Biased Information, Biased Preferences: How Fragmented Media Environments Undermine Support for Redistribution*

Oscar Barrera Rodriguez† Lora Pavlova‡ Li Yang§

This version: September 25, 2026

Abstract

When voters choose what to read about the economy, does the choice polarize what they want and whom they vote for, and does forcing the other side on them backfire? We answer with a pre-registered four-arm experiment on 4,524 German adults using two narratives about the 2022–23 inflation surge. Respondents were either forced to read a left-leaning article on inflation's regressive incidence, offered it, or offered a choice between it and a right-leaning article on growth; the three arms separately identify exposure, selection on willingness to read, and selection on the direction of what is read. Selection is confirmatory: right-leaning respondents choose the growth article at 61–67%, left-leaning at 48–52%. Exposure moves preferences toward the narrative read, and more so when it is chosen than when it is imposed (0.27 versus 0.10 standard deviations on a redistribution index). Policy by policy: the left article moves preferences on wealth, inheritance and corporate taxes; the right article moves universal basic income. The two narratives act on different margins: the left article moves policy preferences but not party preference, while the right article raises support for a far-right platform by 3.8 points, consistent with platforms that bundle fiscal with migration positions. Backfire is narrow: forced exposure lowers UBI support only among the most right-leaning fifth, who nonetheless update their factual beliefs as fully as anyone else. Dissonance operates at the step from fact to policy, not on the fact.

Keywords: media fragmentation; confirmation bias; selective exposure; polarization; redistribution; inflation.

JEL codes: D72, D83, D91, H23, L82.

*We gratefully acknowledge financial support from the Leibniz Science Campus MannheimTaxation. The experiment was pre-registered at the AEA RCT Registry (AEARCTR-0014878) and received ethical approval (ZEW-EC-2024-001). We thank seminar participants in Paris for helpful comments. All errors are our own.

†University College Dublin, School of Politics and International Relations; Democracy Challenged Project (ERC). Email: oscar.barrerarodriguez@ucd.ie.

‡ZEW – Leibniz Centre for European Economic Research, Mannheim. Email: lora.pavlova@zew.de.

§ZEW – Leibniz Centre for European Economic Research, Mannheim. Email: li.yang@zew.de.

1 Introduction

What people believe about inequality shapes how much redistribution they want. In the World Values Survey, 29% of Americans think the poor are trapped in poverty against 60% of Europeans; 30% think income is a matter of luck against 54%; and 60% think the poor are lazy against 26%. The welfare states differ accordingly (Alesina and Glaeser, 2004), and the beliefs are self-sustaining: where effort is believed to pay, low redistribution keeps it necessary (Bénabou and Tirole, 2006). The divide is not only across countries but within them, and it follows the parties: explanations for who is rich and who is poor have polarized along partisan lines in the United States, while beliefs about inequalities politicians do not contest have not (Suhay, 2025). Beliefs about the economy are therefore not a summary of the facts. They are an input to policy, and they are formed largely through the media.

The media environment that forms them has fragmented. When a single medium reached everyone, it moderated: the arrival of broadcast television in the United States reduced the ideological extremism of members of Congress, consistent with a shift in voters' own views toward the center (Campante and Hojman, 2013). The problem is not only that people interpret the same facts differently but that they no longer see the same facts: after a single social-media intervention brought a British scandal to prominence, half of one party's voters had heard of it and a tenth of another's had (YouGov, 2025). Algorithmic curation reduces cross-ideological exposure and hardens prior views (Levy, 2021), and consumption is ideologically segregated online as offline (Gentzkow and Shapiro, 2011, Bakshy et al., 2015). In such an environment the tendency to prefer confirming information, confirmation bias, rooted in the discomfort of dissonance (Festinger, 1957, Lord et al., 1979, Taber and Lodge, 2006), is no longer a quirk of individual psychology but a property of how information is supplied.

This paper asks what that property does to preferences for redistribution and to the vote. Two literatures have studied the pieces separately. Information-provision experiments give respondents facts or narratives and measure the response (Kuziemko et al., 2015, Alesina et al., 2018, Andre et al., 2026, Haaland et al., 2023); they do not let respondents decline the information, so they measure the effect of exposure on a population that did not choose it. Studies of selective exposure observe which outlets people consume (Gentzkow and Shapiro, 2011, Chopra et al., 2024); they do not observe what the same people would have done under forced exposure. Yet the two decisions, what to read and how much to believe, are governed by one trade-off between the value of information and the cost of dissonance, and their interaction is the object of interest: if people avoid uncongenial content, exposure effects are effects on the wrong population; if they update selectively, who reads determines what the aggregate looks like. We write both decisions as consequences of one utility function with two parameters, an accuracy weight and a dissonance cost, and design an experiment in which the same mechanism is tested at the moment of choice and at the moment of updating.

Inflation in Germany is the issue. The 2022-23 surge was the largest in the country's

post-war history and made inflation the most salient economic question of the period. Its distributional incidence is poorly understood, the public's understanding of inflation is still scarce (Binetti et al., 2024). Household expectations respond to narratives (Andre et al., 2026), and it admits two coherent, politically coded readings: that inflation hits the poor hardest and calls for redistribution, or that it calls for growth rather than transfers. Both readings appeared in the German press. That gives the experiment two real narratives and a set of redistributive policies on which they disagree.

We embed a pre-registered randomized experiment in an online survey of 4,524 German adults, fielded between December 2024 and February 2025[1]. Respondents are assigned to four arms. A control group reads nothing or an unrelated text. A forced arm reads a left-leaning article on the regressive incidence of inflation, assembled from published reporting with its sources visible. An optional arm sees the article's headline and may open it. A choice arm sees the headlines of that article and of a right-leaning article on growth, and chooses one. The three arms separate the effect of exposure from selection on the willingness to read and from selection on the direction of what is read. Outcomes are six redistributive policies, a redistribution index, a ranking of four party platforms that bundle fiscal with migration positions, and, at the end of the survey, the factual beliefs on which the left article rests.

Four findings. First, selection is confirmatory: 61 to 67% of respondents who voted for right-of-center parties choose the growth article, against 48 to 52% of those who voted for left-of-center parties, and the choice is predicted by partisanship and media diet, not by measured psychological traits. Second, exposure moves preferences toward the narrative read, and more so when the narrative is chosen than when it is imposed: forced exposure raises the redistribution index by 0.10 standard deviations, voluntary reading by 0.19, and choosing the left article by 0.27, while choosing the right article lowers it by 0.09. The effects are specific to the policies each article names, wealth, inheritance and corporate taxes for the left article, universal basic income for the right, and survive correction for multiple testing. Third, the two narratives are not mirror images. On policy the left article has the larger pull; on party preference only the right article moves anything, raising support for a far-right platform by 3.8 percentage points, while the left article moves policy without moving the vote. We attribute the asymmetry to the bundling of fiscal and migration positions (Roemer, 1998): a respondent persuaded on economics can move only to a platform whose migration stance she accepts. Fourth, backfire exists but is narrow. Forced exposure lowers UBI support among the most right-leaning fifth of respondents by 0.16 standard deviations while raising it among the most left-leaning fifth by 0.13; it moves no other policy against the message. The same right-leaning respondents update their factual belief about low-income inflation as fully as anyone else, three-quarters of the way to the article's figure, the same as every other group. They accept the fact and resist its policy implication.

[1] During the fielding of the experiment, inflation had fallen back to 2%. During that time the campaign for the February 2025 federal election took place.

That last result is what the model was built to detect, and it relocates the mechanism. Dissonance governs which article is chosen; accuracy governs what is believed once it is read; dissonance returns at the step from an accepted fact to a policy position, where it produces backfire on the single most contested policy. This is consistent with the finding that information about inequality moves beliefs more than it moves redistributive preferences (Kuziemko et al., 2015) and with identity-protective reasoning (Kahan, 2013), and it is consistent with the replication evidence that backfire is rare and narrow rather than general (Wood and Porter, 2019, Guess and Coppock, 2020).

The paper contributes to three literatures. To the work on beliefs about inequality and support for redistribution, it adds the realistic ingredient that people choose whether to be informed, and shows that the choice, not the information, carries most of the effect. To the work on the demand for news and selective exposure, it adds a randomized counterfactual, the same content under compulsion, and a redistributive rather than partisan issue. To the work on backfire and reactance, it adds a design that separates the belief from the preference and locates the resistance in the second. Relative to field evidence that media shifts votes (Gerber et al., 2009, Martin and Yurukoglu, 2017) and that conservative media has polarized asymmetrically (Grossmann and Hopkins, 2016), it offers a mechanism for the asymmetry that does not rest on the media's slant but on the structure of the political menu.

Section 2 sets out the model and its predictions. Sections 3 and 4 describe the experiment, the sample and the balance of the randomization. Section 5 states the estimating equations and the inference. Section 6 presents the results, Section 7 the robustness checks, and Section 8 the mechanisms, including the estimates of belief updating and the design of a second wave. Section 9 concludes.

2 Conceptual Framework

2.1 Motivation

A voter who encounters news about the economy makes two decisions, not one. She decides what to read, and then she decides how much of it to believe. The literature on media and political attitudes has mostly studied the second decision, by providing information and measuring the response (Alesina et al., 2018, Andre et al., 2026); the literature on selective exposure has mostly studied the first, by observing which outlets people consume (Gentzkow and Shapiro, 2011, Levy, 2021). Yet the two decisions are governed by the same underlying trade-off, and a design that observes only one of them cannot say how they interact. If people avoid uncongenial content, then the effect of information measured under forced exposure is an effect on a population that would never have read it; if people update selectively, then the composition of who reads is a first-order determinant of what the aggregate response looks like. The framework below writes both decisions as consequences of one utility function with two parameters, so that our four-arm

design can test the same mechanism at the moment of choice and at the moment of updating.

2.2 Related Literature

The idea that beliefs are objects people consume and defend, rather than passive summaries of evidence, enters economics with Akerlof and Dickens (1982), who model workers choosing beliefs about job hazards subject to the psychological cost of contradicting them, and is developed into a general framework by Bénabou and Tirole (2016) and Bénabou (2015): agents trade the accuracy of a belief against its desirability, and information is valued not only for what it enables but for how it feels. Our dissonance-cost term is that trade-off written at the point of exposure. The specific functional form we use, a posterior that minimizes a weighted sum of squared distances to the signal and to the prior, is the directional-motives formulation of Little (2019), whose optimum is a convex combination of the two; it is also, as we note below, isomorphic to Gaussian precision-weighting, which is what makes its parameters interpretable to readers who prefer a Bayesian reading.

On the selection side, Chopra et al. (2024) show experimentally that demand for news reflects both accuracy concerns and belief-confirmation motives, and that the second is substantial. Our first-stage utility is their tension in one line, applied to a redistributive issue rather than a partisan one. The older selective-exposure literature (Gentzkow and Shapiro, 2011, Levy, 2021) documents ideological segregation of consumption; the choice arm of our design produces the same phenomenon under randomization, with the counterfactual of forced exposure observed in the same sample.

On the updating side, two strands matter. Rabin and Schrag (1999) model confirmatory bias as the misreading of a signal in the direction of the current belief, and Fryer et al. (2019) show that when evidence is open to interpretation, agents encode it as prior-consistent and then update as if it were unambiguous; both provide the micro-foundation for the attenuation parameter in Section 2.5. The second strand concerns whether updated beliefs translate into policy preferences. Kuziemko et al. (2015) find that information about inequality moves beliefs substantially but redistributive preferences only modestly, and Alesina et al. (2018) find the response to mobility information is asymmetric by political orientation. This gap between accepting a fact and acting on it is where our evidence on backfire turns out to sit (Section 8.3).

