs rather than missing information.

References

Akbaş, M., Ariely, D., and Yuksel, S. (2019). When is Inequality Fair? An Experiment on the Effect of Procedural Justice and Agency. *Journal of Economic Behavior & Organization*, 161(C):114–127.

Alesina, A. and Fuchs-Schündeln, N. (2007). Good-bye Lenin (or not?): The effect of communism on people's preferences. *American Economic Review*, 97(4):1507–1528.

Alesina, A., Stantcheva, S., and Teso, E. (2018). Intergenerational Mobility and Preferences for Redistribution. *American Economic Review*, 108(2):521–554.

Alesina, A. F., Glaeser, E. L., and Sacerdote, B. (2001). Why doesn't the US have a European-Style Welfare System? *Brookings Papers on Economic Activity*, (1):187–254.

Almås, I., Cappelen, A. W., and Tungodden, B. (2020). Cutthroat Capitalism versus Cuddly Socialism: Are Americans More Meritocratic and Efficiency-Seeking than Scandinavians? *Journal of Political Economy*, 128(5):1753–1788.

Andreoni, J. and Miller, J. (2002). Giving According to GARP: An Experimental Test of the Consistency of Preferences for Altruism. *Econometrica*, 70(2):737–753.

Balafoutas, L., Kocher, M. G., Putterman, L., and Sutter, M. (2013). Equality, Equity and Incentives: An Experiment. *European Economic Review*, 60:32–51.

Barr, A., Miller, L., and Ubeda, P. (2023). Is the Acknowledgment of Earned Entitlement Effect Robust Across Experimental Modes and Populations? *Sociological Methods & Research*, 52(1):209–230.

Bartling, B., Weber, R. A., and Yao, L. (2015). Do Markets Erode Social Responsibility? *Quarterly Journal of Economics*, 130(1):219–266.

Bauer, M., Cassar, A., Chytilová, J., and Henrich, J. (2014). War's Enduring Effects on the Development of Egalitarian Motivations and In-Group Biases. *Psychological Science*, 25(1):47–57.

Bellemare, C., Kröger, S., and Van Soest, A. (2008). Measuring Inequity Aversion in a Heterogeneous Population Using Experimental Decisions and Subjective Probabilities. *Econometrica*, 76(4):815–839.

Bolton, G. E. and Ockenfels, A. (2000). ERC: A Theory of Equity, Reciprocity, and Competition. *American Economic Review*, 91(1):166–193.

Brunnermeier, M. K. and Parker, J. A. (2005). Optimal expectations. *American Economic Review*, 95(4):1092–1118.

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

Bénabou, R. and Tirole, J. (2002). Self-Confidence and Personal Motivation. *The Quarterly Journal of Economics*, 117(3):871–915.

Bénabou, R. and Tirole, J. (2006). Belief in a Just World and Redistributive Politics. *The Quarterly Journal of Economics*, 121(2):699–746. Publisher: Oxford University Press.

Caplin, A. and Leahy, J. V. (2019). Wishful thinking. Technical report, National Bureau of Economic Research.

Cappelen, A. W., Hole, A. D., Sørensen, E. Ø., and Tungodden, B. (2007). The pluralism of fairness ideals: An experimental approach. *American Economic Review*, 97(3):818–827.

Cappelen, A. W., Konow, J., Sørensen, E. Ø., and Tungodden, B. (2013). Just Luck: An Experimental Study of Risk-Taking and Fairness. *American Economic Review*, 103(4):1398–1413.

Cappelen, A. W., Mollerstrom, J., Reme, B.-A., and Tungodden, B. (2022). A Meritocratic Origin of Egalitarian Behaviour. *Economic Journal*, 132(646):2101–2117.

Cappelen, A. W., Sørensen, E. Ø., and Tungodden, B. (2010). Responsibility for What? Fairness and Individual Responsibility. *European Economic Review*, 54(3):429–441.

Cappelen, A. W. and Tungodden, B. (2019). *The Economics of Fairness*. Edward Elgar Publishing Limited.

Cassar, L. and Klein, A. H. (2019). A Matter of Perspective: How Failure Shapes Distributive Preferences. *Management Science*, 65(11):5050–5064.

Charness, G. and Rabin, M. (2002). Understanding Social Preferences with Simple Tests. *Quarterly Journal of Economics*, 117(3):817–869.

Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

Chetty, R., Grusky, D., Hell, M., Hendren, N., Manduca, R., and Narang, J. (2017). The fading American dream: Trends in absolute income mobility since 1940. *Science*, 356(6336):398–406. Publisher: American Association for the Advancement of Science.

Corak, M. (2013). Income Inequality, Equality of Opportunity, and Intergenerational Mobility. *Journal of Economic Perspectives*, 27(3):79–102.

Cruces, G., Perez-Truglia, R., and Tetaz, M. (2013). Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment. *Journal of Public Economics*, 98:100–112.

Danz, D., Vesterlund, L., and Wilson, A. J. (2022). Belief Elicitation and Behavioral Incentive Compatibility. *American Economic Review*, 112(9):2851–2883.

Davidai, S. and Gilovich, T. (2015). Building a more mobile america—one income quintile at a time. *Perspectives on Psychological Science*, 10(1):60–71.

De Tocqueville, A. (1835). *Democracy in America*. Saunders and Otley.

Durante, R., Putterman, L., and Van der Weele, J. (2014). Preferences for redistribution and perception of fairness: An experimental study. *Journal of the European Economic Association*, 12(4):1059–1086.

Durlauf, S. N., Kourtellos, A., and Tan, C. M. (2022). The great Gatsby curve. *Annual Review of Economics*, 14:571–605.

Durlauf, S. N. and Seshadri, A. (2018). Understanding the Great Gatsby curve. *NBER Macroeconomics Annual*, 32:333–393.

Eil, D. and Rao, J. M. (2011). The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself. *American Economic Journal: Microeconomics*, 3(2):114–138.

Engelmann, D. and Strobel, M. (2004). Inequality Aversion, Efficiency, and Maximin Preferences in Simple Distribution Experiments. *American Economic Review*, 94(4):857–869.

Engelmann, J. B., Lebreton, M., Salem-Garcia, N. A., Schwardmann, P., and Van Der Weele, J. J. (2024). Anticipatory Anxiety and Wishful Thinking. *American Economic Review*, 114(4):926–960.

Exley, C. L. and Kessler, J. B. (2024). Equity Concerns are Narrowly Framed. *American Economic Journal: Microeconomics*, 16(2):147–179.

Fehr, D. and Vollmann, M. (2022). Misperceiving Economic Success: Experimental Evidence on Meritocratic Beliefs and Inequality Acceptance. *SSRN Electronic Journal*.

Fehr, E., Epper, T., and Senn, J. (2024). Social preferences and redistributive politics. *Review of Economics and Statistics*, pages 1–45.

Fehr, E., Kirchsteiger, G., and Riedl, A. (1993). Does Fairness Prevent Market Clearing? An Experimental Investigation. *Quarterly Journal of Economics*, 108(2):437–459.

Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868.

Giuliano, P. and Spilimbergo, A. (2025). Aggregate shocks and the formation of preferences and beliefs. *Journal of Economic Literature*, 63(2):542–597.

Hossain, T. and Okui, R. (2013). The Binarized Scoring Rule. *The Review of Economic Studies*, 80(3 (284)):984–1001.

Hoy, C. and Mager, F. (2021). Why are relatively poor people not more supportive of redistribution? evidence from a randomized survey experiment across ten countries. *American Economic Journal: Economic Policy*, 13(4):299–328.

Huffman, D., Raymond, C., and Shvets, J. (2022). Persistent Overconfidence and Biased Memory: Evidence from Managers. *American Economic Review*, 112(10):3141–3175.

Hvidberg, K. B., Kreiner, C. T., and Stantcheva, S. (2023). Social positions and fairness views on inequality. *Review of Economic Studies*, 90(6):3083–3118.

Karadja, M., Mollerstrom, J., and Seim, D. (2017). Richer (and holier) than thou? the effect of relative income improvements on demand for redistribution. *Review of Economics and Statistics*, 99(2):201–212.

Konow, J. (1996). A Positive Theory of Economic Fairness. *Journal of Economic Behavior & Organization*, 31(1):13–35.

Konow, J. (2000). Fair Shares: Accountability and Cognitive Dissonance in Allocation Decisions. *American Economic Review*, 90(4):1072–1091.

Konow, J. (2009). Is Fairness in the Eye of the Beholder? An Impartial Spectator Analysis of Justice. *Social Choice and Welfare*, 33(1):101–127.

Krawczyk, M. (2010). A Glimpse Through the Veil of Ignorance: Equality of Opportunity and Support for Redistribution. *Journal of Public Economics*, 94(1–2):131–141.

Krueger, A. B. (2012). The rise and consequences of inequality in the United States. Speech at the Center for American Progress, January 12.

Kuziemko, I., Norton, M. I., Saez, E., and Stantcheva, S. (2015). How elastic are preferences for redistribution? evidence from randomized survey experiments. *American Economic Review*, 105(4):1478–1508.

Malmendier, U. and Nagel, S. (2011). Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, 126(1):373–416.

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

Malmendier, U. and Shen, L. S. (2024). Scarred consumption. *American Economic Journal: Macroeconomics*, 16(1):322–355.

Mollerstrom, J., Reme, B.-A., and Sørensen, E. Ø. (2015). Luck, Choice and Responsibility: An Experimental Study of Fairness Views. *Journal of Public Economics*, 131:33–40.

Müller, D. and Renes, S. (2021). Fairness Views and Political Preferences: Evidence from a Large and Heterogeneous Sample. *Social Choice and Welfare*, 56(4):679–711.

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

Rabin, M. (1993). Incorporating Fairness into Game Theory and Economics. *American Economic Review*, 83(5):1281–1302.

Roth, C. and Wohlfart, J. (2018). Experienced inequality and preferences for redistribution. *Journal of Public Economics*, 167:251–262.

Schwardmann, P. and van der Weele, J. (2019). Deception and self-deception. *Nature Human Behaviour*, 3(10):1055–1061. Number: 10 Publisher: Nature Publishing Group.

Stantcheva, S. (2021). Understanding tax policy: How do people reason? *Quarterly Journal of Economics*, 136(4):2309–2369.

Sugden, R. and Wang, M. (2020). Equality of Opportunity and the Acceptability of Outcome Inequality. *European Economic Review*, 130:103597.

Zimmermann, F. (2020). The Dynamics of Motivated Beliefs. *American Economic Review*, 110(2):337–363.

A Additional figures and tables: experimental evidence

Table A.1: The belief channel: OLS and instrumental-variables estimates

	(1) OLS	(2) IV: Feedback×Score
Rank-1 belief	0.032***	0.043**
	(0.003)	(0.020)
First-stage F		23.1
N	1,199	1,199

Notes: This table reports OLS and instrumental-variables estimates of the effect of the Rank-1 belief (0 to 100) on the dollar amount allocated to the Rich (2 to 10). The sample includes subjects in the Baseline and Feedback treatments. Column 1 is OLS. Column 2 is two-stage least squares instrumenting the Rank-1 belief with the interaction of feedback and quiz score. Both columns control for feedback, the inequality condition, quiz score, and its square. Robust standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.2: Descriptive statistics by inequality condition (Baseline treatment)

Inequality condition	High inequality			Low inequality			
Variable	N	Mean	SD	N	Mean	SD	Test
Gender	377			396			X2=1.285
... Female	186	49%		198	50%		
... Male	186	49%		190	48%		
... Other	5	1%		7	2%		
... Rather not say	0	0%		1	0%		
Age in years	377	38	13	396	37	12	F=0.709
College degree (0/1)	377	0.7	0.46	396	0.7	0.46	F=0.01
Ethnicity	377			396			X2=4.159
... White	226	60%		234	59%		
... African American	78	21%		101	26%		
... Hispanic	18	5%		14	4%		
... East Asian	22	6%		19	5%		
... Other	30	8%		24	6%		
... Prefer not to say	3	1%		4	1%		
Voted Republican in last election (0/1)	377	0.32	0.47	396	0.36	0.48	F=1.205
Household income last year (USD)	377			396			X2=3.088
... $0 - $30,000	61	16%		76	19%		
... $30,001 - $60,000	93	25%		106	27%		
... $60,001 - $100,000	99	26%		96	24%		
... $100,001 - $150,000	74	20%		75	19%		
... $150,001 or above	39	10%		36	9%		
... Prefer not to say	11	3%		7	2%		
Rank-1 belief	377	41	27	396	34	23	F=14.748***
Share to the Rich	377	48	25	396	41	22	F=18.306***
Quiz score	377	6.5	3	396	6.6	2.9	F=0.199
Time on the quiz (max. 10 minutes)	377	5.2	2.2	396	4.9	2.3	F=3.653
Importance of effort for success in the quiz (0 to 100)	377	41	18	396	40	16	F=1.648
Importance of luck for success in the quiz (0 to 100)	377	18	14	396	18	13	F=0.005
Importance of intelligence for success in the quiz (0 to 100)	377	41	19	396	43	17	F=1.358
Self-reported effort (1 to 10)	377	9	1.6	396	9	1.4	F=0.001
Beliefs about others' effort (1 to 10)	377	8.1	1.8	396	8.1	1.8	F=0.002
Willingness to take risks (0 to 10)	377	6	2.6	396	5.9	2.6	F=0.028
It is fair for Rank 1 to earn more (0 to 10)	377	7	2.6	396	7	2.6	F=0.001
Society should not aim to equalize incomes (1 to 10)	377	5.2	2.7	396	5.2	2.8	F=0
Position in relative income distribution (1 to 5)	377	2.7	0.73	396	2.8	0.71	F=0.057
Chance of being in the highest income group (0 to 100)	377	9.3	15	396	8.6	14	F=0.533
Used AI during the quiz (1 = yes, 2 = no, 3 = never heard of ChatGPT or Claude)	377	2	0.17	396	2	0.1	F=0.084
Thought others used AI (1 = yes, 2 = no)	377	1.9	0.35	396	1.9	0.33	F=0.478

Statistical significance markers: * p<0.05, ** p<0.01, *** p<0.001

Table A.3: Empirical distribution of quiz scores and rank probabilities

Quiz score	Share of subjects (%)	P(Rank 1)	P(Rank 2)	P(Rank 3)	P(Rank 4)	P(Rank 5)
0	2.92	4.0%	4.0%	4.1%	8.3%	79.6%
1	4.50	4.0%	4.0%	5.1%	18.2%	68.7%
2	5.42	4.0%	4.3%	8.1%	28.0%	55.6%
3	6.26	4.1%	5.1%	12.6%	34.1%	44.2%
4	6.17	4.1%	6.4%	18.4%	37.3%	33.7%
5	8.76	4.7%	10.0%	25.4%	36.5%	23.4%
6	9.17	5.8%	15.9%	30.7%	32.5%	15.2%
7	9.67	8.3%	22.9%	33.4%	25.6%	9.8%
8	13.93	14.6%	31.0%	31.9%	16.6%	5.8%
9	15.51	28.9%	37.2%	21.2%	8.2%	4.5%
10	17.68	59.8%	24.0%	7.8%	4.3%	4.0%
Average		20.0%	20.0%	20.0%	20.0%	20.0%

Table A.4: Initial inequality, the Rank-1 belief, and inequality acceptance

	(1) Rank-1 belief	(2) Share to the Rich	(3) Share to the Rich	(4) Share to the Rich
High inequality condition	6.865***		7.225***	5.305**
	(1.788)		(1.689)	(1.629)
Rank-1 belief		0.294***		0.280***
		(0.032)		(0.033)
Constant	34.427***	33.198***	40.788***	31.158***
	(1.248)	(1.469)	(1.179)	(1.589)
Num.Obs.	773	773	773	773
R2	0.019	0.097	0.023	0.109

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001
This table reports OLS regressions for the Baseline treatment. Column 1 regresses the Rank-1 belief on an indicator for the High inequality condition. Columns 2 to 4 regress inequality acceptance, measured by the share to the Rich, on the Rank-1 belief (Column 2), the indicator (Column 3), or both (Column 4). Standard errors in parentheses.

