Historical Legacies and Preferences for Inter-Regional Redistribution*

August 16, 2026

Abstract

Regional inequalities are large, persistent, and politically divisive. Yet public support for inter-regional redistribution remains limited, and the beliefs that shape it are poorly understood. We test whether informing citizens about the historical roots of regional disparities can shift redistribution preferences. In a pre-registered online experiment with 8,000 participants across Brazil, Spain, and the United States, we first provide novel evidence on individuals' beliefs about regional inequality, its persistence, and its causes, using both quantitative measures and open-ended responses analyzed with Large Language Models. We then randomly assign participants to an information treatment summarizing peer-reviewed economic history research on the deep roots of regional inequality. We assess the treatment's effect on self-reported and high-stakes redistributive measures, tracing the chain from belief updating to shifts in responsibility attribution and emotional reactions. Crucially, since historical narratives may trigger fatalism and the belief that nothing can be changed, we implement a second treatment arm that emphasizes the role of public policy in overcoming historical legacies. We further assess how various dimensions may influence belief updating and the causal chain from beliefs to policy support. Our study sheds light on how historical narratives can shape attitudes toward redistribution and how academic research can be communicated without generating backlash.

Keywords: Inequality, Persistence, Redistribution, History, Narratives, Experiment.

JEL Classification: D63; D72; H23; N10; N90.

Study pre-registration: Pre-registration will be completed before the commencement of data collection.

*We thank the Editor and three anonymous reviewers for their constructive comments. We also thank for feedback participants of the LISER Working Group and the online workshop Junior Applied Microeconomists in Europe (JAMiE). All errors are our own.

Proposed timeline:

We plan to complete the research project within 8 months after pre-results review acceptance. The project was approved by the University Pablo de Olavide Ethical Review (Reference 25/6-74) in October 2025. The following timeline highlights the main milestones of the project:

Timeline of the project execution

Tentative Date	**Milestone**
May 2026	Submission pre-results review
	↕ *6 months*
Nov 2026	JDE Stage 1 acceptance
Nov 2026	Pre-registration at AEA RCT Registry
	↕ *3 months*
Feb 2027	Survey 1 (Intervention)
	↕ *2 months*
Apr 2027	Survey 2 (post-treatment)
	↕ *3 months*
Jul 2027	Paper completed and final submission

1 Introduction

Large and persistent differences in economic prosperity across regions are a defining feature of modern economies. In developed countries, GDP per capita in the richest regions (top decile) exceeds that of the poorest by 70% on average; in developing countries, this gap reaches 200% (Floerkemeier et al., 2021). These inequalities have proven highly persistent. In countries as diverse as the United States, Brazil, and Spain, the correlation between regional GDP per capita in 1900 and 2020 exceeds 0.65. Regional inequalities can erode social cohesion (International Monetary Fund, 2019; Wuthnow, 2019), fuel political polarization and populism (Autor et al., 2020; Rodríguez-Pose et al., 2023), and shape the geographic sorting of voters (Rodden, 2019). Despite their scale and consequences, however, public support for inter-regional redistribution remains limited and highly polarized, and the factors that shape citizens' willingness to address territorial disparities are still poorly understood (Beramendi, 2012; Balcells et al., 2015).

This paper provides novel evidence on perceptions of regional inequality and tests whether exposure to historical narratives can shift redistribution preferences. A large literature links beliefs about the origins of income differences to preferences for interpersonal redistribution, and in particular to whether inequalities are attributed to effort or to luck and circumstances (Piketty, 1995; Bénabou and Tirole, 2006; Alesina et al., 2018). Beliefs about why someone is poor imply an attribution of responsibility, which in turn shapes judgements of deservingness and support for transfers. This sequence extends to the territorial domain with the crucial distinction that a region is an aggregate rather than a single agent, so that it is far less obvious whose conduct is being weighed (e.g., its current inhabitants, regional governments or elites, past generations or historical events, impersonal forces, etc.). This indeterminacy gives causal narratives room to shift attribution. The *persistence* literature in economic history has documented that a significant share of regional income variation traces back to historical processes rooted far in the past (Spolaore and Wacziarg, 2013; Nunn, 2020). Informing citizens of this evidence may lead them to view regional poverty as reflecting inherited conditions rather than the behavior or choices of current inhabitants, strengthening the moral case for transfers. Yet the same information could backfire. If regional disparities are historically determined and highly persistent, citizens may conclude they are too entrenched to be addressed by policy, thereby fostering fatalism that paradoxically reduces support for redistribution. Doubts about whether transfers reach their intended beneficiaries, together with the long horizon over which results become visible, could reinforce this conclusion. Understanding this tension between the solidarity-enhancing and fatalism-inducing potential of historical narratives is critical for policy design and science communication.

We conduct a pre-registered online experiment with 8,000 participants from the United States, Spain, and Brazil, to causally identify how information about the historical roots of regional inequality affects redistribution preferences. We combine rich baseline measurement of beliefs, randomized information provision, and incentivized behavioral outcomes. The experimental design has three distinctive features. First, rather than providing statistics on inequality levels or income ranks, we provide engaging narratives about the ultimate causes of regional disparities, drawing directly from peer-reviewed economic history research. Specifically, we produce professional 2:30 minutes animated videos summarizing findings on the historical

origins of regional inequality in each country as well as a neutral placebo video for the control group. Second, to directly test the fatalism hypothesis, we implement a two-tiered treatment. Some participants watch the historical video alone (*T1*) while others the same video plus a one-minute extension emphasizing that "history is not destiny" and that well-designed public policy can overcome historical legacies (*T2*). Third, we measure redistribution preferences through multiple instruments including self-reported attitudes, incentivized lottery decisions with high monetary values, and a real-stakes petition to the government.

The treatment effect is heterogeneous by construction, and this shapes how the results should be interpreted. Because the intervention works by moving a specific belief, its effect should be greater among participants who, before treatment, attributed little importance to distant historical events in explaining current regional inequalities. We elicit such beliefs cardinally before and after the treatments, which allows us to investigate the heterogeneity by initial priors and within-individual belief updating.

We also disentangle the mechanisms through which historical narratives operate. We assess the extent to which individuals shift responsibility attribution away from current residents, whether they report increased empathy and moral obligation, and whether they exhibit fatalism about future convergence. It is particularly on this dimension of fatalism and policy efficacy where the two treatments can operate quite distinctively. Importantly, we analyze open-ended survey responses using Large Language Models (LLMs) to classify the causal narratives respondents invoke when explaining regional inequality and justifying their redistribution preferences. This allows us to observe how people reason about these questions without imposing predefined categories. We complement this evidence with measures on perceptions on historically regional inequality and its persistence since 1900.

Our pre-registered hypotheses predict that exposure to historical information (*T1*) will increase support for regional redistribution relative to control, operating through two channels: (i) decreased attribution of responsibility to current residents of poor regions and more to historical events, and (ii) increased empathy and moral obligation to redistribute resources. However, we also hypothesize that *T1* may increase fatalism about future regional convergence which the second treatment (*T2*) would partially attenuate (see Figure 1). The persistence of the treatment effects (*T1* and *T2*) on both redistribution preferences and each of the proposed channels is evaluated two months later through an obfuscated follow-up survey.

We further expect the treatment effect to be highly heterogeneous beyond the initial priors and distinguish two key stages. First, we investigate what could boost belief updating per se, examining the role of the credibility of the information and the degree to which it can be assimilated by the individual (in terms of complexity and compatibility with their current mental framework of inequality). The second stage concerns the step between beliefs and changes in preferences. Here we examine the moderating role of how personal the content feels, individuals' predisposition to solidarity, the perceived personal cost of redistribution, and trust that redistribution will be effective. We complement these analyses with a machine-learning approach (Chernozhukov et al., 2025), which considers the full set of pre-treatment covariates jointly and can uncover unexpected sources of heterogeneity. These heterogeneity explorations provide additional evidence to our theoretical mechanisms and identify which

populations are more responsive to historical information.

This paper contributes to four strands of the literature. First, we contribute to the literature on redistribution preferences by shifting our attention to inter-regional redistribution.[1] The determinants of interpersonal redistributive preferences have been extensively studied; particularly the role of beliefs about effort versus circumstances (Piketty, 1995; Alesina and Angeletos, 2005; Fong, 2001; Alesina and La Ferrara, 2005). Experimental work shows that providing information about income inequality, social mobility, and the role of luck can shift these beliefs and subsequently alter preferences (Cruces et al., 2013; Karadja et al., 2017; Alesina et al., 2018; Hvidberg et al., 2023). We provide the first large-scale experimental evidence that research-based historical narratives about the origins of regional inequality may shift redistribution preferences, and examine the channels through which they operate. Moreover, we provide novel evidence on citizens' perceptions of historical regional inequality, long-term persistence, and the causal factors behind regional disparities — economic beliefs that remain largely unexplored in the literature.

Second, we contribute to the literature on narratives and the formation of beliefs and on how individuals process novel information (Shiller, 2020; Bénabou et al., 2018; Esposito et al., 2023; Graeber et al., 2024). To understand how citizens reason about economic policies, recent studies emphasize the need to capture the subjective mental models and causal accounts (or narratives) that individuals use to interpret the world (Faia et al., 2024; Stantcheva, 2021; Andre et al., 2024). We contribute to this agenda in two ways. First, we provide experimental evidence on the impact of historical narratives about inequality grounded on scientific research instead of statistics.[2] Second, from a methodological standpoint, we leverage open-ended survey questions analyzed via LLMs to elicit the unprompted narratives of respondents, advancing the quantitative analysis of unstructured economic narratives (Haaland et al., 2024; Andre et al., 2024).

Third, we contribute to the literature on historical legacies and long-term persistence. Research in this area has increased exponentially in the last two decades, emphasizing the role of deep structural factors and path dependence, and profoundly influencing how economists think about development (Spolaore and Wacziarg, 2013; Nunn, 2020; Voth, 2021; Cirone and Pepinsky, 2022).[3] The large body of evidence on long-lasting effects of historical factors together with the high degree of persistence observed in historical data may raise concerns about its implications for policy and on how this evidence could be interpreted by the public (Abad and Maurer, 2021; Conley and Kelly, 2025). Paradoxically, while recognizing the importance of history, individuals may be less willing to share resources if there is no scope for contemporary policy or believe that development paths are pre-determined. We directly explore this possibility by assessing whether information on historical legacies may induce fatalism, rais-

[1] Section 2 discusses some features that distinguish inter-regional from interpersonal redistribution.

[2] Our approach is connected to recent findings that historical accounts can influence contemporary political behavior and institutional satisfaction (Ramos-Toro and Voytas, 2024).

[3] Examples of seminal work include Becker and Woessmann (2009), who connect Protestantism to higher human capital through Bible-reading instruction; Dell (2010), who links forced labor in colonial Peru to modern-day local development; Nunn and Wantchekon (2011), who relate transatlantic slave trade to current mistrust in Africa; and Alesina et al. (2013), who document that suitability of land for plough use in early agriculture contributes to explain current gender norms. This literature has also produced highly influential popular-economics books (Acemoglu and Robinson, 2012; Galor, 2022) and Nobel Prize-winning research (Acemoglu, Johnson, and Robinson, 2024).

ing concerns about future convergence and policy efficacy. Moreover, we examine whether coupling narratives with content on policy agency changes how this information may affect preferences.

Finally, our paper will inform debates about place-based policies and fiscal federalism, where public support is essential for political sustainability (Kline and Moretti, 2014; Neumark and Simpson, 2015; Austin et al., 2018). If historical narratives increase solidarity, this suggests a communication strategy for policymakers and advocates. If they induce fatalism, this identifies a pitfall to avoid. By exploring in depth how different profiles of individuals respond to information, we reveal which populations are persuadable and which require different approaches. This connects to broader questions about the political economy of divided societies and the conditions for solidarity across groups (Alesina and Glaeser, 2004; Shayo, 2009; Chinoy et al., 2026). In sum, our study speaks to how the persistence literature, one of the most influential areas in economic history, can be communicated to the public without inadvertently undermining support for the very policies it might inform.

2 Framework and Related Literature

The way people form their political preferences is highly complex and shaped by their social background, their life experiences, and the information they are exposed to. In the case of preferences for redistribution, the causal model of why the inequality exists to begin with and whether a recipient of a transfer deserves it or not is crucial for policy support. When considering regional inequality, this analogy travels through but becomes much more nuanced: a region is not a person who can bear responsibility the way individuals do, and it also presents several particularities that make it distinct. This section discusses the parallels with the literature on interpersonal redistribution and the factors that may distinguish the regional case.

As mentioned, how people reason about the causes of inequality influences whether support is justified. For instance, societies that attribute income differences to effort and merit show less support for redistribution than those that attribute them to forces beyond the control of a person, such as their family background or structural inequalities (Piketty, 1995; Alesina and Angeletos, 2005; Bénabou and Tirole, 2006; Alesina and La Ferrara, 2005). Beliefs are not entirely fixed and sometimes can be shaped by information. A large experimental literature shows that information about one's position in the income distribution, about rates of intergenerational mobility, or about the weight of circumstances relative to choices can shape both beliefs and the policy preferences that rest on them (Cruces et al., 2013; Karadja et al., 2017; Alesina et al., 2018; Hvidberg et al., 2023; Stantcheva, 2021; Haaland et al., 2023). These concepts translate to the regional case insofar as beliefs about the causes of these inequalities (the mental frame a person relies on) feed into the *deservingness* of regional transfers and thus their support (Gniza and Wrede, 2024). Yet, there are a number of features that are more salient or unique in the regional case and that are worth highlighting. These features inform our research design described in the following sections.

1. An undetermined agent

The responsibility and deservingness of a region are very different from those of a person. When thinking about the dichotomy effort vs. luck of a person, there is a single agent whose agency can be judged or evaluated. When transposing this into a region, there is not a clear agent. One may think that the underlying cause of a region's relative poverty is linked to its current residents collectively, the regional government or elites, maybe past generations, or maybe even nobody at all (just structural forces like bad geography). Hence, there may be a lot of heterogeneity in the mental frame a person has, and this can even switch. This makes our information treatments particularly useful because individuals may not have a solidly defined mental framework, or there can be large variety in the way people respond. Also related to the undetermined agent is the dissociation between who or what is responsible for the economic backwardness and who should close the gap, which implies that these two judgments are involved in the beliefs for regional redistribution.

2. Collective goods and a collective agent

While the main objective of interpersonal redistribution is to increase private consumption to equalize social welfare, a primary goal of inter-regional transfers is to fund public services, infrastructure, and other local public goods. Resources are centralized, then transferred to an intermediary entity (a sub-national government) through often complex arrangements, and then reach the population (Oates, 1999; Boadway and Shah, 2009; Beramendi, 2012). This collective intermediary raises the concern that the resources could be subject to corruption, be captured by local elites, or be wasted before they reach their beneficiaries (Bardhan and Mookherjee, 2005; Daniele and Geys, 2015b; Brollo et al., 2013), which can mediate how individuals respond to information about inequalities (Kuziemko et al., 2015). Therefore, judging that the current inhabitants and the government are blameless about their income gap is not enough to support redistribution. One must trust that funds will be used honestly and efficiently. While interpersonal redistribution is not entirely free of this concern, it is far less central than in the regional case.

3. Identity groups

Regions are often home to ethno-linguistic groups and tend to be a source of personal identity. Regional identities can be a major source of attachment and an impediment to solidarity between territories. Inter-regional transfers have an explicit in-group and out-group structure and thus can be heavily influenced by this territorial identity (Balcells et al., 2015; Enke et al., 2023; Chinoy et al., 2026). All this contributes to a very distinctive political economy of regional redistribution. Moreover, regional boundaries often correlate with other salient differences between individuals. Countries often exhibit regional distinctiveness across ethnic, linguistic, or religious lines. In the case of Spain, these differences have a substantially cultural and linguistic basis; regional inequality in the United States is entangled with the history of slavery and its legacy on preferences and attitudes; and in Brazil the relative underdevelopment of the North-East is tied to race and to a long tradition of stereotyping its population.[4] These coincidences

[4]This variation is one advantage of running the study in three countries: the ethno-linguistic and income cleavages overlap in different directions across them, so the design does not rest on a single configuration of territorial and

of regional boundaries with ethnic, linguistic, and cultural groups mean that a judgment about a region may in part reflect a judgment about a group. This is important because the literature has identified that biases and misperceptions also shape policy support (Alesina et al., 2023).

While regional identity forms where people grow up, populations are mobile. Internal migration separates current residence from family origin, so exposure to a region's history is not confined to those living in it. The three countries of our setting experienced large migrations from poorer to richer regions during the twentieth century, and consequently a substantial share of residents of prosperous regions have roots in disadvantaged ones.

4. Political actors

Regions can be highly effective at influencing national politics. In decentralized countries, regions have governments that can exercise a veto as a bloc, and they carry substantial weight in the national legislature, since territorial representation often gives them more seats than their population share would imply. A recent literature has focused on how regions left behind can contest the territorial settlement and destabilize the political system (Rodríguez-Pose et al., 2023; Autor et al., 2020; Rodden, 2019).

Regional redistribution often operates through explicit territorially defined institutional arrangements, fiscal federalism arrangements, EU cohesion funds, etc. Hence, these transfer systems have named beneficiaries, tend to be visible, and thus are highly contested and politically salient. Many countries exhibit recurrent regional conflict over regional financing and fiscal autonomy (e.g., Spain), the degree of federal transfers (Brazil), or who are net contributors to the system (US). Given the identity dimension discussed above, political controversy over regional redistribution aligns with territorial and, in some countries, nationalist cleavages, rather than with the left–right one characteristic of interpersonal redistribution.

5. Persistence and long horizons

People can move up the social ladder over the course of their lives. Regional convergence is much slower, counted in decades. As mentioned above, across the three countries comprising this study, the correlation between regional GDP per capita in 1900 and 2020 exceeds 0.65. The persistence literature has documented how much of the current variation in regional prosperity can be traced to institutions, shocks, and settlement patterns located centuries in the past (Nunn, 2020; Spolaore and Wacziarg, 2013; Dell, 2010; Naritomi et al., 2012; Oto-Peralías and Romero-Ávila, 2016). This is also observed today, for instance, in Europe, where spatial disparities across metropolitan areas are profound and persistent, and may be increasing (Ehrlich and Overman, 2020). Place-based policies are the first candidates to look at when thinking about addressing this problem, but research on their success provides mixed evidence. Little is known about which policies generate self-sustaining gains over the longer run (Neumark and Simpson, 2015; Ehrlich and Overman, 2020), and where gains do materialize, they can unfold over decades (Kline and Moretti, 2014). Citizens may come to believe that helping poor regions takes too long and requires too many resources, with uncertain success, thereby devel-

ethnic divisions.

oping a sense of fatalism.

3 Research Design

3.1 Theory of change

The main goal of this paper is to investigate whether improving knowledge about the causes of economic disparities across regions influences preferences for redistribution toward poorer geographic areas. As described in Section 2, beliefs about the origin of individual income differences may shape the extent of support for interpersonal redistribution (Alesina and Angeletos, 2005). Individuals may attribute someone's misfortune to her own actions or to forces out of her control. In the case of a region, while the agent is undetermined, a parallel logic may apply. The *deservingness* of inter-regional transfers may hinge on the perceived root cause of why a region is relatively poor, on whether its population is perceived as responsible. The historical legacies literature —which has put emphasis on explaining economic geographic differences within countries— is particularly well suited to affect such beliefs. If people are informed that regional income differences are not the consequence of misbehavior of regional governments nor their population, but are rather due to deep-rooted historical factors, they may be more willing to favor redistribution to poorer regions.

Figure 1 summarizes our theory of change. The intervention consists of two information treatments: *T1*, a 2.5-minute animation video summarizing the results of research on the historical origins of regional economic differences in the respondent's country, and *T2*, the same video plus an additional 1-minute video emphasizing that history is not destiny and that economic policies can make a difference to address regional inequalities.[5].

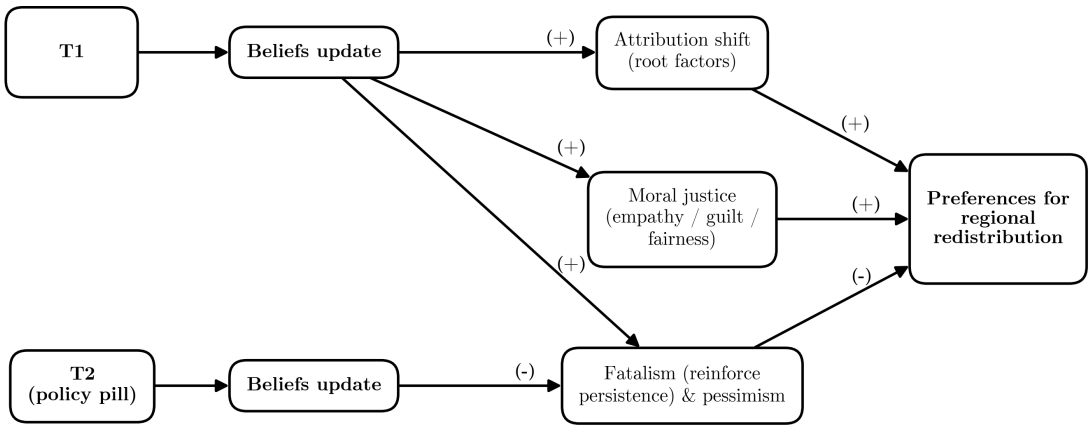

Figure 1: Theory of Change

We expect *T1* to change individuals' beliefs about the importance of historical factors to explain regional inequalities in their country, which may in turn affect preferences for regional

[5]More details in Section 3.3.2

redistribution through three channels. First, *T1* may lead to a responsibility attribution shift towards deep-rooted factors beyond the control of the current population and regional government. This attribution shift should increases support for regional redistribution: if the relative poverty of a region reflects historical accidents, the moral case for transfers is stronger. This is in line with Gniza and Wrede (2024) who argue that lack of responsibility affects the perceived *deservingness* of regions to receive funding. Second, *T1* may generate an emotional reaction of empathy and solidarity, possibly along with a sense of moral responsibility to help poorer regions overcome their historical misfortune. Third, there can also be a backlash effect generated by a sense of fatalism and pessimism. If participants learn that regional income differences are historically rooted and very persistent, then they might consider that they are too difficult to address, so regional redistribution is useless.

The last mechanism, fatalism, is critical in our setting. The concern is that diffusing information about historical legacies may make the public more pessimistic about future convergence and undermine support for redistribution.[6] To explore this channel, we implement a supplementary treatment *T2* highlighting the role that public policies can play. We do not claim that public policy can easily undo the imprints of history, on the contrary, the video highlights how sticky the underlying mechanisms are. Yet, we provide concrete examples of successful experiences (see Section 3.3.2) to understand if mentioning policy agency in this context can matter. Hence, our results can have important consequences for how to disseminate research results among the public, particularly regarding persistence studies. If *T2* significantly increases preferences for redistribution by mitigating fatalism, this would be key for designing the transfer of research results to society.

Finally, in addition to providing new information, *T1* may operate through a saliency channel. By making the historical origins of regional disparities more prominent to the respondent, the information treatment could activate the three previous mechanisms without necessarily altering their prior knowledge. We explain how we address this distinction in Section 3.2.1.

3.2 Hypotheses

3.2.1 Core and secondary hypotheses

We consider three core and six secondary hypotheses. The former refer to the effect of *T1* and *T2* on preferences for regional redistribution (both self-declared and behavioral measures) and the potential heterogeneity by priors. The latter examine the core mechanisms linking the information provision to changes in preferences. These hypotheses are motivated in the previous section and therefore exposed succinctly here.

Core hypotheses: average treatment effects and by priors

C.H1: Exposure to the historical narratives video (*T1*) will increase support for inter-regional redistribution relative to the control group (*CG*).

C.H2: Exposure to the video of historical narratives combined with policy agency (*T2*)

[6] Interestingly, the literature has highlighted that better understanding the historical roots of current disparities is a prerequisite for crafting effective public policies aimed at overcoming these path-dependent obstacles (Nunn, 2020).

will increase support for inter-regional redistribution more than those exposed to historical narratives alone (*T1*).

C.H3: The effect of *T1* on preferences for inter-regional redistribution is decreasing in respondents' pre-treatment priors about the share of current regional income differences attributable to events more than a century ago.

Secondary hypotheses: belief updating

An information treatment can only shift policy preferences if it first shifts the beliefs that the information targets. If treatment moved redistribution preferences without moving beliefs about the historical origins of regional disparities, the effect would be more plausibly attributed to the salience of the topic than to the informational and narrative content of the treatment.

S.H1: Individuals exposed to the historical narratives video (*T1*) will attribute a larger share of current regional income differences to events that occurred more than 100 years ago, compared to individuals not exposed to it (*CG*).

The cardinal item on which *S.H1* is tested is elicited with the same instrument before treatment, immediately after treatment, and again in the follow-up survey, so that updating can be signed at the individual level. We estimate the treatment effect on the post-treatment level while controlling for the baseline elicitation.

Distinguishing genuine belief updating from priming, whereby the treatment merely makes the topic temporarily more salient, is a central challenge in information provision experiments. Our design incorporates two approaches that mitigate this concern (Haaland et al., 2023): prior beliefs are elicited from every respondent before treatment, so that historical persistence is equally salient across arms; and the follow-up survey two months later separates the information provision from a subsequent measurement of the main outcomes. *C.H3* provides a further test: heterogeneity in treatment effects by prior beliefs is commonly interpreted as evidence that effects are driven by genuine changes in beliefs rather than by priming.[7]

Secondary hypotheses: mechanisms

Hypotheses **S.H2-S.H5** relate to the three mechanisms outlined in Figure 1 (attribution shift, moral justice, and fatalism):

S.H2: Individuals exposed to *T1* will attribute more responsibility to historical factors than to the current population for the relative underdevelopment of poor regions, compared to individuals not exposed to it (*CG*).

S.H3: Individuals exposed to *T1* will show stronger feelings of empathy and moral responsibility than individuals not exposed to it (*CG*).

[7] Differentiating genuine updating from agreement with the treatment is a related challenge, which can also be addressed by asking whether the shift depends on respondents' priors. Thus, mere agreement implies similar movement across respondents; genuine revision implies movement concentrated among those whose priors were furthest from the video's message (*C.H3*). In addition, the follow-up survey discriminates between the two on the basis of persistence, since agreement offered to satisfy a perceived expectation should not survive two months, whereas a revised belief may survive. Moreover, the outcome-neutral items on unrelated topics bound generalized acquiescence: if respondents were simply agreeing with what the survey appeared to favor, the pattern should be visible there as well.

S.H4: Individuals exposed to *T1* will express a stronger belief that poor regions will continue to be poor in the future, and a stronger feeling of pessimism, than individuals not exposed to it (*CG*).