Finally, the possibility that counter-attitudinal information entrenches rather than persuades, the backfire effect of Nyhan and Reifler (2010) and Bail et al. (2018), has not fared well in replication: Wood and Porter (2019) and Guess and Coppock (2020) find backfire rare and narrow. Our model predicts exactly that: backfire requires an unusually high dissonance cost on a specific, contested dimension, so it should appear on one policy for one group rather than as a general response. The psychological mechanism we have in mind is identity-protective reasoning (Kahan, 2013, Taber and Lodge, 2006): the fact is accepted, but its policy implication is resisted because it threatens a group-defining position. Explanations for inequality are group-

defining in precisely this sense: the parties supply rival accounts of why some people are poor and others rich, and adherence to those accounts follows party lines (Suhay, 2025).

2.3 Setup: Choosing What to Read

Respondent i holds a prior μ_i about the incidence of inflation, in our application, how much harder inflation hits low-income households, and chooses a source $s \in \{L, R, \text{skip}\}$. Reading s delivers a message $\hat{\theta}_s$: the position the article pushes toward. Utility from reading s is

$$u_i(s) \;=\; \alpha_i \,\mathbb{E}\big[\text{Accuracy} \mid s\big] \;-\; \kappa_i \,\big|\hat{\theta}_s - \mu_i\big| \;+\; \varepsilon_{is}. \tag{1}$$

The first term is the value of information: $\mathbb{E}[\text{Accuracy} \mid s]$ is how much i expects source s to improve her understanding, assessed before reading, and α_i is the weight she places on it. The second term is the dissonance cost: $|\hat{\theta}_s - \mu_i|$ is the distance between the article's message and her prior, and κ_i prices it, in the sense of Festinger (1957) and Akerlof and Dickens (1982). The idiosyncratic term ε_{is} is a standard taste shock. The respondent reads the source with the highest $u_i(s)$; skipping delivers no information and no dissonance.

The same distance enters both terms with opposite signs, and this is the crux of the selection stage. The source one expects to learn most from is often the one disagrees with, because it contains what one has not yet heard; the same source is also the one that stings. When α_i dominates, the respondent seeks out the counter-attitudinal source; when κ_i dominates, she selects the congenial one. The model is agnostic between the two regimes, and Section 6.1 lets the data decide.

2.4 Updating: How Much to Believe

After reading, the respondent forms a posterior. In the motivated-belief reading of Bénabou and Tirole (2016) the posterior is itself a choice: the respondent picks the belief θ that minimizes the loss

$$L_i(\theta) \;=\; \alpha_i \,(\theta - \hat{\theta}_s)^2 \;+\; \kappa_i \,(\theta - \mu_i)^2, \tag{2}$$

trading the cost of being far from the message against the cost of being far from her prior. The first-order condition $2\alpha_i(\theta - \hat{\theta}_s) + 2\kappa_i(\theta - \mu_i) = 0$ gives (Section E)

$$\mu_i^{\text{post}} \;=\; \mu_i + \lambda_i \,(\hat{\theta}_s - \mu_i), \qquad \lambda_i = \frac{\alpha_i}{\alpha_i + \kappa_i} \in (0,1), \qquad \frac{\partial \lambda_i}{\partial \kappa_i} < 0. \tag{3}$$

The posterior is a convex combination of message and prior, and the updating weight λ_i falls with the dissonance cost. Three remarks. First, (3) has the form of Bayesian precision-weighting with α_i in the role of signal precision and κ_i in the role of prior precision; the two readings are observationally equivalent in updating data alone, and we return to what separates them in Section 2.7. Second, only the ratio $\alpha_i : \kappa_i$ is identified from (3), so we normalize $\alpha_i + \kappa_i = 1$ and

read $\lambda_i = \alpha_i$. Third, the model is solved by backward induction: source choice in (1) anticipates the update in (3) and the dissonance it will carry.

2.5 From Beliefs to Preferences, and Backfire

In the baseline, $\lambda_i \in (0, 1)$: a respondent can discount a message but cannot move away from it. Movement away, backfire, requires an additional ingredient, and the evidence in Section 8.3 tells us where it sits. One candidate is skepticism about the messenger: if the respondent treats the article as a persuasive attempt and de-biases it by a share $\rho \geq 0$ of its distance from her prior, the effective weight becomes $\lambda_i^{\text{eff}} = (1 - \rho)\lambda_i$, which turns negative when $\rho > 1$ (Section E); this is Rabin and Schrag (1999)'s misperception for $\rho \leq 1$ and requires source-motive inference beyond it. That mechanism operates on the *belief*. The alternative is that the belief is accepted but its *policy implication* is resisted: a respondent for whom "inflation hits the poor harder" is uncontroversial may still refuse the inference to a universal basic income if that policy is a marker of an out-group. In that case λ_i on the factual belief is high and homogeneous, while the mapping from belief to policy preference is where κ_i bites, and backfire appears on the contested policy only. Our data allow us to distinguish these two accounts, because we observe both the factual belief and the policy preference.

2.6 Predictions

Prediction P1 (Selection)**.** If dissonance costs dominate at the selection stage ($\kappa_i > \alpha_i$ for most respondents), right-leaning respondents are more likely to choose the right-leaning article and left-leaning respondents the left-leaning one.

Prediction P2 (Updating)**.** Conditional on exposure, factual beliefs move toward the article's statistics with weight $\lambda \in (0, 1)$, and preferences move toward the article's position, most strongly on the policies it names. Respondents who did not see a statistic do not update on it.

Prediction P3 (Narrow backfire)**.** Movement against the article under forced exposure is concentrated among respondents whose priors are most distant from it and on the policy for which the dissonance cost is highest, rather than uniform across outcomes. If the belief-based mechanism operates, λ on the factual belief is lower for those respondents; if the preference-based mechanism operates, λ is homogeneous and only the policy preference moves against the message.

2.7 Taking the Model to the Data

The design identifies the model's comparative statics directly. Random assignment to T1, T2 and T3 makes the contrasts in Section 6 causal; the choice arm tests Prediction P1; the forced and voluntary arms test Prediction P2; and the interaction of forced exposure with a pre-treatment

predictor of right-article choice tests Prediction P3. Two objects in the model are, in addition, estimable at the group level from this wave. The two knowledge questions asked at the end of the survey elicit the factual beliefs on which the left article rests, and because the control group received no message, its answer is the counterfactual prior. For any arm, the updating weight is therefore

$$\hat{\lambda} = \frac{\text{posterior} - \text{prior}}{\text{signal} - \text{prior}}, \tag{4}$$

with the posterior the arm's median belief, the prior the control group's, and the signal the statistic the article asserts. Table 8 reports $\hat{\lambda}$ by arm and, within forced exposure, by quintile of right-leaning propensity; the arms that never saw a statistic serve as placebos. What this wave cannot deliver is κ_i at the individual level: that requires eliciting each respondent's prior before treatment, which we did not do to avoid priming, and separating κ from prior precision requires either a measure of confidence or the selection-stage evidence that information can have negative value for a dissonance-averse agent but not for a Bayesian. Both are the object of the follow-up wave described in Section 8.4.

3 Experimental Design

3.1 Sample and Randomization

We embed a pre-registered experiment[2] in an online survey fielded in Germany between 12 December 2024 and 25 February 2025 through the Qualtrics panel, with a monetary completion incentive. Quotas on age, gender and education match population margins; Section 4 compares the sample with the 2021 federal election result to ensure representativeness along this dimension. Qualtrics assigned respondents to four equal groups by block randomization within age×gender×education strata. The control group was split evenly between a pure control (no text) and a placebo group that read a short text on the history of chocolate in Germany; our specifications pool the two, and Section 7 shows that outcomes do not differ between them (Table C.3).

Two attention checks screen inattentive respondents. The first appears immediately after the treatment block. It explains that respondents often answer survey questions without reading them, and then asks "What is your favourite colour?", instructing the respondent to answer *violet* regardless of their actual preference. A respondent who answers anything else is told that the answer is incorrect and asked again, up to three attempts; those who never comply do not proceed and do not appear in the data. Of the analysis sample, 3,640 respondents (80.5%) comply on the first attempt, 771 on the second and 113 on the third. Table 5 re-estimates the main results on the first-attempt compliers alone. A second check is embedded among the redistributive-policy items: formatted like the surrounding 0–10 questions, it instructs respondents to ignore the

[2]AEA RCT Registry AEARCTR-0014878; ethics approval ZEW-EC-2024-001.

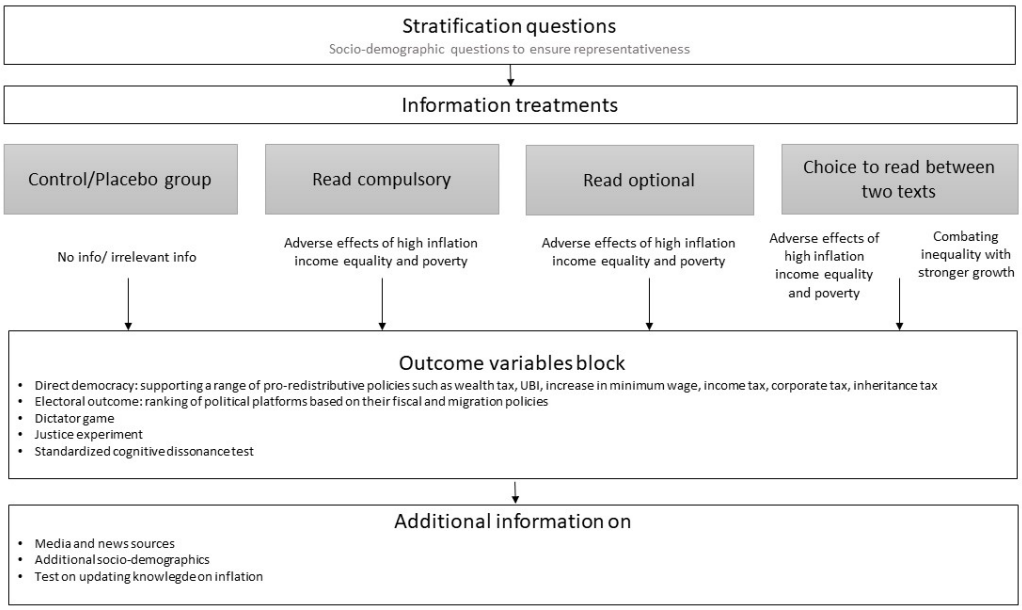

Figure 1. Survey flow. Stratification block, treatment block, outcome block, covariate block, and knowledge questions, in order of administration.

usual response options and select the value "seven".

The analysis sample comprises 4,524 respondents. Of the 4,533 who completed the survey, we drop two with missing education and seven whose federal state is not recorded; the latter are excluded because the federal state is the level at which we cluster standard errors (Section 5.3), leaving all sixteen German states represented. Figure 1 summarizes the survey flow; the questionnaire was administered in German and is reproduced in translation in Section B.

3.2 Treatments

The treatment block follows the stratification questions so that reading, which takes about two minutes, occurs at peak attention. Two texts are used. The *left-leaning article* ("Social gap in inflation at its highest level") documents the regressive incidence of the 2022–23 inflation surge and proposes an excess-profits tax, a wealth tax, and a higher minimum wage. The *right-leaning article* ("Handouts won't save Germany") argues that growth, not social benefits, is the answer and proposes corporate tax cuts. Both are composites of published German reporting with the original sources linked in the text (TAZ, DGB, and the Hans-Böckler Foundation; Welt, Focus, DW, and the Financial Times), so that readers can see each article's political alignment. Full texts are in Section A. The three treated groups differ only in whether, and how, respondents choose what to read:

- **T1 (forced exposure).** Respondents are directly shown the full left-leaning article.
- **T2 (optional exposure).** Respondents see the left-leaning article's headline and a one-sentence description and can choose whether to open it; 87% do.

- **T3 (choice).** Respondents see the headline and description of both articles, in randomized order, and can choose one to read; 57% choose the right-leaning article.

T1 thus identifies the effect of exposure, T2 adds selection on willingness to read, and T3 adds selection on direction (Section 5).