Table A.5: Robustness to rank probabilities above 84%

	(1) Rank-1 belief	(2) Rank-1 belief	(3) Rank-1 belief	(4) Share to the Rich	(5) Share to the Rich	(6) Share to the Rich
High inequality condition	6.865***	4.539**	6.238***			
	(1.788)	(1.613)	(1.681)			
Rank-1 belief				0.294***	0.312***	0.314***
				(0.032)	(0.037)	(0.034)
Constant	34.427***	32.859***	34.035***	33.198***	32.584***	32.676***
	(1.248)	(1.115)	(1.174)	(1.469)	(1.531)	(1.514)
Num.Obs.	773	735	773	773	735	773
R2	0.019	0.011	0.018	0.097	0.088	0.097

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001
This table reports OLS regressions of the Rank-1 belief on an indicator for the High inequality condition (Columns 1 to 3) and of the share to the Rich on the Rank-1 belief (Columns 4 to 6). The sample includes subjects in the Baseline treatment. Columns 1 and 4 use the full sample, and Columns 2 and 5 exclude subjects who stated an impossible probability (above 84%) for any rank. Columns 3 and 6 instead keep these subjects and replace their Rank-1 belief with 0.8 times the stated probability plus 4 percentage points to account for luck in the ranking (a stated 100% becomes 84%). Standard errors in parentheses.

Table A.6: Descriptive statistics by inequality condition (Feedback treatment)

Inequality condition	High inequality			Low inequality			
Variable	N	Mean	SD	N	Mean	SD	Test
Gender	207			219			X2=0.568
... Female	97	47%		102	47%		
... Male	108	52%		113	52%		
... Other	2	1%		4	2%		
Age in years	207	38	12	219	39	13	F=0.969
College degree (0/1)	207	0.65	0.48	219	0.72	0.45	F=2.38
Ethnicity	207			219			X2=2.566
... White	137	66%		139	63%		
... African American	45	22%		52	24%		
... Hispanic	7	3%		12	5%		
... East Asian	4	2%		6	3%		
... Other	14	7%		10	5%		
... Prefer not to say	0	0%		0	0%		
Voted Republican in last election (0/1)	207	0.3	0.46	219	0.4	0.49	F=4.496*
Household income last year (USD)	207			219			X2=5.009
... $0 - $30,000	36	17%		25	11%		
... $30,001 - $60,000	61	29%		58	26%		
... $60,001 - $100,000	54	26%		64	29%		
... $100,001 - $150,000	35	17%		42	19%		
... $150,001 or above	18	9%		27	12%		
... Prefer not to say	3	1%		3	1%		
Rank-1 belief	207	34	27	219	29	22	F=3.364
Share to the Rich	207	47	25	219	42	23	F=4.436*
Quiz score	207	6.5	2.9	219	6.5	2.9	F=0.034
Time on the quiz (max. 10 minutes)	207	5.3	2.3	219	5	2.1	F=2.389
Importance of effort for success in the quiz (0 to 100)	207	39	16	219	40	15	F=0.133
Importance of luck for success in the quiz (0 to 100)	207	20	15	219	20	15	F=0.009
Importance of intelligence for success in the quiz (0 to 100)	207	41	18	219	40	16	F=0.177
Self-reported effort (1 to 10)	207	9	1.5	219	8.8	1.8	F=0.917
Beliefs about others' effort (1 to 10)	207	8.3	1.7	219	8	2	F=2.976
Willingness to take risks (0 to 10)	207	5.7	2.6	219	6.2	2.5	F=3.396
It is fair for Rank 1 to earn more (0 to 10)	207	7.2	2.6	219	7	2.6	F=0.343
Society should not aim to equalize incomes (1 to 10)	207	4.9	2.8	219	5.3	2.8	F=1.589
Position in relative income distribution (1 to 5)	207	2.7	0.75	219	2.8	0.65	F=1.41
Chance of being in the highest income group (0 to 100)	207	8.2	16	219	8.9	15	F=0.235
Used AI during the quiz (1 = yes, 2 = no, 3 = never heard of ChatGPT or Claude)	207	2	0.18	219	2	0.095	F=1.149
Thought others used AI (1 = yes, 2 = no)	207	1.9	0.33	219	1.9	0.35	F=0.588

Statistical significance markers: * p<0.05, ** p<0.01, *** p<0.001

Figure A.1: Share to the Rich by inequality condition

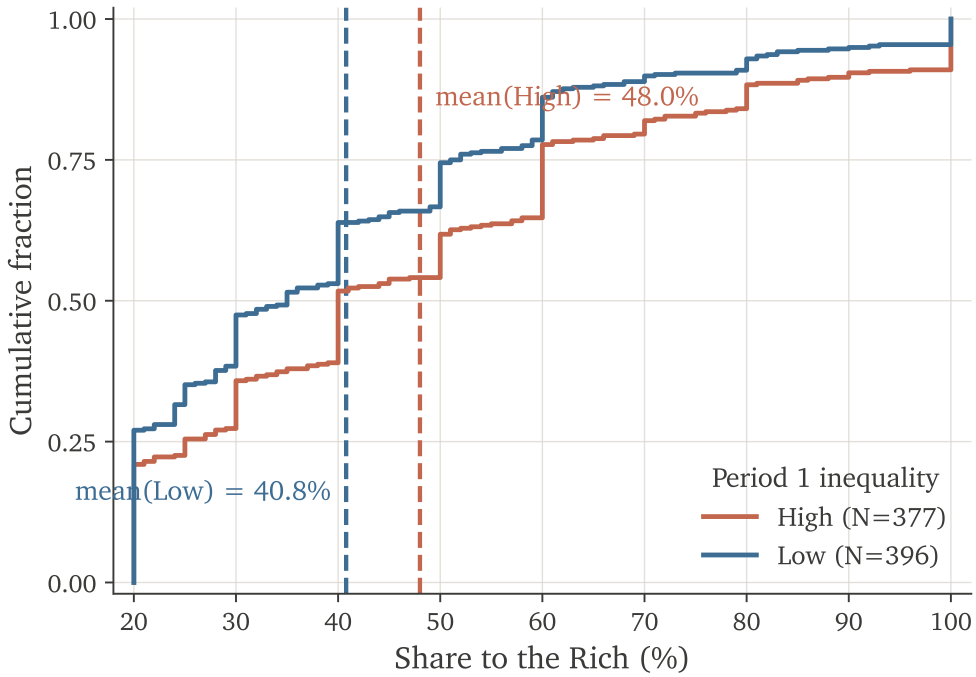

Notes: This figure shows the cumulative distribution of the share to the Rich by inequality condition in the Baseline treatment. The rust line is the High inequality condition, the slate line the Low. Vertical dashed lines mark the mean share to the Rich in each inequality condition. A Kolmogorov–Smirnov test rejects equality of the two distributions ($p < 0.01$).

Figure A.2: Effect of initial inequality on the Rank-1 belief and inequality acceptance: Baseline versus Feedback

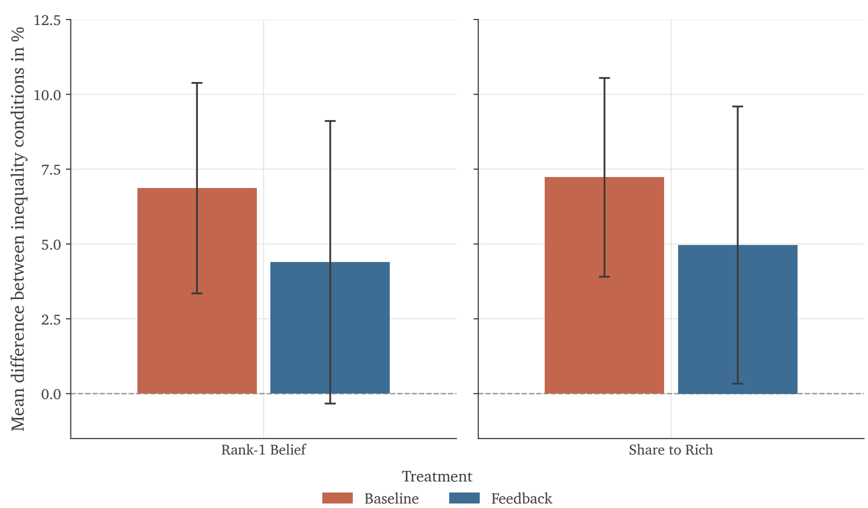

Notes: This figure shows the mean difference between inequality conditions (High minus Low, in percentage points) for the Rank-1 belief (left) and the share to the Rich (right), by treatment. The effect of initial inequality is smaller in the Feedback treatment, where it is significant at the 10% level for the Rank-1 belief and at the 5% level for the share to the Rich. The difference between the effects in the two treatments is not statistically significant. Error bars show 95% confidence intervals.

Figure A.3: Distribution of Rank-1 and Rank-2 beliefs by inequality condition

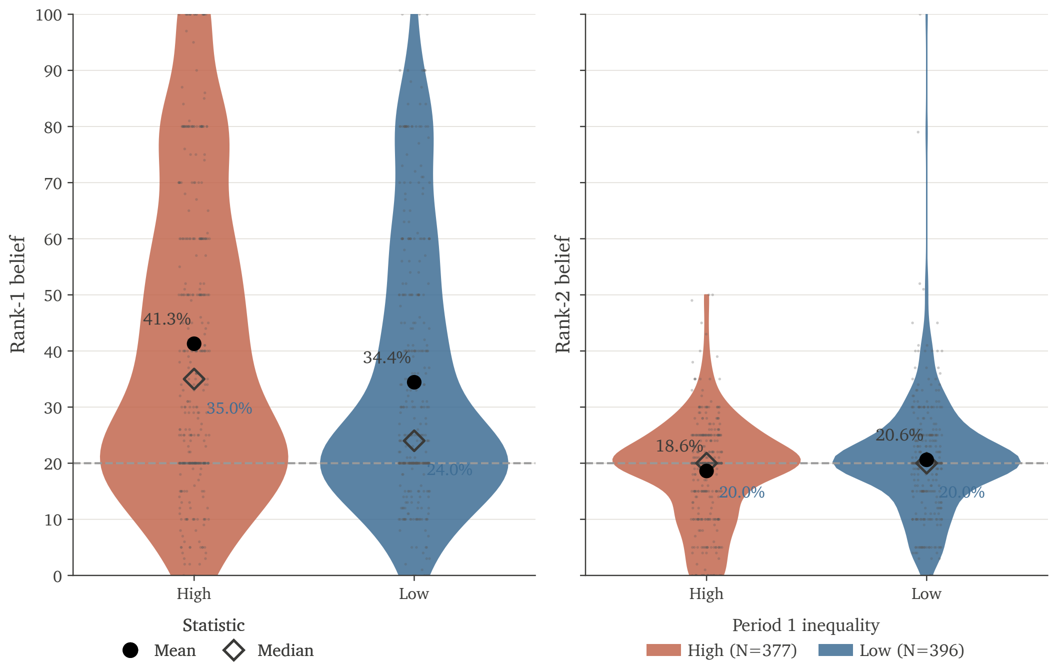

Notes: This figure shows the distribution of Rank-1 and Rank-2 beliefs by inequality condition in the Baseline treatment. Black dots are means and hollow diamonds are medians. The dashed line marks the 20% average objective probability. A Kolmogorov–Smirnov test rejects equality of the Rank-1 belief distributions across inequality conditions ($p < 0.01$), and the High distribution has more mass in the upper tail. See Figure 5 for the average belief for each rank.

Figure A.4: Deviations from the objective probability by quiz score and inequality condition

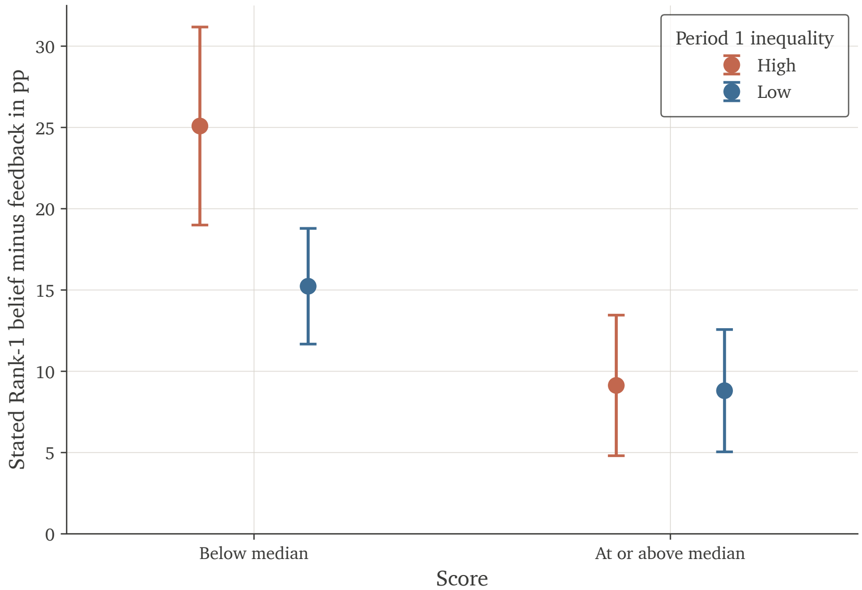

Notes: This figure shows the mean deviation of the stated Rank-1 belief from the objective probability in the Feedback treatment, by inequality condition and quiz score. "Below median" is a quiz score of 6 or fewer correct answers. The deviation is concentrated among below-median performers, for whom it is significantly larger in the High inequality condition than in the Low inequality condition ($p < 0.05$). Among performers at or above the median, the two inequality conditions do not differ significantly. Error bars show 95% confidence intervals.

Figure A.5: Self-reported effort and beliefs about others' effort

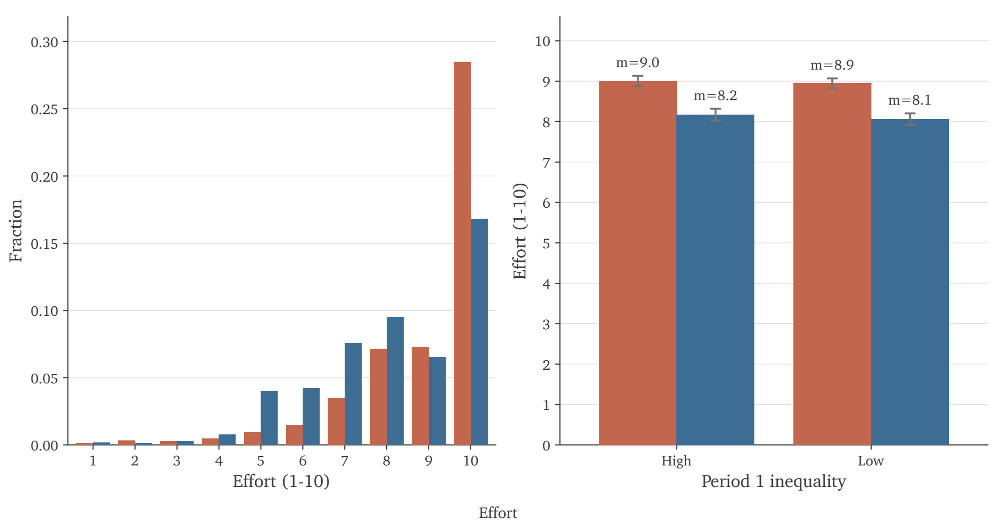

Notes: This figure shows how much effort subjects reported exerting in the quiz and how much effort they thought others exerted, both on a scale from 1 to 10. The left panel shows the distribution of ratings, and the right panel shows mean ratings by inequality condition. The sample includes subjects in the Baseline and Feedback treatments. Error bars show 95% confidence intervals.