While *S.H2*–*S.H4* concern the mechanisms activated by *T1*, the following hypothesis states the mechanism through which *T2* is expected to operate, and is therefore the counterpart to *C.H2*:

S.H5: Individuals exposed to *T2* will be less likely to believe that poor regions will continue to be poor in the future and will show a weaker feeling of pessimism than individuals exposed to *T1* (but still show higher fatalism than the control group).

We expect participants exposed to *T2* to still display a higher degree of fatalism than the control group. This conjecture is based on how both treatments differ. The *T1* content is a 2 minutes and 30 seconds video of an engaging historical narratives about the country of the respondent. As specified in *S.H4*, we expect *T1* to raise fatalism. *T2* still incorporates the same information plus the policy agency pill, but the latter is only 1 minute and not linked to the country of the respondent. Similarly, while the examples of *T2* are still based on academic research, the policy content relies on more cautious language (e.g., *"these examples are not universal guarantees..."*). As a result, we expect the policy agency content of *T2* to only partly offset the raise in fatalism induced by the historical narratives.

Secondary hypotheses: specificity of the treatment effect

The preferences we study are place-based, concerning support for transfers directed to poorer regions. A treatment describing the historical origins of regional disparities could, however, shift redistributive preferences more broadly, making respondents more favorable to redistribution of any kind. The following hypothesis aims to differentiate whether the effect of the information treatment is stronger on preferences for inter-regional redistribution than on those for inter-personal redistribution, measured through a standard item on support for reducing income differences between richer and poorer individuals (Alesina et al., 2023).

S.H6: The effect of *T1* on preferences for inter-regional redistribution will be larger than its effect on support for inter-personal redistribution.[8]

3.2.2 Heterogeneity analysis

This paper aims to improve our understanding of how to communicate academic research to maximize impact while mitigating backlash. We move beyond average effects to analyze how individuals react differently to information. We do so in two distinct ways. First, we specify some concrete hypotheses linked to key stages in the causal path between information and changes in preferences. Second, we employ a machine learning approach based on Chernozhukov et al. (2025), which exploits the full set of pre-treatment covariates jointly to identify the largest treatment effects and characterize the profiles of heterogeneities.

Heterogeneity hypotheses:

[8] If *C.H1* is confirmed but we fail to find support for this hypothesis, it would suggest that the treatment shifts redistributive preferences generally rather than specifically in inter-regional terms.

Updating stage. We first explore what mediates how much a respondent learns and updates her beliefs. New information will update beliefs more if individuals trust the source, they have the cognitive capacity to assimilate the content, or if the narrative is more compatible with their pre-existing mental frameworks. Previous evidence highlights that information treatments may yield highly divergent impacts based on pre-existing mindsets (Alesina et al., 2018; Kuziemko et al., 2015). Hence, the compatibility between the framework of the historical narrative and the causal framework of individuals may be critical.

H.H1: *Trust in the message.* The effect of *T1* will be stronger for participants with higher trust in scientists.

H.H2: *Cognitive capcity.* The effect of *T1* will be stronger for participants with higher levels of education and more complex reasoning.

H.H3: *Frame compatibility.* The effect of *T1* will be stronger for participants who assign more importance to circumstances in explaining interpersonal inequality and structural explanations for regional disparities.

Policy-support stage. We next explore what mediates the influence of belief updating into policy preferences. The extent to which recognizing historical legacies translate into redistribution may depend on the respondent's moral and political commitments, including the geographic radius of their altruism (Enke et al., 2023). Even a respondent who accepts that a region's disadvantage was inherited may not support transfers if they expect the funds to be wasted or captured (Kuziemko et al., 2015; Daniele and Geys, 2015a). Similarly, the policy response may be entirely different if they belief they will be net beneficiaries or if they view it as loss for their in-group (Balcells et al., 2015; Gniza and Wrede, 2024). Individuals from poorer regions are not only those that could benefit the most from more inter-regional redistribution, they are also the ones for which the narrative may feel more personal and resonate more.

H.H4: *Personal resonance/Perceived benefit.* The effect of *T1* will be stronger for participants who come from a relatively poorer region - bottom quartile (i.e., residing there or parents born there).

H.H5: *Moral predisposition.* The effect of *T1* will be larger for individuals with higher moral universalism or leftist ideology (both dimensions will be averaged).

H.H6: *Policy trust.* The effect of *T1* will be larger for individuals with higher trust in the government and on people (both dimensions will be averaged).

H.H7: *Perceived cost.*The effect of *T1* will be smaller for participants who come from a relatively richer region - top quartile (i.e., residing there or parents born there).

While we specify the direction of the expected effects, we acknowledge the limitations of this analysis. Some of these dimensions may correlate with one another (or with the priors on the importance of history, or the update in priors per se), and thus affect how to interpret the results. Hence, we will explore the relationship between all these variables, both in the full sample and by country. Whenever a pair of variables across two distinct families would have a pearson-correlation coefficient above 0.50, we would flag that those heterogeneity results should be taken with caution, as one cannot confidently attribute it to a single dimension

unequivocally.[9] Furthermore, we would assess the robustness of the heterogeneity with partial interactions while controlling for the additional covariate that presented the high correlation, like in a horse-race exercise.

Machine-learning approach. We will use a machine-learning approach to explore heterogeneous results based on the Generic ML Inference framework of Chernozhukov et al. (2025). This method presents several advantages. First, it is agnostic and it does not require us to prespecify which moderators matter or in which direction. Second, it considers all moderators jointly, handling the correlations between them that would confound the results. Third, it can uncover unanticipated sources of heterogeneity that our theory-driven hypotheses may miss.

We will use the full set of pre-treatment covariates as candidate moderators: sociodemographic, political and cultural values, economic beliefs, misperception measures, priors on importance of history, and the content categories and composite measures from the open-ended response on causes of inequality. We will apply the procedure to the pooled treatment (T1 + T2 vs. Control) as the primary specification and then separately for T1 vs. Control and T2 vs. Control. The latter will allow us to examine whether different characteristics predict responsiveness to the historical narrative alone versus the narrative combined with the policy-agency message. The ML approach is exploratory, complementing our theory-driven heterogeneity hypotheses, and may generate new hypotheses for future work.

For each of the three comparisons, we implement the procedure suggested by Chernozhukov et al. (2025) as follows. Since the treatment assignment is done with equal probability in all comparisons (0.5 for *T1* vs. control and for *T2* vs. control; 2/3 for the pooled *T1*+*T2* vs. control comparison), the propensity score is known and constant across all specifications, removing one potential source of estimation error compared to the general case. We randomly split the analysis sample corresponding to each comparison into a main and auxiliary sample of equal size, stratifying by country and treatment so that both subsamples maintain the initial randomization ratios. We use the auxiliary sample to construct machine learning-based proxies for the baseline conditional mean and the conditional average treatment effect, employing two learners: elastic net regression, with the penalty parameter selected by cross-validated mean squared error, as the default one, and random forest as an alternative. We select between them, separately for each split, according to the goodness-of-fit measure Λ proposed by Chernozhukov et al. (2025), favoring the proxy which is the most correlated with the actual treatment effect.

On the main sample, we estimate three objects. First, the best linear predictor (BLP) of the treatment effect of the machine-learning proxy that provides us both an estimate of the average treatment effect as well as a test for whether any observable heterogeneity is present ("HET"). Second, the sorted group average treatment effects (GATES), calculated based on population quintiles ($K = 5$) of the proxy. Third, the classification analysis (CLAN), in which we compare the pre-treatment characteristics of the top and bottom quintiles, so as to find out which covariates are linked with larger or smaller effects. Since all these calculations rely on the arbitrarily chosen split of the sample into the main and auxiliary subsamples, the whole exercise is done

[9]Note that dummies capturing originating from a bottom or top-quartile income region may be highly correlated. Hence, in this case, we also report the estimates excluding the other top (bottom) quartile from the sample.

using $S = 200$ independently randomly chosen splits.[10] The reported statistics include the median of each point estimate across the 200 splits; a confidence interval that takes the median of the per-split lower and upper bounds at the 95% confidence level; and, for each test, a two-sided p-value defined as twice the minimum of the median one-sided p-value against each alternative, taken to be at most one, according to the quantile aggregation procedure defined in Chernozhukov et al. (2025).

Heterogeneity by country. Based on our conceptual framework, we predict that the mechanisms operate similarly in the three countries, and we accordingly expect the main and secondary hypotheses to hold for each of them. Hence, we report country-specific estimates for our core and secondary hypotheses (C.H1–C.H3, S.H1–S.H6), obtained jointly on the pooled sample.[11] Cross-country differences in the estimated effect may arise because respondents differ across countries in characteristics that moderate the treatment, or because the treatment itself operates differently. To explore the potential underlying cause of cross-country differences we do two complementary analyses. First, we add interactions between the treatments and potential moderators (the variables used in the heterogeneity analysis) that could plausibly vary largely across countries. The moderator slopes are constrained to be common across countries, so that a single moderating relationship is applied to all respondents, and we report how far the country-specific coefficients converge once these interactions are included. If the country differences attenuate substantially, the mechanism is common across settings and the differences reflect the countries' sample compositions. If substantial heterogeneity remains, it could reflect substantial heterogeneity on other dimensions (either observable or unobservable) or that the treatment genuinely works differently across settings. Second, we leverage the machine-learning approach for heterogeneity previously discussed and apply it country by country. In doing so, we first aim to understand if the profile that predicts larger treatment effects is similar across countries. Once such profile is identified, we report to which extent the countries differ significantly in the variables flagged as critical for heterogeneity.

3.2.3 Follow-up hypotheses

The follow-up survey allows us to assess whether the effects of the treatments persist after eight weeks, how much they attenuate, and whether there is a differential decay rate. An effect present in both waves is reported as a persistent shift in beliefs and preferences, while an effect observed immediately but not in the second wave is reported as a transitory shift.[12] We also expect the effects of the treatments on preferences for regional redistribution, on belief updating, and on each of the three channels to attenuate over time, but not necessarily at

[10] Chernozhukov et al. (2025) use 200 splits in their microfinance application and 250 in their immunization application, stating that their results do not change when further increases in the number of splits are performed. We choose the lower value as being sufficient due to our sample size.

[11] The precise estimating equation is reported in Section .

[12] The latter should not be read as evidence of priming or other biases. There is an asymmetry in the use of the follow-up worth noting. While finding persistence constitutes strong evidence against priming and experimenter demand effects, since the follow-up is an obfuscated survey, a lack of persistence does not imply the presence of these biases. This is because a null result at two months is equally consistent with genuine decay and with priming or demand effects. The main test to distinguish between them is *C.H3*, since an effect produced by salience should not vary with respondents' priors. The placebo items and the bias-perception index explained below are also used to check for these biases.

the same rate. In particular, we expect the effect of *T2* to decay more than that of *T1*. This prediction is based on the treatments themselves: the historical segment is longer, tells a single narrative, and concerns the respondent's own country, whereas the policy-agency extension lasts a minute, rests on two cases from elsewhere, and depends heavily on trust in government action, which can deteriorate quickly.

Inspired by this discussion, we specify the following two sets of hypotheses:

> ***F.H1*** *(persistence)*: The effects predicted in *C.H1* and *C.H2* and in *S.H1–S.H5* will still be present two months after treatment.
>
> ***F.H2*** *(decay)*: The effects in *F.H1* will be smaller in absolute value than the corresponding immediate effects in the baseline.

The follow-up hypotheses are secondary for several reasons. First, the follow-up sample is expected to be about 65% of the initial sample, which raises the minimum detectable effect from 0.077 to roughly 0.096 SD. Second, the follow-up sample is not itself a randomized sample, so inference there rests on the balance tests, inverse-probability weighting and Lee bounds, which the main-survey estimates do not require. Third, the main survey is richer in information than the follow-up. The regional redistribution index cannot be reconstructed because two of its three components (the petition and the lottery) cannot be introduced again without compromising its obfuscated character. The same applies to the open-ended question used in each mechanism index. Overall, our main research question concerns whether a historical narrative moves preferences, while documenting the drivers of persistence is a distinct and additional question.

3.3 Methodology: Experimental Survey

We conduct an online survey with an embedded information intervention to 8,000 respondents from Brazil, Spain, and the US, three countries with salient regional inequality patterns. The survey flow is represented in Figure 2 and the full questionnaire in English is available in Appendix C.

3.3.1 Survey Structure

The survey is structured in four core blocks: (i) socio-demographics, (ii) economic beliefs, (iii) political and cultural values, and (iv) preferences for redistribution. The information treatment takes place before the last block. An obfuscated follow-up survey is conducted two months after the initial assessment to collect data on the primary outcomes.

Socio-demographic variables. We collect basic socio-demographic variables that serve as control variables. These include age, gender, region/state of residence, education, income, civil status, children, household size, type of area of residence (rural or urban), employment status, and occupation. To test the heterogeneity of the treatment effect by region of family roots, we ask for the region/state of birth of respondents, their parents, and their grandparents.

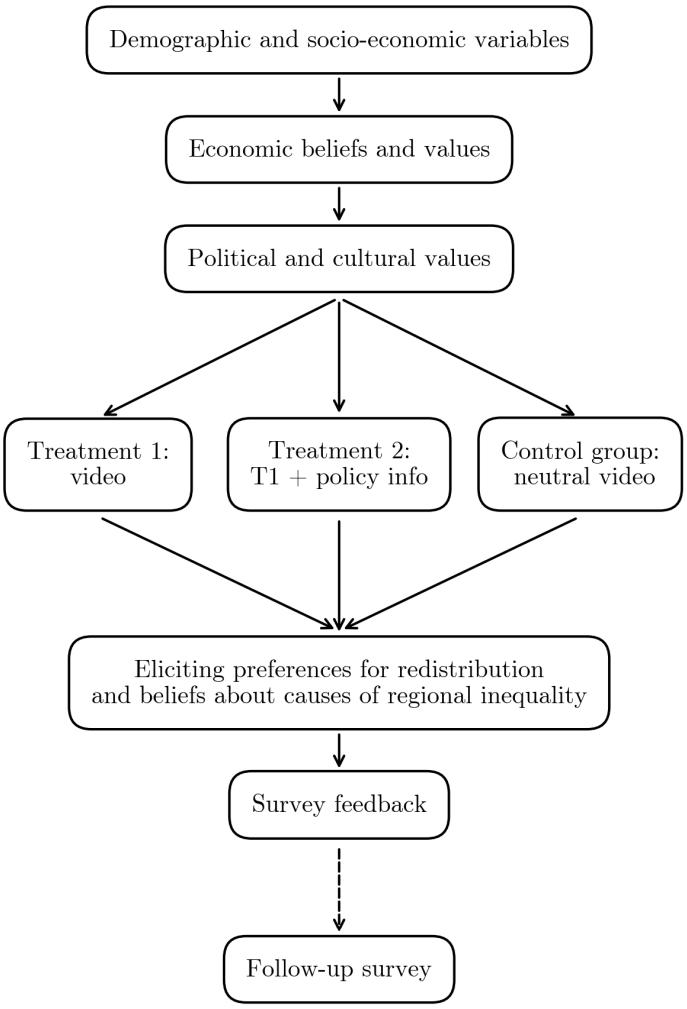

Figure 2: Survey structure

Economic beliefs and values. The block first measures the zero-sum mindset regarding income groups, international trade, and regional differences (Chinoy et al., 2026). It continues with two items to assess respondents' priors on the causes of poverty—specifically, the role of effort versus circumstances (Alesina and Angeletos, 2005).

Next, respondents guess which are the richest and poorest regions today, alongside their relative income, using an interactive bar graph (Figure C.2). Once respondents have given their estimate, an open-ended question asks for the factors explaining current regional income differences. This order mirrors how people encounter the issue outside the survey: they first observe that some regions are richer than others, and only then ask why. LLM analysis of these responses captures individual narratives and development models without researcher-induced priming (Haaland et al., 2024).[13]

We then elicit the respondents' priors about historical legacies by asking them to estimate how much of current income differences across regions is attributable to recent developments and the current characteristics of each territory as opposed to events and processes that took

[13]Sections 4.4–4.5 explain how we plan to exploit and analyze the responses to the open-ended questions.

place a long time ago (more than 100 years ago).

Finally, respondents estimate the country's poorest and richest regions in 1900, alongside their relative income. These items, together with the ones referring to current differences, gauge the perceived persistence of regional income. We also include an explicit question on the extent of persistence between 1900 and 2025, preceded by a brief introduction to the economic concept of persistence.

Political and cultural values. This block aims to capture some political and cultural values of the respondents, including regional identity, political ideology, interpersonal trust, trust in scientists, and trust in the government. These variables are used in the heterogeneity analysis.

Information treatments. Participants are assigned to one of the three experimental arms, as detailed in Section 3.3.2. A landing page explains that they will proceed to a short video after which they will have to answer a few questions about it.

Preferences for redistribution. After the information treatments, we include a set of questions to elicit respondents' preferences for inter-regional redistribution:

1. Whether they favor or oppose rich regions transferring more resources to poor regions to help finance their public services.

2. A real-stakes question asking respondents to support one of the following two policies for a petition to be sent to the national government:

    - The fiscal system should transfer more resources from higher-income regions to lower-income regions, so that the latter can improve their public services and infrastructure.
    - The fiscal system should reduce the transfer of resources from higher-income regions to lower-income regions, so that each region relies mainly on the resources generated within its own territory.

3. A behavioral, real-stakes monetary question in which the respondent decides how to allocate a lottery prize (equivalent to USD 500). The respondent faces two scenarios, one with a split between herself and a low-income participant from her region, and a second between herself and a low-income participant from the poorest region of the country. This lottery question captures willingness to transfer across territorial lines.[14]

This block also measures the underlying mechanisms. First, we include an open-ended question asking respondents to explain why they favor or oppose inter-regional transfers, which allows us to gauge their first-order motivations without priming them with predefined answer options (Ferrario and Stantcheva, 2022; Stantcheva, 2023). Second, we ask respondents

[14] In the two scenarios, we explicitly state the income threshold of the low-income participant, which is set below the At-Risk-of-Poverty (AROP) line. This is defined as 60% of national median equivalized disposable income after social transfers, which stood at €12,220 per year per consumption unit in Spain in 2025 (Instituto Nacional de Estadística, 2025). Since the number of consumption units is at least one, a respondent reporting monthly household income below the threshold necessarily falls below the AROP line. If the respondent already resides in the poorest region, the second scenario will be with respect to the second-poorest region of the country. Given this peculiarity of the question, we present results without respondents from the poorest region as a robustness check (although this only affects roughly 2% of the sample).

to report their feelings when thinking about the poorer regions in their country, specifically measuring empathy or compassion, moral responsibility, and pessimism. Third, to test for a shift in responsibility attribution, we ask whether respondents agree that poor regions still pay the price for disadvantages created long ago, as opposed to being poorer due to the work attitudes and initiative of their population. To capture any potential fatalism induced by the information treatment, we ask how likely it is that the currently poorest regions will remain the poorest 30 years from now. We further measure perceived policy efficacy by asking to what extent well-designed economic policies can generate real economic impacts.

Finally, we ask participants the same question on the relevance of historical legacies, namely to rate the relative contribution of current regional characteristics and historical events (more than 100 years ago) to present-day regional income differences.

3.3.2 Information Treatments

The first video treatment (*T1*) summarizes findings from academic articles on the historical origins of regional economic differences in the countries covered by the survey. For Brazil, the video draws on Naritomi et al. (2012), Musacchio et al. (2014) and Wigton-Jones (2020) to explain how colonial sugar and gold booms established extractive institutions and unequal land distribution, enduringly hindering local governance and human capital accumulation. For Spain, the video draws on Oto-Peralías and Romero-Ávila (2016) and Beltrán Tapia and Martinez-Galarraga (2018) to explain how the pace and conditions of the medieval Christian Reconquest shaped the concentration of land ownership across Spanish regions, depressing human capital accumulation and regional income down to the present day. For the US, the video builds upon Engerman and Sokoloff (2002), Acemoglu et al. (2002), and Bruhn and Gallego (2012) to illustrate how variations in initial endowments led to either smallholder farming or slavery-reliant plantations, generating long-lasting institutional and economic disparities.

The videos are designed to be short (2:30-3:30 minutes), engaging, and accessible. The narrative is structured in four parts. First, it poses the guiding question of what factors explain regional income inequality within the country. Second, it introduces the importance of history to explain these disparities. Third, it presents specific economic history findings tailored to the participant's nationality.[15] Finally, the video concludes with a call to reflect on the importance of understanding the deep roots of regional inequalities.

A second treatment arm (*T2*) exposes participants to the same content as *T1*, followed by a short 1-minute extension stressing policy agency. It illustrates the "history is not destiny" idea (Nunn, 2020; Banerjee and Duflo, 2014) with two specific, well-documented cases in which place-based interventions produced a clear positive effect: the United Kingdom's Regional Selective Assistance program, which Criscuolo et al. (2019) show raised area-level manufacturing employment without simply moving jobs from nearby places, and China's Special Economic Zones, which Lu et al. (2019) show —focusing on zones created in 2005-2008— significantly increased employment and output with limited spillovers to surrounding areas. Both examples are drawn from countries outside our study sample (Brazil, Spain, and the United States), so

[15] That is, Brazilian participants watch the video concerning Brazil, Spanish participants the one concerning Spain, and US participants the one concerning the US.

that the policy-agency message is symmetric across treatment arms.[16]

As exposure to *T2* implies one further minute of video, the comparison between *T2* and *T1* identifies the effect of adding the policy-agency segment as implemented, that is, its content together with the additional exposure it requires. We interpret it accordingly throughout, and we pre-specify a set of checks on treatment receipt and engagement, reported by arm, to bound the alternative channels of attention and effort.[17]

Participants are randomly assigned with equal probability to *T1*, *T2*, or the mock control group. The latter group watches a video matched to *T1* in length (2:30 minutes) and production quality, explaining the mechanics of the bank deposit money multiplier.[18] This placebo video is appropriate for at least two reasons. First, it is genuinely economics-relevant, so that watching it feels consistent within the context of an economics survey. Second, it is entirely unrelated to regional inequality, redistribution, historical causation, or policy effectiveness: it involves no groups, no regions, and no distributional content of any kind, minimizing the risk that the control video itself shifts beliefs relevant to our outcomes.[19]

This design allows us to isolate the effect of the videos' informational content from the effect of time-on-task, engagement, and production quality alone (Haaland et al., 2023). Randomization is stratified by state (in Brazil and the US) and region (in Spain) to ensure that the treatment and control groups are fully balanced across all geographic areas. To maximize respondent attention, all videos—including the mock control—feature high-quality production, utilizing animations by a professional cartoonist and voiceovers by professional actors. The full scripts, drafted by the research team, are provided in Appendix B.

Immediately following the videos, we include questions to assess respondent attention and comprehension. We ask whether they watched the video until the end, along with a factual manipulation check regarding the video's content. To gauge genuine interest, we offer respondents the option to receive a link to the video via email after the survey for re-watching or sharing. Furthermore, in the final feedback section of the survey, we explicitly ask participants to what extent they were persuaded by the video's arguments, providing an additional measure of treatment efficacy.

[16] We acknowledge that we illustrate two examples in which policies appear to have worked, and we could have also discussed examples in which they did not. This is an explicit choice because the purpose of *T2* is not to correct a misbelief or to educate participants on what the consensus in the literature may be, but rather to allow examining a communication strategy that bundles persistence findings with evidence of policy agency. Importantly, the wording is intended to ensure that respondents are not persuaded to believe that historical persistence is easily solved or that the examples portrayed are representative of place-based policy in general.

[17] The checks are the following: (i) whether respondents skipped the video at some point; (ii) their self-reported confirmation that they watched and listened to the end; (iii) performance on the post-video comprehension check; (iv) break-off at the video screen; (v) a direct measure of perceived persuasive intensity; and (vi) the share requesting the video link by email, as a behavioral indicator of engagement. If these checks are comparable across arms, the attention and effort channels are largely excluded and the effect can be attributed to the added content; if not, we report this and qualify the interpretation accordingly. Layered treatments of the kind used here, differing in both content and length, are common in the information provision literature (e.g., Alesina et al., 2018; Kuziemko et al., 2015).

[18] We follow Stantcheva (2023)'s terminology, which distinguishes a *mock* control—receiving unrelated content, primarily to equalize duration and engagement—from an *active* control, which receives different information on the same topic.

[19] We considered and rejected several alternative topics — including AI and employment, climate change, demographic aging, geopolitical risk, and immigration — because each carries a meaningful risk of activating beliefs about policy effectiveness, distributional fairness, or institutional trust, undermining their suitability as a neutral placebo.

3.3.3 Follow-up survey

Participants who take part in the main survey are invited to complete a very short (3-4 minute) obfuscated follow-up survey, taking place two months after the initial survey. The questionnaire contains some socio-demographic questions to give it the appearance of an independent survey. It also contains some other unrelated policy questions for the same purpose. We then include five questions to elicit preferences for redistribution and to measure the three hypothesized mechanisms. The objective of the follow-up survey is to test whether the information treatment has only a short-term or also a medium-term effect on preferences for regional redistribution, and to examine which mechanisms exhibit greater persistence. For instance, it is possible that the effect on moral justice dissipates in two months, while the effect on the responsibility attribution shift does not. Furthermore, it serves as an additional approach to assess whether the effects are influenced by short-lived priming and other related biases, such as social desirability bias and experimenter demand effects.

3.3.4 Quality checks

The survey contains a number of quality and attention checks designed to maximize respondents' attention and detect low-quality answers. First, the landing page emphasizes the importance of reading carefully and answering honestly, indicating that rushed and incoherent responses will be flagged through statistical control methods. We also encourage participants by stressing that they are contributing to our collective knowledge as a society and can learn new facts and express their views.

We first ask socio-demographic questions so that we can analyze whether there is any non-random selection in attrition at baseline based on the topic, although we expect a very low overall level of attrition, as explained below in Section 3.3.5. At the end of block three, we include a self-reported honesty check, placed strategically before the information treatment (Stantcheva, 2023). Right after the treatments, we include factual manipulation checks to verify respondents' comprehension and attention to the videos. In section five (eliciting preferences for redistribution), we include an instructional manipulation check to identify inattentive respondents, alongside two outcome-neutral or placebo tests to ensure that the results obtained can credibly test the specified hypotheses.

To detect careless answering ex-post, we analyze response patterns, specifically flagging respondents who exhibit mechanical answering (using a Response Pattern Index following Meade and Craig (2012)) or who provide nonsensical or very short answers to open-ended questions.

Furthermore, the time spent on the survey can be used to detect low-quality responses. We check the robustness of the results to excluding outliers in the survey duration—specifically, the bottom 5% (to filter out rushed answers) and the top 5% (to filter out respondents who were likely distracted). Crucially, we also monitor the exact time spent on the video treatment pages, flagging respondents who spend significantly less time than the video's duration or who spend an excessively long time.