3.3 Outcomes

All groups answer the same outcome block. *Redistributive-policy preferences* cover six policies, phrased neutrally, e.g. "To what extent do you agree or disagree with the introduction of a wealth tax in Germany that requires high net worth individuals to pay a percentage of their total wealth as tax?" (0–10). Universal basic income and a minimum-wage increase use the same scale; the preferred top income-tax rate for a €280,000 earner is a 0–100 slider anchored on the information that a minimum-wage earner pays about 27%; corporate and inheritance tax are sliders around their current rates. We summarize the six in a redistribution index, the first principal component of the standardized items (Section 5).

Electoral preferences are elicited by ranking four fictitious platforms that each combine a fiscal and a migration position (Table B.1), shown in random order with a blank-vote option; our outcomes are indicators for the platform ranked first.

Cognitive dissonance is measured by three 1–7 items on discomfort with opposing political views, whether disagreement leads one to question one's beliefs, and difficulty reconciling opposing arguments; they enter the propensity model of Section 5.4.

Distributive-justice orientation is elicited by a question in which respondents award a flute to one of three children, the most talented, the most deprived, or the one who inherited it, coded as talent, need, and inheritance (Section B).

Political self-placement (0–10, far left to far right) is elicited at the end of the outcome block rather than before treatment, to avoid priming. It is therefore post-treatment, and we use it only descriptively and in the propensity model of Section D, where results are shown to be unchanged when it and the other post-treatment items are excluded.

The block also contains a dictator game and, at the end of the survey, two knowledge questions on the 2022 inflation rate overall and for low-income singles, the statistics reported in the left-leaning article, and a question on the perceived political slant of the survey. The knowledge questions serve as a manipulation check and as a belief-updating measure (Section 8); wording for all items is in Section B.

3.4 Covariates

After the outcome block we elicit media consumption, main and secondary news source among broadcast and print, online portals, and social media, with outlet-level indicators within each, and socio-demographic characteristics (education, income, marital status, religion, region of

birth, employment, home ownership, parental education and birth region) together with the 2021 federal vote. These variables are elicited after treatment but are predetermined; they serve as controls and, together with the dissonance and justice items, as candidate predictors in Section D.

4 Data, Balance and Descriptives

4.1 Balance and Representativeness

Table C.1 reports means of the stratification variables, socio-economic characteristics, 2021 vote and the justice question by arm, with p-values for the difference between each treatment arm and the control group. Of the 120 pairwise comparisons, four are nominally significant at the 5% level (age 35–44 and Catholic affiliation in T3; age 65+ and household income above €6,000 in T2), roughly the number expected by chance, and none survives the Romano–Wolf correction in the last column. Justice orientations, elicited by the flute question (Sen, 2009), are likewise balanced (43% allocate to the child in need, 28% to the most talented, 21% to the heir). An omnibus multinomial logit of arm assignment on all covariates does not reject balance (likelihood-ratio $\chi^2(162) = 182.7$, $p = 0.13$); arm-specific joint F-tests are marginal for T2 ($p = 0.005$) and T3 ($p = 0.02$), reflecting the same few small differences. Since assignment was block-randomized by the survey provider, these are chance imbalances; all specifications include the full covariate set, and Section 7 shows that results are insensitive to it.

The sample is a quota sample, not a probability sample, and column (5) shows where it departs from the SOEP (2022) population benchmarks. Education and gender match closely. The age distribution is compressed toward the middle: respondents aged 18–24 (10% vs. 17%) and 65+ (12% vs. 21%) are under-represented, and full-time employment is over-represented (54% vs. 37%). Protestants are under-represented (16% vs. 33%). Table C.3 reports estimates reweighted to these margins; the main results are unaffected.

Reported 2021 voting departs from the official result in one dimension: 17.6% of party voters report having voted AfD, against an actual 10.4%, while non-voters are under-represented (15% vs. 23%). Two explanations are plausible and not mutually exclusive. The panel may over-sample AfD supporters. Alternatively, recall of a vote cast three years earlier may drift toward current preference: the survey was fielded between December 2024 and February 2025, during the campaign for the federal election of 23 February 2025 in which the AfD obtained 20.8%, and the sample share lies between the 2021 and 2025 results. Because the 2021 vote enters only as a control and as a descriptive split, neither explanation affects the treatment estimates; for external validity Table C.3 reweights to both benchmarks which incorporates both AFD shares (10,4% for 2021 and 20.8% as for 2025).

4.2 Baseline Preferences for Redistribution

Figure 2 shows the control-group distribution of the six policy outcomes. Three features matter for reading the treatment effects. First, support is not uniform across instruments: a minimum-wage increase draws near-consensus (80% at 6 or above; 35% at the maximum), the wealth tax a clear majority (65%), and universal basic income a divided public (57% at 6 or above, but 35% at 4 or below), so UBI is the most polarized of the three agreement items and the one with the most room to move in either direction. Second, the two slider outcomes carry a large mass at the status quo (30% for corporate tax, 19% for inheritance tax), and the inheritance-tax distribution leans toward *reduction* (45% below the status quo vs. 37% above), so the positive left-article effect on inheritance tax in Table 2 moves respondents from a baseline that favors cutting the tax. Third, the preferred top income-tax rate has a median of 38% (interquartile range 28–50%) against the 27% anchor given for a minimum-wage earner. Because outcomes are standardized to the control group, the standard deviations reported in Figure 2 are the units of the effect sizes in Section 6.

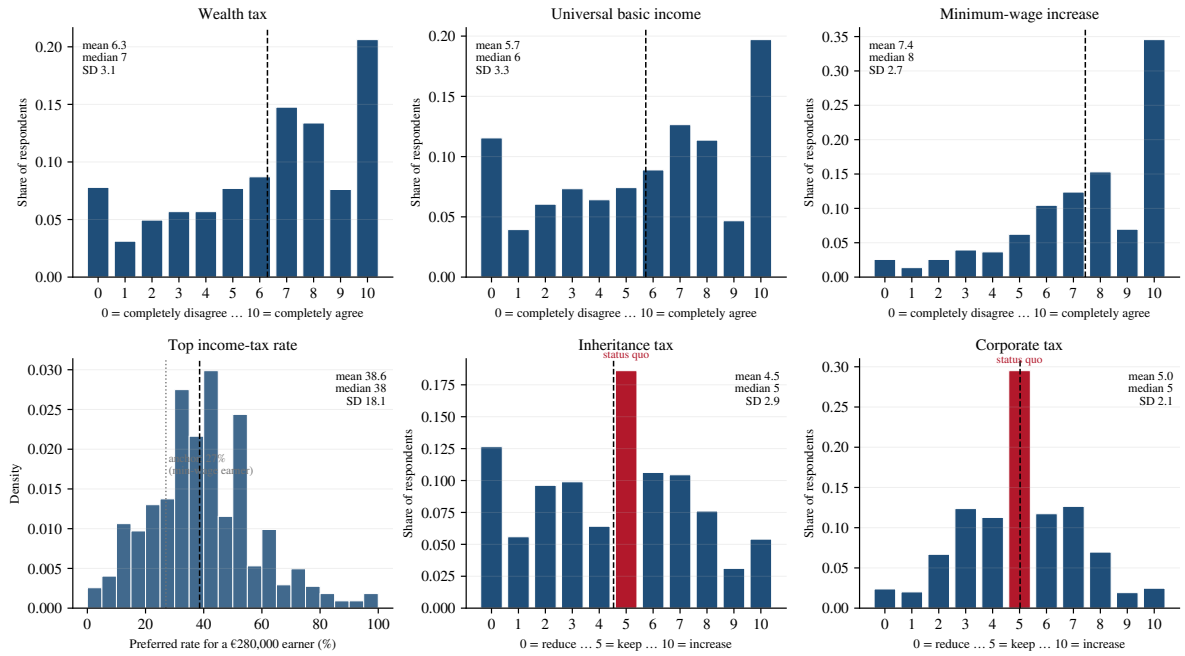

Figure 2. Distribution of the six redistributive-policy outcomes in the control group. Dashed line: mean. The top income-tax question is anchored on the information that a minimum-wage earner pays about 27%; the inheritance- and corporate-tax sliders are centered on the current rate (status quo = 5). N = 1,091.

4.3 Perceived Inflation

At the end of the survey respondents were asked for the 2022 inflation rate overall and for low-income singles, the two statistics on which the left article rests. Figure 3 reports control-group answers by household income. Two facts stand out. Perceptions are, at the median, accurate and

flat across income: 6.9–7.0% in every income group against an official 7.9%, and 7.0–8.2% for low-income singles against the Böckler Foundation's 10.7%. Means tell a different story only because of a thin right tail of implausible answers (3% of respondents report rates above 50%), which is heavier among low-income respondents and generates a spurious income gradient in the mean. We therefore use medians, or means winsorized at 50%, throughout. These two questions also serve as the post-treatment belief measures in Section 8.

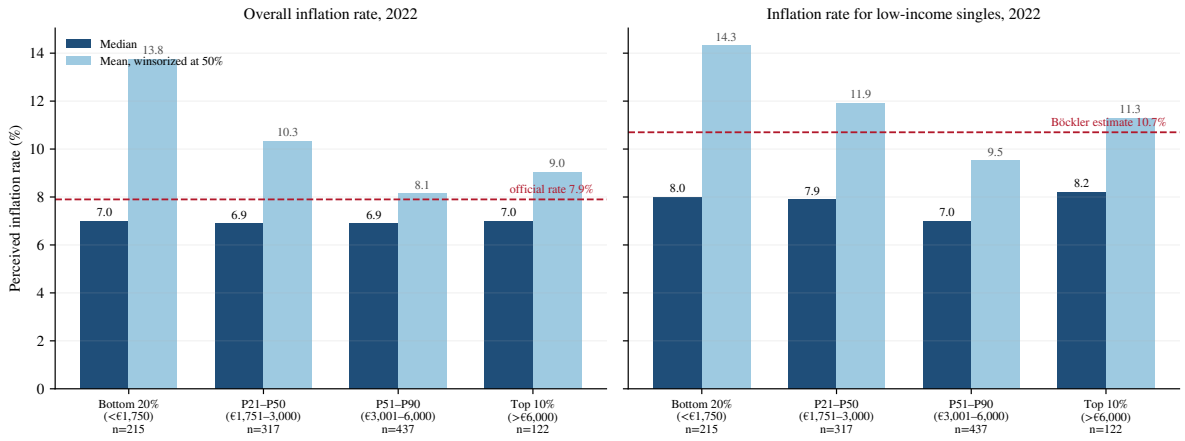

Figure 3. Perceived 2022 inflation by household income, control group. Dark bars: medians; light bars: means after winsorizing at 50%. Dashed lines: the official overall rate (7.9%) and the Böckler Foundation estimate for low-income singles (10.7%) cited in the left article.

5 Empirical Strategy

5.1 Estimating Equations

Assignment to the four arms is random, so differences in outcomes across arms are causal. Our baseline specification estimates intent-to-treat effects by arm,

$$Y_i = \alpha + \beta_1 \, \mathrm{T1}_i + \beta_2 \, \mathrm{T2}_i + \beta_3 \, \mathrm{T3}_i + X_i' \gamma + \varepsilon_i, \tag{5}$$

where Y_i is an outcome, $\mathrm{T1}_i$–$\mathrm{T3}_i$ are assignment indicators with the no-article control as the omitted category, and X_i collects the stratification variables (age, gender, education) and the socio-demographic controls listed in Section 3.4. The coefficients β_1–β_3 are the effects of being *offered* each arm.