Figure A.6: Factors determining success in the quiz

Notes: This figure shows the average shares of effort, luck, and intelligence that subjects thought were needed to succeed in the quiz, by inequality condition. The sample includes subjects in the Baseline treatment.

Figure A.7: Relevant factors in the allocation choice

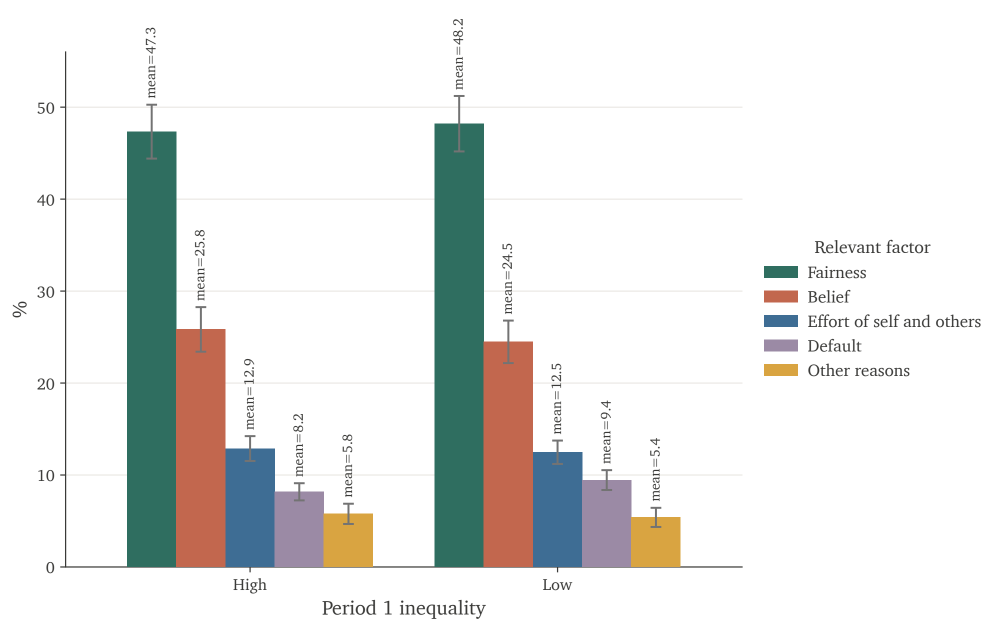

Notes: This figure shows the factors that subjects reported as having influenced their allocation choice, by inequality condition. Subjects were given 100 points to distribute between the factors fairness, belief about being Rich, effort, the default distribution, and other factors, to explain what influenced their decision the most. The sample includes subjects in the Baseline treatment. Error bars show 95% confidence intervals.

Figure A.8: Experimental evidence by subgroup

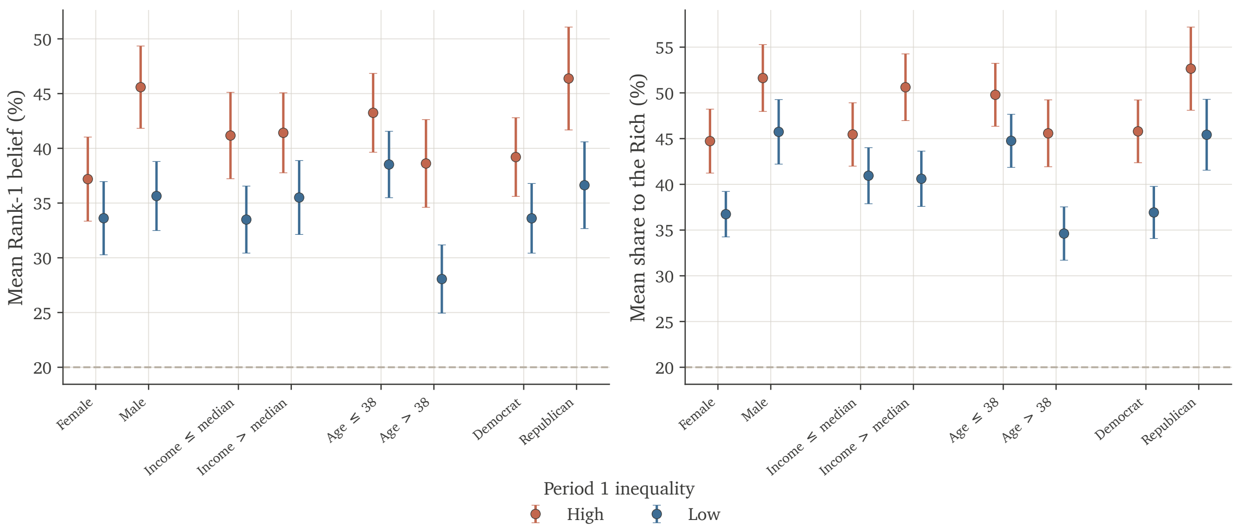

Notes: This figure shows the mean Rank-1 belief (left) and the mean share to the Rich (right), in percent, by inequality condition (High in rust, Low in slate) within each subgroup: gender, household income (split at the sample median), age (split at the US median of 38.5 years), and political orientation (Democratic versus Republican voters). The sample includes subjects in the Baseline treatment. The dashed line marks the 20% average objective probability in the left panel and the share to the Rich that fully equalizes income in the right panel. Error bars show 95% confidence intervals. The estimated effect of initial inequality on both outcomes is positive in every subgroup, and no subgroup interaction with the High inequality condition is significant at the 5% level (Tables A.7 and A.8).

Table A.7: Subgroup analysis: effect of initial inequality on the Rank-1 belief

	(1)	(2)	(3)	(4)	(5)
High inequality condition	6.865***	4.720*	3.804	8.010***	5.677**
	(1.788)	(2.295)	(2.479)	(2.299)	(2.194)
Age >38		−10.465***			
		(2.528)			
High inequality condition × Age >38		5.838			
		(3.598)			
Male			2.336		
			(2.483)		
High inequality condition × Male			6.140+		
			(3.554)		
Income >median				2.802	
				(2.578)	
High inequality condition × Income >median				−3.024	
				(3.663)	
Republican					3.448
					(2.587)
High inequality condition × Republican					4.071
					(3.755)
Constant	34.427***	38.523***	33.306***	33.372***	33.182***
	(1.248)	(1.582)	(1.720)	(1.582)	(1.554)
Num.Obs.	773	773	773	773	773
R2	0.019	0.044	0.034	0.020	0.031

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001

This table reports OLS estimates of the effect of initial inequality on the Rank-1 belief by subgroup in the Baseline treatment. The interaction terms test whether this effect differs by age (above the US median age of 38.5 years), gender (male), household income (above $80,000, about the 2023 US median), and political orientation (Republican). Standard errors in parentheses.

Table A.8: Subgroup analysis: effect of initial inequality on the share to the Rich

	(1)	(2)	(3)	(4)	(5)
High inequality condition	7.225***	5.034*	8.274***	5.087*	7.626***
	(1.689)	(2.167)	(2.322)	(2.166)	(2.064)
Age >38		−10.136***			
		(2.387)			
High inequality condition × Age >38		5.925+			
		(3.398)			
Male			9.514***		
			(2.325)		
High inequality condition × Male			−2.387		
			(3.328)		
Income >median				0.254	
				(2.430)	
High inequality condition × Income >median				5.180	
				(3.452)	
Republican					7.250**
					(2.433)
High inequality condition × Republican					−0.398
					(3.532)
Constant	40.788***	44.755***	36.223***	40.692***	38.170***
	(1.179)	(1.493)	(1.611)	(1.490)	(1.462)
Num.Obs.	773	773	773	773	773
R2	0.023	0.049	0.055	0.029	0.043

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001

This table reports OLS estimates of the effect of initial inequality on inequality acceptance, measured by the share to the Rich, by subgroup in the Baseline treatment. The interaction terms test whether this effect differs by age (above the US median age of 38.5 years), gender (male), household income (above $80,000, about the 2023 US median), and political orientation (Republican). Standard errors in parentheses.

B Robustness of the cross-country experience-expectation gap (Figure 1)

This appendix presents the robustness checks for the cross-country experience-expectation gap in Figure 1, which Section 3 summarizes. The experience-expectation gap holds under alternative measures of inequality (the top-5% and top-1% income shares in place of the Gini coefficient), under alternative measures of mobility (the continuous measure, status persistence as in Corak (2013), and a rags-to-riches transition rate), when South Africa and the Philippines, the two outlier countries, are included, and, for expected mobility, under country-level controls (log GDP per capita, GDP growth, unemployment, V-Dem polyarchy, and years of education). To make the opposite signs directly comparable, every check except the one with country-level controls pairs an actual-mobility panel with an expected-mobility panel measured on the same ladder. Finally, we report the country-level expected-minus-actual gap, the difference between expected and actual mobility.

Outlier exclusion: the full 27-country sample

The experience-expectation gap holds when South Africa and the Philippines are included. Figure B.1 plots the Gini coefficient against the binary measure for the full sample of 27 countries (top row) and after excluding South Africa from both panels and the Philippines from the expected-mobility panel (bottom row). In every panel the sign is unchanged: actual mobility is lower in more unequal countries, while expected mobility is higher.

Figure B.1: Gini and mobility on the binary measure: full sample and outlier exclusion

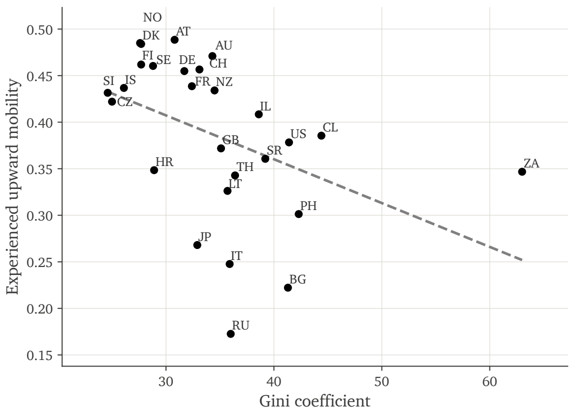

(a) Actual mobility, all 27 countries

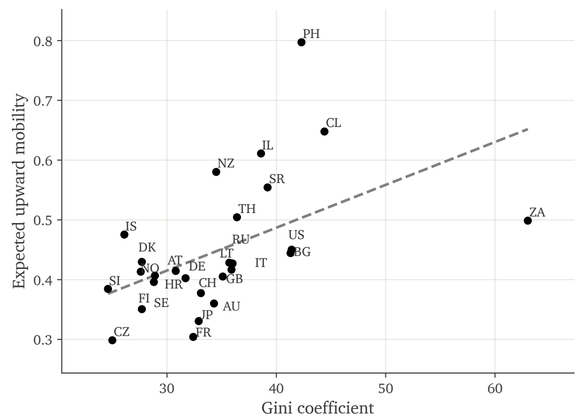

(b) Expected mobility, all 27 countries

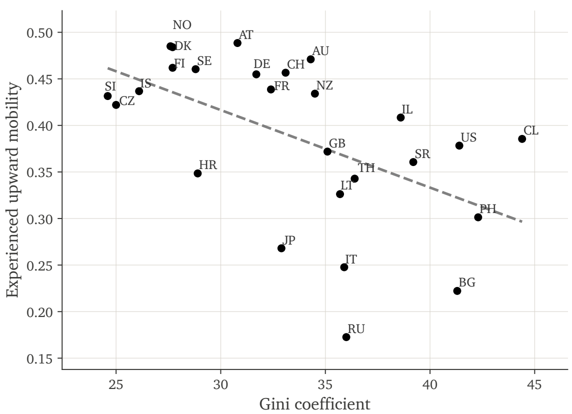

(c) Actual mobility, outliers excluded

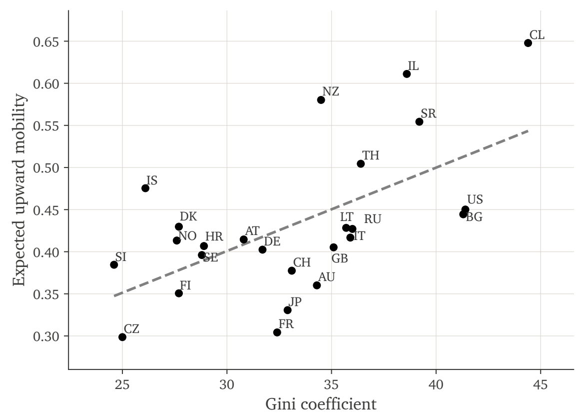

(d) Expected mobility, outliers excluded

Notes: This figure plots the Gini coefficient against actual and expected mobility. Each point is one of the 27 countries in the 2019 International Social Survey Programme (ISSP), a survey that asks respondents to place their parents, themselves today, and themselves in 10 years on a social ladder from 1 to 10. Inequality is the Gini coefficient of disposable income. Actual mobility is the country share of respondents who place themselves above their parents ($v_{41} > v_{42}$); expected mobility is the country share who expect a higher status in 10 years than today ($v_{43} > v_{41}$). The top row retains all 27 countries. The bottom row excludes South Africa from the actual-mobility panel (26 countries) and South Africa and the Philippines from the expected-mobility panel (25 countries). The dashed line in each panel is the linear fit.

Alternative inequality measures: top-5% and top-1% income shares

Measuring inequality by top income shares leaves the experience-expectation gap intact. Figure B.2 replaces the Gini coefficient with the top-5% income share (top row) and the top-1% income share (bottom row). Both signs are unchanged under either share: actual mobility is lower, and expected mobility is higher, in countries where the top income share is larger.