We expect high engagement with the videos and we introduce several ways of measur-

ing attention at this specific stage.[20] We do not attempt to artificially raise attention through warnings or requiring repeated correct answers to a screener. The literature finds such interventions to have limited and inconsistent effects and may increase dropout (Berinsky et al., 2016; Stantcheva, 2023). Instead, we rely on video quality and brevity to keep participants engaged.

Finally, in the feedback section, we ask for the respondents' opinion on whether they felt the survey was biased in any sense, which serves as a check for experimenter demand effects, and we provide them the possibility to enter open-ended comments.[21] We ellaborate how we plan to use this information in Section 3.6.

3.3.5 Sample and statistical power

The sample size is 8,000 individuals, divided equally among Brazilian, Spanish, and US nationals. With two treatment arms and a control group, this results in 2,666 respondents per group. This largely exceeds the 700-respondent threshold recommended by Haaland et al. (2023) for information provision experiments.[22] Our sample size is consistent with other large-scale, multi-country survey studies (e.g., Alesina et al. (2023)). Furthermore, this sample size provides sufficient statistical power to explore potential heterogeneity in treatment effects. Splitting each of the three experimental groups into two branches for heterogeneity analysis still yields over 1,000 respondents per branch.[23]

When conducting follow-up surveys, it is necessary to account for attrition and increase the baseline sample accordingly. Assuming an estimated attrition rate of 35%, our baseline sample of 2,666 respondents per group provides a substantial safety net. We expect approximately 1,733 respondents per group to complete the follow-up survey (i.e., $(1 - 0.35) \times 2,666$).

Overall, our sample size provides the statistical power necessary to detect relatively small effects and our minimum detectable effects (MDE) align well with the literature (more details in Appendix A.3). The MDE for a one-variable outcome under the core and secondary hypotheses is 0.077 of a standardized mean difference, while the MDE for the heterogeneity analysis is 0.109. For the follow-up survey, assuming the 35% attrition rate, the MDE is 0.095. Given that we will use indices for our main independent outcomes, the MDE could reasonable fall between 18%-28% depending on how correlated their components of the index are.[24] These thresholds are best assessed in comparison with comparable information-provision experiments. Table A.4 in the Appendix reports MDE and treatment effects of ten studies on

[20] Following Berinsky et al. (2016), who recommend relying on more than one attention item, we combine three measures recorded for every respondent in every arm: the self-reported video completion item, the factual manipulation check, and time on the video page.

[21] Asking respondents at the end of the survey if they perceived any bias is a standard practice to detect experimenter demand effects (Stantcheva, 2023). It allows us to verify whether the information treatments inadvertently signaled a specific intent, and provides a way to conduct robustness checks by re-estimating treatment effects excluding those who felt the survey was biased.

[22] Haaland et al. (2023) recommend that information provision experiments "should have at least 80 percent power to detect a treatment effect of 15 percent of a standard deviation" (p. 30), which is important for measuring null findings with confidence. This translates to a minimum of 700 respondents per treatment arm.

[23] This calculation assumes that treatment arms are divided equally in the heterogeneity analysis, which is not necessarily the case. However, our large sample size provides flexibility; it is sufficient if each subgroup constitutes at least one-quarter of a treatment arm (approximately 666 participants).

[24] The precise formula underneath this computation is provided in Appendix Section A.3

attitudes to redistribution and other related political outcomes. The average MDE sits around 0.10, and the effects tend to range between 0.10-0.15 SD. Hence, our MDE lies below the range in which comparable effects have been found.

For country-specific estimates, the corresponding MDE is 0.133 SD for a one-variable outcome and potentially in the range 0.097–0.109 SD for our redistribution index. Once again, this aligns with literature. If the analysis would proved underpowered to estimate country-specific effects, we plan to pool T1 and T2 into a single treatment indicator and test whether the historical narrative has any effect relative to the control, which would in turn lower the country-specific MDE by approximately 30%.

3.4 Representativeness of the sample

The survey is administered by the specialized survey firm *Bilendi-Netquest*. The company recruits participants through its online panel, sending targeted messages about new survey opportunities, and continues recruiting participants until the sample size is achieved (or specific quotas respected). We restrict participation to individuals born in the country and aged 18 to 70. We apply quotas on education (% with college degree), age group (% above 40), gender, and state or region of residence.[25] Hence, the sample should be seen not as nationally representative but rather as quota-representative of each country's adult population on these demographic dimensions. This is the standard approach in the literature on information-provision experiments using online panels (Alesina et al., 2018; Stantcheva, 2023; Chinoy et al., 2026). The underlying panel is self-selected: individuals who volunteer for online surveys may differ from the general population, and those who opt into a particular survey may differ from the broader panel. This is a well-known limitation of online panel experiments, which quotas and reweighting only partially alleviate. We present as baseline results unweighted estimates, following the convention in this literature. Additionally, as a robustness we report estimates reweighted to the population joint distribution of the quota variables, computing the weights from the most recent data for each country (2020 American Community Survey ACS 1-year for the United States, 2021 census for Spain, 2022 census for Brazil).

3.5 Attrition in the baseline survey

The firm in charge of the online survey (*Bilendi-Netquest*) is one of the most experienced in the sector, ensuring high-quality standards. Participants are compensated only for completing the survey; bonuses are rewarded for accurate completion; and attention checks are put in place to monitor and remove participants who answer carelessly. The survey firm commits to provide the agreed number of *complete* questionnaires (8,000) full-filling our requested quotas. Additionally, it will provide all incomplete questionnaires that we will use to report baseline attrition rates by treatment arm.[26] Given the aforementioned high-quality standards

[25] Since Spain and Brazil have fewer than 30 regions/states, when presenting country-specific results, we will supplement clustered standard errors with wild bootstrap p-values (Roodman et al., 2019).

[26] As mentioned above, the first block of the questionnaire contains socio-demographic questions which allow us to analyze a potential selection in attrition at baseline, conditional on the participant having answered a sufficient number of questions.

of *Bilendi-Netquest*, the moderate duration of the survey, and the fact that the information interventions appear only midway, we expect negligible dropout at baseline and dropout rates to be orthogonal to the treatments.

3.6 Attrition in the follow-up survey

Individuals will be recontacted to participate in a follow-up survey two months after the initial survey. We expect a recontact rate of about 65%, the figure anticipated by the survey firm and in line with previous studies using a similar recontact period (Stantcheva, 2023). Note that our power calculations already consider this rate of participation.[27] Now, to reduce attrition in the follow-up participants will receive a bonus for completion 25% higher than in the baseline. We specify a detailed protocol to assess and handle attrition regardless of its ultimate level:

1. *Differential attrition.* We regress an indicator for follow-up participation on the treatment indicators and the full set of baseline covariates, testing whether attrition differs across arms and by observable characteristics. The mock control video removes the component of any differential that would otherwise arise simply from having watched a video and from the additional time this takes. A one-minute differential nonetheless remains between $T2$ and the other two arms, and the test is reported for both contrasts separately.

2. *Baseline effects on the follow-up subsample.* As an indirect test of whether the follow-up sample is selected on responsiveness to the treatment, we re-estimate the main-survey treatment effects on the subsample of respondents who complete the follow-up, and compare them with the full-sample estimates.

3. *Inverse-probability weighting.* All follow-up results are reported both unweighted and with inverse-probability weights, estimated from a model of follow-up participation on the treatment indicators and baseline covariates.

4. *Lee bounds.* We report Lee (2005) bounds on the follow-up estimates as a robustness check. We anticipate that attrition at this rate will produce wide bounds, and we will interpret them as worst-case limits rather than as our preferred estimates.

3.7 Social desirability and demand effects

We acknowledge that the survey content and the information treatments may induce experimenter demand effects and social desirability bias. Some respondents may believe that researchers expect them to answer with a pro redistribution stance (demand effects). Similarly, they may be prompted to think that a pro-redistribution stance is what is socially expected from them (social desirability). If the information intervention affect these dimensions, then our treatment estimates would be biased. These concerns would to some extent be alleviated in the obfuscated follow-up survey, particularly as there is no priming on regional inequality nor any information treatment then. That said, to provide some evidence on these threats in

[27] Importantly, our core analyses rely on the baseline survey and thus are not dependent on the degree of participation in the follow-up

the baseline, the feedback block asks all respondents whether they felt the survey was biased in any sense; those answering yes, are then invited to explain in an open-ended question.

From these two items we construct a dummy indicator (reporting bias or not) and a perceived bias index, coded from the open-ended responses with Large Language Models (LLMs). The index is signed, from -2 to $+2$: positive values record the perception that the survey promoted the view that more resources should be redistributed to poorer regions (or a general bias in favor of poorer regions), which is the direction that would bias our estimates upward; negative values record the perception of a slant in the opposite direction, which would bias them downward; and zero covers respondents reporting no bias, a bias unrelated to regional redistribution, or a regional slant whose direction is not stated. The LLM implementation is described in Appendix Section A.6. We rely on three models, assess their coding against human coding, and select the best-performing one. The prompt is reproduced in Section A.7. These measures allow us to test for differences in perceived bias across arms and, if we find them, to run a battery of descriptive tests to investigate the extent to which they may influence our core results.[28] Note, however, that these exercises are only suggestive, as it can happen that individuals do update their priors thanks to the video and their preferences shift (an actual treatment effect) and that simultaneously they perceive the videos slightly biased. This is a limitation that we cannot overcome.

3.8 Country Selection & External Validity

To enhance external validity of our study, we conduct the survey in three countries: Brazil, Spain, and the United States. This selection captures a diversity of institutional and cultural backgrounds, effectively mitigating the single-country bias common in survey experiments.

Despite their socioeconomic differences, these countries share critical characteristics that make them convenient settings. All three exhibit large, persistent spatial inequalities and feature federal or highly decentralized political structures. In these systems, explicit or implicit inter-regional transfers operate mitigating territorial disparities, even if the intensity of the public debate surrounding these transfers varies across the three nations, being -for instance- particularly salient in Spain. Finally, there is a rich literature documenting some of the historical roots of regional inequalities, providing credible content for our information treatments.

Yet, the three countries differ in how regional boundaries relate to other social categories: Spanish regional identities have partly linguistic foundations, while regional disparities in Brazil and the United States are historically entangled with race and with slavery. Because no such category is common to the three countries, a treatment effect that appears in all three is unlikely to operate through any one of them, which is a further reason for the cross-country design and for reporting country-specific estimates as part of the main analysis.

[28] These tests include estimating the core hypotheses with interactions between the treatment indicators and these measures, and reproducing the estimates for subsamples defined by the perceived bias measures.

4 Data & Measurement

4.1 Data collection

The data will be collected by the survey company in two waves: the main survey and an obfuscated follow-up survey two months later. Before conducting the main survey, we will run a pilot survey of 300 individuals (100 in each country) to ensure that the questionnaire is accurately programmed and implemented. If the pilot survey proceeds without issues, the data collection will continue as planned, and all responses will be pooled. If we identify any issues, they will be corrected, and the pilot responses will be discarded.[29] In either case, the final sample size will be 8,000 respondents, as we have budgeted for such eventualities.

The data collected through the questionnaire can be classified into four groups of variables: pure controls, variables for heterogeneity analysis, independent variables, and outcome variables. Table A.1 in the Appendix provides the description and operationalization of all the variables. Next, we focus on explaining our approach to measuring: (i) priors on historical-attribution and miss-perceptions on historical inequality, and (ii) beliefs about regional inequalities and redistribution using open-ended questions.

4.2 Historical-attribution to current regional inequality (priors/beliefs)

Before treatment, all participants are asked to provide an estimate of how important history is at explaining contemporary regional inequalities. We formulate the question as follows:

> *Still thinking about income disparities between regions, in your opinion* ***how much of these differences is due to events and processes that took place a long time ago (more than 100 years ago), as opposed to more recent developments and the current characteristics of each region*** *— its population, its government, its economy?*

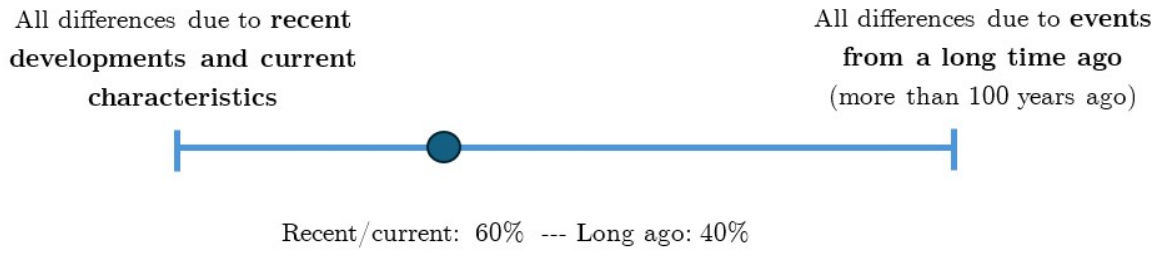

This question serves as our proxy for priors on historical-attribution, and since it is asked before and after treatment, it allows us to measure belief updating for each individual.

Eliciting priors before the treatment has implications for the interpretation of the estimand. As all respondents, in every arm, answer the same pre-treatment questions about regional

[29] We will use this pilot to report, by treatment arm, the distribution of time spent on the video page, the video-player completion rate, and the pass rate on the post-video comprehension check, and we pre-specify that if completion falls below 80% we will adapt the video.

income differences and their causes, the estimand is the effect of the historical narrative on respondents who have already been prompted to think about regional disparities and their origins, and not on respondents encountering the topic for the first time. Doing so in all arms equalizes the salience of historical persistence and is a standard remedy for priming in information-provision designs (Haaland et al., 2023). It affects the estimand, not identification.

4.3 Misperceptions of historical inequality and persistence

Participants are asked to provide an estimate of the ratio between the income of the richest and the poorest region of the country (both today and in 1900). We also present the concept of long-term persistence, and ask individuals to provide a guess. We measure misperceptions on historical inequality and persistence by comparing participants' responses to the actual historical evidence for their countries. Specifically:

Historical inequality misperception:

$$\left| r_{1900}^{actual} - r_{1900}^{belief} \right|$$

where r_{1900}^{actual} denotes the actual ratio between the richest and poorest region in 1900 and r_{1900}^{belief} is the ratio declared by the respondent.

Long-term persistence misperception:

$$\left| c_{1900-2025}^{actual} - c_{1900-2025}^{belief} \right|$$

where $c_{1900-2025}^{actual}$ denotes the actual correlation of regional GDP p.c. 2025–1900 p.c and $c_{1900-2025}^{belief}$ the degree of correlation declared by the respondent.

4.4 Beliefs on regional inequalities from qualitative evidence

To gain insight into respondents' baseline beliefs about the causes of regional economic disparities, we exploit an open-ended question asking participants to explain why some states or regions are poorer than others in the country. Open-ended responses provide a unique window into respondents' salient mental models, without imposing predefined categories or theoretical frames, and mitigating experimenter demand effects (Haaland et al., 2024). We interpret these responses not as fully articulated theories of inequality, but rather revealing which explanations are top-of-mind, whom respondents hold responsible (e.g., markets, governments, individuals), and whether inequality is perceived as a present-day outcome or as the result of longer-run processes. These dimensions are central to understanding preferences for inter-regional redistribution and provide a natural complement to standard quantitative measures.

We analyze the open-ended responses using LLMs, which allow us to systematically and consistently classify qualitative text at scale. The full procedure, including model versions, settings, human coding, validation and robustness, is pre-specified in Appendix A.6.

Each response will be coded into a set of non–mutually exclusive categories capturing the main types of explanations articulated by respondents. The categories are designed to reflect empirically salient and theoretically relevant dimensions of how individuals reason about regional inequality, while avoiding an ex ante emphasis on historical explanations. Specifically, responses are classified according to whether they reference:

- **Economic structure:** labor market conditions, unemployment, sectoral composition (e.g., tourism, agriculture, industry), productivity, firm location, or agglomeration effects.
- **Human capital and demography:** education, skills, innovation, youth out-migration, brain drain, or population aging.
- **Infrastructure and geography:** transportation, connectivity, rurality, peripherality, access to services, or geographic location.
- **Politics, governance, and institutions:** quality of regional government, corruption, mismanagement, centralization, favoritism, or policy choices.
- **Public finance and redistribution:** allocation of public funds, subsidies, transfers, or perceived fiscal imbalances.
- **Culture and attitudes:** work ethic, entrepreneurship, mentality, or effort.
- **Historical or long-run legacies:** direct references to the past, historical events, long-run development paths, or temporal language extending beyond the present.
- **Non-explanatory:** expressions of uncertainty, generic statements ("many factors"), or circular reasoning.
- **Other:** a residual category to capture spontaneous or unanticipated rationales.

Beyond content classification, we construct three composite measures that characterize how respondents reason about inequality. First, we code *agency attribution*, identifying whether inequality is framed as the result of structural forces (markets, geography, institutions), external (imposed by central authorities or outside elites), or regional/individual (reflecting regional choices, culture, or effort). We code these agents as independent dummies. Third, we build a *grievance intensity* indicator that flags responses employing moralized language that frames inequality as unjust, unfair, or characterized by exploitation and zero-sum competition. Finally, we construct a *sophistication index* based on response length (normalized), a dummy for mentioning more than one distinct explanatory category, a dummy for the presence of explicit causal reasoning (indicated by connectives such as "because," "therefore," or "leads to"), and a dummy for the presence of nuance (acknowledgment of complexity and interaction between factors).

These measures serve two main purposes. First, they provide novel descriptive evidence on how individuals conceptualize regional inequality in the absence of informational priming. Second, they allow us to study whether and how the informational treatment interacts with explanatory frames, such as increasing temporal depth, shifting agency attribution, or crowding out purely present-oriented explanations.

Appendix A.4 presents illustrative responses that mimic what could be obtained from participants, varying substantially in length, depth, and content. Table A.6 displays the corresponding classification of these responses according to our coding framework.

A few methodological considerations are worth highlighting. Every response is classified independently by three models: Claude Sonnet 4.5 (Anthropic), GPT-4o (OpenAI) and Mistral Small 3 (Mistral AI). We validate the LLM categorizations on a sample of 600 responses (stratified by country and arm) that are coded by humans. We report inter-coder agreement between human coders, humans against LLMs, and the LLMs across each other. We will report the typical measures of performance such as precision, recall, and F1 for each category and composite measure. We take as our baseline the classification of the model with the highest average F1 on the validation sample.[30]

4.5 Justifications for regional redistribution policy views

We exploit a second open-ended question to understand the *normative reasoning* underlying redistribution preferences. While the beliefs question captures causal mental models, the justifications question reveals the moral, pragmatic, and identity-based arguments that individuals invoke when taking positions on redistribution policy. We analyze these responses using the same LLM-based classification approach described above, validating against human coding on a subset of responses. Each response is coded into non-mutually exclusive categories capturing the main types of justifications. The categories are designed to reflect salient and theoretically relevant dimensions of reasoning about redistribution:

- **Solidarity and justice:** justifications based on the moral or ethical imperative to support disadvantaged regions, ensure equal opportunity, or level the playing field.
- **National unity and cohesion:** justifications based on the principle of national solidarity, equal rights, and the uniform provision of public services regardless of location.
- **Historical redress:** explicitly framing redistribution as compensation for historical legacies or long-run processes that created current disparities (e.g., "richer regions benefited from policies that hurt poorer ones").
- **Autonomy and self-interest:** arguments rooted in the desire to retain and manage own resources, the preference to minimize local tax burdens, or right of self-determination.
- **Moral hazard and incentives:** concerns that transfers create dependency, discourage effort or reform, reward poor governance, or distort incentives for regional development.
- **Distrust and misuse:** skepticism that transferred funds will be used effectively, that policy can achieve change, concerns about corruption, elite capture, or mismanagement.
- **Non-explanatory:** expressions of uncertainty, generic statements, or circular reasoning.
- **Other:** a residual category to capture spontaneous or unanticipated rationales.

[30] When using LLM-constructed variables in regression analysis, we show as robustness the results using the variables based on the three different LLMs.

Beyond content classification, we construct the same composite measures to evaluate the depth and orientation of respondents' policy reasoning: agency attribution, grievance intensity, and sophistication. Agency attribution is adapted to capture whom respondents hold responsible for *enacting* redistribution and it is coded as four separate dummies (e.g, the state, rich regions, citizens collectively, or poor regions themselves), with a separate variable indicating when the person invokes more than one agent.

As before, these measures provide novel descriptive evidence on how individuals justify redistribution preferences revealing the "spontaneous" normative views they hold. Moreover, they allow us to test whether the historical information treatment reshapes justificatory frames. Specifically, whether exposure to historical narratives increases the salience of historical redress arguments or crowds out purely present-oriented or self-interested reasoning. Because the treatment may also change how much respondents write, which mechanically affects category counts and the sophistication index, we will report mean response length by arm and re-estimate all effects on LLM-derived measures controlling for response length. Finally, by linking the beliefs and justifications questions, we can examine the coherence between causal attributions and normative positions: for instance, whether respondents who attribute inequality to historical processes are more likely to invoke historical redress as a justification for redistribution, and whether the treatment strengthens this connection.

Appendix A.5 presents illustrative responses and Table A.7 the results from applying our classification scheme.

4.6 Construction of summary indices

Given the variety of outcomes to test some of the hypotheses, we follow the approach recommended by Anderson (2008) regarding the construction of summary indices. We define a single *primary index of preferences for inter-regional redistribution*, constructed as an equally weighted average of the standardized values of its three components: support for inter-regional transfers, the petition choice, and the incentivized lottery allocation. Standardization is performed within country, using the control group's mean and standard deviation. We prefer equal weighting to the inverse-covariance weighting of Anderson (2008) for two reasons. First, it is transparent: the index is interpretable directly from its components. Second, it is more appropriate here because inverse-covariance weighting assigns greater weight to components that are weakly correlated with the others. Our incentivized lottery allocation is the noisiest of the three measures and is likely to be the least correlated with them, so that procedure would weight it most heavily for reasons of statistical independence rather than of substantive importance. *C.H1* —*C.H3* are tested on the equally weighted index, although we also report the Anderson inverse-covariance index as a robustness check. The individual components are also reported for interpretation purposes, that is, to asses which of them drives or offsets the effect on the index, but no confirmatory claim rests on any one of them.

The mechanism outcomes are also structured into three indices, constructed in the same way. Within each index, components are oriented so that higher values correspond to the direction predicted by the associated hypothesis, so reverse-oriented items enter with the opposite sign. All variables are standardized within country on the control-group mean and SD.

The indices also incorporate measures derived from the LLM classification of open-ended responses on justifications for supporting or opposing redistribution. Rather than letting each LLM-coded measure enter as a separate component, we first average them into a single sub-index, so that the open-ended question does not dominate the index merely because more categories are coded from it. Each mechanism hypothesis is tested on its index. As with the primary index, the individual components are reported for interpretation purposes.

- *Attribution shift*: $\frac{1}{3}\times$ *Poor regions still pay the price* $+\frac{1}{3}\times$ *Work ethic* (reverse-oriented) $+\frac{1}{3}\times$ *LLM sub-index*. The latter is the average of *Agency attribution: poor regions themselves* (-) and *Historical redress*. Higher values indicate greater attribution of regional underdevelopment to historical circumstances rather than to the current population.
- *Moral justice*: $\frac{1}{3}\times$ *Empathy* $+\frac{1}{3}\times$ *Moral responsibility* $+\frac{1}{3}\times$ *LLM sub-index*. The latter is the average of *Solidarity and justice* and *Autonomy and self-interest* (-). Higher values indicate stronger empathy towards, and a stronger sense of moral responsibility for, the situation of poorer regions.
- *Fatalism*: $\frac{1}{4}\times$ *Future persistence* $+\frac{1}{4}\times$ *Pessimism* $+\frac{1}{4}\times$ *Policy effectiveness* (-) $+\frac{1}{4}\times$ *LLM sub-index*. The latter is the average of *Moral hazard and incentives* and *Distrust and misuse*. Higher values indicate a stronger belief that regional disparities are intractable, that is, that poorer regions will remain poor and that policy can do little to alter this.

4.7 Further methodological details

Missing values, null answers and outliers. We expect a very low incidence of missing values given that our survey questions have predefined scales and categories for the answers. Only a limited number of questions allow for null answers (e.g., "I do not know") or open-ended answers, which we may classify in a residual category as explained in Section 4.4. We expect null answers among a negligible proportion of responses and we will check that they are balanced across treatment and control groups. Moreover, the large sample size relative to the number of treatment arms gives us enough margin (in terms of required statistical power) even if the number of null answers is higher than expected. All our relevant outcomes have answers with predefined scales and categories, hence we thus do not expect any outliers.

Multiple hypothesis testing. We apply the false discovery rate (FDR) correction to the individual tests of outcomes of hypotheses belonging to pre-defined families (Anderson, 2008). These are: i) the three components of the primary redistribution index; ii) the components of each of the three mechanism indices, treated as three separate families; iii) the numbered heterogeneity hypotheses (*H.H1* to *H.H7*); and iv) the follow-up hypotheses (*F.H1* and *F.H2*), treated as a single family. Within each family we report sharpened q-values (Benjamini et al., 2006) as implemented by Anderson (2008), alongside unadjusted p-values.

Three elements fall outside FDR correction. The primary index is a single pre-specified test and is not corrected. We apply no correction across the three mechanism indices, since they correspond to distinct hypotheses with distinct predicted signs rather than to interchangeable

outcomes. The additional heterogeneity analysis (generic machine learning and country-level analyses) is no part of the confirmatory.

5 Empirical Analysis of Information Treatment

5.1 Overview of the analysis

Before presenting the estimating equations, we summarize here how the hypotheses developed in Section 3 map onto a hierarchy of tests. Figure 3 illustrates the analysis plan.