Because T2 and T3 record what respondents chose, we also estimate the decomposed specification

$$Y_i = \alpha + \beta_1 \, \mathrm{T1}_i + \beta_{2r} \, \mathrm{T2}_i^{\mathrm{read}} + \beta_{2s} \, \mathrm{T2}_i^{\mathrm{skip}} + \beta_{3\ell} \, \mathrm{T3}_i^{\mathrm{left}} + \beta_{3r} \, \mathrm{T3}_i^{\mathrm{right}} + X_i' \gamma + \varepsilon_i, \tag{6}$$

which splits T2 by whether the article was opened and T3 by which article was chosen. The

interpretation differs across coefficients and we are explicit about it. β_1 is the causal effect of exposure to the left article on a randomly assigned population. β_{2r}, β_{2s}, $\beta_{3\ell}$ and β_{3r} condition on a post-assignment choice and are therefore effects of an exposure-plus-selection bundle: the effect of reading the left article *on the kind of respondent who opens it*, and so on. We do not attempt to purge selection from these coefficients, there is no instrument for the reading choice, and we do not need to. The bundle is the object of interest, because it is what a fragmented media environment delivers, and β_1 provides the benchmark of the same content without selection. The contrast between β_1 and $\beta_{3\ell}$, identical content read under compulsion versus by choice, is the central comparison of the paper. The symmetry test $H_0 : \beta_{3\ell} + \beta_{3r} = 0$, that the two articles have effects of equal size and opposite sign on those who choose them, is a clustered Wald test on (6).

5.2 Outcomes and Scaling

The six policy outcomes are standardized to the control group (mean zero, standard deviation one), so coefficients are effect sizes in control-group standard-deviation units and are comparable across instruments with different scales. The redistribution index is the first principal component of the six standardized items, rescaled the same way; it explains 37.7% of their variance with all loadings positive. Platform outcomes are indicators for ranking a platform first, estimated by linear probability models with coefficients in percentage points; probit average marginal effects are reported as a check.

5.3 Inference

Standard errors are clustered at the federal-state level ($G = 16$). Because asymptotic cluster-robust inference is unreliable with this few clusters, every headline coefficient is accompanied by a wild-cluster bootstrap-t p-value with the null imposed and Webb six-point weights (Cameron et al., 2008); Table 5 reports both. For the six policy outcomes we control the family-wise error rate within each treatment across the six tests using the Romano and Wolf (2005) step-down procedure with 1,000 cluster-bootstrap replications, and report unadjusted and adjusted p-values side by side. **[TODO: State the pre-registered families and hypotheses (AEARCTR-0014878) and mark which analyses below are exploratory: the propensity-based heterogeneity and the λ estimates were not pre-specified in this form.]**

The registration (AEARCTR-0014878), published on 3 December 2024 before fielding began, specifies preferences for redistributive policies as the primary outcome and states two directional hypotheses: that information on the regressive incidence of inflation increases support for redistribution, and that information on the growth costs of redistribution decreases it. The pre-registration does not specify multiple-testing families; we define them as the six policy outcomes within each treatment, matching the registered primary outcome. We treat the intent-to-treat estimates of each treatment on these outcomes as the confirmatory tests.

The following analyses are exploratory: the redistribution index; the decomposition of T2 by whether respondents read the article and of T3 by which article they chose; the propensity-score heterogeneity analysis, including the λ estimates; the LASSO models of article choice; and the use of the cognitive-dissonance items as predictor.

5.4 Heterogeneity and Backfire

Prediction P3 concerns respondents for whom the left article is incongruent. We proxy congruence with $\hat{p}_i$, the predicted probability of choosing the right article, obtained in two steps: an ℓ_1-penalized logit of right-article choice on 77 baseline predictors (stratification and socio-demographic variables, media consumption, the three cognitive-dissonance items and the justice question), estimated on the choice arm with the penalty chosen by 5-fold cross-validation, retains 16 predictors; an unpenalized logit on those 16, again estimated on the choice arm, is then applied to every respondent (Section D). The fitted score has a cross-validated AUC of 0.64: it predicts choice, but coarsely, so it should be read as a noisy proxy for ideological alignment. Because $\hat{p}_i$ is a function of covariates and T1 is randomly assigned, the interaction $\mathrm{T1}_i \times \hat{p}_i$ in (6), and the corresponding quintile interactions, estimate heterogeneous treatment effects without further assumptions. Two caveats. The attitudinal predictors were elicited after treatment; they are balanced across the randomized arms (Table C.1), and results are unchanged when they are dropped from the score. And measurement error in $\hat{p}_i$ attenuates the interaction, so the backfire estimates are conservative.

5.5 Belief Updating

The knowledge questions at the end of the survey elicit the factual beliefs on which the left article rests, and the control group's answer is the counterfactual prior under random assignment. For each arm we estimate the updating weight of (3) as $\hat{\lambda} = (\text{posterior} - \text{prior})/(\text{signal} - \text{prior})$, with medians in place of means because a thin right tail of implausible answers dominates the latter, and cluster-bootstrap confidence intervals. Arms that never saw a statistic serve as placebos. $\hat{\lambda}$ is a group-level quantity; individual dissonance costs are not identified in this wave (Section 2.7).

6 Results

We present results in the order of the three predictions: who chooses which article (Prediction P1), how exposure moves preferences (Prediction P2), and where forced exposure backfires (Prediction P3). Throughout, policy outcomes and the index are in control-group standard-deviation units and platform outcomes in percentage points.

6.1 Selective Exposure

In the choice arm, 57% of respondents open the right (growth) article. Choice is confirmatory (Figure 4): among respondents who voted for a right-of-center party in 2021 the share is 67% (CDU/CSU), 65% (AfD) and 61% (FDP); among left-of-center voters it is 52% (Linke), 49% (SPD) and 48% (Greens). The share rises by 5 percentage points per step of political self-placement ($p < 0.001$) and is unaffected by the order in which the two headlines were shown (a recency effect of about 3 points). This is Prediction P1: at the point of choice, dissonance costs outweigh the informational pull of the unfamiliar source.

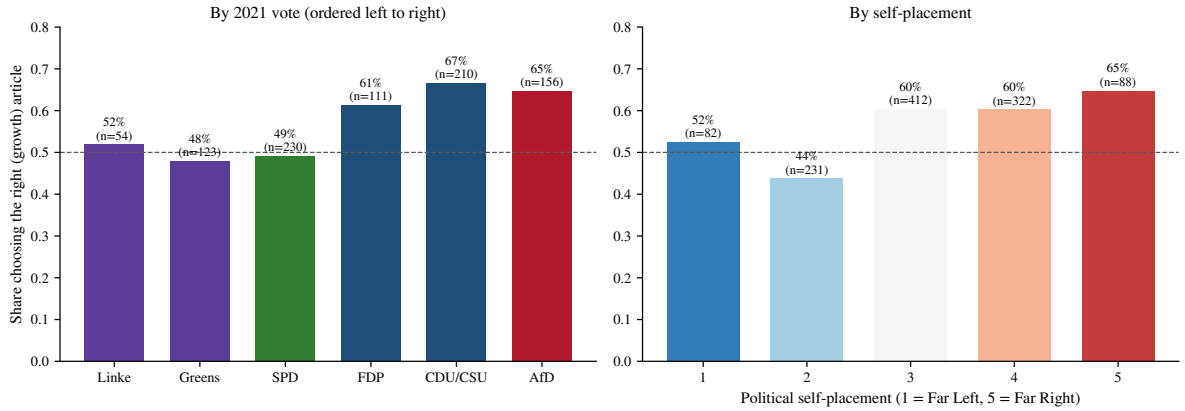

Figure 4. Article choice in the choice arm (T3). Share of respondents choosing the right (growth) article by 2021 vote and by political self-placement. Right-leaning respondents choose the congenial article more often.

Who selects is predicted by partisanship and media diet, not by psychological disposition. Of 77 candidate predictors, the LASSO retains 16 (Table D.1): readers of *Bild*, *Welt* and Yahoo News, full-time employees and CDU/CSU voters are more likely to choose the growth article; SPD voters, respondents who place themselves on the left, those aged 25–34, and users of Instagram, Snapchat and TikTok are less likely. None of the three cognitive-dissonance items and neither the justice question nor education is retained. The fitted model predicts choice only coarsely (cross-validated AUC 0.64), which is itself informative: most of the variation in what people choose to read is not explained by who they are on observables.

Engagement with the content does not depend on whether it was chosen. In the optional arm 87% open the article. Asked afterwards whether they would share it, 42% of forced readers and 45% of voluntary readers say yes, as do 46% of those who chose the left article in T3; among those who chose the right article the share is 36%.

6.2 Effects on the Redistribution Index

Table 1 reports effects on the redistribution index; Figure 5 plots the decomposed specification. Every arm that delivers the left article raises support for redistribution. As intent-to-treat effects, forced exposure raises the index by 0.10 SD, the optional arm by 0.16 SD and the choice arm,

which mixes readers of both articles, by 0.06 SD (column 1). The decomposition (column 2) orders the channels: forced exposure 0.10 SD, voluntary reading 0.19 SD, choosing the left article 0.27 SD, all significant at 5% or better, while choosing the right article lowers the index by 0.09 SD ($p = 0.025$). Respondents who skip the article in T2 are not different from the control group (−0.09 SD, $p = 0.26$). The equal-weight index gives the same ordering (column 3).

Table 1. Treatment effects on the redistribution index (control-group SD units)

	Redistribution index (PC1)		Equal-weight z-index
	(1) Pooled arms	(2) Decomposed arms	(3) Decomposed arms
T1: Forced (left article)	0.103**	0.103**	0.094**
	(0.043)	(0.043)	(0.043)
T2: Optional	0.154***		
	(0.025)		
T2: Read (left article)		0.190***	0.176***
		(0.026)	(0.026)
T2: Skipped		-0.093	-0.126
		(0.082)	(0.088)
T3: Choice	0.062**		
	(0.032)		
T3: Chose left		0.265***	0.251***
		(0.039)	(0.041)
T3: Chose right		-0.093**	-0.114***
		(0.041)	(0.042)
Observations	4,524	4,524	4,524
Adjusted R^2	0.126	0.136	0.132

Notes: The redistribution index is the first principal component of the six standardized policy outcomes (PC1 explains 37.7% of the variance; all six loadings positive, range 0.28 - 0.51), rescaled to mean 0 and SD 1 in the control group; column (3) applies the same rescaling to the equal-weighted average of the six z-scores. Column (1) pools each arm; columns (2) and (3) decompose T2 by whether the respondent read or skipped the article and T3 by which article was chosen. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

Two qualifications matter for how much weight each coefficient can bear. First, the voluntary and directional effects are robust to few-cluster inference: the wild-cluster bootstrap p-values for T2-read and T3-left are below 0.001 (Table 5, Panel C). The forced-exposure effect is not: its bootstrap p-value is 0.12, and that of T3-right is 0.08. We therefore treat the ordering, forced < voluntary < chosen, as the finding, and the T1 and T3-right index coefficients as suggestive. Second, among respondents who passed the attention check on the first attempt the point estimates are unchanged but T3-right loses precision (−0.11 SD, $p = 0.11$; Table 5, Panel A).

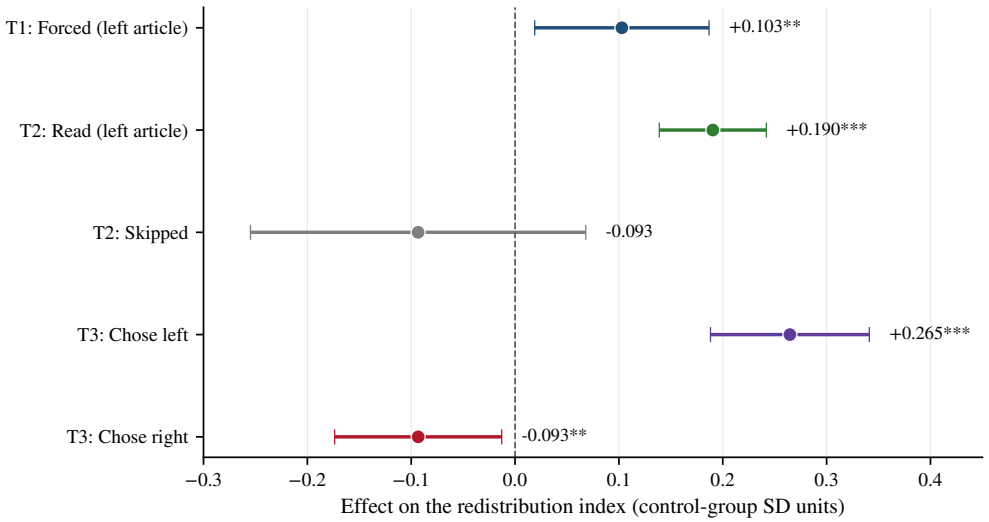

Figure 5. Treatment effects on the redistribution index. Point estimates and 95% confidence intervals from the decomposed specification (Table 1, column 2). Standard errors clustered at the federal-state level ($G = 17$).