Figure B.2: Top income shares and mobility on the binary measure

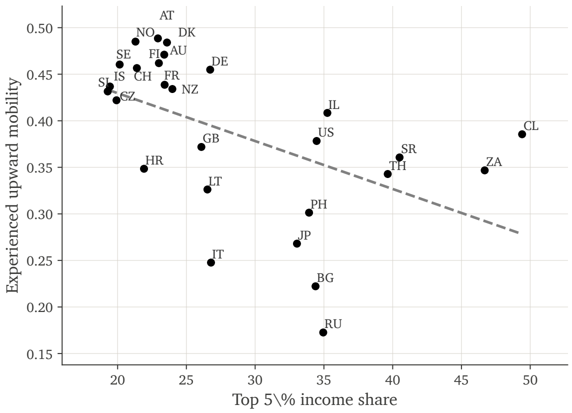

(a) Actual mobility, top-5% share

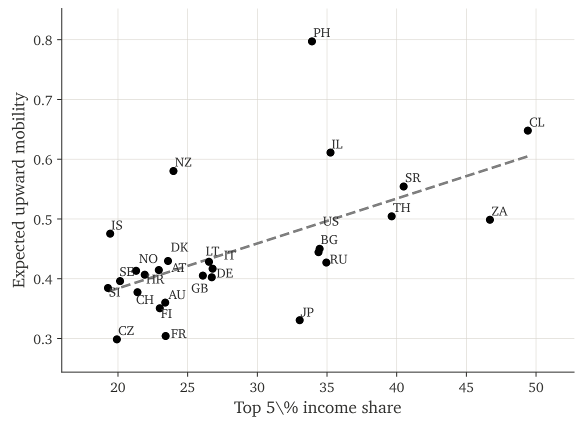

(b) Expected mobility, top-5% share

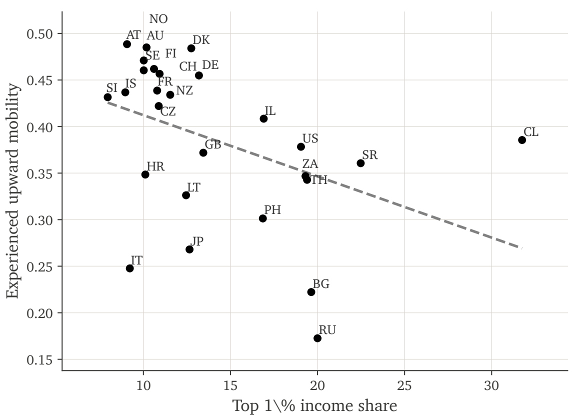

(c) Actual mobility, top-1% share

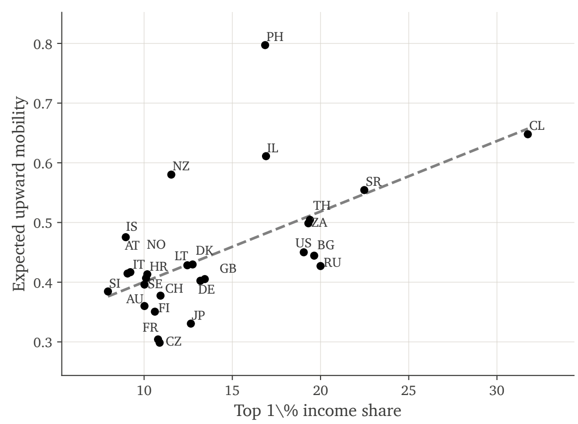

(d) Expected mobility, top-1% share

Notes: This figure plots top income shares against actual and expected mobility. Each point is one of the 27 ISSP 2019 countries. Inequality is measured by the top-5% income share (top row) and the top-1% income share (bottom row) from the World Inequality Database (WID), the fraction of national income accruing to the highest 5% and the highest 1% of earners respectively. Actual mobility is the country share of respondents who place themselves above their parents ($v_{41} > v_{42}$); expected mobility is the country share who expect a higher status in 10 years than today ($v_{43} > v_{41}$). The dashed line in each panel is the linear fit.

Alternative mobility measures

The experience-expectation gap also holds under alternative measures of mobility. Figure B.3 replaces the binary measure with the continuous measure, the country mean of $v_{41} - v_{42}$

for actual mobility and of $v_{43} - v_{41}$ for expected mobility. Figure B.4 measures mobility by status persistence in the sense of Corak (2013), the within-country slope of current status on family-of-origin status, where a larger slope indicates less mobility. Finally, Figure B.5 measures it by a rags-to-riches transition rate, the share of all respondents who start on rungs 1 to 3 of the ladder and reach rungs 8 to 10. In each case the actual measure is associated with lower mobility in more unequal countries while the expected measure is associated with higher mobility.

Figure B.3: Gini and mobility on the continuous measure

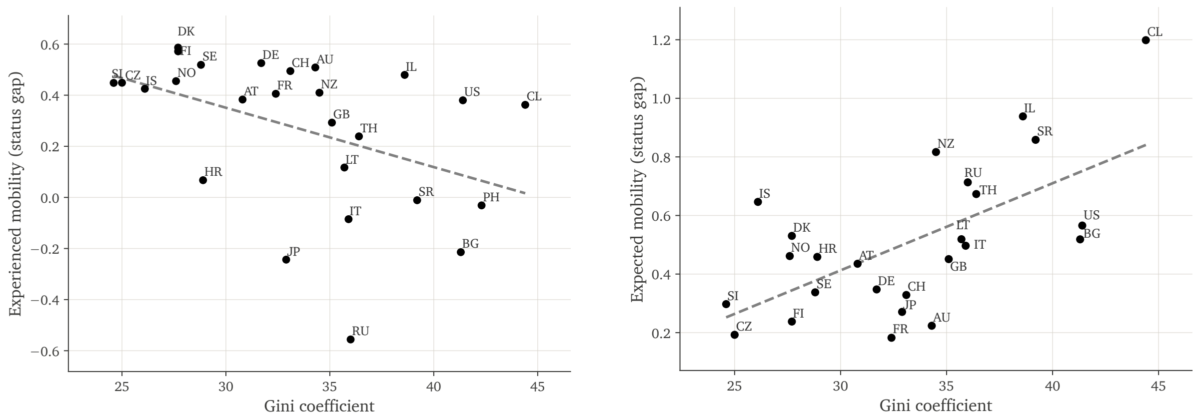

(a) Actual mobility **(b) Expected mobility**

Notes: This figure plots the Gini coefficient against actual and expected mobility. Each point is an ISSP 2019 country. The actual-mobility panel excludes South Africa (26 countries), and the expected-mobility panel excludes South Africa and the Philippines (25 countries). Inequality is the Gini coefficient of disposable income. Mobility is the continuous measure on the 1 to 10 ladder: actual mobility is the country mean of $v_{41} - v_{42}$ (current status minus family-of-origin status); expected mobility is the country mean of $v_{43} - v_{41}$ (expected status in 10 years minus current status). The dashed line in each panel is the linear fit.

Figure B.4: Gini and status persistence

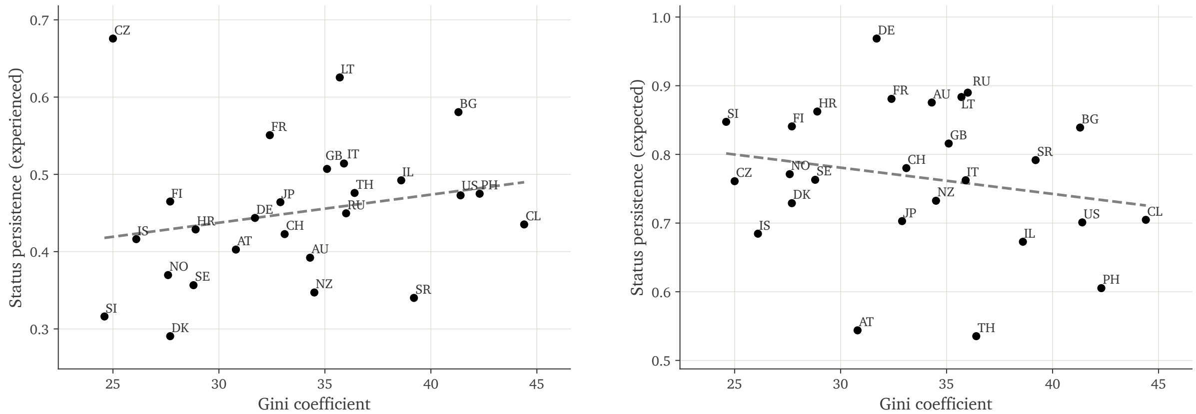

(a) Actual mobility **(b) Expected mobility**

Notes: This figure plots the Gini coefficient against actual and expected mobility. Each point is an ISSP 2019 country, excluding South Africa (26 countries in each panel). Inequality is the Gini coefficient of disposable income. Mobility is measured by status persistence in the sense of Corak (2013): the within-country slope from a regression of next status on prior status, where a larger slope indicates that status persists more strongly and hence less mobility. The actual panel uses the slope of current status (v_{41}) on family-of-origin status (v_{42}); the expected panel uses the slope of expected status (v_{43}) on current status (v_{41}). The dashed line in each panel is the linear fit.

Figure B.5: Gini and the rags-to-riches transition rate

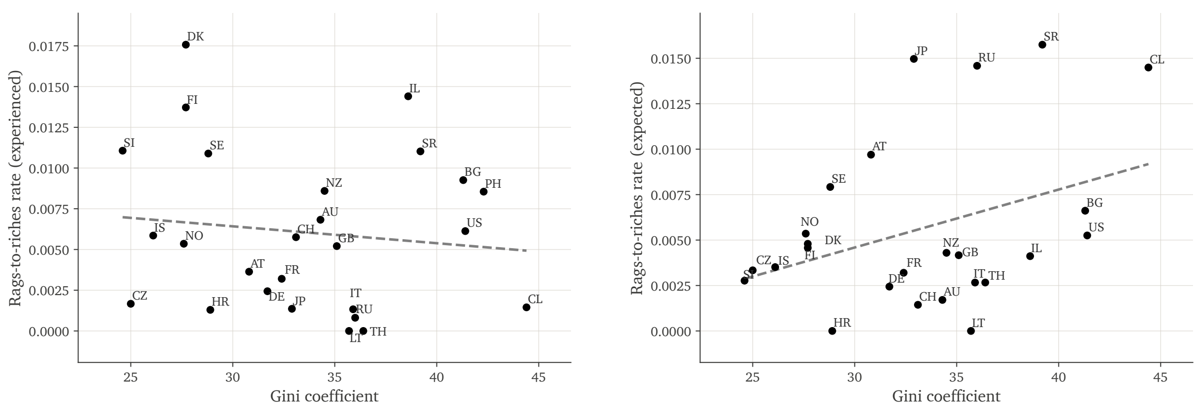

(a) Actual mobility **(b) Expected mobility**

Notes: This figure plots the Gini coefficient against actual and expected mobility. Each point is an ISSP 2019 country. The actual-mobility panel excludes South Africa (26 countries), and the expected-mobility panel excludes South Africa and the Philippines (25 countries). Inequality is the Gini coefficient of disposable income. Mobility is measured by a rags-to-riches transition rate, the share of all respondents in a country who start on rungs 1 to 3 of the 1 to 10 ladder and reach rungs 8 to 10, in the spirit of the bottom-to-top transition probabilities of Chetty and coauthors. The actual panel uses transitions from family-of-origin status (v_{42}) to current status (v_{41}); the expected panel uses transitions from current status (v_{41}) to expected status (v_{43}). The dashed line in each panel is the linear fit.

Country-level controls

The positive association between inequality and expected mobility is robust to country-level controls. We regress the country-mean probability of expecting to move up on the Gini coefficient, adding controls in sequence (Figure B.6 and Table B.1). The coefficient is 0.007 ($p < 0.05$) in the raw bivariate regression and 0.010 ($p < 0.01$) under the full control set of log GDP per capita, 10-year average GDP growth, unemployment, V-Dem polyarchy, and country-mean years of education. The R^2 is 0.26 in the bivariate regression and 0.58 under the full set, with $N = 27$ countries (26 once controls requiring complete data are included).

Figure B.6: Gini coefficient on expected upward mobility across control sets

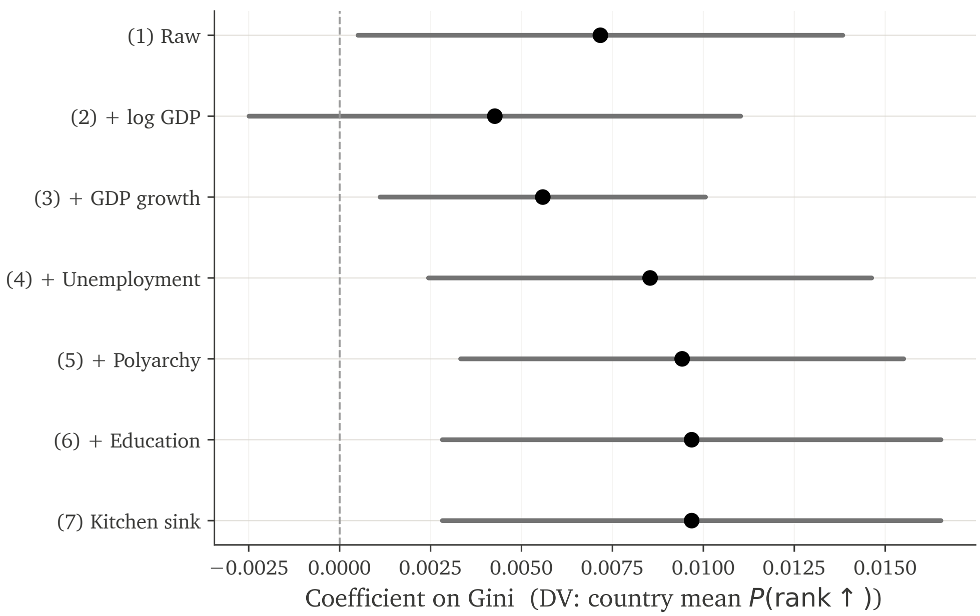

Notes: This figure shows coefficients from country-level regressions of expected upward mobility on the Gini coefficient. The dependent variable is the country-mean probability that a respondent expects a higher status in 10 years than today, across the 27 ISSP 2019 countries (including South Africa and the Philippines). Each marker is the estimated coefficient on the Gini coefficient of disposable income from a country-level OLS regression. Horizontal error bars show 95% confidence intervals based on heteroskedasticity-robust (HC1) standard errors. The specifications add controls in sequence, from the raw bivariate regression to the full set of log GDP per capita, 10-year average GDP growth, unemployment, V-Dem polyarchy, and country-mean years of education. Table B.1 reports the full regressions.

Table B.1: Gini coefficient and expected upward mobility: country-level controls

	Dependent variable: country-mean $P(\text{rank} \uparrow)$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gini	0.007**	0.004	0.006**	0.009***	0.009***	0.010***	0.010***
	(0.003)	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)
log GDP per capita		-0.046	0.009	-0.013	-0.051	-0.060	-0.060
		(0.032)	(0.025)	(0.033)	(0.053)	(0.075)	(0.075)
GDP growth (10y avg)			6.272**	4.127	3.512	3.369	3.369
			(2.631)	(2.664)	(2.735)	(2.999)	(2.999)
Unemployment				-0.009	-0.013*	-0.014	-0.014
				(0.006)	(0.007)	(0.009)	(0.009)
Polyarchy					0.186	0.193	0.193
					(0.146)	(0.160)	(0.160)
Years of education						0.004	0.004
						(0.019)	(0.019)
Observations	27	27	26	26	26	26	26
R^2	0.258	0.349	0.497	0.547	0.576	0.577	0.577

Notes: This table reports country-level OLS regressions of expected upward mobility on the Gini coefficient for the 27 ISSP 2019 countries (including South Africa and the Philippines). The dependent variable is the country-mean probability that a respondent expects a higher status in 10 years than today. Column 1 is the raw bivariate regression. Columns 2 to 6 add, in sequence, log GDP per capita (World Bank, 2019), 10-year average GDP growth (Maddison), the country mean of experienced unemployment (the unemployment rate during respondents' formative years, ages 18 to 25), V-Dem polyarchy (2010–2019 mean), and country-mean years of education in the ISSP sample. Column 7 uses the full control set and is identical to Column 6, which already includes every control. Columns 3 to 7 include 26 countries because 10-year average GDP growth is unavailable for Suriname. Heteroskedasticity-robust (HC1) standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

The expected-minus-actual gap

On average, expected mobility exceeds actual mobility at the country level, and the expected-minus-actual gap is larger in more unequal countries. Figure B.7 plots each country's expected mobility against its actual mobility, on the binary measure in Panel (a) and the continuous measure in Panel (b). Points above the 45-degree line are countries whose respondents expect more upward mobility than they have realized. The comparison is imperfect, because expected mobility refers to the next 10 years while actual mobility refers to the move relative to one's parents. Sixteen of the 27 countries lie above the line on the binary measure and 18 on the continuous measure. The mean gap is 0.063 on the binary measure and 0.318 on the continuous measure, and the correlation between the gap and the Gini coefficient is 0.61 on the binary measure and 0.59 on the continuous measure (both $p < 0.01$). On the binary measure, the gap is largest in the Philippines, followed by Chile and Russia, and smallest in France, followed by Czechia and Finland. The full country ranking and correlations appear in Table B.2.