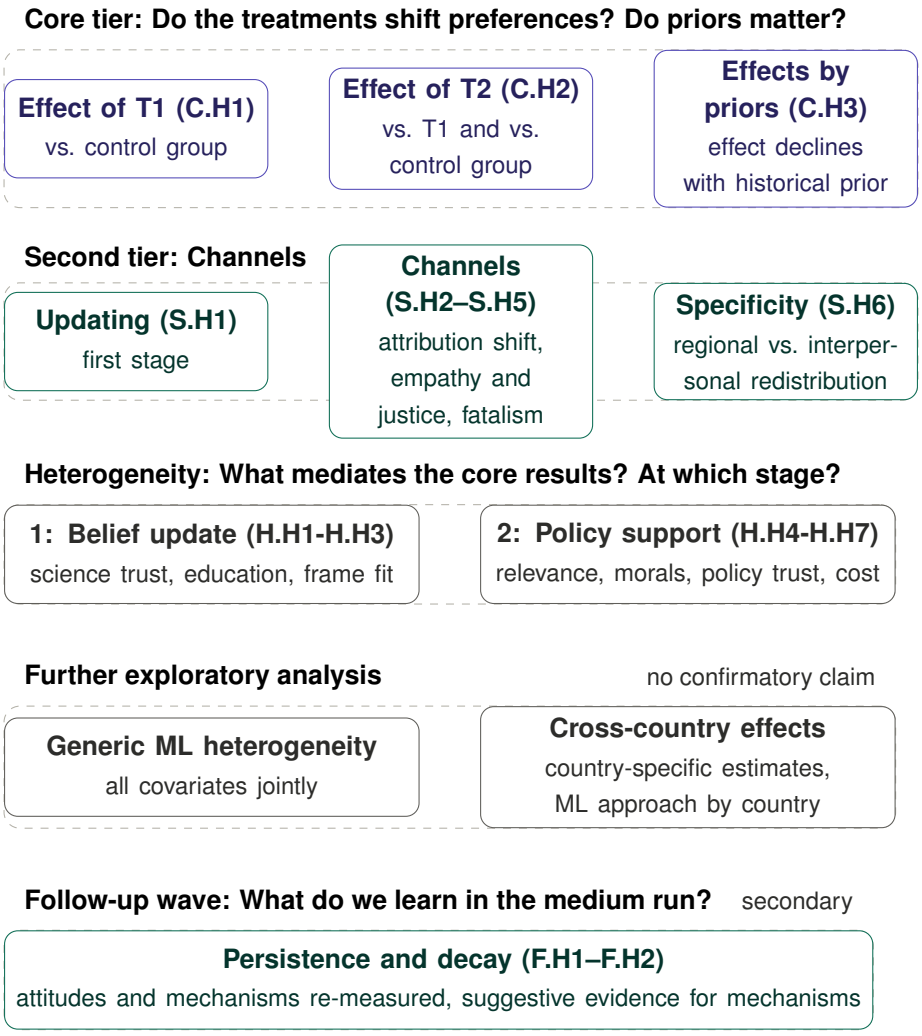

Figure 3: Hierarchy of the analysis plan

The plan rests on a small confirmatory core: the effect of the historical-legacies treatment on a single primary index of preferences for inter-regional redistribution (*C.H1*), the differential effect of the second treatment that adds the policy-agency extension (*C.H2*), and the moderation of the treatment effect by respondents' priors (*C.H3*). All three are tested on the same index of redistribution as confirmatory.[31]

[31] All the analyses that rely on indices as independent variables are also carried out separately for each of their components, treating them as a family outcomes for adjusting inference.

A secondary tier examines whether respondents update the belief the treatment targets (*S.H1*), the mechanisms through which any effect on redistribution preferences operates (attribution shift, empathy, fatalism: *S.H2–S.H5*), and whether the effect is specific to the regional dimension rather than reflecting a general shift in redistributive preferences (*S.H6*).

We consider a set of pre-specified heterogeneity hypotheses, organized around two stages of our theory of change: factors influecing of how much respondents may update (*H.H1–H.H3*) and factors influencing of how updated beliefs translate into political support (*H.H4–H.H7*). We complement these with two exploratory analyses. First, a data-driven machine-learning analysis that searches for heterogeneity across the full set of pre-treatment covariates jointly. Second, a cross-country analyses estimating country-specific effects and repeating the machine learning approach by country.

Finally, we examine the persistence and decay of the effects in the two-month follow-up (*F.H1–F.H2*). Section 5.2 states the corresponding estimating equations, and Table 3 lists each hypothesis with its predicted sign and the family within which multiple-inference correction is applied.

5.2 Estimating equations for baseline survey

We test our hypotheses using four specifications. Equation 1 is our baseline and is used for the two main hypotheses on preferences for inter-regional redistribution (*C.H1*, *C.H2*) and for the four mechanism hypotheses (*S.H2–S.H5*). Equation 2 adds the pre-treatment prior and its interactions with the treatments, and is used to test *C.H3*. Equation 3 tests the first stage on beliefs (*S.H1*), and Equation 4 tests the specificity of the treatment effect (*S.H6*).

We begin with Equation 1:

$$Y_{irc} = \alpha + \beta_1 T1_{irc} + \beta_2 T2_{irc} + X'_{irc}\Omega + \lambda_{rc} + \varepsilon_{irc} \tag{1}$$

where Y_{irc} is the outcome variable of individual i in region r in country c. $T1_{irc}$ and $T2_{irc}$ are dummies capturing exposure to *T1* and *T2* respectively. X_{irc} is a vector of control variables that includes age, woman, education, income, civil status, number of children, and urban area. λ_{rc} is a set of 93 state/region fixed effects for Brazil (26), Spain (17), and the US (50), and ε_{irc} denotes the error term. We estimate the equation by ordinary least squares (OLS), with standard errors clustered at the region level (Abadie et al., 2023). Importantly, β_1 and β_2 are average treatment effects which could differ drastically over the distribution of priors, as shown in other settings (Arias et al., 2022; Cox et al., 2024).

C.H3 states that the effect of *T1* declines with the respondent's prior belief about the share of current regional income differences attributable to events more than 100 years ago. It is estimated on the primary redistribution index as:

$$Y_{irc} = \alpha + \beta_1 T1_{irc} + \beta_2 T2_{irc} + \gamma \widetilde{P}_{irc} + \delta_1 (T1_{irc} \times \widetilde{P}_{irc}) + \delta_2 (T2_{irc} \times \widetilde{P}_{irc}) + X'_{irc}\Omega + \lambda_{rc} + \varepsilon_{irc} \tag{2}$$

where $\widetilde{P}_{irc}$ is the baseline cardinal prior, centered on its country-specific mean. The coeffi-

cient of interest is δ_1, predicted to be negative.[32] Centering implies that β_1 and β_2 are treatment effects evaluated at the country mean prior. These are close to, but not numerically identical to, the unconditional average effects estimated in Equation 1, and *C.H1* and *C.H2* are tested on the latter. The baseline prior is elicited with a forced-response slider, so Equation 2 is estimated on the same sample as Equation 1. We also report the marginal effect of each treatment across the observed range of the prior, $\beta_1 + \delta_1 \widetilde{P}$ with its confidence interval, together with the empirical distribution of the prior. We present descriptive evidence of the distribution of priors across the three countries.

S.H1 tests whether the treatment shifts the belief it targets. The outcome is the cardinal item elicited immediately after treatment (i.e., share of current regional disparities attributed to events happening more than 100 years ago), and we control for the same measure elicited at baseline:

$$P_{irc}^{post} = \alpha + \beta_1 T1_{irc} + \beta_2 T2_{irc} + \gamma P_{irc}^{pre} + X'_{irc}\Omega + \lambda_{rc} + \varepsilon_{irc} \tag{3}$$

with $\beta_1 > 0$ as the coefficient of interest. Conditioning on the baseline elicitation is what allows β_1 to be interpreted as belief updating rather than as a difference in post-treatment belief levels across arms: the estimate is the effect of the treatment on where a respondent ends up relative to where that respondent started. It also improves precision, as the two measurements of the same item are correlated. We will additionally report a version of this specification including the interactions between the treatments and the baseline prior, as in Equation 2, to assess whether updating is concentrated among respondents whose priors lie furthest below the content of the treatment. Since the prior is bounded above, respondents with high baseline values have limited scope to move upward irrespective of treatment, which mechanically favors a negative interaction. We therefore treat this coefficient as descriptive.

S.H6 compares the effect of *T1* on preferences for *inter-regional* redistribution with its effect on preferences for *inter-personal* redistribution. Both outcomes are standardized within country on the control group, and the two are stacked so that each respondent contributes two observations:

$$Y_{isrc} = \alpha + \beta_1 T1_{irc} + \beta_2 T2_{irc} + \theta R_s + \pi_1 (T1_{irc} \times R_s) + \pi_2 (T2_{irc} \times R_s) + X'_{irc}\Omega + \lambda_{rc} + \varepsilon_{isrc} \tag{4}$$

where s indexes the two outcomes and R_s equals one for the inter-regional index and zero for the inter-personal item. The coefficient of interest is π_1, predicted to be positive. Clustering at the region level accommodates the correlation between the two observations of a given respondent, since respondents are nested within regions.

[32] We also report a version in which the prior enters as dummies for quartiles of its distribution, computed on the pooled sample across the three countries and with the lowest quartile omitted. This allows the reader to assess whether the linear form adequately summarizes the relationship. Should the linear coefficient be indistinguishable from zero while the quartile estimates are monotonically ordered, we will report this as suggestive of non-linear heterogeneity rather than as support for *C.H3*..

The heterogeneity hypotheses *H.H1–H.H7* are tested on the primary redistribution index using Equation 5:

$$Y_{irc} = \alpha + \beta_1 T1_{irc} + \beta_2 T2_{irc} + \beta_3 IND_{irc} + \beta_4 (T1_{irc} \times IND_{irc}) + \beta_5 (T2_{irc} \times IND_{irc}) + X'_{irc}\Omega + \lambda_{rc} + \varepsilon_{irc} \tag{5}$$

where IND_{irc} is the dimension of interest, either a single variable or the simple average of two, and ranges between 0 and 1 for every moderator. The coefficients β_4 and β_5 capture, respectively for *T1* and *T2*, the average difference in treatment effects between individuals with the highest and lowest score on the relevant indicator.

Country-specific effects are estimated jointly through the following equation:

$$Y_{irc} = \sum_c \beta_{1c} \left(T1_{irc} \times C_c\right) + \sum_c \beta_{2c} \left(T2_{irc} \times C_c\right) + X'_{irc}\Omega_c + \lambda_{rc} + \varepsilon_{irc} \tag{6}$$

where C_c are country indicators and the controls are allowed to enter with country-specific coefficients, so that the point estimates coincide with those from three separate regressions. Country main effects are absorbed by the region fixed effects λ_{rc}, since regions are nested within countries. For each country we report β_{1c}, together with the contrast $\beta_{2c} - \beta_{1c}$ and its standard error.

Table 1 summarizes the mapping of each hypothesis to its corresponding outcome variables (as defined in Table A.1) and the coefficients of interest. The main hypotheses are evaluated on a single index of support for inter-regional redistribution, and each mechanism hypothesis on one of three indices. The individual components of each index, their predicted signs, and the families over which multiple-testing correction is applied are reported in Table A.2. These components are reported for interpretation purposes, that is, to understand what drives or offsets the effect of the treatments, and no confirmatory claim rests on any one of them.

Table 1 (and Table 2 and Table 3 below) serves as our pre-registration commitment: every hypothesis, outcome variable, and coefficient of interest is specified here before data collection begins. Exploratory analyses conducted beyond this pre-specified list will be clearly labeled as such in the results paper.

Table 1: Main treatment effects

Hypothesis	**Eq.**	**Coef. of interest**	**Outcome variables**
Main hypotheses: Preferences for inter-regional redistribution			
C.H1	(1)	$\beta_1 > 0$	Primary redistribution index: Support for regional redistribution; Petition pro-redistribution; Lottery

Table 1. Main treatment effects (continued)

Hypothesis	**Eq.**	**Coef. of interest**	**Outcome variables**
C.H2	(1)	$\beta_2 - \beta_1 > 0$	Primary redistribution index: Support for regional redistribution; Petition pro-redistribution; Lottery
C.H3 (Heterogeneity by priors)	(2)	$\delta_1 < 0$	Primary redistribution index: Support for regional redistribution; Petition pro-redistribution; Lottery
Secondary hypotheses: Belief updating			
S.H1 (First stage)	(3)	$\beta_1 > 0$	Share attributed to events more than 100 years ago (post-treatment)
Secondary hypotheses: Mechanisms			
S.H2 (Attribution shift)	(1)	$\beta_1 > 0$	Attribution index: Poor regions still pay the price; Work ethic (reverse-oriented); LLM sub-index[1]
S.H3 (Moral justice)	(1)	$\beta_1 > 0$	Moral justice index: Empathy; Moral responsibility; LLM sub-index[1]
S.H4 (Fatalism)	(1)	$\beta_1 > 0$	Fatalism index: Future persistence; Pessimism; Policy effectiveness (reverse-oriented)
S.H5 (Fatalism mitigation)	(1)	$\beta_2 - \beta_1 < 0$, $\beta_1 > \beta_2 > 0$	Fatalism index: Future persistence; Pessimism; Policy effectiveness (reverse-oriented); LLM sub-index[1]
Secondary hypotheses: Specificity			
S.H6 (Specificity)	(4)	$\pi_1 > 0$	Primary redistribution index vs. support for interpersonal redistribution

Notes: All indices are equally weighted averages of standardized components, oriented so that higher values correspond to the direction predicted by the associated hypothesis. Individual components are also reported, with multiple-testing correction applied within the families defined in Section 4.6. [1]The LLM sub-index aggregates measures derived from the classification of post-treatment open-ended responses justifying support or opposition to redistribution. Indices and variable definitions are available in Table A.1

Table 2: Heterogeneous treatment effects

Hypothesis	IND_{irc}	**Outcome variables**
Belief updating factors		

Table 2. Heterogeneous treatment effects (continued)

Hypothesis	IND_{irc}	**Outcome variables (effect sign)**
H.H1 (Credibility of the message)	Trust in scientists dummy	Primary redistribution index (+)
H.H2 (Cognitive capacity)	College education dummy; LLM sophistication-index	Primary redistribution index (+)
H.H3 (Narrative compatibility)	Dummy (1) if both "poor: lack of effort" and "rich: hard work" are equal to 0; LLM agency-attribution structural dummy	Primary redistribution index (+)
Policy action factors		
H.H4 (Personal resonance/ Perceived benefit)	Dummy for residing in bottom-quartile income region or family roots there	Primary redistribution index (+)
H.H5 (Moral predisposition)	Leftist ideology, Universalism	Primary redistribution index (+)
H.H6 (Perceived cost)	Dummy for residing in top-quartile income region or family roots there	Primary redistribution index (-)
H.H7 (Policy trust)	Trust in Gov, Interpersonal trust	Primary redistribution index (+)

Notes: The coefficient β_4 in equation 5 captures the differential effect of $T1$ for individuals with a given heterogeneity dimension ($IND = 1$ relative to $IND = 0$). The expected sign for each outcome is indicated in the table. Variable definitions in Table A.1.

5.3 Estimating equations for follow-up survey

The follow-up survey allows us to test the medium-term persistence and decay of the treatment effects, which is the object of hypotheses *F.H1* and *F.H2*. Concerning *F.H1* (persistence), the hypotheses *C.H1*, *C.H2*, *S.H2–S.H5* are re-estimated with Equation 1 and *S.H1* with Equation 3.

F.H2 (decay) is estimated with the two-wave analogue of the stacked specification we use for S.*H6*, on the balanced panel of respondents who complete both waves:

$$Y_{ircw} = \alpha + \beta_1 T1_{irc} + \beta_2 T2_{irc} + \gamma W_w + \delta_1 (T1_{irc} \times W_w) + \delta_2 (T2_{irc} \times W_w) + \mathbf{X}'_{irc}\theta + \lambda_{rc} + \varepsilon_{ircw}, \quad (7)$$

where w indexes the wave and W_w equals one for the follow-up. Both waves are standardized using the control-group mean and standard deviation of the main survey within each country, and standard errors are clustered at the region level. For the belief-updating outcome

(*S.H1*), the baseline priors elicitation is included in $\mathbf{X}_i$, so that the treatment effect is estimated conditional on the pre-treatment prior, as in the main survey. β_1 and β_2 are the immediate effects of the two arms relative to the control group, and $\beta_1 + \delta_1$ and $\beta_2 + \delta_2$ the effects two months later. Under *F.H2* each effect moves toward zero between waves, so δ_1 and δ_2 therefore carry the sign opposite to that of the corresponding immediate effect.

The incremental effect of *T2* over *T1* tested in *C.H2* is $\beta_2 - \beta_1$ immediately and $(\beta_2 + \delta_2) - (\beta_1 + \delta_1)$ at two months, so its decay is $\delta_2 - \delta_1$. This is the least precisely estimated quantity in the follow-up analysis because both coefficients are interaction terms, and we will interpret it accordingly.

The decay parameters are not directly comparable across outcomes whose immediate effects differ in size. To compare how far the first-stage effect and each mechanism shrink in relative terms, we also report the retained share of the effect, $(\beta_1 + \delta_1)/\beta_1$, where β_1 is distinguishable from zero. This allows us to asses, for example, whether the affective mechanisms fade faster than the cognitive ones. We report the retained shares alongside one another for that purpose, but as exploratory evidence and without a directional prediction.[33]

Table 3 summarizes, for each hypothesis re-estimated two months later, the outcome variable used and the coefficients that carry the persistence and decay predictions.

Table 3: Hypotheses testing in the follow-up survey

Hypothesis re-estimated	**Outcome variable in the follow-up**	***F.H1*** (persistence) Eq 1 and 3	***F.H2*** (decay) Eq 7
C.H1	Support for regional redistribution	$\beta_1 > 0$	$\delta_1 < 0$
C.H2	Support for regional redistribution	$(\beta_2 - \beta_1) > 0$	$\delta_2 - \delta_1 < 0$
S.H1 (First stage)	Share attributed to events more than 100 years ago	$\beta_1 > 0$	$\delta_1 < 0$
S.H2 (Attribution shift)	Attribution index	$\beta_1 > 0$	$\delta_1 < 0$
S.H3 (Moral justice)	Moral justice index	$\beta_1 > 0$	$\delta_1 < 0$
S.H4 (Fatalism)	Fatalism index	$\beta_1 > 0$	$\delta_1 < 0$
S.H5 (Fatalism mitigation)	Fatalism index	$(\beta_2 - \beta_1) < 0$	$\delta_2 - \delta_1 > 0$

[33] The decay of the effect may also vary by respondent characteristics, including their priors and regional identity. However interesting this analysis can be, these hypotheses require a "treatment" × "characteristic" × "wave" interaction estimated on the roughly 65% of respondents, and the resulting minimum detectable effects would probably be larger than any effect we could plausibly expect. We will report the estimates descriptively and label them as exploratory.

Table 3. Hypotheses testing in the follow-up survey (continued)

Hypothesis re-estimated	**Outcome variable in the follow-up**	***F.H1*** (persistence)	***F.H2*** (decay)

Notes: F.H1 is tested on the follow-up wave alone, using the same specifications as the corresponding main-survey hypotheses. *F.H2* is tested on the two-wave stacked specification in Equation 7, estimated on the balanced panel of respondents completing both waves, with standard errors clustered at the region level. β denotes the immediate effect and $\beta + \delta$ the effect two months later, so δ is the change between waves. The petition, the lottery and the open-ended question are one-off instruments and cannot be repeated, so the primary redistribution index is not available in the follow-up. The attribution, moral justice and fatalism indices are constructed without the LLM subindex. Variable definitions in Table A.1.

6 Interpreting Results

6.1 Descriptive Evidence on Beliefs about Regional Inequality and Redistribution

Before turning to the treatment effects, we present descriptive evidence on four related dimensions: (i) the distribution of priors on historical-attribution to contemporary regional inequality, (ii) how accurately respondents perceive historical regional inequality and long-term persistence, (iii) which beliefs they hold about the causes of regional inequality, and (iv) what justifications they provide for supporting or opposing inter-regional redistribution. This exploratory analysis sheds light on the mental framework of respondents and establishes the scope for belief updating through our informational intervention. Throughout this section, we analyze the priors on history, misspercceptions, and beliefs about specific causes using the full sample (as these questions precede treatment), while restricting the analysis of justifications to individuals from the control group to avoid contamination by treatment effects.[34]

6.1.1 Beliefs about the importance of historical events

We first present the distribution of the baseline measure of priors on historical-legacies, namely which share of current regional inequalities the respondent attributes to historical events rather than to current regional characteristics or contemporary factors. We do so separately by country, which allows to visualize how far beliefs can increase across the different populations. Similarly, we perform a descriptive exercise regressing these priors on a large set of variables, including socio-demogaphics, political and cultural variables, and the measures of misperceptions on historical inequality and persistence. This exercise will show how much of the variation in priors is explained by our baseline observables, as well as which covariates carry the highest explanatory power.[35]

[34] We examine treatment effects on justificatory frames in Section 6, where we explore how historical information may shape not only policy preferences but the reasoning underlying those preferences.

[35] We also apply other methods that could educate which variables to select over others such as Lasso regression.

6.1.2 Misperceptions of historical inequality and persistence

Prevalence of misperceptions. We first present the distribution of misperceptions regarding historical regional inequality by country, documenting the extent and direction of the gaps (i.e., whether respondents tend to over- or underestimate historical inequality and persistence). We expect substantial systematic misperceptions, consistent with evidence from the interpersonal inequality literature (Cruces et al., 2013).

Correlates of misperceptions. To identify the main determinants of misperceptions, we regress both measures on the full set of sociodemographic controls alongside variables capturing political and cultural values.[36] We expect education to be a major predictor, with more educated respondents exhibiting smaller misperceptions. We also anticipate that left-leaning respondents may overestimate historical inequality and persistence while right-leaning respondents underestimate them, consistent with motivated reasoning. Regional income may correlate negatively with inequality perceptions to the extent that citizens in poorer regions may perceive current inequality as larger.

Respondents who underestimate historical inequality and persistence are particularly relevant to the heterogeneity analysis, as the historical treatment is most likely to make inequality more salient and contribute update their beliefs. We use the misperception measures in the machine-learning approach for heterogeneity.

6.1.3 Beliefs about the causes of regional inequality

Prevalence of explanatory frames. First, we examine the share of respondents invoking each category (economic structure, human capital, infrastructure, politics and governance, public finance, culture, historical legacies, non-explanatory), separately by country. We anticipate that economic structure and politics/governance will be the most frequently mentioned frames across all three countries, while historical legacies will be a minority frame. This pattern would directly motivate the information treatment by establishing limited baseline historical consciousness. We also examine the distribution of the three composite measures of reasoning: agency attribution, grievance intensity, and sophistication.[37]

Individuals may combine categories in systematic ways that reveal coherent mental models of inequality. To explore this, we can present a heatmap of pairwise co-occurrence rates, revealing whether certain frames cluster into bundles (e.g., a "structural" bundle combining economic, human capital, and infrastructure explanations; a "political" bundle combining governance and public finance). The composite measures can also feature in this type of analysis.

Correlates of explanatory frames. To shed light on who reasons how, we will present regression results examining the correlates of each content category and composite measures. We will present tables regressing each variable on the sociodemographic, political and cultural variables, together with the measures of misperceptions. We expect the historical-legacies category to be positively correlated with education, left-wing ideology, structural attribution, and

[36] Controls include gender, age, education, regional income, left-right ideology, regional vs. national identity, and trust in government. See Table A.1 in the Appendix for complete definitions.

[37] In the Appendix, we can show these distributions according to key respondent characteristics (e.g., region income, education, etc.).

lower misperceptions on persistence, while negatively correlated with the belief that poverty reflects lack of individual effort and individual attribution.

6.1.4 Justifications for redistribution preferences

The second open-ended question asks respondents to explain why they support or oppose inter-regional redistribution. We restrict this analysis to the control group to capture baseline normative reasoning uncontaminated by treatment.

Prevalence of justificatory frames. As before, we can display the share of control-group respondents invoking each justification category and the three composite measures, separately by country. We anticipate that solidarity/justice to be the most common frames among supporters, while moral hazard and distrust to dominate among opponents. We expect historical redress (i.e., a explicit framing of redistribution as compensation for past wrongs) to be rare in the control group, paralleling the limited historical consciousness in beliefs.

Correlates of justification frames. We will examine the profile of individuals with different justification schemes, presenting regression results for the correlates of each justification category. Importantly, we can control or not for stated redistribution preferences. The latter allow us to examine whether, among respondents with similar policy views, certain characteristics predict different justificatory frames.

6.1.5 Coherence between beliefs and justifications

Finally, we examine the extent to which causal beliefs about inequality may align with normative justifications for redistribution. If respondents hold coherent worldviews, those who attribute inequality to historical processes should be more likely to justify redistribution as historical redress; those who emphasize governance failures may invoke distrust; those who emphasize culture may invoke moral hazard.

We can first present a heatmap of associations between belief categories and justification categories in the control group. We will highlight the key associations (historical beliefs $\leftrightarrow$ historical redress; governance beliefs $\leftrightarrow$ distrust; culture beliefs $\leftrightarrow$ moral hazard) and discuss any unexpected patterns.

We can test whether invoking a particular belief category predicts the corresponding justification category, controlling for respondent characteristics. If beliefs and justifications are linked, this suggests that interventions targeting causal beliefs (like our historical information treatment) may reshape normative reasoning as well.

6.2 Causal Effects of Information Treatments

6.2.1 Main results

If learning about the historical roots of regional inequality increases support for inter-regional redistribution, it will demonstrate that providing information about the origins of underdevelopment actively shapes contemporary policy preferences. This would imply that historical persistence studies have an important role to play in public narratives by shaping citizens'

economic attitudes. Furthermore, if the extended policy information treatment generates an additional positive effect on redistribution preferences, it will provide critical insights into how academic research should be disseminated. As argued by Nunn (2020), communicating findings from the historical legacies literature must leave room for human agency to figure out policies that address structural obstacles. The evidence from the information intervention would inform the design of science communication strategies.[38]

Conversely, finding precisely estimated null effects would be equally informative. A null result would suggest that preferences for inter-regional redistribution are deeply entrenched and highly inelastic to new information regarding the origins of inequality. This would align with recent evidence indicating that while factual information can successfully alter perceptions of an economic issue, it often fails to change policy preferences (e.g., Kuziemko et al. (2015)). Furthermore, a lack of movement in policy support might imply that citizens' redistributive motives are driven strictly by contemporary material self-interest or in-group bias, rather than by considerations of historical injustice.

It is also possible to observe a negative effect from the baseline historical treatment, which would mean a backlash effect driven by a sense of fatalism. If exposed individuals conclude that historical legacies are impossible to overcome, they may view redistributive policies as a futile waste of resources. In this scenario, comparing the baseline effect of historical narratives with the extended treatment, which includes the policy content, becomes highly relevant. If the effects of this bundled message are different, this will indicate that adding policy-agency message can offset the backlash generated by the historical narrative alone.

Beyond estimating average treatment effects, a central contribution of our project is to provide evidence on the mechanisms through which historical narratives operate. By evaluating shifts in the attribution of responsibility, empathy, and pessimism, we can identify the specific cognitive and emotional channels that translate factual information into policy preferences. Crucially, assessing these mechanisms provides actionable insights for the design of tailored research communication without unintentionally generating fatalistic backlash.