6.3 Narrative-Specific Effects

The index conceals which policies move. Table 2 and Figure 6 show that each article moves the instruments it names. Readers of the left article, which proposes a wealth tax, an excess-profits tax and a higher minimum wage, raise their support for the wealth tax (0.16 SD when read voluntarily, 0.24 SD when chosen), the inheritance tax (0.16 and 0.21 SD) and the corporate tax (0.20 and 0.23 SD). These six effects survive Romano–Wolf correction across the six outcomes ($p \leq 0.035$; Panel B). Readers of the right article, which attacks the citizen's allowance and argues for tax cuts, lower their support for universal basic income by 0.20 SD (Romano–Wolf $p = 0.068$) and for a minimum-wage increase by 0.11 SD ($p = 0.125$ after correction). No arm moves the preferred top income-tax rate. Forced exposure moves the corporate tax (0.16 SD, corrected $p = 0.066$) and little else.

The skippers in T2 are the mirror image of the readers. They are not moved by an article they did not open, yet they report lower support for UBI (−0.18 SD) and the minimum wage (−0.18 SD) than the control group. Since assignment to T2 is random, this is selection: respondents who decline the left article are those least sympathetic to the transfers it defends.

6.4 Two Asymmetries: Policy Moves Left, Politics Moves Right

If the two articles were mirror images, the effect of choosing left and the effect of choosing right would sum to zero. Table 3 rejects this in two opposite directions. On *policy*, the left article dominates: the sum is +0.17 SD on the index ($p = 0.005$), +0.22 on the wealth tax ($p = 0.001$) and +0.24 on the inheritance tax ($p < 0.001$), while on UBI the right article dominates (−0.16, $p = 0.054$) and the minimum wage is symmetric. On *party* preference the asymmetry runs the

Table 2. Treatment effects on six redistributive-policy outcomes

	Wealth tax (1)	UBI (2)	Min. wage (3)	Income tax (4)	Inheritance tax (5)	Corporate tax (6)
Panel A: OLS coefficients (control-group SD units)						
T1: Forced (left article)	0.051	-0.022	0.049*	0.003	0.087**	0.165***
	(0.034)	(0.026)	(0.027)	(0.063)	(0.040)	(0.043)
T2: Read (left article)	0.164***	-0.015	0.088***	0.025	0.163***	0.198***
	(0.033)	(0.023)	(0.030)	(0.032)	(0.042)	(0.040)
T2: Skipped	0.056	-0.185*	-0.176**	-0.046	-0.104	0.000
	(0.074)	(0.109)	(0.089)	(0.069)	(0.078)	(0.081)
T3: Chose left	0.235***	0.036	0.117*	0.058	0.214***	0.227***
	(0.038)	(0.053)	(0.062)	(0.066)	(0.032)	(0.047)
T3: Chose right	-0.020	-0.198***	-0.107***	-0.037	0.024	-0.073
	(0.038)	(0.047)	(0.034)	(0.030)	(0.046)	(0.048)
Panel B: Multiple-testing adjustment (p-values)						
T1 Forced – unadjusted p	0.135	0.396	0.073	0.956	0.029	0.000
Romano–Wolf p	0.425	0.672	0.344	0.953	0.278	**0.061**
T2 Read – unadjusted p	0.000	0.527	0.003	0.432	0.000	0.000
Romano–Wolf p	**0.003**	0.571	**0.027**	0.571	**0.010**	**0.003**
T3 Left – unadjusted p	0.000	0.490	0.058	0.374	0.000	0.000
Romano–Wolf p	**0.015**	0.650	0.325	0.650	**0.012**	**0.034**
T3 Right – unadjusted p	0.595	0.000	0.001	0.205	0.606	0.125
Romano–Wolf p	0.815	**0.051**	0.113	0.523	0.815	0.523
Control-group mean (raw scale)	6.29	5.73	7.44	38.59	4.53	5.02
Control-group SD (raw scale)	3.13	3.34	2.66	18.14	2.87	2.08
Observations	4,524	4,524	4,524	4,524	4,524	4,524
Adjusted R^2	0.125	0.092	0.146	0.059	0.130	0.064

Notes: Outcomes are standardized to the control group (mean 0, SD 1), so coefficients are effect sizes in control-group standard-deviation units. Raw scales: wealth tax, UBI and minimum wage are 0–10 agreement scales; income tax is the preferred top rate in percent (0–100); inheritance and corporate taxes are slider responses (higher = more redistribution); raw control-group means and SDs are reported in the footer. Panel B reports, for each treatment, the unadjusted p-value and the Romano–Wolf (2005) step-down p-value controlling the family-wise error rate across the six outcomes (1,000 cluster-bootstrap replications; bold: adjusted $p < 0.10$); p-values are invariant to the standardization. T2 Skipped is omitted from Panel B because of its small cell ($n = 149$). Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

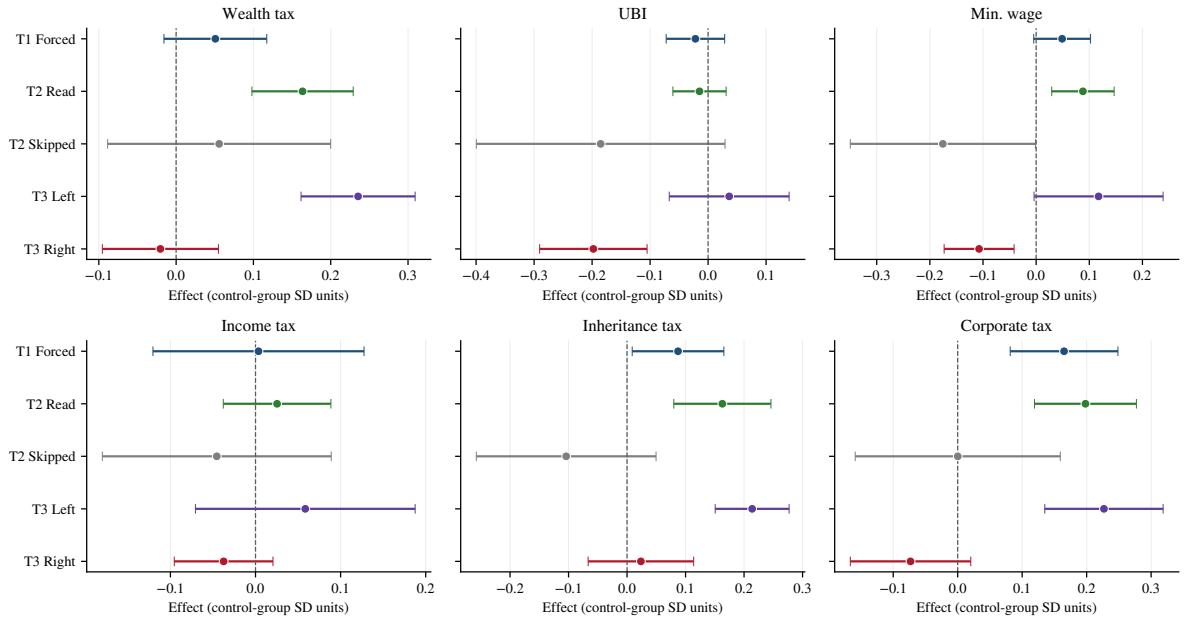

Figure 6. Treatment effects on six redistributive-policy outcomes. Coefficients from Table 2, Panel A, in control-group standard-deviation units; 95% confidence intervals.

other way. Table 4 and Figure 7 report effects on ranking each platform first. Choosing the right article raises the probability of ranking the far-right platform first by 3.8 points ($p < 0.01$), and forced exposure to the left article raises it by 2.5 points ($p < 0.05$). Choosing the left article does not raise support for either left platform (2.3 and 1.8 points, neither significant); its only effect on party preference is to lower the ranking of the Center-Right platform by 3.9 points ($p < 0.05$), an effect voluntary readers in T2 share (−3.1 points). On the far-right platform the sum of the two T3 effects is +5.0 points ($p = 0.053$).

The left article therefore changes what people want without changing who they would vote for; the right article changes both. Section 8.2 interprets this through the bundling of fiscal and migration positions in the platforms.

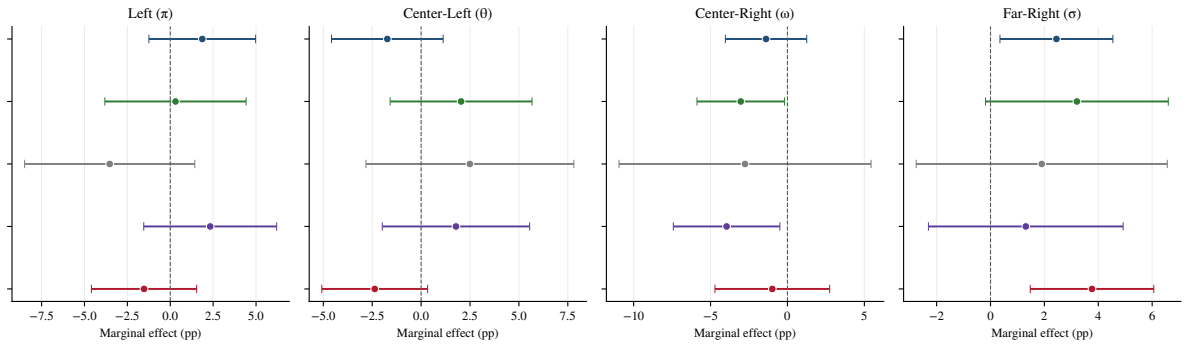

Figure 7. Marginal effects on political-platform choice. Linear-probability estimates in percentage points (Table 4); 95% confidence intervals.

Table 3. Symmetry tests: do the left and right articles have mirror-image effects?

Outcome	$\hat{\beta}_L$	$\hat{\beta}_R$	$\hat{\beta}_L + \hat{\beta}_R$	(SE)	F	p
Redistribution index	0.265	-0.093	0.171	(0.062)	7.66	0.006***
Wealth tax	0.235	-0.020	0.215	(0.065)	10.85	0.001***
UBI	0.036	-0.198	-0.161	(0.085)	3.61	0.057*
Min. wage	0.117	-0.107	0.010	(0.081)	0.02	0.901
Income tax	0.058	-0.037	0.021	(0.083)	0.06	0.801
Inheritance tax	0.214	0.024	0.238	(0.066)	13.15	0.000***
Corporate tax	0.227	-0.073	0.154	(0.087)	3.13	0.077*
Left (π)	2.33	-1.53	0.80	(3.02)	0.07	0.791
Center-Left (θ)	1.78	-2.37	-0.59	(2.25)	0.07	0.793
Center-Right (ω)	-3.96	-0.98	-4.94	(3.16)	2.45	0.118
Far-Right (σ)	1.31	3.77	5.08	(2.61)	3.77	0.052*
Restrictive-migration bundle ($\omega \cup \sigma$)	-2.65	2.78	0.14	(2.82)	0.00	0.962

Notes: Test of $H_0 : \beta_L + \beta_R = 0$, i.e., that choosing the left and the right article have equal-magnitude effects of opposite sign, from the decomposed specification. A significantly positive sum indicates the left article dominates (net leftward asymmetry); a significantly negative sum that the right article dominates. Index and policy outcomes in control-group SD units; platform outcomes in percentage points. F is the clustered Wald statistic with one restriction. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