Figure B.7: Expected versus actual mobility by country

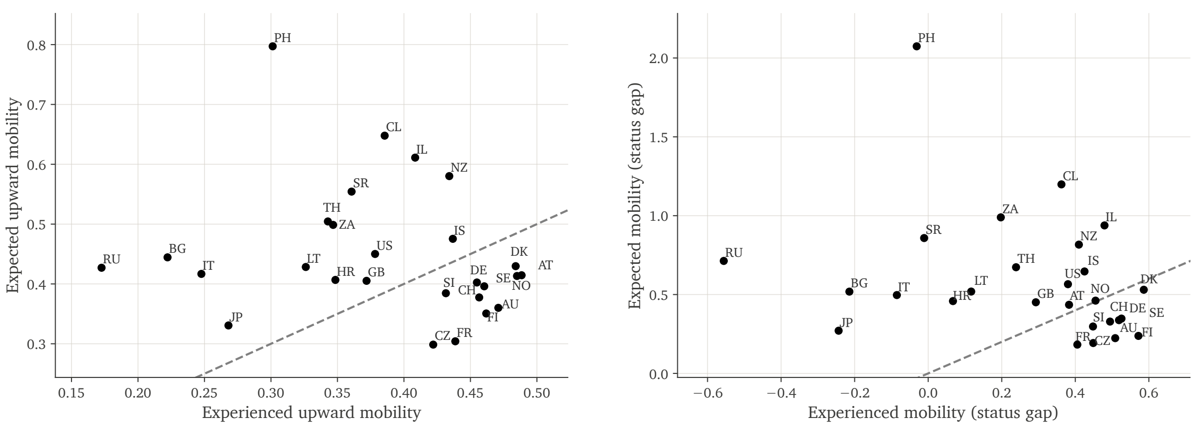

(a) Binary measure **(b) Continuous measure**

Notes: This figure plots expected mobility (vertical axis) against actual mobility (horizontal axis). Each point is one of the 27 ISSP 2019 countries (including South Africa and the Philippines). The dashed line is the 45-degree line; a country above it is one whose respondents expect more upward mobility over the next 10 years than they have realized relative to their parents. Panel (a) uses the binary measure: the country share with $v_{41} > v_{42}$ (actual) and with $v_{43} > v_{41}$ (expected). Panel (b) uses the continuous measure: the country mean of $v_{41} - v_{42}$ (actual) and of $v_{43} - v_{41}$ (expected). The mean expected-minus-actual gap is 0.063 on the binary measure and 0.318 on the continuous measure, and the gap is larger in more unequal countries: its correlation with the Gini coefficient is 0.61 on the binary measure and 0.59 on the continuous measure (both $p < 0.01$).

Table B.2: The expected-minus-actual gap by country, ranked

		Binary measure			*Continuous measure*		
Country	Gini	Actual	Expected	Gap	Actual	Expected	Gap
PHL	42.3	0.301	0.797	+0.496	-0.031	+2.074	+2.105
CHL	44.4	0.386	0.648	+0.262	+0.362	+1.199	+0.836
RUS	36.0	0.173	0.427	+0.254	-0.556	+0.713	+1.269
BGR	41.3	0.222	0.444	+0.222	-0.214	+0.519	+0.733
ISR	38.6	0.408	0.611	+0.203	+0.479	+0.938	+0.459
SUR	39.2	0.361	0.554	+0.194	-0.011	+0.858	+0.869
ITA	35.9	0.248	0.417	+0.169	-0.085	+0.497	+0.582
THA	36.4	0.343	0.504	+0.162	+0.239	+0.673	+0.434
ZAF	63.0	0.347	0.499	+0.152	+0.197	+0.989	+0.792
NZL	34.5	0.434	0.580	+0.146	+0.410	+0.817	+0.407
LTU	35.7	0.326	0.428	+0.102	+0.117	+0.519	+0.402
USA	41.4	0.378	0.450	+0.072	+0.380	+0.566	+0.186
JPN	32.9	0.268	0.331	+0.063	-0.244	+0.271	+0.514
HRV	28.9	0.348	0.407	+0.058	+0.067	+0.459	+0.391
ISL	26.1	0.437	0.475	+0.039	+0.425	+0.646	+0.221
GBR	35.1	0.372	0.405	+0.033	+0.293	+0.451	+0.158
SVN	24.6	0.432	0.385	-0.047	+0.448	+0.297	-0.151
DEU	31.7	0.455	0.402	-0.052	+0.526	+0.348	-0.178
DNK	27.7	0.484	0.430	-0.054	+0.586	+0.530	-0.056
SWE	28.8	0.460	0.396	-0.064	+0.519	+0.338	-0.181
NOR	27.6	0.485	0.413	-0.072	+0.455	+0.461	+0.006
AUT	30.8	0.488	0.415	-0.074	+0.383	+0.435	+0.052
CHE	33.1	0.457	0.377	-0.079	+0.494	+0.329	-0.166
AUS	34.3	0.471	0.360	-0.111	+0.509	+0.224	-0.285
FIN	27.7	0.462	0.351	-0.111	+0.572	+0.238	-0.334
CZE	25.0	0.422	0.299	-0.123	+0.449	+0.193	-0.256
FRA	32.4	0.439	0.304	-0.134	+0.406	+0.182	-0.223
Mean	—	0.385	0.449	+0.063	+0.266	+0.584	+0.318

Notes: This table reports actual mobility, expected mobility, and the expected-minus-actual gap for the 27 ISSP 2019 countries (including South Africa and the Philippines), sorted in descending order of the gap on the binary measure. The binary measure is the country mean of $\mathbf{1}\{v_{41} > v_{42}\}$ (actual) and $\mathbf{1}\{v_{43} > v_{41}\}$ (expected). The continuous measure is the country mean of $v_{41} - v_{42}$ (actual) and $v_{43} - v_{41}$ (expected). Gap = expected minus actual. A positive gap means that the country's respondents expect more upward mobility over the next 10 years than they have realized relative to their parents. Pearson correlations of the gap with country inequality: corr(binary gap, Gini) = +0.61; corr(binary gap, top-5%) = +0.70; corr(continuous gap, Gini) = +0.59; corr(continuous gap, top-5%) = +0.65. All four are significant at $p < 0.01$.

C Additional figures and tables: survey evidence (ISSP and GSS)

Table C.1: Descriptive statistics for main variables of interest in ISSP dataset

	More equal countries (Gini below the median)					More unequal countries (Gini above the median)				
	Mean	SD	Min	Max	N	Mean	SD	Min	Max	N
Current status	5.73	1.68	1.00	10.00	12433	5.09	1.86	1.00	10.00	11628
Family-of-origin status	5.30	1.78	1.00	10.00	12433	4.93	1.99	1.00	10.00	11628
Expected status (in 10 years)	6.03	1.90	1.00	10.00	12433	6.02	2.11	1.00	10.00	11628
Rich should pay higher taxes than poor	3.96	0.76	1.00	5.00	12433	3.95	0.86	1.00	5.00	11628
Age	46.19	10.83	26.00	65.00	12433	44.73	10.89	26.00	65.00	11628
Male	0.48	0.50	0.00	1.00	12415	0.46	0.50	0.00	1.00	11621
Years of education	14.58	4.02	0.00	57.00	12228	11.87	4.31	0.00	52.00	11237
Experienced inequality	0.21	0.05	0.10	0.39	12433	0.34	0.09	0.15	0.52	11628
Experienced unemployment	5.90	2.98	0.37	14.53	11152	8.52	5.92	0.87	26.75	10509

This table reports descriptive statistics for ISSP 2019 respondents aged 26 to 65 in more equal countries (Gini below the median) and more unequal countries (Gini above the median). The more equal countries are Australia, Austria, Croatia, Czechia, Denmark, Finland, France, Germany, Iceland, New Zealand, Norway, Russia, Slovenia, Sweden, and Switzerland. The more unequal countries are Bulgaria, Chile, Israel, Italy, Japan, Lithuania, the Philippines, South Africa, Suriname, Thailand, the UK, and the US. Experienced inequality and experienced unemployment are constructed as in Roth and Wohlfahrt (2018).

C.1 Experienced inequality: full results and robustness

This appendix presents the supporting material for the experienced-inequality analysis of Section 3.2. Table C.2 reports the regression estimates for the four specifications of Figure 2, and Table C.3 reports the coefficient on every regressor other than the fixed effects in those specifications. All coefficients discussed below are replicated in Python and Stata (`reghdfe`) to four decimal places.

Functional form. Figure C.1 residualizes both experienced inequality and expected mobility on the individual controls and the country and subregion-by-cohort fixed effects, and then bins the residuals. The binned means lie close to the linear fit.

Per-rung decomposition. Figure C.2 estimates the association between experienced inequality and the probability of expecting each rung separately. The point estimates are positive for rungs 8 to 10, whereas most lower rungs have negative or zero point estimates. Each per-rung coefficient is imprecisely estimated. The above-mean test below therefore pools the rungs.

Mapping rungs to income, and the above-mean test. Figure C.3 maps each ladder rung to the mean within-country household income of respondents at that rung, locating the country-mean income between rungs 7 and 8. Figure C.4 then shows that, as the framework predicts, experienced inequality is associated with a higher probability of expecting an above-mean-income position. The at-or-below-mean coefficient mirrors it by construction.

Heterogeneity. The coefficient is positive in every age group, and it is concentrated among lower-status respondents, turning negative only in the small top-status bin (Figure C.5). The household-income panel covers only respondents who report their income. In that sample, all five quintile coefficients are smaller than the Specification (3) estimate, and none

Table C.2: **Experienced inequality and expected mobility**

	Dependent variable: Expected mobility			
	(1)	(2)	(3)	(4)
Experienced inequality	0.466***	0.335***	0.288***	0.252**
	(0.135)	(0.113)	(0.104)	(0.111)
Individual controls		Yes	Yes	Yes
Country FE			Yes	Yes
Subregion×cohort FE			Yes	Yes
Experience controls				Yes
N	22,558	21,879	21,879	19,152

Notes: This table reports regressions of expected mobility on experienced inequality. The dependent variable is expected mobility, the respondent's expected status in 10 years minus their current status, each on the 1 to 10 social ladder. Experienced inequality, the Roth–Wohlfahrt measure, is the average top-5% income share in the respondent's country during the formative years (ages 18 to 25), standardized within the sample. The top-5% income share is the share of national income received by the top 5% of earners, and a higher share means that the top 5% earn more relative to mean income. The sample includes ISSP 2019 respondents aged 29 or older. Individual controls are age, an indicator for female, and years of education. Experience controls are family-of-origin status and the six country-level macroeconomic experiences (unemployment, GDP per capita growth, CPI inflation, electoral democracy, banking-or-currency-crisis exposure, and rare-disaster exposure), each averaged over the same formative years. Table C.3 reports the coefficients on each control. Standard errors are clustered at the country level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

is statistically significant. When the top-1% or top-10% income share replaces the top-5% income share, the coefficient remains positive and statistically significant at the 5% level, with or without the six macroeconomic experiences as additional controls (Figure C.6).

Placebo windows. Figure C.7 re-estimates Specification (3) using inequality experienced in windows other than the formative years. Inequality experienced at ages 0 to 7, before the formative years, and at ages 34 to 41, after them, has no significant association with expected mobility. Only inequality experienced during the 18 to 25 window does. The coefficient on the Malmendier–Nagel lifetime-weighted average (Malmendier and Nagel, 2011), which also covers the formative years, is larger than the 18 to 25 coefficient, but its confidence interval includes zero.

C.2 Additional survey evidence: status distributions and the experienced-inequality measure

Figure C.1: Residualized binned scatter plot of expected mobility on experienced inequality

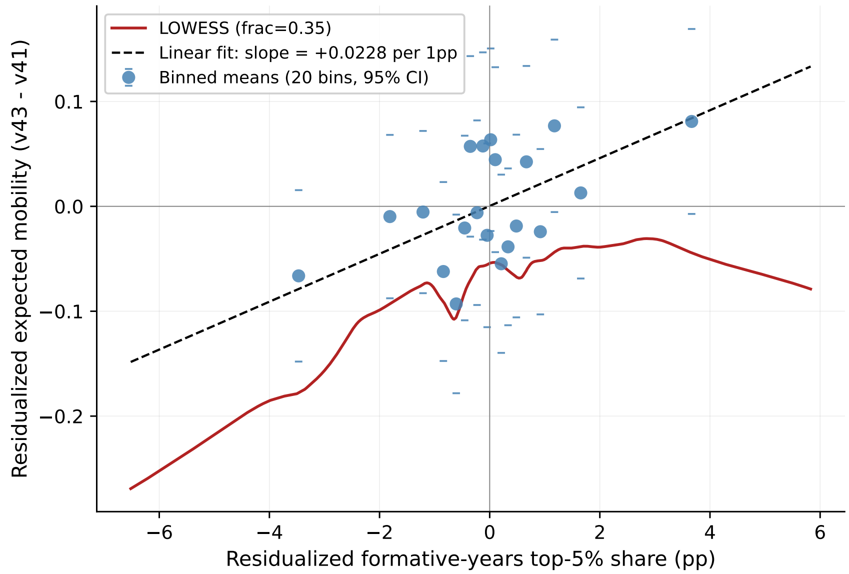

Notes: This figure shows a binned scatter plot of residualized expected mobility against residualized experienced inequality. Both variables are residualized on country fixed effects, subregion-by-cohort fixed effects, age, an indicator for female, and years of education. Blue points are means within 20 equal-sized bins of the residualized regressor. The red curve is a LOWESS smoother, and the dashed line is the linear fit on the residuals. The plot excludes observations in the top and bottom 1% of the residualized regressor.