6.2.2 Heterogeneity

The heterogeneity analysis addresses pivotal questions regarding how best to communicate research to diverse audiences. For instance, exploring treatment effect variations based on trust in scientists or low trust in the government can shed light on the limits of academic dissemination in low-trust environments, informing strategies to reach more skeptical segments of the public. Understanding how historical narratives may influence differently individuals with a moral predisposition for solidarity, or who perceived themselves as net beneficiaries or contributors brings new insights to the political economy literature (Chinoy et al., 2026; Cappelen et al., 2025). Importantly, exploring heterogeneity with non-deterministic approaches (Chernozhukov et al., 2025) can reveal a variety of profiles and help us understand why certain groups react to research dissemination in a given way.

[38] We acknowledge that a more nuanced view on the effectiveness of policies, providing both examples of success and failures could generate different conclusions. Manipulating the precise content and framing of this type of interventions should be object of further investigation.

6.2.3 Country-specific results

While the information treatments aim to be as comparable as possible, the historical narratives for Brazil, Spain, and the United States differ in content, political sensitivity, and likely familiarity to participants. Similarly, these countries differ in many dimensions that could mediate the effects even of a homogeneous information treatment, such as their level of development, their institutional setting, and key socio-demographics (e.g., ethnicity, education) and their distribution across space. We investigate country-specific results and try to explain potential differences both with additional empirical evidence and contextual knowledge. Consistent effects across the three countries would support the claim that the mechanisms we identify are not specific to one national narrative alone and that they may hold in different settings. If we find divergent effects, our analysis adding partial interactions and the machine learning approach will help us understand to which extent these may be driven by compositional differences in observable characteristics across countries. Either finding, consistent or divergent effects, coupled with our heterogeneity analyses, will contribute to our understanding of when narratives of historical legacies may shape preferences over regional policy.

7 Concluding remarks

Regional inequality remains one of the most persistent and politically salient forms of economic disparity. While many countries operate systems of inter-regional transfers or place-based policies, public support for them is often contested and polarized. Understanding the beliefs and narratives that shape citizens' attitudes toward territorial redistribution is therefore essential for both their sustainability, political stability, and the economic development of poorer regions.

Our study provides the first experimental evidence on whether historical narratives can shift preferences for inter-regional redistribution, while also investigating the channels through which they may operate. We conduct a pre-registered experiment with 8,000 participants across the United States, Spain, and Brazil. We measure redistribution preferences using both self-reported attitudes and incentivized behavioral outcomes and the potential mechanism with a rich battery of measures capturing attribution shift and empathy responses. Given the risk that historical narratives may fuel fatalism and the belief that things are immutable, we examine whether complementing the information treatment with a policy agency message can counteract this effect. Beyond the experimental analysis, we provide novel descriptive evidence on how citizens perceive regional inequality, its persistence over time, and its ultimate causes, dimensions of economic beliefs so far largely unexplored. We use Large Language Models (LLMs) to analyze open-ended survey responses, revealing the causal narratives people invoke when reasoning about territorial disparities.

Our findings have strong implications for both policy making and science communication. In particular, they shed light on whether informing citizens about the historical legacies of regional inequalities strengthens or undermines public support for redistribution, and under which conditions such information may backfire.

Administrative information

Funding: This project has been supported by the Fundación Ramon Areces through the programm "XXIII Concurso Nacional de Ayudas a la Investigación en Ciencias Sociales [Project: "Historical legacies, preferences for redistribution and territorial cohesion"].

Institutional Review Board: The project has been approved by the University Pablo de Olavide Ethical Review. Reference 25/6-74.

Declaration of interest: None.

Acknowledgments: We thank participants at LISER's Working Group and the online workshop Junior Applied Microeconomists in Europe (JAMiE) for their comments.

References

Abad, L. A. and Maurer, N. (2021). History never really says goodbye: a critical review of the persistence literature. *Journal of Historical Political Economy*, 1(1):31–68.

Abadie, A., Athey, S., Imbens, G. W., and Wooldridge, J. M. (2023). When should you adjust standard errors for clustering? *The Quarterly Journal of Economics*, 138(1):1–35.

Acemoglu, D., Johnson, S., and Robinson, J. A. (2002). Reversal of fortune: Geography and institutions in the making of the modern world income distribution. *The Quarterly Journal of Economics*, 117(4):1231–1294.

Acemoglu, D. and Robinson, J. (2012). *Why nations fail: The origins of power, prosperity and poverty*. Profile, London.

Alesina, A. and Angeletos, G.-M. (2005). Fairness and redistribution. *American Economic Review*, 95(4):960–980.

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

Alesina, A. and Glaeser, E. L. (2004). *Fighting poverty in the US and Europe: A world of difference*. Oxford University Press.

Alesina, A. and La Ferrara, E. (2005). Preferences for redistribution in the land of opportunities. *Journal of public Economics*, 89(5-6):897–931.

Alesina, A., Miano, A., and Stantcheva, S. (2023). Immigration and redistribution. *The Review of Economic Studies*, 90(1):1–39.

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

Anderson, M. L. (2008). Multiple inference and gender differences in the effects of early intervention: A reevaluation of the abecedarian, perry preschool, and early training projects. *Journal of the American statistical Association*, 103(484):1481–1495.

Andre, P., Haaland, I., Roth, C., Wiederholt, M., and Wohlfart, J. (2024). Narratives about the macroeconomy. Working Paper 426, SAFE.

Arias, E., Larreguy, H., Marshall, J., and Querubin, P. (2022). Priors rule: When do malfeasance revelations help or hurt incumbent parties? *Journal of the European Economic Association*, 20(4):1433–1477.

Austin, B. A., Glaeser, E. L., and Summers, L. H. (2018). Jobs for the heartland: Place-based policies in 21st century america. Technical report, National Bureau of Economic Research.

Autor, D., Dorn, D., Hanson, G., and Majlesi, K. (2020). Importing political polarization? the electoral consequences of rising trade exposure. *American Economic Review*, 110(10):3139–83.

Balcells, L., Fernández-Albertos, J., and Kuo, A. (2015). Preferences for inter-regional redistribution. *Comparative Political Studies*, 48(10):1318–1351.

Banerjee, A. and Duflo, E. (2014). Under the thumb of history? political institutions and the scope for action. Working Paper 19848, National Bureau of Economic Research.

Bardhan, P. and Mookherjee, D. (2005). Decentralizing antipoverty program delivery in developing countries. *Journal of Public Economics*, 89(4):675–704. Cornell - ISPE Conference on Public Finance and Development.

Becker, S. O. and Woessmann, L. (2009). Was weber wrong? a human capital theory of protestant economic history. *The quarterly journal of economics*, 124(2):531–596.

Beltrán Tapia, F. J. and Martinez-Galarraga, J. (2018). Inequality and education in pre-industrial economies: Evidence from Spain. *Explorations in Economic History*, 69:81–101.

Bénabou, R., Falk, A., and Tirole, J. (2018). Narratives, imperatives, and moral persuasion. Working Paper 24798, National Bureau of Economic Research.

Bénabou, R. and Tirole, J. (2006). Belief in a just world and redistributive politics. *The Quarterly Journal of Economics*, 121(2):699–746.

Benjamini, Y., Krieger, A. M., and Yekutieli, D. (2006). Adaptive linear step-up procedures that control the false discovery rate. *Biometrika*, 93(3):491–507.

Beramendi, P. (2012). *The political geography of inequality: Regions and redistribution*. Cambridge University Press.

Berinsky, A. J., Margolis, M. F., and Sances, M. W. (2016). Can we turn shirkers into workers? *Journal of Experimental Social Psychology*, 66:20–28.

Boadway, R. and Shah, A. (2009). *Fiscal Federalism: Principles and Practice of Multiorder Governance*. Cambridge University Press, Cambridge.

Brollo, F., Nannicini, T., Perotti, R., and Tabellini, G. (2013). The political resource curse. *American Economic Review*, 103(5):1759–1796.

Bruhn, M. and Gallego, F. A. (2012). Good, bad, and ugly colonial activities: Do they matter for economic development? *The Review of Economics and Statistics*, 94(2):433–461.

Cappelen, A. W., Enke, B., and Tungodden, B. (2025). Universalism: Global evidence. *American Economic Review*, 115(1):44–77.

Chernozhukov, V., Demirer, M., Duflo, E., and Fernández-Val, I. (2025). Fisher–schultz lecture: Generic machine learning inference on heterogeneous treatment effects in randomized experiments, with an application to immunization in india. *Econometrica*, 93(4):1121–1164.

Chinoy, S., Nunn, N., Sequeira, S., and Stantcheva, S. (2026). Zero-sum thinking and the roots of us political differences. *American Economic Review*, 116(3):1052–1096.

Cirone, A. and Pepinsky, T. B. (2022). Historical persistence. *Annual Review of Political Science*, 25:241–259.

Conley, T. G. and Kelly, M. (2025). The standard errors of persistence. *Journal of International Economics*, 153:104027.

Cox, L., Eyzaguirre, S., Gallego, F. A., and García, M. (2024). Punishing mayors who fail the test: How do voters respond to information about educational outcomes? *Journal of Development Economics*, 171:103315.

Criscuolo, C., Martin, R., Overman, H. G., and Van Reenen, J. (2019). Some causal effects of an industrial policy. *American Economic Review*, 109(1):48–85.

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.

Daniele, G. and Geys, B. (2015a). Interpersonal trust and welfare state support. *European Journal of Political Economy*, 39:1–12.

Daniele, G. and Geys, B. (2015b). Organised crime, institutions and political quality: Empirical evidence from italian municipalities. *The Economic Journal*, 125(586):F233–F255.

Dell, M. (2010). The persistent effects of peru's mining mita. *Econometrica*, 78(6):1863–1903.

Ehrlich, M. V. and Overman, H. G. (2020). Place-based policies and spatial disparities across european cities. *Journal of Economic Perspectives*, 34(3):128–149.

Engerman, S. L. and Sokoloff, K. L. (2002). Factor endowments, inequality, and paths of development among new world economies. *Economía*, 3(1):41–88.

Enke, B., Rodríguez-Padilla, R., and Zimmermann, F. (2023). Moral universalism and the structure of ideology. *The Review of economic studies*, 90(4):1934–1962.

Esposito, E., Rotesi, T., Saia, A., and Thoenig, M. (2023). Reconciliation narratives. *The American Economic Review*, 113(6):1461–1504.

Faia, E., Fuster, A., Pezone, V., and Zafar, B. (2024). Biases in information selection and processing: Survey evidence from the pandemic. *Review of Economics and Statistics*, 106(3):829–847.

Ferrario, B. and Stantcheva, S. (2022). Eliciting people's first-order concerns: Text analysis of open-ended survey questions. In *AEA papers and proceedings*, volume 112, pages 163–169. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Floerkemeier, H., Spatafora, N., and Venables, A. (2021). Regional disparities, growth, and inclusiveness. Technical report, International Monetary Fund.

Fong, C. (2001). Social preferences, self-interest, and the demand for redistribution. *Journal of Public economics*, 82(2):225–246.

Galor, O. (2022). The journey of humanity: The origins of wealth and inequality. ny, dutton.

Gniza, J. and Wrede, M. (2024). Public acceptance of regional redistribution in germany: A survey experiment on the perceived deservingness of regions. *Journal of Social Policy*, 53(4):1052–1072.

Graeber, T., Roth, C., and Zimmermann, F. (2024). Stories, statistics, and memory. *The Quarterly Journal of Economics*, 139(4):2181–2225.

Haaland, I., Roth, C., Stantcheva, S., and Wohlfart, J. (2024). Understanding economic behavior using open-ended survey data. (32421).

Haaland, I., Roth, C., and Wohlfart, J. (2023). Designing information provision experiments. *Journal of Economic Literature*, 61(1):3–40.

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.

Instituto Nacional de Estadística (2025). Encuesta de condiciones de vida. https://ine.es/dyngs/INEbase/en/operacion.htm?c=Estadistica_C&cid=1254736176807&menu=resultados&idp=1254735976608. Accessed 14 August 2026.

International Monetary Fund (2019). Closer together or further apart? subnational regional disparities and adjustment in advanced economies. In *World Economic Outlook*, chapter 2.

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.

Kline, P. and Moretti, E. (2014). Local economic development, agglomeration economies, and the big push: 100 years of evidence from the tennessee valley authority. *The Quarterly journal of economics*, 129(1):275–331.

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.

Lee, D. S. (2005). Training, wages, and sample selection: Estimating sharp bounds on treatment effects.

Lu, Y., Wang, J., and Zhu, L. (2019). Place-based policies, creation, and agglomeration economies: Evidence from china's economic zone program. *American Economic Journal: Economic Policy*, 11(3):325–60.

Meade, A. W. and Craig, S. B. (2012). Identifying careless responses in survey data. *Psychological methods*, 17(3):437.

Musacchio, A., Fritscher, A. M., and Viarengo, M. (2014). Colonial institutions, trade shocks, and the diffusion of elementary education in brazil, 1889–1930. *The Journal of Economic History*, 74:730–766.

Naritomi, J., Soares, R. R., and Assunção, J. J. (2012). Institutional development and colonial heritage within brazil. *The journal of economic history*, 72(2):393–422.

Neumark, D. and Simpson, H. (2015). Place-based policies. In *Handbook of regional and urban economics*, volume 5, pages 1197–1287. Elsevier.

Nunn, N. (2020). The historical roots of economic development. *Science*, 367(6485):eaaz9986.

Nunn, N. and Wantchekon, L. (2011). The slave trade and the origins of mistrust in africa. *American Economic Review*, 101(7):3221–3252.

Oates, W. E. (1999). An essay on fiscal federalism. *Journal of Economic Literature*, 37(3):1120–1149.

Oto-Peralías, D. and Romero-Ávila, D. (2016). The economic consequences of the Spanish Reconquest: The long-term effects of Medieval conquest and colonization. *Journal of Economic Growth*, 21(4):409–464.

Piketty, T. (1995). Social mobility and redistributive politics. *The Quarterly Journal of Economics*, 110(3):551–584.

Ramos-Toro, D. and Voytas, E. (2024). Historical narratives and political behavior in the us. Working paper.

Rodden, J. A. (2019). *Why cities lose: The deep roots of the urban-rural political divide*. Hachette UK.

Rodríguez-Pose, A., Terrero-Dávila, J., and Lee, N. (2023). Left-behind versus unequal places: interpersonal inequality, economic decline and the rise of populism in the usa and europe. *Journal of Economic Geography*, 23(5):951–977.

Roodman, D., Nielsen, M. Ø., MacKinnon, J. G., and Webb, M. D. (2019). Fast and wild: Bootstrap inference in stata using boottest. *The Stata Journal*, 19(1):4–60.

Shayo, M. (2009). A model of social identity with an application to political economy: Nation, class, and redistribution. *American Political science review*, 103(2):147–174.

Shiller, R. J. (2020). Narrative economics: How stories go viral and drive major economic events.

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

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

Stantcheva, S. (2023). How to run surveys: A guide to creating your own identifying variation and revealing the invisible. *Annual Review of Economics*, 15(1):205–234.

Voth, H.-J. (2021). Persistence–myth and mystery. In *The handbook of historical economics*, pages 243–267. Elsevier.

Wigton-Jones, E. (2020). Legacies of inequality: the case of brazil. *Journal of Economic Growth*, 25(4):455–501.

Wuthnow, R. (2019). *The left behind: Decline and rage in small-town America*.

Appendix

Historical Legacies and Preferences for Inter-Regional Redistribution

Martín Fernández & Daniel Oto-Peralías

Table of Contents

A Additional information

A.1 List of variables

Table A.1: List of variables

Variable name	Description and operationalization
Composite indices	
Primary redistribution index	Equally weighted averages of components standardized within country on the control-group mean and SD of the variables *support for regional redistribution*, *petition pro-redistribution* and *lottery*.
Attribution shift index	$\frac{1}{3}\times$ *Poor regions still pay the price* $+\frac{1}{3}\times$ *Work ethic* (-) $+\frac{1}{3}\times$ *LLM sub-index*. The latter is the average of *Agency attribution: poor regions themselves* (-) and *Historical redress*. All variables standardized within country on the control-group mean and SD. For the follow-up survey, the index only contains the first two components.
Moral justice index	$\frac{1}{3}\times$ *Empathy* $+\frac{1}{3}\times$ *Moral responsibility* $+\frac{1}{3}\times$ *LLM sub-index*. The latter is the average of *Solidarity and justice* and *Autonomy and self-interest* (-). All variables standardized within country on the control-group mean and SD. For the follow-up survey, the index only contains the first two components.
Fatalism index	$\frac{1}{4}\times$ *Future persistence* $+\frac{1}{4}\times$ *Pessimism* $+\frac{1}{4}\times$ *Policy effectiveness* (-) $+\frac{1}{4}\times$ *LLM sub-index*. The latter is the average of *Moral hazard and incentives* and *Distrust and misuse*. All variables standardized within country on the control-group mean and SD. For the follow-up survey, the index only contains the first two components.
Outcome variables	
Support for regional redistribution	Dummy (1) if respondent supports or strongly supports that rich regions should transfer more resources to poor regions to help finance their public services.
Reasons to support/oppose redistribution	LLM-generated categorizations of open-ended responses explaining the respondent's reasons to support or oppose regional redistribution, including non-mutually exclusive topic indicators (solidarity and justice, national unity, historical redress, autonomy and self-interest, moral hazard, distrust and misuse) and composite measures: agency-attribution dummies (state, rich regions, citizens collectively, poor regions themselves), grievance intensity, and sophistication.

Table A.1. List of variables (continued)

Variable name	*Description and operationalization*
Share attributed to history	Indicator capturing the share (0-100) of current regional income differences attributed to events more than 100 years ago. Elicited at baseline, post-treatment, and in the follow-up.
Future persistence	Dummy (1) if respondent thinks it is very likely that the poorest regions will still be the poorest in 30 years.
Petition pro-redistribution	Dummy variable capturing support for a petition to the national government in favor of regional redistribution.
Empathy, moral responsibility, pessimism	Set of three dummy variables capturing whether participant feels somewhat or a lot of empathy, moral responsibility, and pessimism when thinking about poorer regions.
Work ethic	Dummy (1) if respondent agrees or strongly agrees that poor regions are poorer due to the work attitudes and initiative of their population.
Poor regions still pay the price	Dummy (1) if respondent agrees or strongly agrees that poor regions suffer from disadvantages created long ago.
Policy effectiveness	Dummy (1) if respondent answers ”a great deal” or ”a lot” that well-designed economic policies can promote the development of left-behind areas.
Lottery	Difference between the amount allocated to a poor participant from the poorest region and to a poor participant from the respondent’s region. For respondents residing in the poorest region, the former refers to another poor region.
Support for inter-personal redistribution	Respondent’s view on whether the government should reduce income differences between rich and poor people, on a 0-10 scale.
Independent variables	
T1	Dummy (1) if exposed to information treatment 1 (research summary video).
T2	Dummy (1) if exposed to treatment 2 (research video + policy information extension).
Pure controls	

Table A.1. List of variables (continued)

Variable name	*Description and operationalization*
Age	Respondents' age.
Woman	Dummy (1) if respondent is female.
State/region of residence	Set of dummy variables (26 for Brazil, 17 for Spain, and 50 for the US) for state or region of residence of the respondent.
Education	Set of 4 dummy variables for highest level of studies completed: primary or less, lower secondary, high school or vocational degree, college degree or above .
Income	Average monthly household income per adult (interval class mark). For the top, open-ended class, the upper bound is set at 1.5 times the lower bound.
Civil status	Set of 4 dummy variables: single, married/*de facto*, separated/divorced, or widowed.
Children	Number of children.
Household size	Number of people living in the household.
Type of area	Set of 5 dummy variables based on settlement population size ($< 5k$ to $> 250k$).
Employment status	Set of 7 dummy variables for current employment category.
Controls for heterogeneity analysis	
Strong regional identity	Dummy (1) if respondent identifies more with their region than their nation.
Leftist Ideology	Dummy (1) if respondent has left-wing ideology (< 5 on 0–10 scale).
Family roots in poor regions	Dummy (1) if parents or ≥ 2 grandparents were born in a bottom-quartile income region.
Interpersonal trust	Dummy (1) if average of two questions of interpersonal trust > 50.
Trust: Gov/Scientists	Two dummy variables for high trust in government and scientists.
Universalism	Trust gap between local neighbors and residents in other regions (0–100 scale)

Table A.1. List of variables (continued)

Variable name	*Description and operationalization*
Zero-sum mindset	Average of three questions regarding zero-sum thinking (1–5 scale).
Poor: lack of effort	Dummy (1) if respondent blames poverty on lack of personal effort.
Rich: hard work	Dummy (1) if respondent attributes wealth to hard work and effort.
Inferred persistence: strong	Dummy (1) if respondent believes poorest/richest regions are the same now as in 1900.
Historical inequality misperception	$\lvert r_{1900}^{actual} - r_{1900}^{belief} \rvert$, where r_{1900}^{actual} denotes the actual ratio between trhe richest and poorest region in 1900 and r_{1900}^{belief} the ratio reports by the respondent. Standardized to 0-1.
Long-term persistence misperception	$\lvert c_{1900-2025}^{actual} - c_{1900-2025}^{actual} \rvert$, where $c_{1900-2025}^{actual}$ is the correlation between GDP per capita levels in 1900 and 2025 across regions and $c_{1900-2025}^{actual}$ the correlation reported by the respondent - Standardized to 0-1.
Reasons of regional inequality	LLM-generated categorizations of open-ended responses explaining the respondent's reasons for regional income differences, including non-mutually exclusive topic indicators (economic structure, human capital, infrastructure and geography, politics and institutions, public finance and redistribution, culture and attitudes, and historical legacies) and composite measures (agency attribution, grievance intensity, and sophistication).
Placebo outcomes and quality-check variables	
AI job effect	Expected effect of artificial intelligence on future employment, from 1 (very possitive) to 5 (very negative).
Climate change	Extent to which climate change is already having a negative economic impact in the respondent's country, from 1 (not at all) to 5 (a great deal)
Honesty check	Dummy (1) if respondent confirms having devoted full attention and answered truthfully.
Instructional manipulation check	Dummy (1) if respondent selects exactly the two options indicated in the instruction ("a little" and "a lot").
Video comprehension check	Dummy (1) if respondent selects the factually correct statement about the video's content and no incorrect statement.

Table A.1. List of variables (continued)

Variable name	*Description and operationalization*
Self-reported completion	Dummy (1) if respondent watched and listened to the video until the end.
Video link requested	Dummy (1) if respondent asks to receive a link to the video by email after the survey.
Persuaded	Dummy (1) if respondent is very persuaded by the video's arguments.
Time on video page	Time (in seconds) spent on the video screen, recorded as paradata. Used to identify respondents who advance before the video ends and to measure break-off at the video screen.
Careless responding	Dummy (1) if respondent is flagged by the Response Pattern Index or provides nonsensical or very short answers to the open-ended questions.
Perceived bias	Dummy (1) if respondent reports feeling that the survey was biased in some sense.
Perceived bias index	Index from -2 to 2 capturing the extent to which the respondent read the survey as promoting redistribution toward poorer regions or on the contrary, favoring rich regions and less redistribution. Respondents who report no bias take the value 0. Open-ended responses from those who report bias are classified by a LLM using the prompt in Appendix Section A.7

Table A.2: Index components and predicted signs

Component	***T1*** β_1	***T2–T1*** $\beta_2 - \beta_1$
Panel A. Primary redistribution index (*C.H1*, *C.H2*, *C.H3*)		
Support for regional redistribution	+	+
Petition pro-redistribution	+	+
Lottery	+	+
Panel B. Attribution index (*S.H2*)		
Poor regions still pay the price	+	
Work ethic	−	
LLM sub-index: *Agency attribution: poor regions themselves* (-) and *Historical redress*	+	
Panel C. Moral justice index (*S.H3*)		
Empathy	+	
Moral responsibility	+	
LLM sub-index: *Solidarity and justice* and *Autonomy and self-interest* (-)	+	
Panel D. Fatalism index (*S.H4*, *S.H5*)		
Future persistence	+	−
Pessimism	+	−
Policy effectiveness	−	+
LLM sub-index: *Moral hazard and incentives* and *Distrust and misuse*.	+	−

Table A.2. Index components and predicted signs (continued)

Component	***T1***	***T2–T1***
	β_1	$\beta_2 - \beta_1$

Notes: Components are the individual items entering each index, standardized within country on the control-group mean and standard deviation. *Work ethic* and *Policy effectiveness* enter their indices with a negative sign, so that all components point in the direction predicted by the hypothesis. The predicted signs reported here are stated for interpretation and do not constitute confirmatory tests: the confirmatory tests are conducted on the indices, as specified in Table 1. Each panel defines a family for multiple-testing correction, within which we report sharpened q-values (Benjamini et al., 2006) as implemented by Anderson (2008), alongside unadjusted p-values. Blank cells for $\beta_2 - \beta_1$ indicate that no directional prediction is made. Variable definitions in Table A.1

A.2 Additional figures

Figure A.1: Long-term persistence in income per capita, 1900-2020

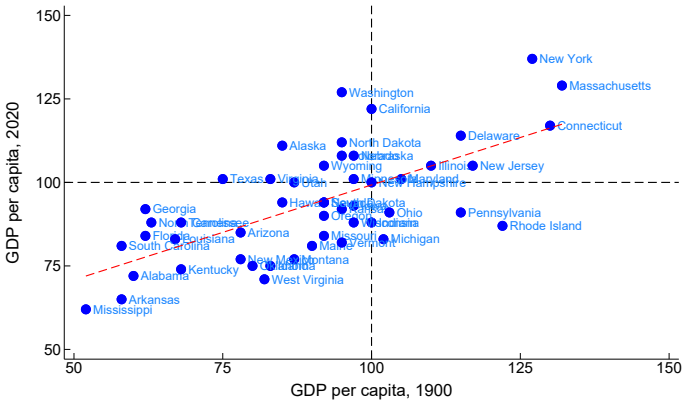

(a) United States

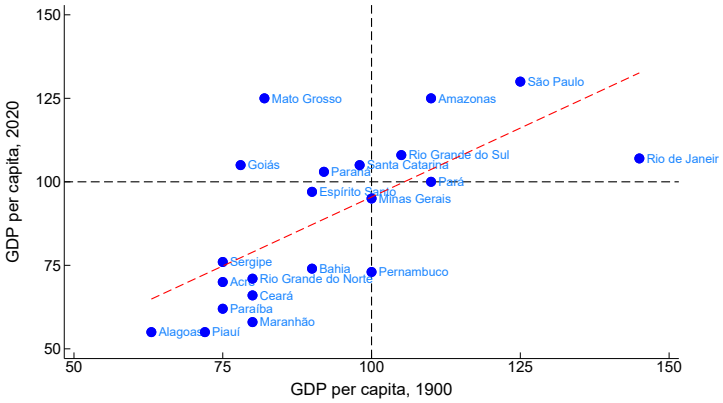

(b) Brazil

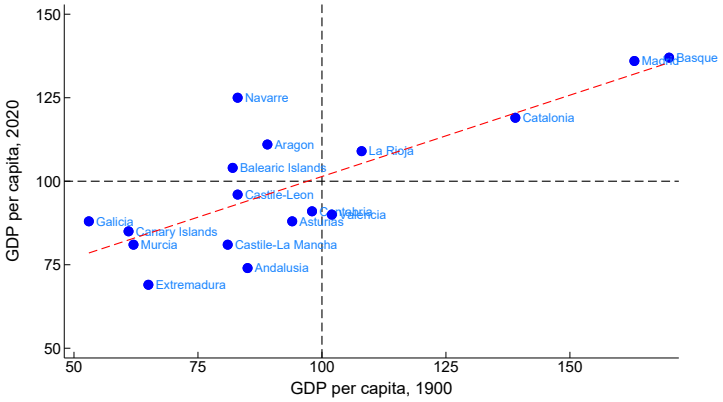

(c) Spain

Notes: Sources for GDP per capita in 1900 are Easterlin (1960) for the US, Buescu (1981) for Brazil, Diez-Minguela et al. (2018) for Spain. For the year 2020, the US Bureau of Economic Analysis (BEA), the Brazilian Institute of Geography and Statistics (IBGE), and the Spanish National Institute of Statistics (INE).