Table 4. Marginal effects on political-platform choice

	Top-ranked platform				Bundle
	Left (π)	Center-Left (θ)	Center-Right (ω)	Far-Right (σ)	Restrictive migration
	(1)	(2)	(3)	(4)	(5)
T1: Forced (left article)	1.87	-1.73	-1.39	2.45**	1.06
	(1.59)	(1.46)	(1.35)	(1.07)	(1.16)
T2: Read (left article)	0.30	2.04	-3.04**	3.21*	0.17
	(2.10)	(1.85)	(1.45)	(1.73)	(1.26)
T2: Skipped	-3.53	2.49	-2.76	1.90	-0.86
	(2.53)	(2.71)	(4.18)	(2.38)	(3.98)
T3: Chose left	2.33	1.78	-3.96**	1.31	-2.65
	(1.98)	(1.92)	(1.77)	(1.84)	(1.61)
T3: Chose right	-1.53	-2.37*	-0.98	3.77***	2.78
	(1.57)	(1.38)	(1.90)	(1.17)	(1.69)
Control-group share (%)	11.5	21.8	31.8	25.6	57.4
Observations	4,524	4,524	4,524	4,524	4,524

Notes: Linear probability models; coefficients are marginal effects in percentage points on the probability of ranking the platform first. Platforms bundle a fiscal and a migration position: π = open borders and radical redistribution; θ = inclusive migration and progressive taxation; ω = selective skilled migration and moderate tax cuts; σ = strict migration controls and corporate tax cuts. Column (5) is an indicator for top-ranking either restrictive-migration platform. Probit average marginal effects are reported in Table 5. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

6.5 Backfire

Prediction P3 asks whether forced exposure moves the respondents least disposed to the left article *against* it. Figure 8 interacts forced exposure with the right-leaning propensity $\hat{p}_i$. On the index and on every tax-the-rich policy the interaction is zero: right-leaning respondents forced to read the article do not lower their support for redistribution as a whole. On UBI it is negative, −0.35 SD per unit of $\hat{p}_i$ ($p = 0.059$; $p = 0.018$ and 0.013 under the alternative penalty rules reported in Table C.4). Figure 9 shows the shape of the effect: forced exposure raises UBI support by 0.13 SD in the most left-leaning fifth ($p < 0.01$) and lowers it by 0.16 SD in the most right-leaning fifth ($p < 0.01$), with the middle quintiles near zero. The same article pushes the two ends of the distribution apart on the one policy that is most contested, and nowhere else. The far-right platform shows no such heterogeneity ($p = 0.55$).

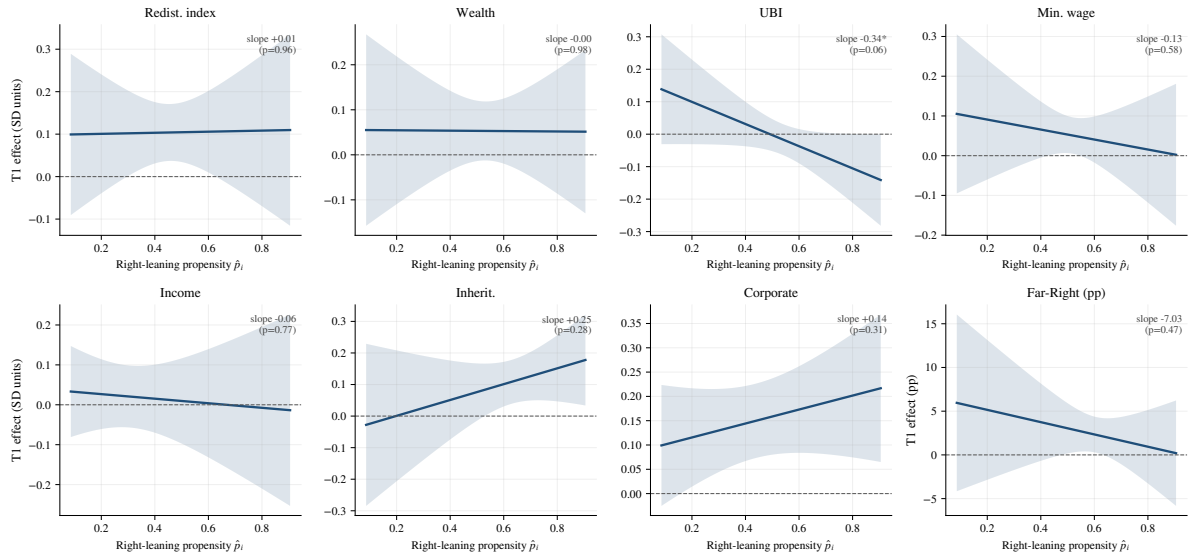

Figure 8. Backfire (I): the forced-exposure effect as a continuous function of right-leaning propensity. Predicted effect of T1 at each value of $\hat{p}_i$, with 95% confidence bands, from a linear interaction model that also includes the other treatment arms and controls. Outcomes in control-group SD units except the far-right platform (percentage points). The annotated slope is the coefficient on T1 × $\hat{p}_i$; a negative slope on a pro-redistribution outcome indicates that right-leaning respondents move *against* the article's message.

Backfire is not disbelief. Table 8 estimates the updating weight λ from the knowledge questions. Readers of the left article move their median belief about inflation for low-income households from the control group's 7.9% to 10.0%, three-quarters of the way to the article's 10.7% ($\hat{\lambda} = 0.75$, CI [0.74, 0.82]), whether forced, voluntary or chosen; skippers and right-article readers, who never saw the figure, do not move ($\hat{\lambda} = 0.00$ and 0.04). Within forced exposure $\hat{\lambda}$ is flat across right-leaning propensity, 0.74–0.82 from the most left- to the most right-leaning quintile. The respondents whose UBI preference moves against the article are the ones who accept its central fact as fully as anyone else. Section 8.3 draws out what this implies for the mechanism.

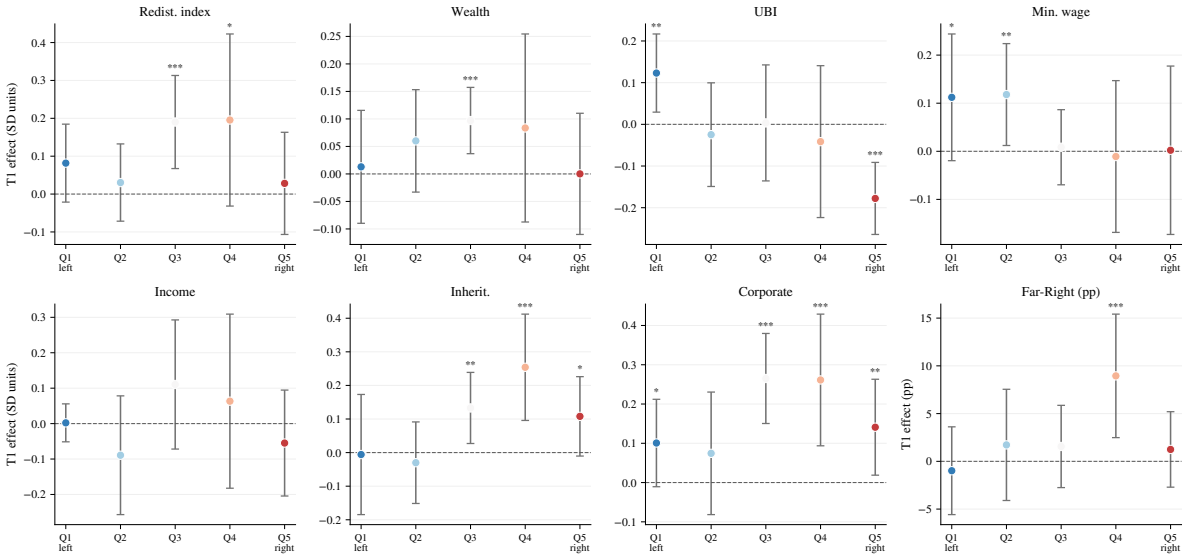

Figure 9. Backfire (II): the forced-exposure effect by quintile of right-leaning propensity. Coefficients on T1 $\times Q_k$ with 95% confidence intervals, Q1 = most left-leaning, Q5 = most right-leaning; stars mark $p < 0.10$. Same specification as Figure 8 with quintile indicators replacing the linear interaction.

6.6 Spillover to the Migration-Laden Bundle

Because the platforms bundle a fiscal with a migration position, the party effects above are also effects on migration-laden choices. Reading the left article lowers the probability of ranking first the Center-Right platform, selective skilled migration with moderate tax cuts, by 3.9 points when chosen and 3.1 points when read voluntarily, in both the linear-probability and probit specifications (Table 5, Panel B). The effect on the broader restrictive-migration bundle ($\omega \cup \sigma$) has the same sign but is imprecise (-2.6 points, $p > 0.1$; Figure 10). A purely economic narrative thus shifts a platform defined jointly by fiscal and migration positions; because the two dimensions are aligned across platforms by design, we do not attribute the shift to migration attitudes specifically.

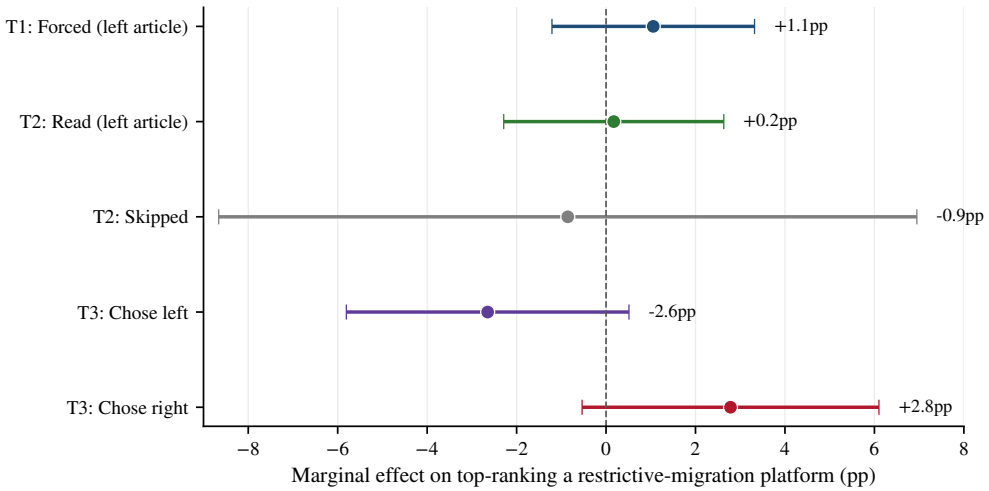

Figure 10. Economic narrative and the migration-laden bundle. Marginal effects (pp) on top-ranking a restrictive-migration platform (ω or σ), Table 4 column 5; 95% confidence intervals.

7 Robustness

We report robustness checks in two tables. Table 5 asks whether the results depend on how carefully respondents engaged with the survey; Table 6 asks whether they depend on how we estimate and how we compute standard errors. Table C.3 in Section C adds the placebo control, the order of the headlines, and reweighting.

Inattentive respondents. Restricting the sample to the 3,640 respondents who complied with the attention check on the first attempt leaves the point estimates unchanged (Table 5, Panel A): forced exposure 0.13 SD, voluntary reading 0.20 SD, choosing the left article 0.26 SD. The effect of choosing the right article keeps its size (−0.11 SD) but loses precision ($p = 0.11$) in the smaller sample. Policy by policy, the left-article effects on wealth, inheritance and corporate taxes and the right-article effects on UBI and the minimum wage are the same in both samples. Because respondents who never complied were screened out before completing the survey, this is a check on immediate versus prompted compliance rather than on attentive versus inattentive respondents, which makes the stability less surprising than it might appear.

Experimenter demand. Because both articles are politically coded, respondents may infer what the survey is looking for and answer accordingly. The closing question asks whether the survey felt politically slanted (Section B.9). Table 5, Panel B, shows that respondents do notice: relative to a control group in which 85.0% report no slant, reading the left article raises the share reporting a left-leaning survey by 4.1 to 4.9 percentage points, and choosing the right article lowers it by 2.6 points. The perception is real but small, and it does not drive the results. Among the 3,746 respondents who report no slant, the treatment effects on the index are *larger* than in the full sample (T1 +0.14 against +0.10 SD; T2-read +0.24 against +0.19; T3-left +0.28

against +0.27), and controlling for perceived slant leaves every coefficient within 0.01 SD of its baseline value. If demand effects were generating the estimates, both checks would move them toward zero; both move them away from it.