Figure C.2: Experienced inequality and the probability of each expected rung

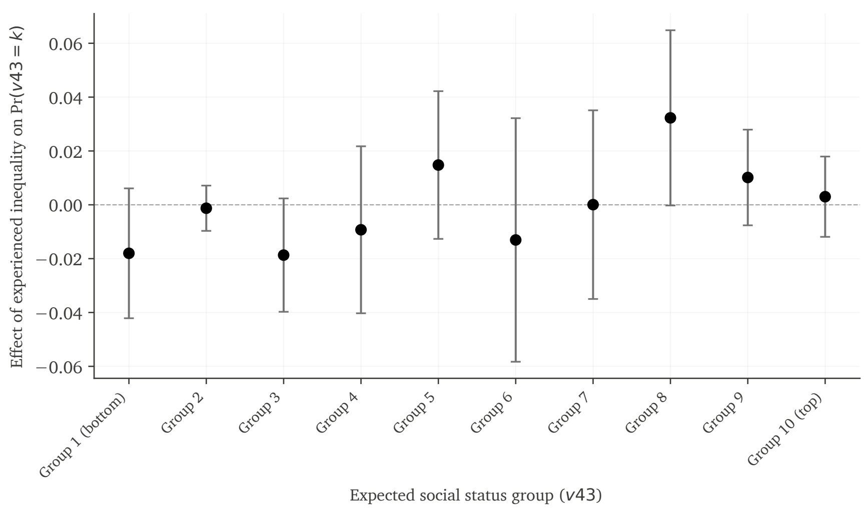

Notes: This figure shows coefficients on standardized experienced inequality, each from a separate linear-probability regression of an indicator for expecting rung k. Each regression includes the individual controls (age, an indicator for female, and years of education), current status (v_{41}), the experienced unemployment rate, and country and birth-cohort fixed effects. Error bars show 95% confidence intervals based on standard errors clustered at the country level.

Table C.3: Experienced inequality and expected mobility: all coefficients

	Dependent variable: Expected mobility			
	(1)	(2)	(3)	(4)
Experienced inequality	0.466***	0.335***	0.288***	0.252**
	(0.135)	(0.113)	(0.104)	(0.111)
Age	—	-0.035***	-0.047***	-0.044***
		(0.003)	(0.007)	(0.007)
Female	—	-0.014	-0.002	0.002
		(0.019)	(0.024)	(0.023)
Years of education	—	-0.033***	-0.012**	-0.005
		(0.012)	(0.005)	(0.006)
Family-of-origin status	—	—	—	-0.044***
				(0.009)
Experienced unemployment rate	—	—	—	0.049
				(0.036)
Experienced GDP per capita growth	—	—	—	0.046*
				(0.024)
Experienced CPI inflation	—	—	—	-0.014
				(0.023)
Experienced electoral democracy	—	—	—	-0.079
				(0.048)
Experienced banking/currency crisis	—	—	—	-0.005
				(0.014)
Experienced rare disaster	—	—	—	0.015
				(0.017)
Country FE			Yes	Yes
Subregion×cohort FE			Yes	Yes
N	22,558	21,879	21,879	19,152

Notes: This table reports the coefficient and standard error on every regressor other than the fixed effects in each column of Table C.2. The dependent variable is expected mobility, the respondent's expected status in 10 years minus their current status, each on the 1 to 10 social ladder. Experienced inequality is the average top-5% income share in the respondent's country during the formative years (ages 18 to 25), standardized within the sample. The six macroeconomic experiences in Column 4 are constructed the same way, averaged over the same formative years and then standardized. The sample includes ISSP 2019 respondents aged 29 or older. Standard errors are clustered at the country level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Figure C.3: Mapping current status to within-country household income

(a) Mean log(income / country mean) **(b) Share above country-mean income**

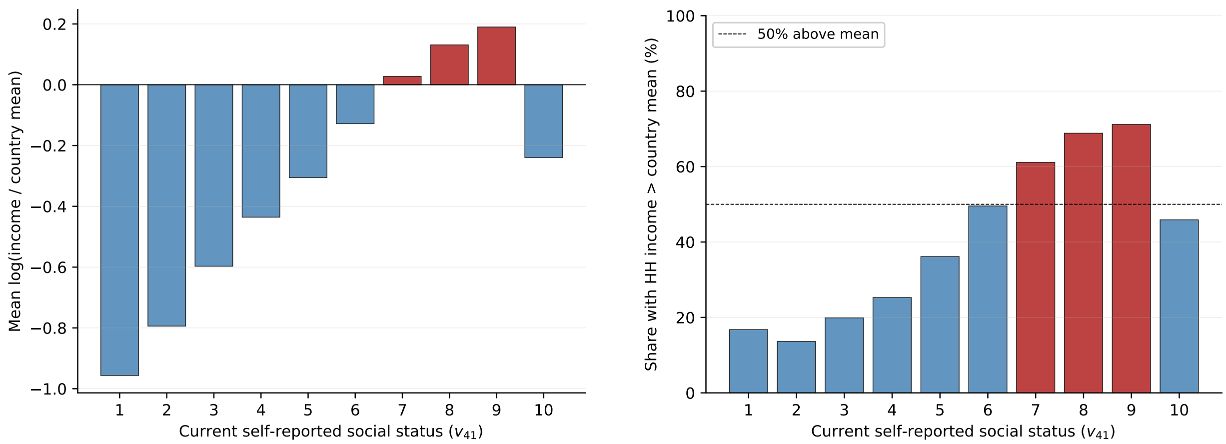

Notes: This figure shows the within-country household income position at each rung of current status v_{41}. Panel (a) shows the mean of log(income/country mean) within each rung, where zero marks the country mean. Panel (b) shows the share of respondents at each rung whose household income exceeds the country mean. Rungs 1 to 7 sit at or below the country mean, and rungs 8 and above sit clearly above it.

Figure C.4: Experienced inequality and the probability of an above-mean versus at-or-below-mean expected position

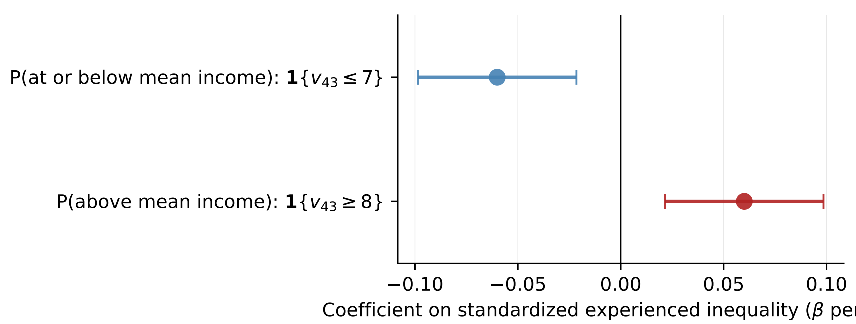

Notes: This figure shows the coefficients on standardized experienced inequality from separate linear-probability regressions for an above-country-mean expected position ($v_{43} \geq 8$) and an at-or-below-mean position ($v_{43} \leq 7$). The two coefficients are equal and opposite by construction. Both regressions include the controls and fixed effects of Specification (3), plus current status (v_{41}) and the experienced unemployment rate. Error bars show 95% confidence intervals based on standard errors clustered at the country level.

Figure C.5: Experienced inequality and expected mobility: heterogeneity by age, income, and status

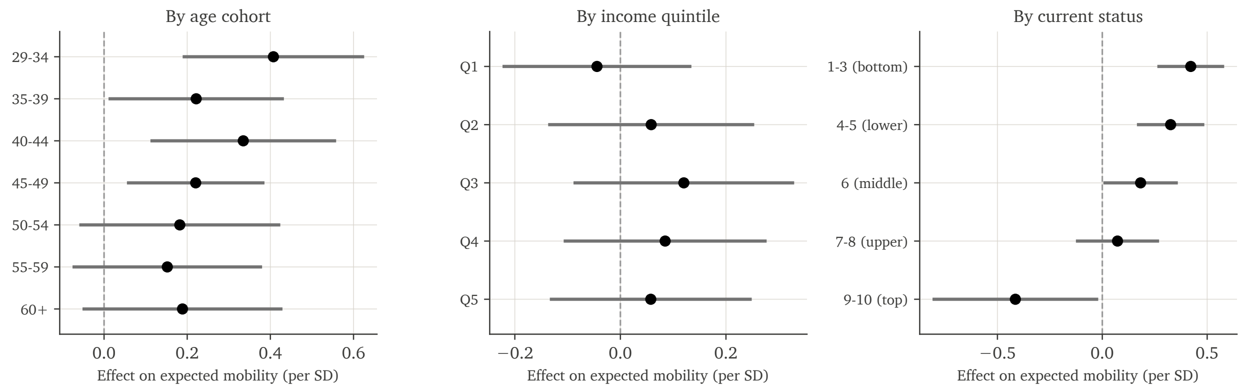

Notes: This figure shows the within-bin coefficient on standardized experienced inequality (per standard deviation) on expected mobility ($v_{43} - v_{41}$), from a full-sample interaction model that lets the slope vary across the bins of each dimension while keeping the fixed effects and individual controls of Specification (3). The left panel shows seven age groups, the middle panel within-country household-income quintiles, and the right panel current social status (v_{41}). The top-status bin in the right panel is small ($N = 673$). The sample includes ISSP 2019 respondents aged 29 or older. Error bars show 95% confidence intervals based on standard errors clustered at the country level.

Figure C.6: Experienced inequality and expected mobility under alternative top income shares

(a) Without the six macroeconomic experiences

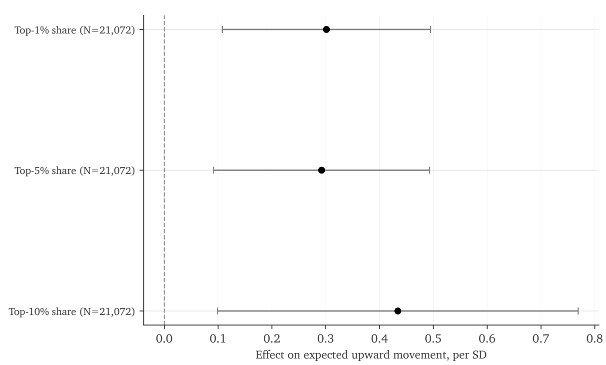

(b) With the six macroeconomic experiences

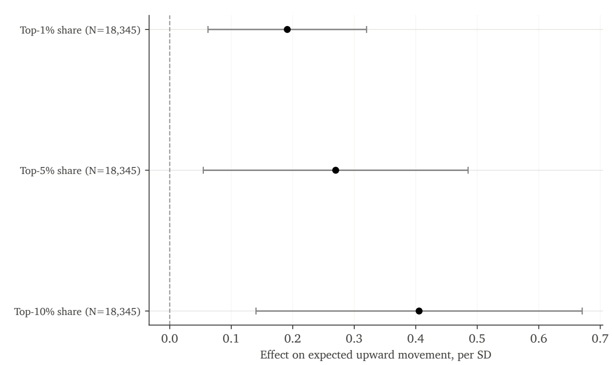

Notes: This figure shows the coefficient on the top-1%, top-5%, or top-10% share averaged over the formative years and standardized. Each coefficient is estimated separately on a common sample with the controls and fixed effects of Specification (3). Panel (b) also controls for the six macroeconomic experiences. Error bars show 95% confidence intervals based on standard errors clustered at the country level.

Figure C.7: Experienced inequality and expected mobility across alternative experience windows

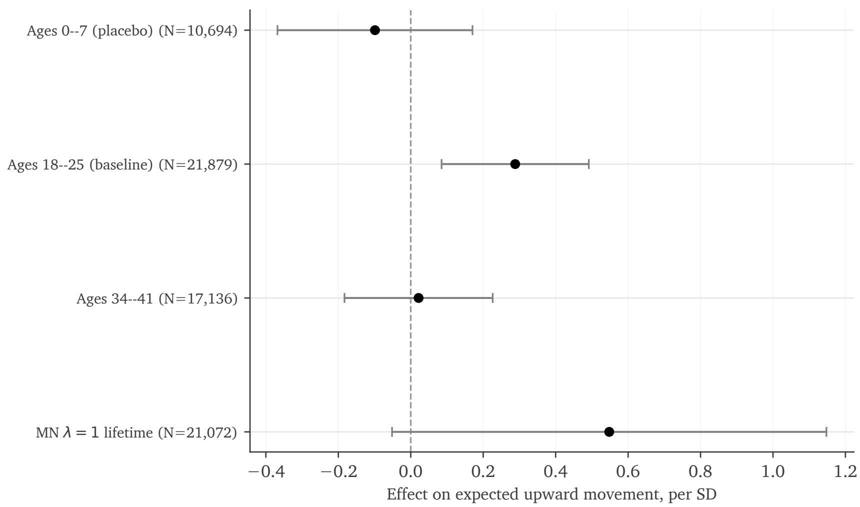

Notes: This figure shows the coefficient on experienced inequality for alternative experience windows. Each row is a separate regression with the controls and fixed effects of Specification (3) but a different definition of the experienced-inequality regressor: the ages 18 to 25 window, a pre-formative placebo (ages 0 to 7), a post-formative window (ages 34 to 41), and a Malmendier–Nagel lifetime-weighted average. All regressors are standardized. Sample sizes differ across rows because each window regressor is available for a different set of respondents. Error bars show 95% confidence intervals based on standard errors clustered at the country level.

Figure C.8: Distribution of expected social status, more equal versus more unequal countries

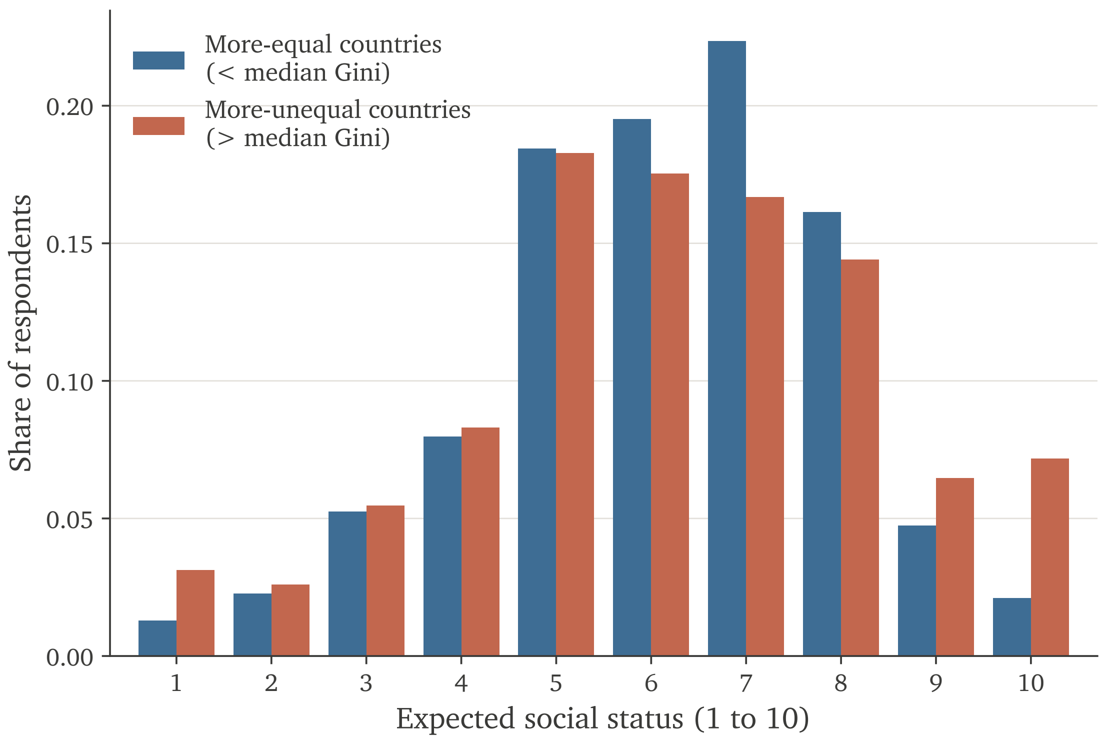

Notes: This figure shows the distribution of expected social status in 10 years (1 to 10) for countries with a Gini coefficient below the median (blue) and above it (red).