A.3 Power calculations and benchmarking

This section presents the main power calculations for our study and a comparison to other studies doing information interventions on redistribution or other political attitudes. For the sake of the example, let's consider our key binary measure of support for inter-regional redistribution. Based on existing survey evidence, let's assume a conservative baseline mean of 0.25 and hence standard deviation 0.433 ($\sqrt{p(1-p)}$). Following Haaland et al. (2023), we consider the ability to detect treatment effects expressed as fractions of the standard deviation. Hence,in our setting a treatment effect equal to 10% of a standard deviations (SD) corresponds to a change of approximately 0.043 in probability terms (4.3 percentage points).

Table A.3 below reports what would be the minimum detectable effect (MDE) for different sample sizes, assuming a conventional significance level of 5% and statistical power of 80%, and three equal-size groups per arm (two treatments and one control) as in our study.

Our planned sample of 8,000 respondents, approximately 2,666 per arm, presents a MDE for our binary outcomes of approximately 7.7% SD using the pooled sample. Our sample would imply a per-country MDE of approximately 13.3% SD. By combining several outcomes into an index, one can reduce the MDE by a factor directly related to the correlation between the components. For example, an index that is the equally weighted mean of k variable with average pairwise correlation ρ, the variance of the index is [1 + (k - 1)ρ]/k, so with k = 3 and ρ between 0.3 and 0.5 the pooled MDE would fall to approximately 5.6%-0.6.3% SD, in other words, 18%-27% lower than the MDE of 0.077 reported for specific binary outcomes. The per-country MDE would fall to approximately 9.7%-10.9 SD.%

Table A.3: Minimum Detectable Effect (MDE) by total sample size

Total N	n per arm	MDE (pp)	MDE (% baseline)	MDE (% std dev)
2,000	666	0.0665	26.59	15.35
4,000	1,333	0.0470	18.80	10.85
6,000	2,000	0.0384	15.34	8.86
8,000	2,666	0.0332	13.29	7.67
10,000	3,333	0.0297	11.89	6.86

Note: Calculations consider a three-arm design with equal size ($n = N/3$ per arm), $\alpha = 0.05$ (two-tailed), power $= 0.80$, and baseline mean $p_0 = 0.25$ with standard deviation $\sigma = 0.43$.

Our MDE of 7.7% SD for specific outcomes in the pooled sample (and 13.3% per-country) is in the range of similar studies in the literature as shown in Table A.4. The average across 10 related studies is 11% for the pooled sample and 13% for country-specific effects. The estimated effects in these papers tend to range between 10%-20% SD, well above our MDE. This comparison shows that our sample size is well suited to detect even small effects and carry out cross-country comparisons and heterogeneity analyses.

Table A.4: Treatment effects in comparable information-provision experiments

Paper	Journal	Topic	Outcome	Treatment	Countries	Sample	Arms	MDE pooled (SD)	MDE country (SD)	Main effect (SD)
This paper	—	Regional redistribution	Self-reported + lottery + petition	Video (3–4 min)	3	8,000	3	0.077	0.133	N/A
Kuziemko et al. (2015)	AER	Redistribution / taxation	Support raising tax on millionaires	Interactive text	1	4,000	2	0.089	0.089	≈0.11
Alesina et al. (2018)	AER	Mobility / redistribution	Support for equality of opportunity policies	Text + Animation	5	12,663	2	0.056	0.125	≈0.06
Hoy and Mager (2021)	AEJ:EP	Inequality / redistribution	Concern about inequality (gap too large)	Text (income position)	10	30,000	2	0.032	0.102	≈0.08
Stantcheva (2021)	QJE	Taxation	Support for tax policy index	Video	1	5,140	4	0.15	0.15	≈0.20
Alesina et al. (2023)	ReStud	Immigration / redistribution	Support for redistribution	Text + video	6	22,006	4	0.053	0.129	≈0.10
Cruces et al. (2013)	JPubE	Inequality / redistribution	Support for redistribution	Text (income rank)	1	1,100	2	0.168	0.168	≈0.15
Karadja et al. (2017)	REStat	Inequality / redistribution	Demand for redistribution (index)	Text (income rank)	1	≈1,000	2	0.177	0.177	≈0.10
Sands and de Kadt (2020)	Nature	Inequality / redistribution	Petition to tax wealthy (field)	Field (exposure to car)	1	1,489	2	0.145	0.145	≈0.22
Lergetporer et al. (2020)	JPubE	Education / redistribution	Support for education policies (index)	Text + graph	1	7,380	2	0.065	0.065	≈0.05
Grigorieff et al. (2020)	Demography	Immigration attitudes	Attitudes toward immigrants (index)	Text (5 statistics)	1	≈2,000	2	0.125	0.125	≈0.10

A.4 Classification of beliefs about regional inequality

This appendix provides detailed documentation of our framework for classifying open-ended responses about the causes of regional inequality. We present: (1) a summary of content categories, (2) detailed definitions of composite measures, (3) illustrative example responses, (4) the classification of these examples, and (5) the LLM prompts used for automated coding.

A.4.1 Summary of content categories

Table A.5 summarizes the non-mutually exclusive content categories used to classify beliefs about the causes of regional inequality.

Table A.5: Content categories for beliefs about regional inequality

Category	Definition and examples
Economic structure	References to labor market conditions, unemployment, sectoral composition (e.g., tourism, agriculture, industry), productivity, firm location, or agglomeration effects (e.g., "big cities attract more companies," "tourism drives the coastal economy," "factories closed and nothing replaced them").
Human capital & demography	References to education, skills, innovation, youth out-migration, brain drain, or population aging (e.g., "young people leave for opportunities elsewhere," "better universities produce qualified workers," "brain drain hurts poorer regions").
Infrastructure & geography	References to transportation, connectivity, rurality, peripherality, access to services, or geographic location (e.g., "regions near ports are richer," "no highways or trains here," "everything is far away").
Politics, governance & institutions	References to quality of regional government, corruption, mismanagement, centralization, favoritism, or policy choices (e.g., "corrupt politicians steal everything," "bad management for decades," "the capital gets all the attention").
Public finance & redistribution	References to allocation of public funds, subsidies, transfers, or perceived fiscal imbalances (e.g., "some regions get more money per person," "the financing system is unfair," "EU funds don't reach us").
Culture & attitudes	References to work ethic, entrepreneurship, mentality, or effort as drivers of regional differences (e.g., "some people work harder," "entrepreneurial culture matters," "attitudes toward risk-taking differ").
Historical & long-run legacies	Direct references to the past, historical events, long-run development paths, or temporal language extending beyond the present (e.g., "regions that industrialized early had advantages," "decisions made decades ago still matter," "colonial history shaped development").
Non-explanatory	Expressions of uncertainty, generic statements, or circular reasoning offering no substantive causal account (e.g., "I don't know," "many factors," "it's complicated," "some places are just richer").
Other	Residual category for substantive explanations not fitting above categories. When coded, a brief description is recorded.

A.4.2 Additional measures: agency, grievance, sophistication

Beyond content classification, we construct additional measures that clarify how respondents reason about inequality. First, we code *agency attribution*, recording whether inequality is attributed to some of these agents: (i) structural factors (e.g., markets, geography, institutions),

(ii) external factors (e.g., central authorities or outside elites), and (iii) internal/regional factors (e.g., local governance, culture, effort, or choices made by regional actors). Every agent mentioned is coded as a separate dummy, and an auxiliary variable flags respondents mentioning more than one agent. We do not classify a single 'dominant' agent, because there is no clear way to establish comprehensive rules (e.g., an agent could be presented as critical but receive less attention in the explanation than a second one). Moreover, separate dummies preserve what the respondent actually said and remain interpretable.[39] Second, we build a *grievance intensity* indicator that flags responses employing moralized language that frames inequality as unjust, unfair, or characterized by exploitation and zero-sum competition. The goal of this variable is to capture how emotionally load the explanations are. Third, we construct a *sophistication index* based on the average of several components: (i) response length (standardized), (ii) a dummy for invoking more than one reason, (iii) a dummy for the presence of explicit causal reasoning (indicated by using connectives such as 'because of this,' 'as a consequence of this,' or 'this has led to'), and (iv) a dummy for complexity and nuance (acknowledging complexity, interactions and dependencies.)

Agency attribution. We classify whom respondents implicitly or explicitly hold responsible for regional inequality with the following dummies:

- *Structural (ST):* Inequality framed as the result of impersonal forces—markets, geography, economic dynamics, agglomeration, or path dependence.
- *External (EX):* Inequality attributed to actors outside the region—central government, national elites, other regions, or EU institutions.
- *Individual/Regional (IN):* Inequality attributed to factors within the region's control—local governance, culture, attitudes, effort, or choices made by regional actors.
- *Unclear (U):* No clear attribution of agency in the response.

This measure captures variation in how respondents conceptualize responsibility for regional disparities, which may predict attitudes toward redistribution.

Grievance intensity. A binary indicator (0/1) that flags responses employing moralized language framing inequality as unjust, unfair, exploitative, or characterized by zero-sum competition. We code 1 when responses contain:

- Explicit claims of injustice, exploitation, or victimization (e.g., "they take everything," "we are forgotten," "it's theft")
- Strong moral condemnation of current arrangements (e.g., "it's outrageous," "completely unfair," "a scandal")
- Zero-sum framing of inter-regional relations (e.g., "their gain is our loss," "they benefit at our expense")

[39] As a validation check, we will report the correlation between each agency-attribution dummy and the content categories to which it relates most closely. Large discrepancy would be reported and investigated.

This measure captures the emotional and moral intensity of beliefs, which may predict policy attitudes and responsiveness to framing.

Sophistication index. An index based on the average of several dimensions of response quality and depth:

- *Length:* Word count (which will be standardized).
- *Multiple causes:* dummy for invoking more than one justification.
- *Causal reasoning:* Presence of explicit logical connectives (e.g., "because," "therefore," "as a result," "this leads to")
- *Nuance:* Acknowledgment of complexity or interactions between factors.

A.4.3 Illustrative responses

The following fictional responses illustrate the range of explanations participants might provide, varying in length, analytical depth, and substantive content. Demographic characteristics are included to show the diversity of the sample.

Response 1 (Female, 28, College, Center-left, Above-average region)

> Honestly, I think it has a lot to do with investment in infrastructure and education. Regions like Madrid and Catalonia have better universities, more companies, and attract talented people from everywhere. Also, tourism plays a huge role—coastal areas and cities with cultural heritage get more money. Poorer regions maybe didn't diversify their economy enough when industries like mining or agriculture declined. Politicians also matter—some regional governments are just better at attracting investment and managing resources than others.

Response 2 (Male, 52, High school, Center-right, Below-average region)

> Because Madrid gets everything and the rest of us get nothing. It's that simple. All the big companies have headquarters there, all the government money goes there. We had factories here that closed and nobody cared. Now young people leave because there are no jobs.

Response 3 (Female, 41, Postgraduate, Extreme left, Average region)

> The unequal distribution of wealth reflects decades of neoliberal policies that have concentrated resources in major urban centers while abandoning rural and peripheral regions. The lack of industrial policy, the precariousness of labor, and austerity measures have devastated regions that depended on traditional industries. Additionally, the constitutional framework created autonomous communities with very different fiscal capacities, perpetuating structural inequalities. The right has historically favored Madrid and certain areas while neglecting others, particularly those with strong regional identities.

Response 4 (Male, 23, College, Center-left, Above-average region)

> Geography and location matter a lot. Regions near ports or with good connections to Europe are richer. Also education levels and innovation. And tourism obviously.

Response 5 (Female, 67, High school, Center-right, Below-average region)

> Some regions have always been richer, I think it's just how things are. The climate is better in some places, the land is more fertile. Barcelona and Madrid are big cities with everything, of course they have more money. In my town there's nothing, the young people all leave.

Response 6 (Male, 35, College, Extreme right, Average region)

> Too many autonomous communities, too much bureaucracy, too much waste. Catalonia and the Basque Country get special treatment and financing while places like Extremadura are forgotten. The nationalists always complaining but getting everything. Also, decades of socialist governments in some regions that destroyed the economy with corruption and bad management. We need less politics and more work.

Response 7 (Female, 29, Postgraduate, Center-left, Above-average region)

> There are multiple factors: access to quality education, investment in R&D and innovation, industrial diversification, infrastructure connectivity, demographic dynamics (brain drain from poorer regions), institutional quality at the regional level, and differential access to EU funds. Some regions successfully transitioned from traditional economies to service and knowledge-based economies while others remained dependent on declining sectors. The financing system for autonomous communities also creates disparities in public service provision.

Response 8 (Male, 44, High school, Center-left, Below-average region)

> The government forgets about us. All the investments go to Madrid and Barcelona. Here we don't have good roads, the hospital is far away, young people can't find work. Politicians promise things in elections but then nothing changes.

Response 9 (Female, 31, Postgraduate, Extreme left, Below-average region)

> Spain's regional inequalities are the product of centuries of centralism that favored Castile and Madrid, followed by an incomplete decentralization that maintained structural advantages for central regions. The extraction of resources from peripheral and rural areas to benefit industrial and financial centers continues under capitalism. Neoliberal policies since the 1980s destroyed traditional industries without providing alternatives, creating permanent unemployment and depopulation in many regions. The current financing system perpetuates these inequalities by providing insufficient resources for public services in poorer communities.

Response 10 (Male, 19, High school, Extreme left, Below-average region)

> Capitalism exploits some regions for the benefit of others. The rich areas extract resources and labor from poor areas and give nothing back. It's a system designed to keep inequality.

Response 11 (Female, 33, College, Center-left, Average region)

> I think historical industrialization patterns matter—regions that industrialized early had advantages. Now it's about services, technology, tourism, and having educated people. The financing system for autonomous communities isn't equal, some get more money per person than others. And probably decisions made long ago about where to put universities, factories, train lines etc. still affect things today.

Response 12 (Female, 21, High school, Center-left, Average region)

> Idk, probably because some places have more tourism and businesses? And better universities maybe. Madrid is the capital so obviously everything happens there.

A.4.4 Classification of illustrative responses

Table A.6 presents the classification of the twelve illustrative responses according to our coding framework.

Table A.6: Classification of illustrative beliefs about regional inequality

	Response number											
	1	2	3	4	5	6	7	8	9	10	11	12
Content categories (1 = present)												
Economic structure	1	1	1	0	0	0	1	0	1	1	1	1
Human capital	1	1	0	1	0	0	1	0	0	0	1	1
Infrastructure & geography	1	0	0	1	1	0	1	1	0	0	0	0
Politics & governance	1	1	1	0	0	1	1	1	1	0	0	0
Public finance	0	1	1	0	0	1	1	0	1	0	1	0
Culture & attitudes	0	0	0	0	0	1	0	0	0	0	0	0
Historical legacies	0	0	1	0	0	0	0	0	1	0	1	0
Non-explanatory	0	0	0	0	1	0	0	0	0	0	0	1
Other	0	0	0	0	0	0	0	0	0	1[a]	0	0
Composite measures												
Agency attribution	ST, IN	EX	ST	ST	EX	IN	ST	EX	ST	ST	ST	EX
Grievance intensity	0	1	1	0	0	1	0	1	1	1	0	0
Sophistication (0–1)	0.8	0.4	1	0.4	0.2	0.6	1	0.2	1	0.2	0.6	0.2

Notes: Agency attribution: ST = Structural; EX = External; IN = Individual/Regional; U = Unclear.
[a] Response 10 coded as Other: "Systemic/Marxist"—capitalist exploitation frame not fitting standard categories.

A.4.5 LLM classification prompts

Here we presents the prompts used for LLM-based classification of open-ended responses about beliefs. We provide separate prompts for content classification and composite measure construction.

Content classification prompt

```
You are a research assistant helping to code open-ended survey responses
about beliefs regarding regional economic inequality.

TASK: Classify the following response according to which explanatory
categories are present. A response may invoke multiple categories
(non-mutually exclusive coding).

CATEGORIES:

1. ECONOMIC_STRUCTURE: References to labor market conditions, unemployment,
   sectoral composition (tourism, agriculture, industry), productivity,
   firm location, or agglomeration effects.

2. HUMAN_CAPITAL: References to education, skills, innovation,
   out-migration, brain drain, or demographic factors like aging.

3. INFRASTRUCTURE_GEOGRAPHY: References to transportation, connectivity,
   rurality, peripherality, access to services, or geographic location
   (ports, borders, climate).

4. POLITICS_GOVERNANCE: References to quality of regional government,
   corruption, mismanagement, centralization, favoritism, or policy choices
   by regional or national authorities.

5. PUBLIC_FINANCE: References to allocation of public funds, subsidies,
   transfers, fiscal imbalances, or funding systems for regions.

6. CULTURE_ATTITUDES: References to work ethic, entrepreneurship,
   mentality, effort, or cultural factors as drivers of differences.

7. HISTORICAL_LEGACIES: Direct references to past events, historical
   periods, long-run development paths, or temporal language extending
   beyond the present (e.g., "industrialized early," "decades ago,"
   "colonial history").

8. NON_EXPLANATORY: Vague, circular, or uninformative responses
   (e.g., "I don't know," "many factors," "it's complicated" without
   elaboration, "some places are just richer").

9. OTHER: Substantive explanation not fitting above categories.
   If present, provide brief description.
```

```
RESPONSE TO CLASSIFY:
"""
[INSERT RESPONSE TEXT HERE]
"""

OUTPUT FORMAT (JSON):
{
  "ECONOMIC_STRUCTURE": 0 or 1,
  "HUMAN_CAPITAL": 0 or 1,
  "INFRASTRUCTURE_GEOGRAPHY": 0 or 1,
  "POLITICS_GOVERNANCE": 0 or 1,
  "PUBLIC_FINANCE": 0 or 1,
  "CULTURE_ATTITUDES": 0 or 1,
  "HISTORICAL_LEGACIES": 0 or 1,
  "NON_EXPLANATORY": 0 or 1,
  "OTHER": 0 or 1,
  "OTHER_DESCRIPTION": "brief description if OTHER=1, else null"
}
```

Composite measures prompt

```
You are a research assistant helping to code open-ended survey responses
about beliefs regarding regional economic inequality.

TASK: For the following response, code these composite measures:

1. AGENCY_ATTRIBUTION (ST, EX, IN, or UN):
Who or what is responsible for regional inequality?
  Code EACH of the following independently as 0 or 1.
  Code 1 for every agent the response invokes; do NOT choose a dominant one.
   - AGENT_CAUSES_ST = Structural: impersonal forces
   (markets, geography, agglomeration)
   - AGENT_CAUSES_EX = External: actors outside the region
   (central government, elites)
   - AGENT_CAUSES_IN = Individual/Regional: factors within regional control
   (governance,
         culture, choices)
   - AGENT_CAUSES_UN = Unclear/not specified

   If multiple agents mentioned, list all (e.g., "ST,EX").
   Use the dummy AGENCY_CAUSES_MX (0,1) to code when more than on justification
   is invoked.

3. GRIEVANCE_INTENSITY_CAUSES (0 or 1):
   Code 1 if the response uses moralized language framing inequality
   as unjust, exploitative, or zero-sum.

   Code 1 for:
```

```
   - Claims of injustice, exploitation, victimization
   - Strong moral condemnation ("outrageous," "unfair," "they take")
   - Zero-sum framing ("their gain is our loss")

   Code 0 for:
   - Neutral or analytical descriptions
   - Matter-of-fact statements without moral charge

3. SOPHISTICATION_CAUSES:
   The average of four dimensions: word count (as a fraction of 200, the maximum allowed
   Register independently:
   - CAUSES_LENGTH: answer wordcount  (as a fraction of 200).
   - CAUSES_MULTIPLE: invoking more than one category, 0 or 1.
   - CAUSES_NUANCE: use of causal reasoning, 0 or 1.


RESPONSE TO CLASSIFY:
"""
[INSERT RESPONSE TEXT HERE]
"""

OUTPUT FORMAT (JSON):
{
  "AGENCY_CAUSES": "ST" or "EX" or "IN" or "UN",
  "AGENT_CAUSES_ST":  0 or 1,
  "AGENT_CAUSES_EX":  0 or 1,
  "AGENT_CAUSES_IN":  0 or 1,
  "AGENT_CAUSES_UN":  0 or 1,
  "AGENT_CAUSES_MX":  0 or 1,
  "GRIEVANCE_INTENSITY_CAUSES": 0 or 1,
  "SOPHISTICATION_CAUSES": 0-1
}
```

A.5 Classification of redistribution justifications

This appendix provides detailed documentation of our framework for classifying open-ended responses about redistribution preferences. We present: (1) a summary of content categories, (2) detailed definitions of composite measures, (3) illustrative example responses, (4) the classification of these examples, and (5) the LLM prompts used for automated coding.

A.5.1 Summary of content categories

Table A.7 summarizes the non-mutually exclusive content categories used to classify redistribution justifications.

Table A.7: Content categories for redistribution justifications

Category	Definition and examples
Solidarity & justice	Justifications based on the moral imperative to support disadvantaged regions, ensuring equal opportunity, or leveling the playing field. Must use explicitly moral or ethical framing (e.g., "it's only fair," "they deserve support," "it's the right thing to do").
National unity & cohesion	Justifications based on shared national identity, common citizenship, equal rights across the territory, or uniform provision of public services regardless of location (e.g., "we are all part of the same country," "citizens everywhere deserve the same services").
Historical redress	Explicit framing of redistribution as compensation for past wrongs, historical policies that favored certain regions, or long-run processes that created current disparities. Must reference the past, not merely current conditions (e.g., "historical investments were concentrated elsewhere," "past policies left these regions behind," "paying back a historical debt").
Autonomy & self-interest	Arguments rooted in the right to self-determination, the desire to retain and manage the region's own resources, fiscal autonomy, or the preference to minimize local tax burdens (e.g., "each region should keep what it earns," "we shouldn't have to pay for others," "our money should stay here").
Moral hazard & incentives	Concerns that transfers create dependency, discourage effort or reform, reward poor governance, or distort incentives for regional development (e.g., "it makes them lazy," "why improve if money keeps coming," "bad incentives," "rewards failure").
Distrust & misuse	Skepticism that transferred funds will be used effectively—concerns about corruption, waste, elite capture, or mismanagement in recipient regions (e.g., "the money disappears," "politicians steal it," "nothing ever changes despite the transfers").
Non-explanatory	Expressions of uncertainty, generic statements, tautological reasoning, or responses that do not articulate a substantive justification (e.g., "I don't know," "it's complicated," "just because," "there are many factors").
Other	Residual category for substantive justifications not fitting above categories. When coded, a brief description is recorded (e.g., efficiency arguments, conditional support).

A.5.2 Additional measures: agency, grievance, sophistication

Beyond content classification, we construct additional measures to characterize how respondents reason about redistribution. These measures parallel those used for beliefs about regional inequality, with adaptations for normative reasoning.

Agency attribution. We classify whom respondents implicitly or explicitly hold responsible for redistribution using the following dummies:

- *State/national government (S):* Redistribution framed as a central government policy responsibility or national-level decision.
- *Rich regions specifically (R):* Wealthy regions framed as bearing particular obligations to support poorer areas.
- *Citizens collectively (C):* Redistribution framed as a shared civic duty or collective responsibility of the national community.
- *Poor regions themselves (P):* Emphasis on self-help, local solutions, or the responsibility of disadvantaged regions to address their own problems.
- *Unclear/not specified (U):* No clear attribution of agency in the response.

Responses may be coded with multiple attributions if different agents are mentioned. This measure captures variation in how respondents conceptualize the locus of responsibility for addressing regional disparities.

Grievance intensity. A binary indicator (0/1) that flags responses employing moralized language framing redistribution or inequality as unjust, unfair, exploitative, or characterized by zero-sum competition between regions. We code 1 when responses contain:

- Explicit claims of injustice, exploitation, or victimization (e.g., "we are being exploited," "it's theft," "they take advantage")
- Strong moral condemnation of current arrangements (e.g., "it's outrageous," "completely unfair," "a scandal")
- Zero-sum framing of inter-regional relations (e.g., "their gain is our loss," "they benefit at our expense")

This measure captures the emotional and moral intensity of redistribution attitudes, which may predict behavioral responses beyond stated policy preferences.

Sophistication index. An index based on the average of several dimensions of response quality and depth:

- *Length:* Word count (which will be standardized).
- *Multiple causes:* dummy for invoking more than one justification.

- *Causal reasoning:* Presence of explicit logical connectives (e.g., "because," "therefore," "as a result," "this leads to")
- *Nuance:* Acknowledgment of complexity or interactions between factors.

A.5.3 Illustrative responses

The following fictional responses illustrate the range of justifications participants might provide, varying in length, stated position (support vs. opposition), analytical depth, and substantive content. Demographic characteristics are included to show the diversity of the sample.

Response 1 (Female, 34, College, Center-left, Below-average region) —*Strongly support*

> We are all citizens of the same country. Richer regions have a duty to help those less fortunate—it's simply the right thing to do.

Response 2 (Male, 45, Postgraduate, Center-left, Average region) —*Support*

> The wealth of certain regions today is not just the result of their own efforts. Historical policies concentrated investment and industry in some areas while neglecting others. Infrastructure, universities, factories—all of this was built with national resources but placed in specific locations. Redistribution is a way of correcting these historical imbalances. The poorer regions didn't choose to be left behind; they were left behind by decisions made decades ago.

Response 3 (Female, 38, College, Center-right, Above-average region) —*Support*

> It makes economic sense. When some regions fall too far behind, you get migration pressures, social problems, and political instability that affect everyone. Investing in poorer regions isn't charity—it's smart policy that benefits the whole country in the long run. A more balanced economy is a stronger economy.