Binary outcomes. The platform outcomes are estimated by linear probability models in the main text. Probit average marginal effects (Table 6, Panel A) reproduce them: choosing the left article lowers the probability of ranking the Center-Right platform first by 3.9 points under either estimator, voluntary reading by 3.1 points under either, and the imprecise effects on the restrictive-migration bundle are imprecise under both.

Few clusters. With 16 state clusters, conventional cluster-robust inference over-rejects. Table 6, Panel B, reports wild-cluster bootstrap-t p-values next to the cluster-robust ones for every arm and outcome. The results the paper leans on survive: voluntary reading and choosing the left article remain significant at 1% on the index and on the wealth, inheritance and corporate taxes, and choosing the right article remains significant on UBI ($p = 0.002$) and the minimum wage ($p = 0.018$). Two coefficients do not survive at conventional levels: forced exposure on the index ($p = 0.14$) and choosing the right article on the index ($p = 0.08$). We have accordingly presented those two as suggestive throughout, and the ordering of the three channels,which does not depend on them,as the finding.

Placebo control, article order, and reweighting. Table C.3 addresses three features of the design. First, the control group pools a pure control with a placebo that read an unrelated text; the placebo indicator is zero on the index (0.03 SD, $p = 0.56$) and on every policy, and using the pure control alone as the comparison group leaves the treatment effects unchanged. Second, the two headlines in the choice arm were shown in random order; the order shifts the share choosing the right article by 3 points ($p = 0.16$), and the effect of choosing the left article is positive and significant whichever headline came first (0.20 and 0.35 SD). Third, the sample is a quota sample: it meets its gender and education targets but under-represents respondents aged 18–24 by 6.6 and those aged 65 or over by 8.4 percentage points, and it over-represents AfD voters (17.6% of party voters against 10.4% in 2021). We therefore re-estimate with raking weights matching the SOEP (2022) margins for gender, age and education, and separately with weights that additionally set the AfD share to the 2021 result (10.4%) or to the February 2025 result (20.8%). The latter two are a sensitivity analysis with respect to the AfD share rather than a correction to a known electorate: the survey elicits the 2021 vote, so the first is appropriate if recall is accurate and the second if recall drifted toward contemporaneous preference during the campaign in which the survey was fielded; together they bracket the plausible range. Voluntary reading and choosing the left article are invariant across all schemes (0.18–0.19 and 0.26–0.28 SD), which also indicates that these effects do not vary appreciably with age, gender or education,weighting and controlling would otherwise diverge. Forced exposure ranges from

Table 5. Robustness: respondent attention and experimenter demand

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Attention-check robustness						
	Redistribution index		T3 Left		T3 Right	
	Full sample	AC-passers	Full	AC-passers	Full	AC-passers
T1: Forced (left article)	0.103**	0.125***				
	(0.043)	(0.029)				
T2: Read (left article)	0.190***	0.202***				
	(0.026)	(0.029)				
T2: Skipped	-0.093	-0.060				
	(0.082)	(0.103)				
T3: Chose left	0.265***	0.258***				
	(0.039)	(0.051)				
T3: Chose right	-0.093**	-0.107				
	(0.041)	(0.066)				
Observations	4,524	3,640				
Wealth tax			0.235***	0.212***	-0.020	-0.027
			(0.038)	(0.057)	(0.038)	(0.058)
UBI			0.036	0.053	-0.198***	-0.190***
			(0.053)	(0.051)	(0.047)	(0.072)
Min. wage			0.117*	0.093*	-0.107***	-0.141***
			(0.062)	(0.049)	(0.034)	(0.044)
Income tax			0.058	0.049	-0.037	-0.047
			(0.066)	(0.046)	(0.030)	(0.031)
Inheritance tax			0.214***	0.220***	0.024	0.009
			(0.032)	(0.039)	(0.046)	(0.046)
Corporate tax			0.227***	0.236***	-0.073	-0.072
			(0.047)	(0.057)	(0.048)	(0.073)
Panel B: Experimenter demand — perceived slant of the survey						
	Perceived slant (pp)				Redistribution index	
	Leaned left	Leaned right	Unbiased		Unbiased only	Bias control
T1: Forced (left article)	4.14**	-0.25	-3.89*		0.141***	0.107**
	(1.66)	(1.20)	(2.27)		(0.044)	(0.045)
T2: Read (left article)	4.92**	-0.49	-4.42*		0.235***	0.195***
	(2.01)	(1.23)	(2.68)		(0.035)	(0.027)
T2: Skipped	3.14	-3.84**	0.70		-0.052	-0.094
	(2.64)	(1.65)	(3.19)		(0.089)	(0.080)
T3: Chose left	3.77**	0.57	-4.34*		0.281***	0.269***
	(1.59)	(1.37)	(2.37)		(0.034)	(0.039)
T3: Chose right	-2.62*	1.60	1.02		-0.100**	-0.094**
	(1.43)	(1.11)	(1.82)		(0.046)	(0.041)
Control-group mean	10.4	4.7	85.0			
Observations	4,524	4,524	4,524		3,746	4,524

Notes: Panel A re-estimates the main specifications on the 3,640 respondents who passed the attention check on the first attempt; columns (1)–(2) report the redistribution index by arm and columns (3)–(6) the effects of choosing each article on the six policies. Panel B uses the closing question asking whether respondents perceived the survey as politically slanted (Section B.9): columns (1)–(3) report treatment effects in percentage points on each response, and columns (5)–(6) re-estimate the index specification on the subsample who perceived no slant and on the full sample controlling for perceived slant. Estimates larger in column (5) than in the baseline indicate that experimenter demand attenuates rather than generates the results. Outcomes are standardized to the control group (mean 0, SD 1), so coefficients are effect sizes in control-group standard-deviation units. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

Table 6. Robustness: functional form and inference with few clusters

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Binary outcomes — linear probability model versus probit (percentage points)							
	Restrictive-migration bundle		Center-Right (ω)				
	LPM	Probit AME	LPM	Probit AME			
T1: Forced (left article)	1.06	1.11	-1.39	-1.44			
	(1.16)	(1.19)	(1.35)	(1.33)			
T2: Read (left article)	0.17	0.20	-3.04**	-3.04**			
	(1.26)	(1.26)	(1.45)	(1.42)			
T2: Skipped	-0.86	-0.87	-2.76	-2.75			
	(3.98)	(3.90)	(4.18)	(4.30)			
T3: Chose left	-2.65	-2.59	-3.96**	-4.00**			
	(1.61)	(1.58)	(1.77)	(1.81)			
T3: Chose right	2.78	2.81	-0.98	-1.03			
	(1.69)	(1.74)	(1.90)	(1.88)			
Observations	4,524	4,524	4,524	4,524			
Panel B: Inference with few clusters — wild-cluster bootstrap p-values (federal states, $G = 16$)							
	Index	Wealth	UBI	Min. wage	Income	Inherit.	Corporate
T1 Forced – cluster-robust p	0.016	0.135	0.396	0.073	0.956	0.029	0.000
wild-cluster bootstrap p	0.135	0.267	0.387	0.106	0.962	**0.064**	**0.006**
T2 Read – cluster-robust p	0.000	0.000	0.527	0.003	0.432	0.000	0.000
wild-cluster bootstrap p	**0.000**	**0.002**	0.532	**0.020**	0.459	**0.010**	**0.000**
T2 Skipped – cluster-robust p	0.257	0.449	0.090	0.048	0.507	0.184	0.999
wild-cluster bootstrap p	0.323	0.484	0.181	**0.099**	0.514	0.301	0.999
T3 Left – cluster-robust p	0.000	0.000	0.490	0.058	0.374	0.000	0.000
wild-cluster bootstrap p	**0.000**	**0.000**	0.500	**0.088**	0.462	**0.000**	**0.001**
T3 Right – cluster-robust p	0.023	0.595	0.000	0.001	0.205	0.606	0.125
wild-cluster bootstrap p	**0.083**	0.605	**0.002**	**0.018**	0.210	0.624	0.203

Notes: Panel A reports marginal effects in percentage points on top-ranking a restrictive-migration platform (ω or σ) and the Center-Right platform, from linear probability models and from probit average marginal effects (delta-method standard errors; the probit converged in all cases). Panel B compares the cluster-robust p-value with a wild-cluster bootstrap-t p-value (Cameron, Gelbach and Miller 2008; Webb six-point weights; null imposed; 4,999 replications; bold: $p < 0.10$), which is robust to the small number of clusters; outcomes are in control-group SD units. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

0.09 to 0.15 SD. The one sensitivity is the effect of choosing the right article: −0.09 to −0.10 SD under the demographic and 2021 weights, but −0.04 SD and insignificant at the higher AfD share.

8 Mechanisms and Interpretation

8.1 Taking Stock of the Model

Table 7 sets seven findings against three readings of the model: accuracy only, in which $\kappa_i = 0$; dissonance only, in which κ_i governs both the choice of what to read and the updating that follows; and the hybrid, in which the two forces operate at different points. Each row is a fact and each column asks whether a reading can produce it. The single-force readings fail on identifiable facts rather than on the evidence as a whole: accuracy alone cannot produce five of the seven, dissonance alone two. The hybrid accommodates all seven, but it nests the other two and so cannot fit worse by construction; what the table isolates is which restrictions the data reject, not that the unrestricted model survives.

Pure Bayesian updating fails at the selection stage and at the ordering of effects (rows 1, 3 and 5). It has no reason to expect confirmatory choice, and it predicts that respondents who select the congenial article, whose priors are already close to its message, should move *less* than a random population forced to read it, because they have less distance to cover. They move more: 0.27 SD against 0.10 SD on the index. It also predicts that the two articles, both informative, should produce a small net difference; the observed difference is 0.36 SD. Nor can it account for the respondents who decline the article (row 4). A skipper receives no signal and should be indistinguishable from the control group, yet those who decline report 0.19 SD less support for a universal basic income and 0.18 SD less for a higher minimum wage. Because assignment to T2 is random, this is selection on the preference itself, and selection on the value of information is exactly what an accuracy-only reading has no room for. Pure dissonance avoidance fails instead at the belief stage (rows 6 and 7). It predicts that respondents for whom the left article is incongruent should discount its facts; Table 8 shows that the most right-leaning fifth of forced readers update their factual belief by as much as the most left-leaning fifth ($\hat{\lambda} = 0.81$ versus 0.74). It also predicts backfire wherever the message is incongruent, whereas the movement against the article is confined to one policy and one quintile. The hybrid survives because it lets the two forces operate at different points. Dissonance governs the choice of what to read (Prediction P1); accuracy governs what is believed once read (Prediction P2, with $\hat{\lambda} \approx 0.75$ homogeneous across ideology); and dissonance returns at the step from an accepted fact to a policy position, where it produces the narrow backfire on UBI (Prediction P3).

That last observation sharpens the model rather than confirming it. As written in Section 2.4, the dissonance cost κ_i attaches to the distance between the message and the prior. The evidence says the relevant distance is not that between a statistic and a prior about the statistic,no one's

identity is threatened by an inflation rate,but that between a policy implication and a policy position. The dissonance cost is object-specific: near zero for the fact, high for the inference. Section E states the two-tier version of the model this implies, in which factual beliefs update with weight $\alpha/(\alpha + \kappa_{\text{fact}})$ and the mapping from belief to policy carries κ_{policy}; the data are consistent with $\kappa_{\text{fact}} \approx 0$ and κ_{policy} large only for UBI.

8.2 Why Policy Moves Left but Politics Moves Right

The left article raises desired redistribution by 0.2–0.3 SD among those who read it, yet raises support for neither left platform (Table 4: +2.3 and +1.8 points, neither significant, neither robust). Its only robust effect on party preference is negative: voluntary readers are 3.1 points less likely to rank the Center-Right platform first (bootstrap $p = 0.04$), those who chose the article 3.9 points (bootstrap $p = 0.07$). The right article, by contrast, raises far-right support by 3.8 points (bootstrap $p = 0.005$) and lowers Center-Left support by 2.4 points. Persuasion on economics converts into votes on one side only.