Figure C.9: Distribution of current social status, more equal versus more unequal countries

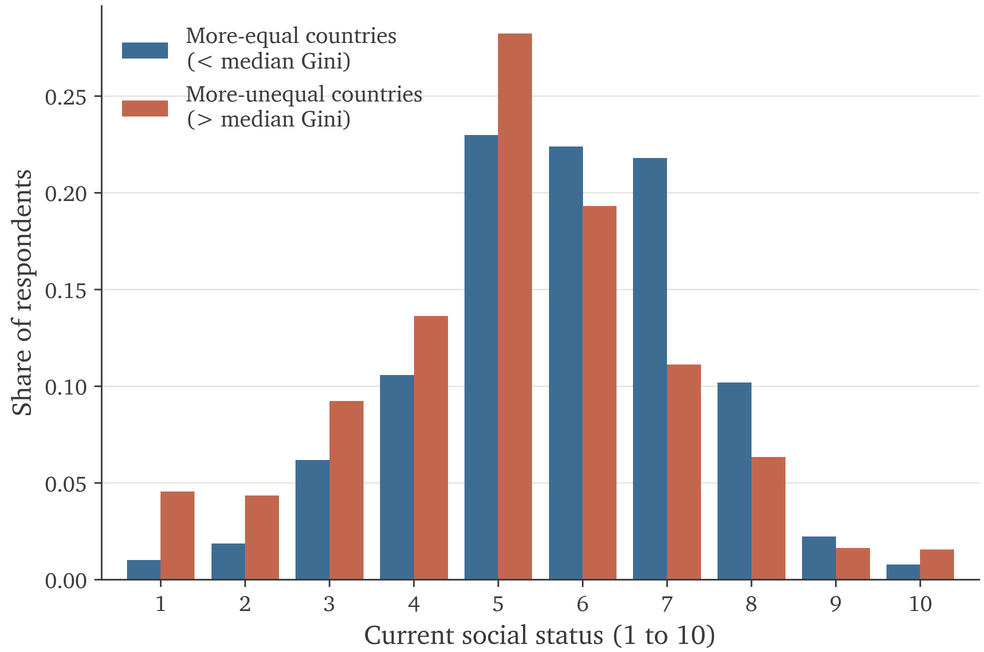

Notes: This figure shows the distribution of current social status (1 to 10) for countries with a Gini coefficient below the median (blue) and above it (red).

Table C.4: Logistic regressions of social status categories, more equal versus more unequal countries

	(1) Expected, rungs 1 to 6	(2) Expected, rungs 7 to 8	(3) Expected, rungs 9 to 10	(4) Current, rungs 1 to 5	(5) Current, rungs 6 to 8	(6) Current, rungs 9 to 10
Gini above the median	0.062**	−0.277***	0.551***	0.603***	−0.603***	−0.047
	(0.021)	(0.023)	(0.038)	(0.021)	(0.021)	(0.061)
Constant	0.361***	−0.651***	−2.618***	−0.197***	0.072***	−3.437***
	(0.015)	(0.015)	(0.029)	(0.015)	(0.015)	(0.042)
Num.Obs.	36 038	36 038	36 038	36 038	36 038	36 038
Log.Lik.	−24 299.094	−22 352.122	−10 743.910	−24 539.588	−24 375.749	−4917.472
F	8.422	147.027	213.611	796.004	788.258	0.595

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001
This table reports logistic regressions of indicators for expected and current social status categories on an indicator for more unequal countries (Gini above the median). Columns 1 to 3 use expected status and Columns 4 to 6 use current status, and the column headers give the rungs of the social ladder in each category. The sample includes ISSP 2019 respondents of all ages in the 27 countries. Standard errors in parentheses.

C.3 Expected status and happiness (General Social Survey)

Table C.5 shows that higher expected status is associated with higher general happiness, controlling for current status (Column 1). Columns 2 to 7 and Figure C.11 split the sample by current status and by income group. The association is significant at the 1% level for respondents with lower current status and for low-income respondents, and none of the other four subgroup coefficients on expected status is significant at the 10% level.

Table C.5: Expected status and general happiness, by current status (columns 2 to 4) and by income group (columns 5 to 7)

	(1) all	(2) lower	(3) middle	(4) upper	(5) low inc.	(6) middle inc.	(7) high inc.
Expected status	0.034**	0.053**	0.019	0.023	0.054***	0.001	0.026
	(0.011)	(0.018)	(0.019)	(0.022)	(0.016)	(0.020)	(0.026)
Current status	0.020	0.008		0.013	0.006	0.034	0.025
	(0.013)	(0.030)		(0.038)	(0.018)	(0.024)	(0.030)
Education	0.002	0.002	0.002	0.001	0.006	−0.015	0.012
	(0.007)	(0.013)	(0.011)	(0.012)	(0.011)	(0.012)	(0.013)
Income	0.013***	0.009	0.013*	0.018**	0.006	0.005	0.020
	(0.003)	(0.006)	(0.005)	(0.006)	(0.006)	(0.026)	(0.018)
Gender	0.013	0.043	0.025	−0.024	0.054	0.034	−0.039
	(0.034)	(0.066)	(0.053)	(0.059)	(0.061)	(0.057)	(0.057)
Constant	1.304***	1.266***	1.507***	1.414***	1.220***	1.778***	1.093*
	(0.120)	(0.225)	(0.213)	(0.349)	(0.209)	(0.532)	(0.452)
Num.Obs.	1596	440	629	527	513	540	543
R2	0.041	0.036	0.014	0.027	0.040	0.007	0.021

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001

This table reports OLS regressions of general happiness on expected status (rank in 10 years) for all respondents (Column 1), by current status (Columns 2 to 4), and by income group (Columns 5 to 7). General happiness is measured on a 3-point scale: 1 (not too happy), 2 (pretty happy), and 3 (very happy). The controls are current status, education, income, and gender. On the 10-rung ladder (10 = top), the current-status groups are rungs 1 to 5 in Column 2, rung 6 in Column 3, and rungs 7 to 10 in Column 4. The income groups have total family income below $40,000 in Column 5, $40,000 to $89,999 in Column 6, and $90,000 or more in Column 7. Column 3 has no current-status coefficient because every respondent in that group is on rung 6. The sample includes US respondents to the 2021 General Social Survey. Standard errors in parentheses.

Figure C.10: Distribution of experienced inequality by country

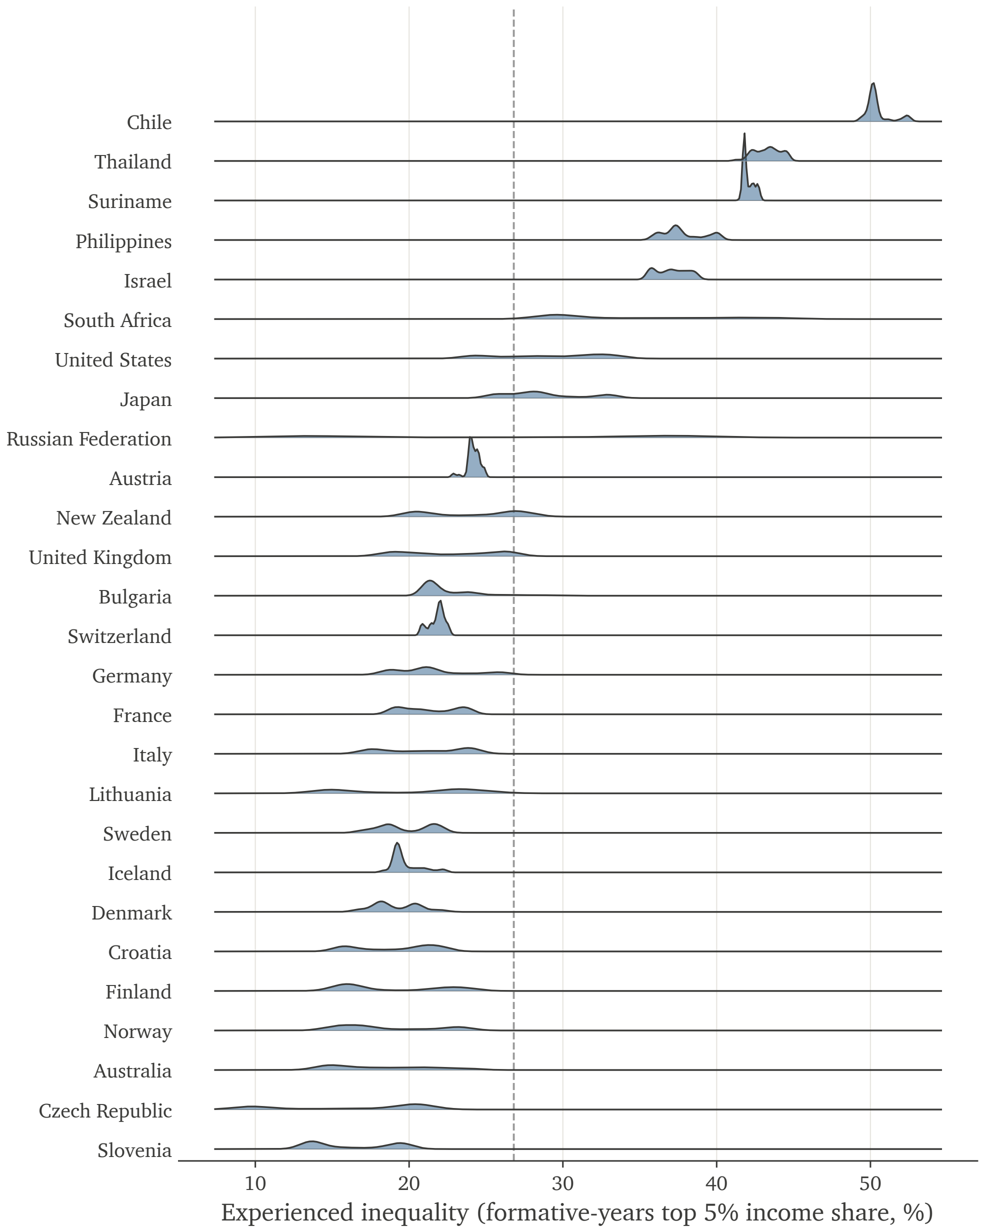

Notes: This figure shows the distribution of experienced inequality by country, constructed as in Roth and Wohlfart (2018): the average top-5% income share in respondent i's country during i's formative years (ages 18 to 25).

Figure C.11: Expected status and general happiness

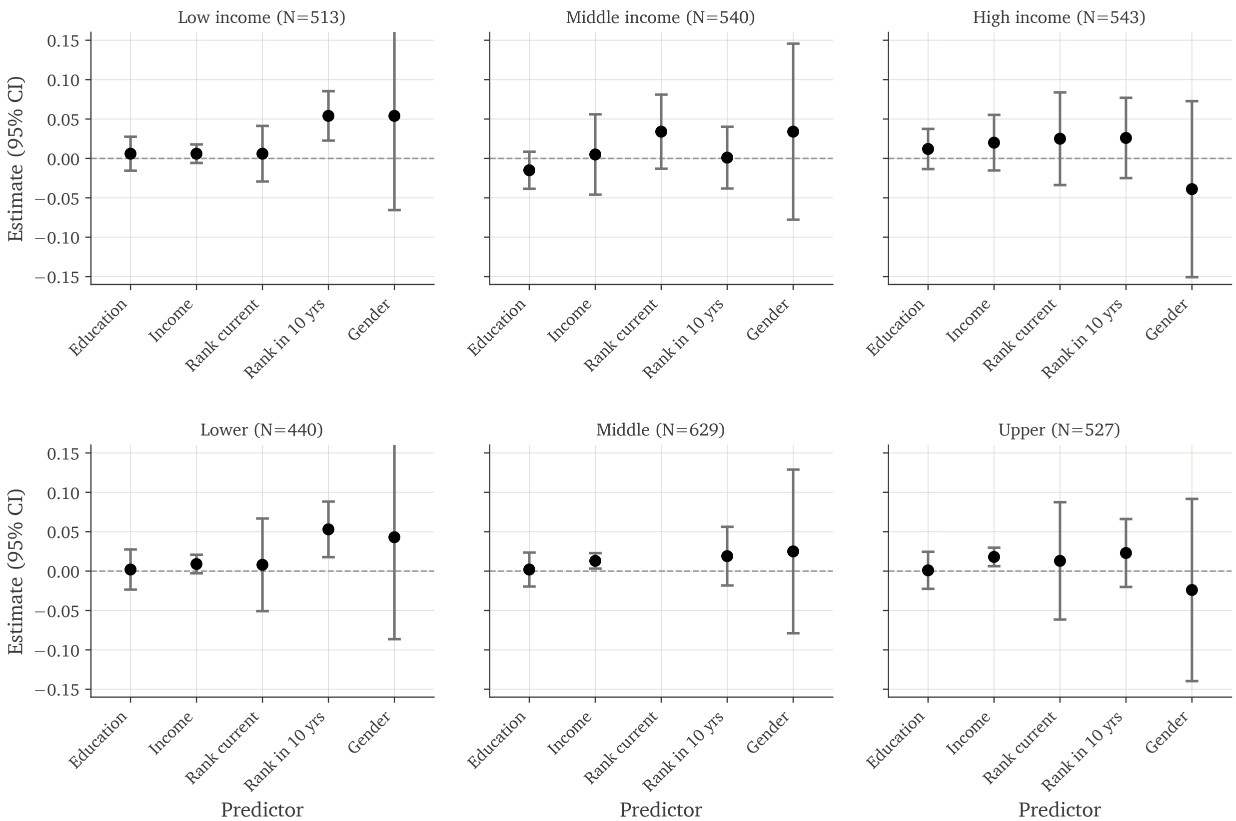

Notes: This figure shows coefficients from OLS regressions of general happiness (1 not too happy, 2 pretty happy, 3 very happy) on expected status ("Rank in 10 yrs"), current status ("Rank current"), education, income, and gender. The regressions are estimated separately by income group (top row) and by current status (bottom row). The middle current-status panel has no current-status coefficient because all respondents in that group place themselves at the same point of the 10-point status scale. Error bars show 95% confidence intervals. The sample includes US respondents to the 2021 General Social Survey.

D Additional details on the experimental design

Variations of the Baseline treatment. The Baseline treatment has two variations based on when subjects complete the quiz. In our conceptual framework, subjects learn about the income allocation choice only after their objective probabilities of achieving each rank are fixed, and they take these probabilities as given when choosing the allocation. The first variation of the Baseline treatment, which we refer to as the *Quiz-First* variation, closely follows this framework. In this variation, subjects complete two example quiz questions and all 10 main quiz questions *before* the belief elicitation in Period 1 and the income allocation choice in Period 2 (see Figure 4).

In the second variation, which we refer to as the *Quiz-Last* variation, we test whether our results are robust to the timing of the effort choice. Here, subjects first complete two example quiz questions and state their beliefs, then make the Period 2 income allocation choice, and finally complete the 10 quiz questions (see Figure D.1). This allows subjects to jointly determine their preferred inequality and effort level when making the allocation choice.