Response 4 (Male, 29, College, Center-left, Above-average region) —*Support*

> I support transfers in principle because basic services like healthcare and education should be available everywhere. But I worry about how the money is spent. If it just goes to inefficient bureaucracies or gets lost to corruption, it doesn't actually help anyone. So I support redistribution, but with accountability mechanisms.

Response 5 (Female, 52, High school, Center, Average region) —*Neither support nor oppose*

> It's complicated. On one hand, solidarity is important. On the other hand, regions should also take responsibility for their own development. I can see arguments on both sides and don't have a strong opinion either way.

Response 6 (Male, 41, College, Center-right, Above-average region) —*Oppose*

Each region should be responsible for its own affairs. If a region is poorly managed, why should others pay for it? This creates bad incentives. Regions that work hard and govern well shouldn't have to subsidize those that don't. Let each region keep what it earns and learn to manage its own resources.

Response 7 (Female, 56, High school, Center-right, Above-average region) —*Strongly oppose*

I've seen how these transfers work in practice. The money disappears into the pockets of local politicians and their friends. Nothing ever changes in those regions no matter how much money you send. Until there's real accountability and an end to corruption, I don't support sending more resources that will just be wasted.

Response 8 (Male, 33, College, Center-left, Below-average region) —*Strongly support*

I live in one of the poorer regions, so obviously I benefit from these transfers. But beyond my personal situation, I genuinely believe it's fair. We pay the same national taxes, we're part of the same country, and we deserve the same quality of public services. The current system already favors richer regions in so many ways—in infrastructure, in jobs, in opportunities. Redistribution just levels the playing field a little.

Response 9 (Male, 48, High school, Extreme right, Average region) —*Oppose*

Handouts don't help anyone. People need to work for what they get.

Response 10 (Female, 22, High school, Center, Average region) —*Neither support nor oppose*

I don't really know. There are probably many factors to consider. It's just my opinion.

Response 11 (Female, 61, College, Center-left, Below-average region) —*Support*

During the dictatorship, all the major investments went to the capital and a few industrial centers. The rest of the country was deliberately ignored. My region lost two generations to emigration because there were no jobs, no future. Now people act surprised that we're poor? The state owes us. This isn't charity—it's paying back a historical debt.

Response 12 (Male, 39, Postgraduate, Center-right, Above-average region) —*Strongly oppose*

Why should my taxes go to regions that consistently vote for corrupt politicians and then complain about being poor? They made their choices. We made ours. We invested in education, attracted businesses, built a functioning administration. They could have done the same. At some point, people need to face the consequences of their own decisions instead of expecting others to bail them out forever.

A.5.4 Classification of illustrative responses

Table A.8 presents the classification of the twelve illustrative responses according to our coding framework.

Table A.8: Classification of illustrative redistribution justifications

	Response number											
	1	2	3	4	5	6	7	8	9	10	11	12
Stated position												
Support/Oppose	S+	S	S	S	N	O	O+	S+	O	N	S	O+
Content categories (1 = present)												
Solidarity & justice	1	1	0	0	1	0	0	1	0	0	0	0
National unity	1	0	0	1	0	0	0	1	0	0	0	0
Historical redress	0	1	0	0	0	0	0	0	0	0	1	0
Autonomy & self-interest	0	0	0	0	0	1	0	1	0	0	0	1
Moral hazard	0	0	0	0	0	1	0	0	1	0	0	1
Distrust & misuse	0	0	0	1	0	0	1	0	0	0	0	1
Non-explanatory	0	0	0	0	0	0	0	0	0	1	0	0
Other	0	0	1[a]	0	0	0	0	0	0	0	0	0
Additional measures												
Agency attribution	RR	SG	SG	SG	UN	PR	PR	SG	PR	UN	SG	PR
Grievance intensity	0	0	0	0	0	0	0	0	0	0	1	1
Sophistication (0–1)	0.4	1	0.8	0.8	0.6	0.8	0.8	0.8	0.2	0.2	0.8	0.8

Notes: Position: S+ = Strongly support; S = Support; N = Neither; O = Oppose; O+ = Strongly oppose. Agency attribution: S = State/national government; R = Rich regions; C = Citizens collectively; P = Poor regions themselves; U = Unclear.

[a] Response 3 coded as Other: "Efficiency/consequentialist"—pragmatic argument about collective economic benefits.

A.5.5 LLM classification prompts

This section presents the prompts used for LLM-based classification of open-ended redistribution justifications. We provide separate prompts for content classification and composite measure construction.

Content classification prompt

```
You are a research assistant helping to code open-ended survey responses
about inter-regional redistribution preferences.

TASK: Classify the following response according to which justificatory
categories are present. A response may invoke multiple categories
(non-mutually exclusive coding).

CATEGORIES:

1. SOLIDARITY_JUSTICE: Justifications based on the moral imperative to
   support disadvantaged regions, ensuring equal opportunity, or leveling
   the playing field. Must use explicitly moral or ethical framing
   (e.g., "it's only fair," "they deserve," "it's the right thing").

2. NATIONAL_UNITY: Justifications based on shared national identity,
   common citizenship, equal rights across the territory, or uniform
   provision of public services (e.g., "we are all one country,"
   "same rights everywhere").

3. HISTORICAL_REDRESS: Explicit framing of redistribution as compensation
   for past wrongs, historical policies, or long-run processes that
   created current disparities. Must reference the past, not just current
   conditions (e.g., "historical investments," "past policies," "decades
   of neglect," "paying back a debt").

4. AUTONOMY_SELF_INTEREST: Arguments based on regional self-determination,
   desire to retain own resources, fiscal autonomy, or minimizing tax
   burden (e.g., "keep what we earn," "not our responsibility,"
   "our money").

5. MORAL_HAZARD: Concerns that transfers create dependency, discourage
   effort or reform, reward poor governance, or distort incentives
   (e.g., "makes them lazy," "bad incentives," "why improve,"
   "rewards failure").

6. DISTRUST_MISUSE: Skepticism about effective use of funds|corruption,
   waste, elite capture, or mismanagement (e.g., "money disappears,"
   "politicians steal," "nothing changes").

7. NON_EXPLANATORY: Vague, circular, tautological, or uninformative
   responses (e.g., "I don't know," "just because," "many factors,"
```

```
   "it's complicated" without elaboration).

8. OTHER: Substantive justification not fitting above categories.
   If present, provide brief description.

RESPONSE TO CLASSIFY:
"""
[INSERT RESPONSE TEXT HERE]
"""

STATED POSITION (for context only|classify based on justification content):
[INSERT: Strongly support / Support / Neither / Oppose / Strongly oppose]

OUTPUT FORMAT (JSON):
{
  "SOLIDARITY_JUSTICE": 0 or 1,
  "NATIONAL_UNITY": 0 or 1,
  "HISTORICAL_REDRESS": 0 or 1,
  "AUTONOMY_SELF_INTEREST": 0 or 1,
  "MORAL_HAZARD": 0 or 1,
  "DISTRUST_MISUSE": 0 or 1,
  "NON_EXPLANATORY": 0 or 1,
  "OTHER": 0 or 1,
  "OTHER_DESCRIPTION": "brief description if OTHER=1, else null"
}
```

Composite measures prompt

```
You are a research assistant helping to code open-ended survey responses
about inter-regional redistribution preferences.

TASK: For the following response, code these three composite measures:

1. AGENCY_ATTRIBUTION_JUST (S, R, C, P, or U):
   Who is responsible for redistribution (or for not redistributing)?
   Code EACH of the following independently as 0 or 1. Code 1 for every
agent the response invokes; do NOT choose a dominant one.

   - SG = State/national government
   - RR = Rich regions specifically
   - CC = Citizens collectively
   - PR = Poor regions themselves (self-help emphasis)
   - UN = Unclear/not specified

   If no agent is identifiable, code all four as 0 except UN. If multiple agents
   mentioned, list all (e.g., "SG,PR"). Use the dummy AGENCY_JUST_MX (0,1) to code
   when more than one justification is invoked.

2. GRIEVANCE_INTENSITY_JUST (0 or 1):
```

```
   Code 1 if the response uses moralized language framing redistribution
   or inequality as unjust, exploitative, or zero-sum.

   Code 1 for:
   - Claims of injustice, exploitation, victimization
   - Strong moral condemnation ("outrageous," "scandal," "theft")
   - Zero-sum framing ("their gain is our loss")

   Code 0 for:
   - Neutral or matter-of-fact descriptions
   - Mild preferences without moral intensity

3. SOPHISTICATION_JUST:
   The average of four dimensions: word count (as a fraction of 200, the maximum allowed
   - JUST_LENGTH: answer wordcount  (as a fraction of 200).
   - JUST_MULTIPLE: invoking more than one category, 0 or 1.
   - JUST_NUANCE: use of causal reasoning, 0 or 1.

RESPONSE TO CLASSIFY:
"""
[INSERT RESPONSE TEXT HERE]
"""

OUTPUT FORMAT (JSON):
{
  "AGENCY_ATTRIBUTION_JUST": "SG" or "RR" or "CC" or "PR" or "UN" (or combination),
  "AGENCY_JUST_SG":  0 or 1,
 "AGENCY_JUST_RR":  0 or 1,
  "AGENCY_JUST_CC":  0 or 1,
  "AGENCY_JUST_PR":  0 or 1,
  "AGENCY_JUST_UN":  0 or 1,
 "AGENCY_JUST_MX":  0 or 1,
  "GRIEVANCE_INTENSITY_JUST": 0 or 1,
  "SOPHISTICATION_JUST": 0-1
}
```

A.6 Implementation notes for LLM classification

The following notes apply to LLM-based classification of both beliefs and justifications.

- *Models and settings.* Every response is classified independently by three models: Claude Sonnet 4.5 (Anthropic), GPT-4o (OpenAI) and Mistral Small 3 (Mistral AI). All are run at temperature 0 to minimize variation across runs. The exact model versions and the dates the code is ran will be recorded and reported in the replication package.
- *Batch processing.* Responses are processed individually rather than in batches, so that the code assigned to one response cannot depend on the others in context or on the order of processing.
- *Language.* For responses in languages other than English, we will translate the answers before applying our prompts. As a robustness check, we also report results from translating the prompts rather than the responses.
- *Validation sample.* We validate the LLM coding against human coding on 600 responses for each open-ended question (200 per country), drawn at random but stratified by country and arm.
- *Human coding procedure.* Each validation response is coded independently by two human coders who are blind to treatment assignment. The human coders receive instructions identical to the model prompts reproduced in Sections A.4.5 and A.5.5 .
- *Agreement statistics.* We report (i) agreement between the two human coders; (ii) agreement between each model and each human coder; and (iii) agreement across the three LLMs.
- *Performance and model selection.* For each content category and composite measures (of each open-ended question), we report recall, precision and F1 (i.e., the harmonic mean of the two). We select as our baseline classification for descriptive and regressions the model that yields the highest average F1 across all dimensions on the validation sample.
- *Replication package.* The replication will incorporate all the prompts, category definitions, coding instructions, model version and settings. Importantly, all anonymized responses to the open-ended questions will be available in the replication package (not just those of the validation but the entire dataset).

A.7 Measuring perceived bias with LLMs

This section describes the prompt used for LLM-based classification of the open-ended responses on bias of the survey.

Perceived bias classification prompt

```
You are research assistant helping to code open-eded respones from a suvey
about regional inequality and redistribution preferences.
Respondents were asked whether they felt the survey was biased in some sense.

TASK: Classify the following respones according to how strongly the respondent
indicates that the survey was biased in favor/against poor regions, or in favor/
against rich regions.  If the respondent declares the survey promoted the view that
resources should be redistributed to poorer regions, or that it favored poorer
regions or greater support for them, that should take positive values.

Assign an integer from -2 to 2:

\begin{itemize}
\item $+2$ = The respondent states emphatically that the survey was biased in favor
of poorer regions, or accuses it of pushing an agenda in favor of redistribution
toward them. Typical markers are explicit accusations of manipulation, propaganda or
an agenda; strong or repeated language; or an explicit claim that the survey was
designed to produce a pro-redistribution answer.

$+1$ = The respondent indicates that the survey favored poorer regions,
promoted greater redistribution toward them, or was designed to elicit
support for helping them, but states this indirectly, in passing,
or in measured terms.

\item $0$ = None of the directional cases apply. Code 0 in each of the
following situations:
(i) the respondent reports a bias unrelated to regional inequality or regional
redistribution --- for example, that the survey was too long or
repetitive, that the questions were confusing or leading in general,
that the survey was political, that the demographic questions
were intrusive, or that the response options were unbalanced;
(ii) the respondent indicates that the survey took a position
on regional inequality or regional redistribution but does not
make clear which position;
(iii) the respondent describes no bias at all;
(iv) the response is empty, unintelligible, or does not address the question.

\item $-1$ = The respondent indicates that the survey leaned against
redistribution toward poorer regions, favored richer regions, or was
designed to elicit opposition to helping poorer regions, but states
this indirectly, in passing, or in measured terms.
```

```
\item $-2$ = The respondent states emphatically that the survey was
biased in favor of richer regions or against poorer ones, or accuses it
of pushing an agenda against redistribution toward poorer regions.
\end{itemize}
```

Judge only what the respondent says. Do not infer a position from the topic of the survey, and do not reward or penalize a response for being articulate. A respondent may be correct or incorrect about the survey; that is irrelevant to the coding.

```
RESPONSE TO CLASSIFY:
"""
[INSERT RESPONSE TEXT HERE]
"""

OUTPUT FORMAT (JSON):
{
 "BIAS_SCORE": <integer -2, -1, 0, 1, 2>
 }
```

A.8 Data protection & Consent

The study is run by researchers based in the European Union and is therefore subject to the GDPR for all three country samples; the Brazilian sample is additionally covered by the LGPD, whose requirements the procedure below meets. Daniel Oto-Peralías and Martín Fernández are joint controllers. Bilendi-Netquest acts as processor under a written data processing agreement and transmits only pseudonymous panel identifiers; the research team never receives names, e-mail addresses, or contact details from the panel. Data are collected under Art. 6(1)(a) consent and, for the political-values block and the petition, under the explicit consent of Art. 9(2)(a), with the research safeguards of Art. 89(1) applied throughout. The linkage key between waves is held separately from the response data and destroyed once the second wave is matched, at which point the dataset is anonymous; respondents are informed *in advance* that erasure requests cannot be honoured after that point, since the data can no longer be traced to them. Only the anonymised dataset is deposited in a journal's replication archive, and open-ended responses are screened for inadvertently identifying content before deposit. Respondents are informed of their rights of access, rectification, erasure, restriction, objection and portability, before the data is fully anonymized.

Extended Landing Page for the Survey

A study on public opinion about economic and political issues

This survey is part of an academic research project carried out by Martín Fernández (Luxembourg Institute of Socio-Economic Research) and Daniel Oto-Peralías (University Pablo Olavide of Seville), who are jointly responsible for the data collected. It takes about 20 minutes.

What the study is about. We study what people know and think about the economy of their country and about public policy. We do not provide extensive details to avoid influencing your answers.

What you will be asked. Background questions about you and your family's regions of origin; your views on political, social and economic questions, including your political position; and your preferences about public policy. You will watch a short video. Some questions ask you to write in your own words.

Re-contacting. You may be invited to participate in a further survey of about [10] minutes as part of this study. Your answers to it would be linked to your answers today through a pseudoanonimzed code, never through personal information. Taking part in it will be entirely up to you at the time. You will be again asked for consent.

Your data and your rights. Your name is never recorded. The survey company identifies you to us only by a code. Your answers are processed on the basis of your consent. Because your political views are involved, we ask for your explicit consent to that separately below. Your answers are seen by the research team and processed by our survey provider, Bilendi-Netquest. Text you write in your own words may be analysed by automated language-processing services operating outside the European Economic Area, under contractual data-protection safeguards. The code linking the two surveys is stored separately and destroyed

once the two waves are matched, and no later than [DATE]. From that point the data are fully anonymous. Anonymous data will be published so that other researchers can verify our results, as scientific journals require. Once the data are anonymous we can no longer find your answers, so please note that we cannot delete them after that point. Until then you may ask us for a copy of your data, or ask us to correct, delete, restrict or object to the processing of it, or withdraw your consent — withdrawing does not affect what was lawful beforehand. Write to [INSTITUTIONAL ADDRESS]. Our data protection officer can be reached at [DPO ADDRESS].

Taking part. Participation is voluntary and you may stop at any time. [Compensation statement — confirm with the panel provider.] Some questions concern historical and political topics that people feel strongly about; if you would prefer not to continue at any stage, you may simply close the survey.

Please read carefully and answer honestly: rushed or inconsistent answers are detected by statistical checks, and the value of the research depends on genuine responses.

You may take a screenshot of this page for your own record.

- $[X]$ I am 18 or older, I have read the above, and I agree to take part and to my answers being used as described.
- $[X]$ I explicitly consent to my answers about my political views and preferences being used for this research.
- $[X]$ I agree to be invited to the short follow-up survey and to my answers being linked by code.

[Start the survey] [I do not wish to take part]

B Intervention Scripts

INFORMATION TREATMENT

Animated Videos — Video Scripts

Treatment 1 (≈2:30 minutes)

Opening (≈10 seconds): The Big Question

Why are some regions in your country much poorer than others, generation after generation?

Differences in culture? Education? Politics?

Core Idea (≈40 seconds): The Power of History

Economic historians suggest the answer lies much deeper, with roots in the past.

Some regions enjoyed early advantages like fertile land and trade, which helped them develop.

Others faced severe barriers like isolation, conflict, and exploitation, making progress much harder.

Once these economic gaps emerged, they tended to persist, passing down across generations through local institutions, infrastructure, and culture.

Economists call this persistence: today's economic map is largely a reflection of the past.

Country-Specific Content (≈90 seconds): Concrete Example

[Insert country-specific script here — see below]

Closing (≈10 seconds): Call to Reflection

History has shaped the present — often in ways we don't even notice.

Treatment 2 (≈3:30 minutes)

Treatment 1 content as above, followed by the extension below.

Extension (≈60 seconds): *From Persistence to Possibilities*

However, history does not dictate the future, as the following well-documented cases illustrate.

In the United Kingdom, a government program offered investment subsidies to manufacturing firms in disadvantaged regions. Economists who studied the program found it significantly increased employment in those regions, without simply moving jobs from nearby places.

In China, special economic zones have been a successful development strategy. Focusing on zones created in the mid-2000s, recent research found they significantly increased employment and output, without displacing nearby activity.

These aren't universal guarantees, but they show that, under the right conditions, regional policy can lift local economies.

Understanding where regional inequalities come from is the first step toward asking what we can do about it.

Country-Specific Scripts

Spain (≈90 seconds)

Let's take the case of Spain. Two economic history studies trace part of today's regional inequality back hundreds of years—to the long struggle of the so-called Reconquista.

For centuries, Christian armies pushed south, reclaiming land from Muslim rule. How that land was distributed shaped each region's future.

In the north and center, the Reconquest progressed gradually. Individual settlers and the Crown led the resettlement, and land was distributed more equally.

By contrast, southern conquests happened much faster, as the Crown relied on nobles and military orders, rewarded with vast estates that concentrated land and power.

Researchers point to two consequences. First, in regions where land was more concentrated, day-laborers were still less likely to read and write centuries later, in the mid-1800s—lacking both the resources and the incentive to learn.

Relatedly, when Spain industrialized, the rich in unequal regions had little reason to invest beyond agriculture, while ordinary people lacked the resources to join new industries.

Today, these historical patterns still show up in income, education, and job opportunities: regions where land ended up more concentrated are nowadays poorer.

The United States (≈90 seconds)

Let's take the case of the US. Economic historians trace today's income gaps across states back centuries, to the country's colonial past.

European settlers' strategies followed geography. The South's fertile land and warm climate favored plantation crops like tobacco, rice, and cotton—most profitable on large estates using enslaved labor.

In the North, less suitable soil and climate favored smaller, family-run farms and more diversified economies.

These beginnings produced very different societies: land, wealth, and political power grew highly concentrated in plantation states, while smallholder states developed more equal land distribution, broader political participation, and better governance.

These early inequalities shaped local institutions: areas more reliant on slavery invested less in public schooling and had lower economic mobility—institutional gaps that persisted even after slavery was abolished.

Northern states industrialized earlier: manufacturing, railroads, and urbanization raised wages and attracted both investment and skilled workers, while the South stayed dependent on agriculture, and by the early twentieth century, incomes there lagged far behind the rest of the country.

Colonial settlement, slavery, institutions, and industrialization combined to shape the economic map we see today.

Brazil (≈90 seconds)

Let's take the case of Brazil. Economic historians trace today's regional inequalities to colonial choices that shaped who owned the land, who held power, and what the state provided.

When the Portuguese arrived in 1500, the warmer and rain-fed northeastern coast became sugar plantations worked by enslaved labor, producing a wealthy planter elite and a large, landless underclass. The interior later saw gold rushes and cattle ranching. The southeast, much later, boomed with coffee.

Each colonial boom left its own institutional imprint. Sugar-cycle areas became oligarchic and land-concentrated—a pattern still visible today. Gold-rush areas developed a heavily inefficient, rent-seeking state that still shows weaker governance and access to justice.

Where elites controlled the land, they also controlled who got educated, who had political voice, and who benefited from public goods. Recent research confirms that municipalities with historically unequal land distribution today spend less on welfare, health, and education, remain more dependent on low-wage agriculture, and have lower human development across the board.

The northeast, heir to the sugar economy, remains today poorer than the south and southeast.

Active control – Neutral video

Opening (≈ 24 seconds): The Big Question, per country

The United States

In the United States, there are over nineteen trillion dollars sitting in bank accounts. But if you tried to collect every physical dollar bill and coin in the country — every single one — you'd gather barely more than a tenth of that. Where did all the actual money go?

Spain / Eurozone

Across the eurozone, there are over ten trillion euros sitting in bank accounts. But if you tried to collect every physical euro banknote and coin in circulation — every single one — you'd gather only about fifteen percent of that. Where did all the actual money go?

Brazil

In Brazil, there are more than seven trillion reais sitting in bank accounts. But if you tried to collect every physical real note and coin in the country — every single one — you'd gather barely five percent of that. Where did all the actual money go? [40]

Core Idea (≈41 seconds): The basic mechanism

Economists have a simple answer: here's what actually happens.

Imagine you deposit a thousand [dollars/euros/reais] into your account. Banks are generally required to keep only a small portion of each deposit on hand — say, ten percent, just as an example.

The rest, they lend out to other people and businesses: a car loan, a mortgage, a loan to open a small shop.

So of your thousand [dollars/euros/reais], the bank might keep a hundred in reserve, and lend out the remaining nine hundred to someone else.

The multiplier (≈72 seconds):

That nine hundred doesn't just vanish—it gets spent, and deposited again, this time by someone else's bank, which again keeps ten percent and lends out the rest: eight hundred and ten.

And the cycle repeats again, and again, at a smaller scale each time.

By the time this plays out, your deposit can turn into several times its value across the banking system—though only the original currency ever existed.

That's why the numbers don't match physical cash.

[40] The monetary aggregate used in each country's script corresponds to the closest available measure of household-accessible bank deposits (including savings) in that country's official classification: M1 for the US (which has included savings deposits since May 2020) and the Eurozone (overnight deposits), and M2 for Brazil (where savings deposits are classified separately from M1)

Economists call this the money multiplier.

Because of this, the total money in an economy isn't fixed—it grows and shrinks through the ordinary process of depositing and borrowing, no printing required.

The overwhelming majority of money in a modern economy exists purely as electronic entries—numbers in a database, updated every time someone gets paid, pays rent, or takes out a loan.

Closing (≈10 seconds)

So next time you wonder where all that money went — now you know: it never needed to be there at all.

C Questionnaire

Historical Legacies and Preferences for Inter-Regional Redistribution

Martín Fernández & Daniel Oto-Peralías

August 2026

SPANISH QUESTIONNAIRE

INDEX

LANDING PAGE

This is a survey conducted for academic purposes by researchers from European universities. It will take approximately **20 minutes** to complete. Your responses will be used **exclusively for research purposes**. You will be compensated if you complete the full survey and your responses pass our quality checks. These checks use sophisticated statistical methods to detect incoherent and rushed responses. It is very important for the validity of our research that you **answer honestly** and **read the questions carefully**. By participating, you are contributing to the progress of scientific knowledge in our society. The survey also gives you an opportunity to learn new information and express your own views.

Please note:

- *Your participation in this study is completely voluntary.*
- *Your name will never be recorded.*
- *The results may include summary data, but you will never be identified.*
- *The anonymous data may be shared with other researchers.*
- *You can take a photo or screenshot of this page for your own record.*
- *If you have any questions, feel free to contact us at xxxx@gmail.com.*

Please, select:

- Yes, I would like to take part in this study, and I confirm that I am 18 or older.
- No, I would not like to participate.

SOCIO-DEMOGRAPHIC BLOCK

- **Year of birth:**
  - _____ [Survey ends if year is 2008 or later, or 1955 or earlier]

- **Gender:**
  - Male
  - Female
  - Other / Prefer not to say

- **Province of birth:**
  - [Drop-down menu] [Survey ends if "Outside Spain"]

- **Province of residence:**
  - [drop-down menu]

- **Highest level of studies completed:**
  - No schooling completed
  - Primary school
  - Lower secondary school
  - High school
  - Vocational degree
  - College degree
  - Master's degree or above

- **Over the last year, what was the average monthly income of your household (after taxes and transfers)?**
  - Less than 1,000 €
  - 1,000 € – 1,999 €
  - 2,000 € – 2,999 €
  - 3,000 € – 3,999 €
  - 4,000 € – 4,999 €
  - More than 5,000 €

- **Civil status:**

  - Single, never married
  - Married or de facto couple
  - Separated or divorced
  - Widowed

- **How many children do you have?**
  - 0
  - 1
  - 2
  - 3
  - 4 or more

- **Number of people living in your household:**
  - Adults (18+): _____
  - Children (under 18): _____

- **What type of area do you live in?**
  - A small town (fewer than 5,000 people)
  - A middle-size town (5,000 – 20,000 people)
  - A large town (20,001 – 50,000 people)
  - A small city or its suburbs (50,001 – 250,000 people)
  - A large city or its suburbs (more than 250,000 people)

- **Select your current employment situation:**
  - Unemployed
  - Employed by someone else
  - Entrepreneur or self-employed
  - Unpaid domestic or caregiving work (e.g., childcare, elder care)
  - Retired or pensioner
  - Student
  - Other: _____

- **Your parents' regions of birth:**
  - Father: [drop-down menu]
  - Mother: [drop-down menu]

- **Your grandparents' regions of birth:**
  - [2×2 matrix with drop-down menus]
- **Regarding your feelings of identity or belonging, which of the following best describes you?**
  - I feel only from my region
  - I feel more from my region than Spanish
  - I feel equally from my region as Spanish
  - I feel more Spanish than from my region
  - I feel only Spanish

ECONOMIC BLOCK

- **To which extent you agree or disagree with the following statements:**

  [Strongly agree; Agree; Neither agree nor disagree; Disagree; Strongly disagree]

  - In Spain, there are different income groups. When one group becomes wealthier, it usually comes at the expense of other groups.
  - In international trade, if one country earns more, it generally means that the other country earns less.
  - Regarding the regions in Spain, if one region does better economically, this generally comes at the expense of the other regions.