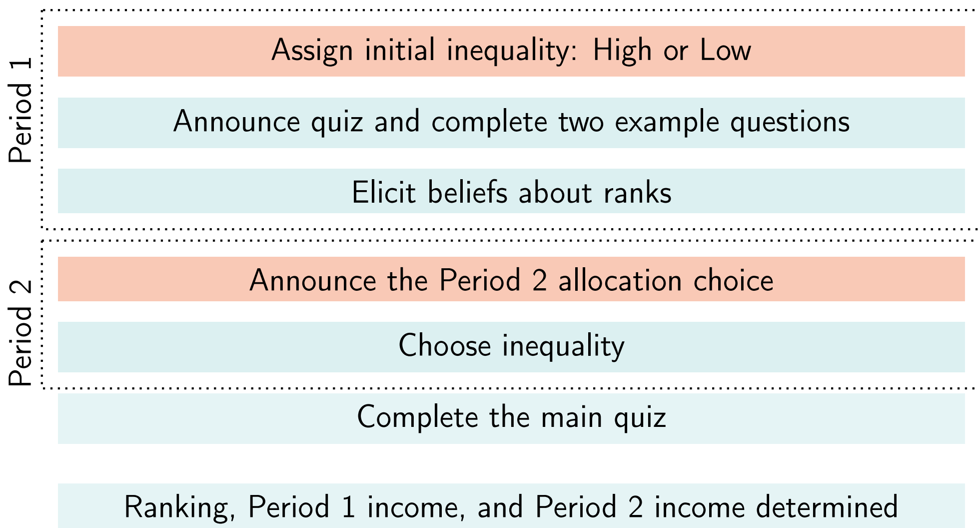

Figure D.1: Experimental design: Baseline treatment, Quiz-Last variation

In our main analysis, we pool both variations of the Baseline treatment as specified in our pre-analysis plan. As a robustness check, we replicate Table A.4 while also controlling for whether a subject was in the the Quiz-Last treatment in Table D.1. All columns confirm that the treatment effect of the High inequality treatment is robust to the inclusion of this indicator variable, which indicates that it is not driven by one of the two treatment variations. Column (1) shows that the effect on Rank-1 belief is similar to the coefficient in the Column (1) of the main specification (5.521 compared to 6.865). The indicator for the treatment variation is significant with Rank-1 belief as the dependent variable, implying that overall beliefs are lower when the full quiz is completed after the belief elicitation. The treatment variation did not have an impact on Share to Rich, see Column (2)-(4).

Table D.1: Treatment effects on inequality acceptance and Rank-1 belief - Quiz-Last vs. Quiz-First

	(1) Rank-1 belief	(2) Share to Rich	(3) Share to Rich	(4) Share to Rich
High Period 1 inequality	5.521*		7.272**	5.720*
	(2.568)		(2.441)	(2.340)
Quiz-Last	−6.825**	−0.400	−0.968	0.951
	(2.485)	(2.963)	(2.362)	(2.269)
Rank-1 belief x Quiz-Last		0.029		
		(0.065)		
Rank-1 belief		0.282***		0.281***
		(0.046)		(0.033)
Constant	37.943***	33.355***	41.286***	30.618***
	(1.783)	(2.194)	(1.696)	(2.043)
Num.Obs.	773	773	773	773
R2	0.032	0.097	0.024	0.109

+ p <0.1, * p <0.05, ** p <0.01, *** p <0.001
OLS regression estimates of treatment effects on inequality acceptance (share allocated to Rich) and Rank-1 belief. Column (1) reports coefficients for Rank-1 belief; column (2)-(4) report coefficients for the share allocated to the rich. Standard errors are reported in parentheses. Sample size and R^2 are shown for each specification.

E Proofs

E.1 Proof of median $P(Rank2)$ result

To quantify what improved performance would predict, we use the empirical rank probabilities for each quiz score from our data (see Table A.3). Consider a subject who correctly answers 7 questions (the median score). Their objective probabilities are P(Rank 1) = 8.3% and P(Rank 2) = 22.9%. If this subject overestimated their performance in the High inequality condition, believing they would answer 8 questions correctly instead of 7, their subjective probabilities would become P(Rank 1) = 14.6% and P(Rank 2) = 31.0%. There is a substantial increase in the probability of achieving each of the top 2 ranks.

Aggregating across all subjects, if everyone believes their performance improves by one correct answer in the High inequality condition while everybody else's performance is constant on average, the average Rank-1 belief would increase by 7.3 percentage points and the average Rank-2 belief would increase by 4.5 percentage points. The median Rank-1 belief would increase by 6.3 percentage points (from 8.3% to 14.6%) and the median Rank-2 belief would increase by 8.1 percentage points (from 22.9% to 31.0%). Even if subjects believe their performance improves by two correct answers, we reach similar conclusions: the average Rank-2 belief would increase by 5.3 percentage points and the median Rank-2 belief would increase by 14.3 percentage points (from 22.9% to 37.2%). The prediction of increased Rank-2 beliefs is robust: the average would remain essentially flat only under the extreme assumption that most subjects believe their performance improves by three or more questions (which would require median performers to believe they scored 10/10), yet even then the median Rank-2 belief would still increase.

We can generalize this observation beyond the specific empirical distribution. Even without assuming subjects know the exact performance distribution, we can prove that, if subjects believe their own score increases while others' scores remain unchanged, the median Rank-2 belief increases under mild assumptions about the believed performance distribution:

Assumption 1: $F(s_{med} + t) < 3/4$, where s_{med} is the median score and t is the believed performance improvement.

This condition requires that the median performer, after believed improvement, remains below the 75th percentile of the original performance distribution.

Given Assumption 1, we can establish the following result: If all individuals believe their own score increases by t units while others' scores remain unchanged, then the median of the distribution of $P(\text{Rank 2})$ beliefs must increase. The intuition is straightforward: for someone to achieve Rank 2, they need exactly one opponent to outperform them and three to underperform. This probability follows an inverse U-shaped curve as a function of one's position in the distribution: it is low for both very weak performers (who rarely beat three opponents) and very strong performers (who usually beat all four). The probability peaks at the 75th percentile. Since the median performer sits at the 50th percentile, improving their believed performance moves them closer to this peak, increasing their P(Rank 2). As long as they do not move past the 75th percentile (Assumption 1), this relationship holds.

Each individual's probability of achieving rank k given score s_i is:

$$P(\text{Rank } k \mid s_i) = C_4^{k-1}[1 - F(s_i)]^{k-1}[F(s_i)]^{5-k}$$

For an individual with score s, the change in $P(\text{Rank 2})$ when believing their score increases from s to $s + t$ is:

$$\Delta P(\text{Rank 2} \mid s, t) = 4\left[[1 - F(s + t)][F(s + t)]^3 - [1 - F(s)][F(s)]^3\right]$$

Define $g(x) = x^3(1 - x)$. The derivative $g'(x) = 3x^2 - 4x^3 = x^2(3 - 4x)$ is positive for $x < 3/4$ and negative for $x > 3/4$, and $g(x)$ is therefore strictly increasing on $[0, 3/4]$ and strictly decreasing on $[3/4, 1]$. For the median performer with score s_{med}, we know that $F(s_{med}) \leq F(s_{med} + t) < 3/4$ by Assumption 1.

Since both $F(s_{med})$ and $F(s_{med} + t)$ lie in the region where g is strictly increasing, and since $F(s_{med} + t) \geq F(s_{med})$, we can conclude:

$$g(F(s_{med} + t)) \geq g(F(s_{med}))$$

Therefore $\Delta P(\text{Rank 2} \mid s_{med}, t) \geq 0$, meaning the median performer's $P(\text{Rank 2})$ belief weakly increases under performance optimism, and the increase becomes strict if we instead assume that $F(s_{med} + t) > F(s_{med})$. Since this individual represents the median of the distribution, the median of all $P(\text{Rank 2})$ beliefs must increase.

F Conceptual framework with an anticipated Period 2 allocation choice

The conceptual framework in Section 2 assumes subjects form their Period 1 beliefs without anticipating the Period 2 allocation choice. This matches the Baseline treatment, in which this choice is introduced only after beliefs are stated. In many real-world settings, however, individuals form beliefs knowing in advance that those beliefs will guide a consequential downstream decision, for example voters anticipating their electoral choices or investors anticipating their portfolio allocations. This appendix shows that the main results of the conceptual framework, strict overoptimism about Rank 1 and a positive inequality–overoptimism link, continue to hold when individuals fully anticipate the Period 2 allocation choice.

F.1 Setup

We retain all primitives of the conceptual framework in Section 2 and add a Period 2 allocation choice whose consequence individuals fully anticipate. There are n individuals, n ranked positions, one Rich position ($k = 1$) paying y_R, and $n - 1$ Poor positions each paying y_P. Inequality $\pi \in [1/n, 1]$ is the share of total Period 1 income Y going to Rank 1, which implies $y_R = \pi Y$ and $y_P = (1 - \pi)Y/(n - 1)$ and leaves the Period 1 mean income $\bar{y} = Y/n$ independent of π. Each individual i has objective probability p_{0ij} of position j with $\sum_j p_{0ij} = 1$.

Period 2. Each individual i chooses an allocation $s_i \in [0, Y_2]$, the dollar amount of a Period 2 budget Y_2 that goes to Rank 1, with the remainder $Y_2 - s_i$ split equally across the $n - 1$ Poor positions. Under random dictator, one of the n individuals is drawn uniformly and this individual's choice is implemented for the entire group. Rank 1 receives the implemented allocation $\bar{s}$ and each of Ranks $2, \ldots, n$ receives $(Y_2 - \bar{s})/(n - 1)$.

Payment structure. Individuals are paid their Period 1 income with probability $1 - \lambda$ and their Period 2 income with probability $\lambda \in (0, 1)$. The two periods are exclusive substitutes: introducing Period 2 reweights the payment lottery rather than adding to it.

Optimal Period 2 allocation. Conditional on subjective belief μ_{i1} about Rank 1, the privately optimal s_i is the Bayesian choice that maximizes expected Period 2 consumption utility under μ_{i1}. Under standard specifications, the optimum is

$$s_i^*(\mu_{i1}) \;=\; \mu_{i1}\, Y_2. \tag{4}$$

Well-being. Individuals choose $\boldsymbol{\mu}_i$ to maximize anticipatory utility net of decision cost. Without an anticipated allocation choice, well-being is given by Equation 1 of Section 2. With an anticipated allocation choice, the well-being function incorporates Period 2 anticipatory utility and uses a decision-cost coefficient $\alpha^A > 0$:

$$W_i^A(\boldsymbol{\mu}_i) \;=\; \sum_{j=1}^{n} \mu_{ij}\left[(1-\lambda)\,y_j + \lambda\,y_j^{(2)}\right] \;-\; \alpha^A \sum_{j=1}^{n}\bigl(\mu_{ij} - p_{0ij}\bigr)^2 \quad \text{s.t.} \;\sum_{j=1}^{n} \mu_{ij} = 1, \tag{5}$$

where $y_1^{(2)} = \bar{s}$ and $y_j^{(2)} = (Y_2 - \bar{s})/(n-1)$ for $j > 1$ denote Period 2 income for each rank.[19]

F.2 Optimal belief with an anticipated allocation choice

Applying the Lagrangian to Equation 5 as in Section 2, with each individual taking the implemented allocation $\bar{s}$ as given, the first-order condition for each position j is

$$(1-\lambda)\,y_j + \lambda\,y_j^{(2)} - 2\alpha^A\bigl(\mu_{ij} - p_{0ij}\bigr) \;=\; \nu,$$

where ν is the multiplier on the adding-up constraint (denoted λ in Section 2). Summing over j and using $\sum_j \mu_{ij} = \sum_j p_{0ij} = 1$, together with $\sum_j y_j = n\bar{y}$ and $\sum_j y_j^{(2)} = Y_2$, yields $\nu = (1-\lambda)\bar{y} + \lambda Y_2/n$. The optimal distortion for position j is

$$\mu_{ij}^* - p_{0ij} \;=\; \frac{(1-\lambda)\bigl(y_j - \bar{y}\bigr) + \lambda\bigl(y_j^{(2)} - Y_2/n\bigr)}{2\alpha^A}. \tag{6}$$

This generalizes the distortion in Equation 3 of Section 2. In the numerator, the first term is the Period 1 income premium relative to the mean, reweighted by the probability that Period 1 is the paid period, and the second term is the analogous Period 2 income premium.

Symmetric equilibrium. The Period 2 income premium depends on the implemented allocation $\bar{s}$, which itself depends on the population's beliefs. Under random dictator, $\bar{s}$ equals a uniformly drawn individual's choice. In the symmetric equilibrium where all individuals share Rank-1 belief μ_R, this gives $\bar{s} = \mu_R Y_2$. Substituting into Equation 6 for Rank 1 and using $y_1^{(2)} = \bar{s} = \mu_R Y_2$ together with the uniform objective probability $p_{0i1} = 1/n$:

$$\mu_R^* - \tfrac{1}{n} \;=\; \frac{(1-\lambda)(y_R - \bar{y}) + \lambda Y_2\bigl(\mu_R^* - 1/n\bigr)}{2\alpha^A}. \tag{7}$$

Solving the fixed point:

$$\mu_R^* - \tfrac{1}{n} \;=\; \frac{(1-\lambda)\bigl(y_R - \bar{y}\bigr)}{2\alpha^A - \lambda Y_2}. \tag{8}$$

For the symmetric equilibrium to be interior, the denominator must be positive: $2\alpha^A > \lambda Y_2$.

[19] The decision-cost coefficient α^A in the specification with an anticipated allocation choice may differ from its value in Section 2 because individuals internalize the realized utility loss from optimizing the allocation under a distorted belief. The robustness results below assume $2\alpha^A > \lambda Y_2$.

F.3 Persistence of overoptimism and the inequality–overoptimism link

The main results of the conceptual framework continue to hold with an anticipated allocation choice.

Proposition 1 (Persistence of overoptimism)**.** *For any $\lambda \in (0,1)$ and $y_R > \bar{y}$, the symmetric equilibrium with an anticipated allocation choice has $\mu_R^* > 1/n$. Individuals remain strictly overoptimistic about being Rank 1.*

Proof. The numerator of Equation 8, $(1-\lambda)(y_R-\bar{y})$, is strictly positive, and the denominator is positive by the interior-equilibrium condition. Therefore $\mu_R^* - 1/n > 0$. The underlying logic is identical to that of Section 2: at $\mu_R = 1/n$, the marginal anticipatory benefit of raising μ_R is $(1-\lambda)(y_R-\bar{y}) > 0$ (first order), while the marginal decision cost is $2\alpha^A(\mu_R - 1/n) = 0$ (second order). The asymmetry from Brunnermeier and Parker (2005) that drives the result in Section 2 is unchanged. □

Proposition 2 (Inequality–overoptimism link)**.** *Holding total Period 1 income Y and Period 2 parameters fixed,*

$$\frac{\partial(\mu_R^* - 1/n)}{\partial \pi} = \frac{(1-\lambda)\,Y}{2\alpha^A - \lambda Y_2} > 0,$$

so higher Period 1 inequality raises the equilibrium Rank-1 belief. The implemented Period 2 allocation $\bar{s} = \mu_R^* Y_2$ also rises in π: higher Period 1 inequality propagates through the belief channel to more unequal Period 2 incomes.

Proof. Substituting $y_R - \bar{y} = (\pi - 1/n)Y$ into Equation 8 and differentiating with respect to π gives the first claim, and $\bar{s} = \mu_R^* Y_2$ in symmetric equilibrium gives the second. □

Both main results of Section 2 therefore continue to hold. Individuals exhibit strict overoptimism about Rank 1, and higher Period 1 inequality raises both the equilibrium belief and the implemented Period 2 allocation. These comparative statics in π are qualitatively the same whether subjects form beliefs without anticipating the Period 2 allocation choice, as in the Baseline treatment, or while fully anticipating it.