- **In your opinion, what plays the biggest role in why someone is poor?**

  - Lack of personal effort
  - Circumstances beyond their control

- **What plays the biggest role in why someone is rich?**

  - Hard work and effort
  - Privileged background or advantages

- **Now, let's explore how income levels vary across regions in Spain. Which regions do you believe are currently the poorest and the richest?**

  - The poorest: [drop-down menu]
  - The richest: [drop-down menu]

- **Based on your perception, how does the income per capita (i.e., per person) of the richest region compare to that of the poorest region in 2025?** [*Select one of the bars*]

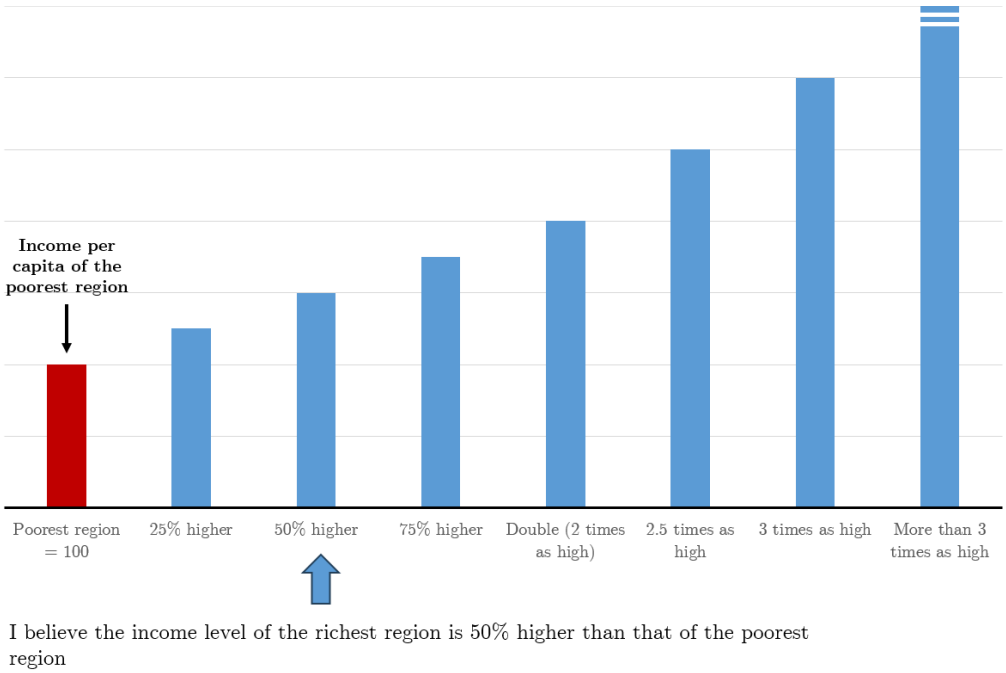

Figure C.2: Incomes 2025

- **Thinking about these differences in regional income, what factors do you think explain why some regions in Spain are richer than others?** You can write as much detail as you wish. Your response is very valuable to our research.

  [Open-ended question]

- Still thinking about income differences across regions in Spain, in your opinion **how much of these differences is due to events and processes that took place a long time ago (more than 100 years ago), as opposed to more recent developments and the current characteristics of each region** — its population, its government, its economy? [*Move the slider to indicate your view*]

  Note to programmers: Handle position starts with no value until first click

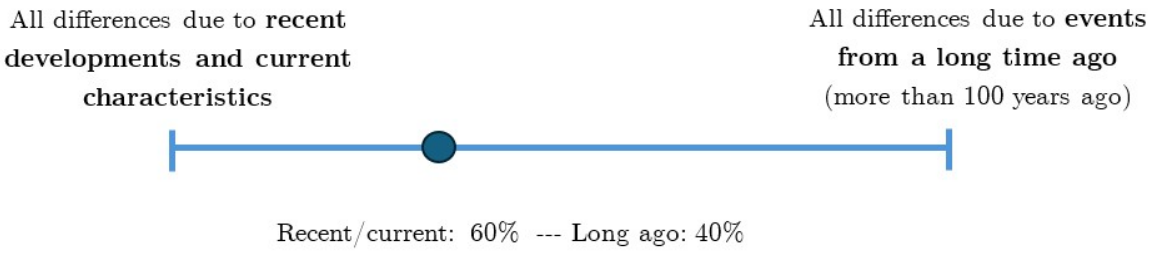

- **Let's now look back to the year 1900. Which regions do you believe were the poorest and the richest in Spain back then?**

  – The poorest in 1900: [drop-down menu]

– The richest in 1900: [drop-down menu]

- **How do you think the income per capita of the richest region compared to that of the poorest region in 1900?** [*Select one of the bars*]

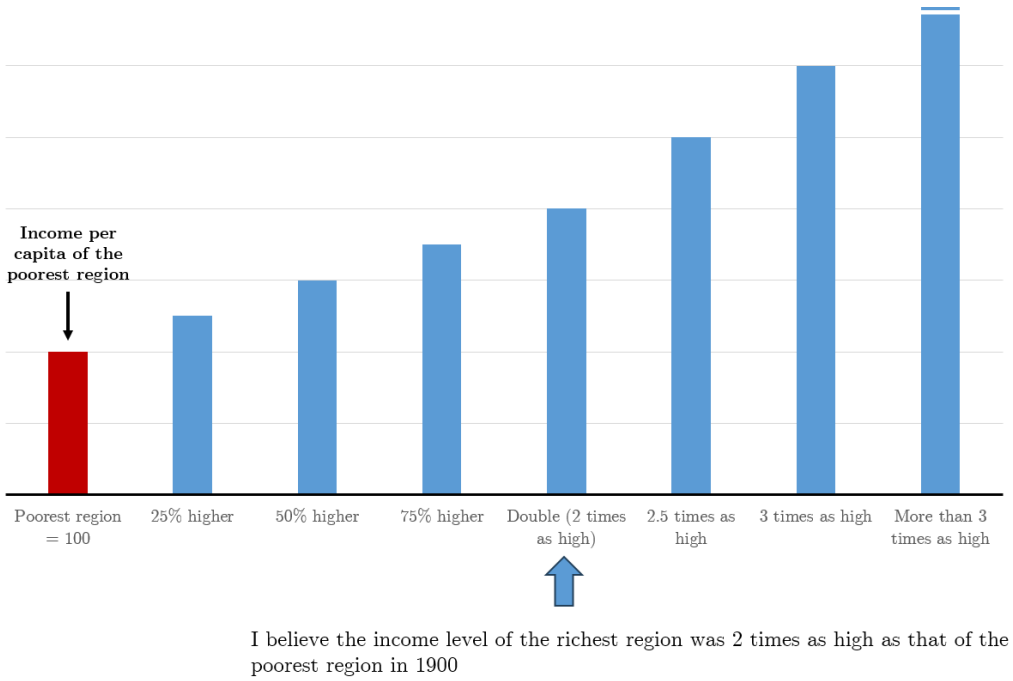

Figure C.3: Incomes 1900

- Economists use the term ***persistence*** to describe how much historical differences in regional income still shape present-day outcomes. Below are three scenarios:

  – **High Persistence**: Regions that were richer in 1900 are still rich today. Poorer regions in 1900 remain poor today.

  – **No Persistence**: No relationship between past and present.

  – **Negative Persistence (Reversal)**: Regions that were poorer in 1900 are now rich, and regions that were richer in 1900 have declined.

  Based on your intuition, how persistent do you think regional income differences have been in Spain from 1900 to 2025? [*Use the slider to answer*]

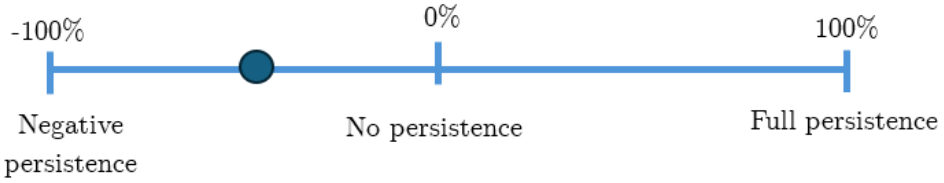

POLITICAL AND CULTURAL QUESTIONS

- **In politics, people often refer to 'the left' and 'the right'. On a scale from 0 to 10, where 0 is "farthest to the left" and 10 is "farthest to the right," where would you place yourself?**

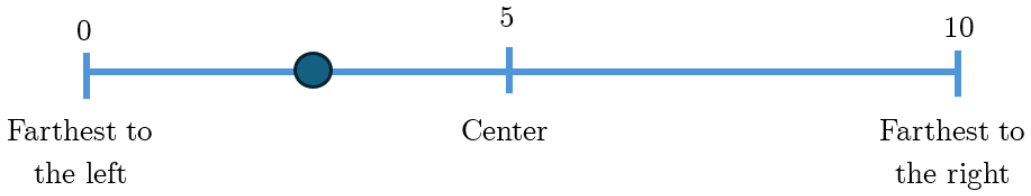

- **On a scale of 0 to 100 (where 0 is "*No trust at all*" and 100 is "*Complete trust*"), how much do you trust each of the following people?**
  - A person you do not know personally who lives in your local neighborhood:
  - A person you do not know personally who lives in a different region of the country:

- **How much do you trust the government?**
  - 1 (= Don't trust at all)
  - 2 (= Don't trust a lot)
  - 3 (= Trust somewhat)
  - 4 (= Trust a lot)
  - 5 (= Trust completely)

- **How much do you trust scientists?**
  - 1 (= Don't trust at all)
  - 2 (= Don't trust a lot)
  - 3 (= Trust somewhat)
  - 4 (= Trust a lot)
  - 5 (= Trust completely)

- **Our study relies on participants devoting their full attention and answering truthfully. Your compensation will not be affected by your answer to this question: honestly, should we use your responses?**
  - Yes, I have devoted full attention and answered truthfully.
  - No, I have not devoted full attention or answered truthfully.

TREATMENT[41]

I. Main treatment

We will now show you a 2 minute 30 second video (with sound) summarizing recent academic research on the causes of regional income differences in [Spain/the US/Brasil]. Please make sure your sound is on and pay close attention to the information presented. After the video, you will be asked a few basic questions about its content to help us assess whether it was viewed and understood. Please watch the video in its entirety before proceeding. When you are ready, click "Next" to continue.

- **Were you able to watch and listen to the video until the end?**
  - Yes
  - No, there was a technical problem
  - No, I skipped part of the video

- **Would you like to receive a link to the video by email after the survey, so you can watch it again or share it?**
  - Yes, send me the link by email
  - No

- **Based on what was said in the video, select ALL the correct answers:**
  - ☐ Recent research in economics shows that the economic prosperity of regions tends to change very quickly over time.
  - ☐ The Spanish Civil War has had major implications for regional economic development in the last 50 years.
  - ☐ Historical events and processes taking place a long time ago are important to explain current regional income differences.
  - ☐ Southern regions are poorer because of the heat-induced effect on the ability to work.

II. Alternative treatment

We will now show you a 3 minute 30 second video (with sound) summarizing recent academic research on the causes of regional income differences in [Spain/the US/Brasil]. Please make sure your sound is on and pay close attention to the information presented. After the video, you will be asked a few basic questions about its content to help us assess whether it was viewed and understood. Please watch the video in its entirety before proceeding. When you are ready, click "Next" to continue.

[41] The survey design does not force respondents to watch the full video before proceeding to the next survey page, as this information (time spent on the page) is valuable to identify those who skip the video too soon.

- **Were you able to watch and listen to the video until the end?**
  - Yes
  - No, there was a technical problem
  - No, I skipped part of the video

- **Would you like to receive a link to the video by email after the survey, so you can watch it again or share it?**
  - Yes, send me the link by email
  - No

- **Based on <u>what was said in the video</u>, select ALL the correct answers:**
  - ☐ ...(*the same as above for each country*).
  - ☐ ...(*the same as above for each country*).
  - ☐ ...(*the same as above for each country*).
  - ☐ ...(*the same as above for each country*).
  - ☐ Well-documented public policies in other countries have been shown to increase employment in disadvantaged regions.

III. Control group

We will now show you a 2 minute 30 second video (with sound) explaining a concept in economics, the money multiplier, how it works and why it is important. Please make sure your sound is on and pay close attention to the information presented. After the video, you will be asked a few basic questions about its content to help us assess whether it was viewed and understood. Please watch the video in its entirety before proceeding. When you are ready, click ”Next” to continue.

- **Were you able to watch and listen to the video until the end?**
  - Yes
  - No, there was a technical problem
  - No, I skipped part of the video

- **Would you like to receive a link to the video by email after the survey, so you can watch it again or share it?**
  - Yes, send me the link by email
  - No

- **Based on <u>what was said in the video</u>, select ALL the correct answers:**

  ☐ The amount of money in the economy is backed by gold reserves in the Central Bank *[US: Federal Reserve*].

  ☐ Banks create new money mainly by printing physical banknotes.

  ☐ When a bank lends out part of a deposit, that money can be redeposited and lent again, expanding the amount of money in the economy.

  ☐ Banks are required to keep 100% of deposits in reserve at all times.

PREFERENCES FOR REDISTRIBUTION AND OTHER OUTCOMES

We would now like to ask your views on a number of topics related to the economy and public policy in Spain.

- **Do you think that Artificial Intelligence will have a negative or a positive effect on employment in the future?** (*Outcome-neutral test*)
  - Very positive
  - Positive
  - Neither positive nor negative
  - Negative
  - Very negative

- **Do you think income differences between rich regions and poor regions in Spain are:**
  - Not a problem at all
  - A small problem
  - A moderate problem
  - A serious problem
  - A very serious problem

- **To what extent do you support or oppose rich regions transferring more resources to poor regions to help finance their public services and infrastructure?**
  - Strongly support
  - Somewhat support
  - Neither support nor oppose
  - Somewhat oppose
  - Strongly oppose

- **Please, explain the main reasons behind your previous answer. Write as much as you would like, your response will be crucial to better understand participants' opinion:**

  [Open-ended question]

- **How likely is it that the poorest regions in Spain today will still be the poorest in 30 years?**
  - Very likely
  - Likely

  - Unlikely
  - Very unlikely

- **To show that you are attentive and responding carefully the questions, please select both "a little" and "a lot" in the following list. Please, select the ones that apply:**
  - Not at all
  - A little
  - Moderately
  - A lot
  - A great deal

- **When thinking about poorer regions in your country, to what extent do you feel…**

  [A lot, Somewhat, Very Little, Not at All]
  - Empathy or compassion
  - A sense of moral responsibility
  - Pessimism regarding their future development

- **To what extent do you agree or disagree that poor regions…**

  [Strongly agree; Agree; Neither agree nor disagree; Disagree; Strongly disagree]
  - Still pay the price for disadvantages created long ago.
  - Are poorer due to the work attitudes and initiative of their population.

- **To what extent do you think that climate change is already having a negative economic impact in your country?** (*Outcome-neutral test*)
  - Not at all
  - A little
  - Moderately
  - A lot
  - A great deal

- **Some people think that the government (at the local, regional, or national level) should not concern itself with income differences between rich and poor people. Others think that the government should do everything in its power to reduce income inequality. Where would you place your own view?** [*Move the slider to indicate your view*]

Note to programmers: Handle position starts with no value until first click. Range: 0 - 10.

- **To what extent do you think that well-designed economic policies of the government are able to:**

  [A great deal; A lot; A moderate amount; A little; Not at all]

  a) Improve the level of education.
  b) Reduce the amount of people in risk of poverty.
  c) Promote the development of left-behind areas.

- Thinking once more about the income differences between regions in Spain, in your opinion **how much of these differences is due to events and processes that took place a long time ago (more than 100 years ago), as opposed to more recent developments and the current characteristics of each region** — its population, its government, its economy? [*Move the slider to indicate your view*]

  Note to programmers: Handle position starts with no value until first click

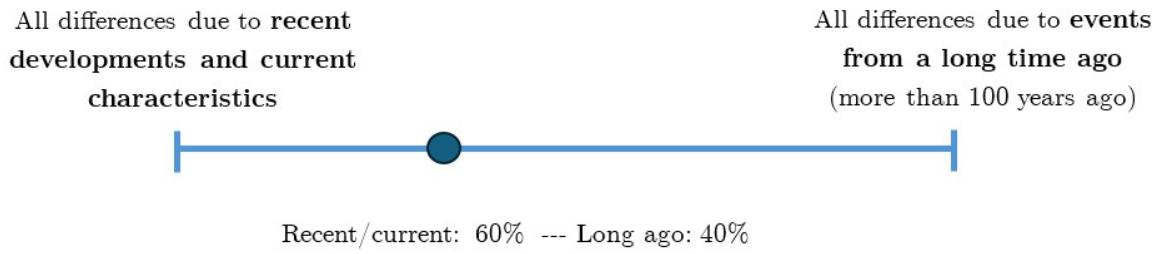

- **In this final part of the survey, we would like to ask you about a petition that we will send to the national government. When the study is complete, the research team will send a letter to the Ministry of Finance informing them which of the following two policies is supported by the majority of a large sample of Spaniards who took part in this survey, and we will publish the result on the project's website. Please indicate which of the two you would support** (naturally, your answer will remain anonymous)**:**

a) The fiscal system should transfer more resources from higher-income regions to lower-income regions, so that the latter can improve their public services and infrastructure.

b) The fiscal system should reduce the transfer of resources from higher-income regions to lower-income regions, so that each region relies mainly on the resources generated within its own territory.

- **Your answer will be included, anonymously, in the overall count that we will send to the Ministry of Finance.**

  Note to programmers: this question must appear on a separate screen, after the respondent has submitted the previous answer, with no option to return to the previous page.

  ☐ Please exclude my answer from the count sent to the Ministry

500€ lottery:

Screen 1 (Instructions)

Instructions

By completing this survey, you are entered into a lottery to win **500€** (paid in equivalent reward points). If you win, you must decide how to share this prize with another anonymous participant in this survey. In both scenarios below, the other participant is someone whose **household income is below 1,000€ per month** (after taxes and transfers).

How it works:

- We will present you with two scenarios (Scenario A and Scenario B).
- If you win, the computer will **randomly pick one** of these two scenarios to be your final, binding choice.
- Because you don't know which one the computer will pick, you should treat **both** as if they were real.

Screen 2 (Scenario A/B – Randomized)

Scenario: Sharing with someone from YOUR region

If you win the 500€, and the computer selects this scenario, how would you like to split the money between yourself and a **randomly selected participant who lives in your region and whose household income is below 1,000€ per month**?

- **Amount to keep for yourself:** € [_____]
- **Amount to give to the other participant:** € [_____]

(Note: The total must sum to 500€)

Screen 3 (Scenario A/B – Randomized) [42]

This is shown for respondents who do not reside in the poorest region (Extremadura)

Scenario: Sharing with someone from the POOREST region (Extremadura)

If you win the 500€, and the computer selects this scenario, how would you like to split the money between yourself and a **randomly selected participant who lives in Extremadura (the poorest region in Spain) and whose household income is below 1,000€ per month**?

[42] Payments are executed after data collection closes. If no participant satisfies both criteria—the designated recipient region and household income below the threshold—the amount is paid to a participant satisfying the income criterion from the next-poorest region. This contingency is unlikely to arise but is specified in advance to ensure the incentive is implementable in all cases. It is not disclosed to respondents, so that the recipient's region is unambiguous at the point of decision.

- **Amount to keep for yourself:** € [_____]
- **Amount to give to the other participant:** € [_____]

(Note: The total must sum to 500€)

Screen 3 (B) (Scenario A/B – Randomized)

This is shown for respondents who reside in the poorest region (Extremadura)

Scenario: Sharing with someone from a poor region (Andalusia)

If you win the 500€, and the computer selects this scenario, how would you like to split the money between yourself and a **randomly selected participant who lives in Andalusia (one of the poorest regions in Spain) and whose household income is below 1,000€ per month**?

- **Amount to keep for yourself:** € [_____]
- **Amount to give to the other participant:** € [_____]

(Note: The total must sum to 500€)

FEEDBACK

- **To what extent have you been persuaded by the video's arguments?**
  - Very persuaded
  - Somewhat persuaded
  - Slightly persuaded
  - Not at all persuaded

- **Do you feel that this survey was biased in some sense?**
  - Yes $\Rightarrow$ If yes: please clarify.
  - No, I do not feel it was biased

- **The survey has ended. Many thanks for participating. You can now enter any comments, thoughts or suggestions about the survey in the field below.**

FOLLOW-UP SURVEY
(obfuscated)

Landing page

This is a short survey with an academic-research goal to gain insights into people's views about economic and political issues. By dedicating 3 or 4 minutes of your time, you contribute to our collective knowledge as society. All answers you give will be fully confidential. We will not ask for information related to your identity. You may withdraw from the study or request the deletion of your data at any time via xxx@gmail.com.

If you are at least 18 years old and freely consent to participate in this study, please click **Next** to start the survey.

Demographic questions

- **What is your age?**
  - 18–24
  - 25–34
  - 35–54
  - 55–70
  - > 70

- **Gender:**
  - Male
  - Female
  - Other / Prefer not to say

- **What is your civil status?**
  - Single. Never married
  - Married or *de facto* couple
  - Separated or divorced
  - Widowed

- **What is your education level?**
  - No schooling completed
  - Primary school
  - Lower secondary school

- High school
- Vocational degree
- College degree
- Master's degree or above

- **Please describe your current occupation/job** (Example: photographer, dental assistant, firefighter, cook, painting worker, financial analyst, etc.). In case you are currently not working, refer to your last job please." _____

Economic policy views

- **How do you think Artificial Intelligence can affect your job?**
  - It can assist and complement my work.
  - It could partially or fully replace my role in the future
  - It does not affect my job
  - I don't know / Not applicable

- **Imagine for a moment that you win a million euros in the lottery. How likely is it that you will spend the money you win on the following things?**

  [4. Very likely – 3. Somewhat likely – 2. Slightly likely – 1. Not at all likely]

  - Quit work and live comfortably: [drop-down menu]
  - Start a business: [drop-down menu]
  - Help family members: [drop-down menu]
  - Donate to charitable causes: [drop-down menu]

- **To what extent do you support or oppose a greater role for the State in the economy?**
  - Strongly support
  - Somewhat support
  - Somewhat opposed
  - Strongly opposed

- **Do you think income differences between rich regions and poor regions in Spain are:**
  - Not a problem at all
  - A small problem
  - A moderate problem

  - A serious problem
  - A very serious problem

- **To what extent do you support or oppose rich regions transferring more resources to poor regions to help finance their public services and infrastructure?**
  - Strongly support
  - Somewhat support
  - Neither support nor oppose
  - Somewhat oppose
  - Strongly oppose

- **Out of every 100 people living in Spain, how many do you think were born in another country?**

- **To what extent do you agree or disagree with the following statement: "Immigration is essential for the long-term growth and stability of our national economy."**
  - Strongly agree
  - Somewhat agree
  - Neither agree nor disagree
  - Somewhat disagree
  - Strongly disagree

- **How likely is it that the poorest regions in Spain today will still be the poorest in 30 years?**
  - Very likely
  - Likely
  - Unlikely
  - Very unlikely

- **When thinking about poorer regions in your country, to what extent do you feel…**
  [A lot, Somewhat, Very Little, Not at All]
  - Empathy or compassion
  - A sense of moral responsibility
  - Pessimism regarding their future development

- **To what extent do you agree or disagree that poor regions…**

  [Strongly agree; Agree; Neither agree nor disagree; Disagree; Strongly disagree]

  - Still pay the price for disadvantages created long ago.
  - Are poorer due to the work attitudes and initiative of their population.

- **To what extent do you think that well-designed economic policies of the government are able to:**

  [A great deal; A lot; A moderate amount; A little; Not at all]

  - Improve the level of education.
  - Reduce the amount of people in risk of poverty.
  - Promote the development of left-behind areas.

- Thinking once more about the income differences between regions in Spain, in your opinion **how much of these differences is due to events and processes that took place a long time ago (more than 100 years ago), as opposed to more recent developments and the current characteristics of each region** — its population, its government, its economy? [*Move the slider to indicate your view*]

  Note to programmers: Handle position starts with no value until first click

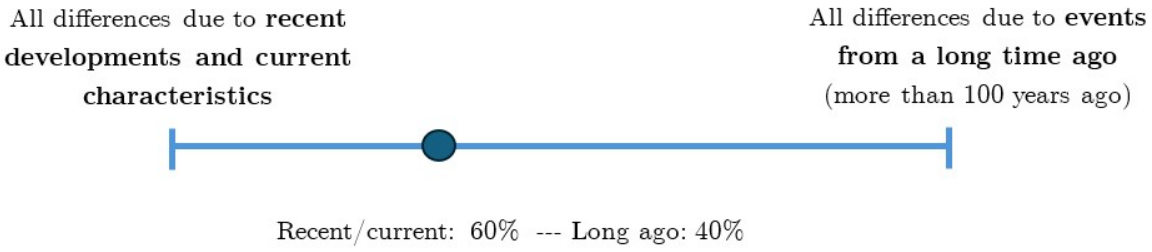

Appendix References

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

Buescu, M. (1981). Regional inequalities in brazil during the second half of the nineteenth century. In *Disparities in Economic Development since the Industrial Revolution*, pages 349–358. Springer.

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(1):100–112.

Diez-Minguela, A., Martinez-Galarraga, J., and Tirado-Fabregat, D. A. (2018). *Regional inequality in Spain*. Springer.

Easterlin, R. (1960). Interregional differences in per capita income, population, and total income, 1840-1950. *Trends in the American economy in the nineteenth century*, pages 73–140.

Grigorieff, A., Roth, C., and Ubfal, D. (2020). Does information change attitudes toward immigrants? *Demography*, 57(3):1117–1143.

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.

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

Lergetporer, P., Werner, K., and Woessmann, L. (2020). Educational inequality and public policy preferences: Evidence from representative survey experiments. *Journal of Public Economics*, 188:104226.

Nunn, N. (2020). The historical roots of economic development. *Science*, 367(6485):eaaz9986.

Sands, M. L. and de Kadt, D. (2020). Local exposure to inequality raises support of people of low wealth for taxing the wealthy. *Nature*, 586(7781):257–261.

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