Roemer (1998) explains why. When political competition has a second, non-economic dimension, parties offer bundles, and a voter persuaded on the economic dimension can move only to a platform whose position on the second dimension she accepts. Our platforms bundle fiscal with migration positions by design: the two left platforms (π, θ) pair redistribution with open or inclusive migration; the two right platforms (ω, σ) pair tax cuts with selective or restrictive migration (Table B.1). A respondent persuaded by the left article that inflation is regressive and that redistribution is warranted faces, on the left, platforms that also require accepting a migration stance she may reject; the closest available move is away from the Center-Right bundle without arriving at a left one, which is what we observe. A respondent persuaded by the right article faces no such tension: lower redistribution and restrictive migration sit in the same bundle. The right-wing narrative is politically productive because its fiscal and identity components are aligned for its audience; the left-wing narrative is not, because its audience is split on the second dimension. This is the sense in which policy preferences move left while party preferences move right.

One further effect fits the same account. Forced exposure to the left article raises far-right support by 2.5 points (bootstrap $p = 0.01$), an effect that does not vary with right-leaning propensity. A narrative of economic grievance that does not carry the reader to a left bundle benefits the platform that combines grievance with restriction. We flag this as consistent with the bundling account rather than as evidence for it; separating the fiscal from the migration component requires platforms that vary the two independently, which ours do not.

8.3 Beliefs Move, Preferences Resist

Table 8 estimates the updating weight λ of (3) from the two knowledge questions, using the control group's median answer as the prior under random assignment. Readers of the left

Table 7. Which reading of the model survives the evidence?

What the experiment shows	Estimate (control-group SD units unless noted)	Can this reading produce the finding? (A) Accuracy only $\kappa_i = 0$	(B) Dissonance only κ_i both stages	(C) Hybrid κ_i chooses, α_i believes
Stage 1: what people choose to read (Prediction P1)				
Choice is confirmatory: readers pick the article that agrees with them	Right-leaning choose the growth article 61–67%, left-leaning 48–52%; +5 pp per step of self-placement, $p < 0.001$ (Figure 4)	✗	✓	✓
Stage 2: what reading does to preferences (Prediction P2)				
Forced exposure moves preferences, but weakly and unevenly	Index $+0.103^{**}$, corporate tax $+0.165^{***}$, inheritance $+0.087^{**}$; wealth +0.051 and UBI −0.022 n.s. (Table 1, Table 2)	(✓)	(✓)	✓
Choosing beats accepting beats being made to read — for identical content	Index: chose left $+0.265^{***}$ > read $+0.190^{***}$ > forced $+0.103^{**}$ (Table 1, col. 2)	✗	✓	✓
Respondents who decline the article already differ from the control group	Index −0.093 n.s., but UBI -0.185^{*} and min. wage -0.176^{**} (Table 2)	✗	✓	✓
The two articles are far from mirror images, and which one dominates flips by policy	Gap of 0.358 SD on the index ($+0.265^{***}$ against -0.093^{**}); $\hat{\beta}_L + \hat{\beta}_R = +0.171^{***}$ on the index but -0.161^{*} on UBI (Table 1, Table 3)	✗	(✓)	✓
Stage 3: where the resistance sits (Prediction P3)				
Backfire appears on one policy, for one group, and nowhere else	UBI: T1 × $\hat{p}_i = -0.341^{*}$; Q5 -0.178^{***} against Q1 $+0.123^{**}$. Index interaction +0.013 n.s.; far-right platform n.s. (Table C.4)	✗	✗	✓
The contested fact is accepted as fully by opponents as by supporters	$\hat{\lambda} = 0.75$ in every arm that saw the figure, 0.74–0.82 across $\hat{p}_i$ quintiles; the two arms that did not, 0.00 and 0.04 (Table 8)	✓	✗	✓[†]
Findings the reading cannot produce (✗)		**5 of 7**	**2 of 7**	**none**

Notes: Each row is a finding; each column asks whether one reading of (1)–(3) can produce it. ✓ consistent; (✓) right sign, wrong magnitude; ✗ inconsistent. (A) sets $\kappa_i = 0$ and fails at the selection stage and on the ordering of effects (rows 1, 3, 5); (B) lets κ_i govern both stages and fails on the homogeneity of factual updating (rows 6, 7); Section 8.1 works through both. Because (C) nests (A) and (B), it cannot fit worse by construction: the informative content is *which* restrictions the data reject, not that the unrestricted model survives.

[†] Row 7 fits (C) only once the dissonance cost is moved from the fact to the fact-to-policy step ($\kappa^{\text{fact}} \approx 0$, κ^{policy} large on the contested policy) — a refinement the evidence forces on the model, not a prediction of it (Section E).

Estimates are OLS coefficients from (6) in control-group SD units, taken from the tables cited rather than the running text; row 4 is descriptive ($n = 149$, selected). Romano–Wolf p-values are in Table 2, Panel B, and wild-cluster bootstrap p-values in Table 6, Panel B: row 2's index effect survives cluster-robust but not bootstrap inference ($p = 0.135$), and is read as suggestive throughout. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

article move their median estimate of inflation for low-income households from 7.9% to 10.0%, three-quarters of the way to the article's 10.7% ($\hat{\lambda} = 0.75$ in each case, with bootstrap intervals inside [0.74, 0.82]), whether the article was imposed (T1), opened (T2) or chosen (T3). The two arms that never saw the figure do not move: skippers ($\hat{\lambda} = 0.00$) and right-article readers (0.04). On the overall 2022 rate, which both articles state, left-article readers move 90–100% of the way; right-article readers, whose article mentions the figure once before turning to growth, move 10%. The shortfall from $\hat{\lambda} = 1$ on the low-income figure reflects rounding: the modal answer is 10, the article says 10.7.

The result that matters is the profile across ideology. Within forced exposure, $\hat{\lambda}$ is 0.74, 0.75, 0.82, 0.74 and 0.81 from the most left-leaning to the most right-leaning quintile. The respondents whose UBI preference moves against the article are the respondents who accept its central statistic most completely. This rules out the account of backfire in which the reader discounts the message as propaganda ($\rho > 1$ in Section 2.5): a reader who disbelieved the article would not report its number. What remains is the account in which the fact is accepted and its policy inference resisted, the pattern Kuziemko et al. (2015) document for information about inequality, where beliefs move substantially and redistributive preferences barely, and the pattern identity-protective reasoning predicts (Kahan, 2013): a fact is cheap to concede, a policy position that marks group membership is not. UBI is the natural locus, because it is the policy the right article attacks by name and the one on which the control group is most divided (Figure 2).

Two limits on this evidence. The knowledge questions ask what the rate *was*; a respondent can report a figure she remembers from the article without endorsing it, so $\hat{\lambda}$ conflates learning with acceptance. The placebo arms show the answers respond to exposure, and the homogeneity across ideology shows that whatever is measured does not vary with the respondent's politics, but the distinction between recalled and believed cannot be closed with these questions. And $\hat{\lambda}$ is a group-level median statistic; it identifies how far an arm moves, not how far any respondent does.

8.4 Identification Caveats and the Next Wave

Four caveats bound what this wave can claim. First, κ and prior precision are observationally equivalent in updating data: a respondent who moves little may be dissonance-averse or simply confident. What separates them is that a confident Bayesian never finds information aversive, whereas a dissonance-averse respondent does; the 13% who decline a free article in T2 and the confirmatory choice in T3 are evidence of aversion, but they are not a measure of it. Second, the moderator $\hat{p}_i$ is built partly from attitudinal items elicited after treatment; they are balanced across arms and the results are unchanged without them (Section D), but a moderator measured before treatment would remove the question. Third, $\hat{\lambda}$ is estimated at the group level from questions that conflate recall and belief. Fourth, the platforms bundle fiscal and migration

Table 8. Belief updating: group-level estimates of the updating weight λ

	N	Prior	Posterior	Signal	$\hat{\lambda}$	95% CI
Panel A: By arm (prior = control-group median)						
Inflation for low-income singles (left article: 10.7%); prior = 7.9						
T1: Forced (left article)	1,150	7.9	10.0	10.7	+0.75	[+0.74, +0.79]
T2: Read (left article)	1,002	7.9	10.0	10.7	+0.75	[+0.74, +0.82]
T2: Skipped	149	7.9	7.9	10.7	+0.00	[-0.43, +0.04]
T3: Chose left	490	7.9	10.0	10.7	+0.75	[+0.74, +0.82]
T3: Chose right	643	7.9	8.0	10.7	+0.04	[+0.00, +0.21]
Overall 2022 inflation (both articles: 7.9%); prior = 6.9						
T1: Forced (left article)	1,150	6.9	7.9	7.9	+1.00	[+0.89, +1.00]
T2: Read (left article)	1,002	6.9	7.8	7.9	+0.90	[+0.89, +0.90]
T2: Skipped	149	6.9	6.9	7.9	+0.00	[-0.90, +0.00]
T3: Chose left	490	6.9	7.8	7.9	+0.90	[+0.89, +1.00]
T3: Chose right	643	6.9	7.0	7.9	+0.10	[+0.00, +0.50]
Panel B: T1 by quintile of right-leaning propensity (prior = control median within quintile)						
Inflation for low-income singles (left article: 10.7%)						
Q1 (most left-leaning)	234	8.0	10.0	10.7	+0.74	[+0.67, +0.89]
Q2	256	7.9	10.0	10.7	+0.75	[+0.64, +0.85]
Q3	220	7.9	10.2	10.7	+0.82	[+0.74, +0.86]
Q4	221	8.0	10.0	10.7	+0.74	[+0.74, +0.85]
Q5 (most right-leaning)	219	7.0	10.0	10.7	+0.81	[+0.74, +0.85]
Overall 2022 inflation (both articles: 7.9%)						
Q1 (most left-leaning)	234	7.0	7.9	7.9	+1.00	[+0.89, +1.00]
Q2	256	6.9	7.8	7.9	+0.90	[+0.89, +1.00]
Q3	220	7.0	7.9	7.9	+1.00	[+1.00, +1.00]
Q4	221	6.9	7.9	7.9	+1.00	[+0.89, +1.00]
Q5 (most right-leaning)	219	6.9	7.9	7.9	+1.00	[+0.89, +1.00]

Notes: Group-level updating weights $\hat{\lambda}$ = (posterior − prior)/(signal − prior), where the posterior is the median answer to the end-of-survey knowledge question in the arm, the prior is the median answer in the control group (the counterfactual prior under random assignment), and the signal is the statistic asserted by the article. Medians are used because a thin right tail of implausible answers dominates means. 95% confidence intervals from a cluster bootstrap over federal states (999 replications; 499 in Panel B). T2 Skipped and T3 Right serve as placebos: neither group saw the low-income figure. Right-leaning propensity is the post-LASSO score of Table D.1.

positions, so the party effects cannot be decomposed.

A second wave removes each of these. Priors are elicited before treatment, incentivized, as the respondent's own belief about the low-income inflation rate together with a confidence interval, so that prior precision is observed and λ_i and κ_i are identified per respondent from (3) under the normalization $\alpha_i + \kappa_i = 1$. A willingness-to-pay module before treatment elicits how much each respondent would pay to read, or to avoid reading, the congenial and the uncongenial article, which identifies the selection-stage dissonance cost in money and tests directly whether the value of information is negative for some respondents. Beliefs and preferences are elicited for two or three policies, so that κ_{policy} can vary across them and the prediction that it is highest for UBI can be tested rather than inferred. And the platform menu crosses fiscal and migration positions in a 2×2 design, so that the bundling account of Section 8.2 becomes a test of whether the left article moves the fiscal dimension when the migration dimension is held fixed. The design is pre-registered before fielding.

9 Conclusion

We asked whether the choice of what to read about the economy polarizes preferences for redistribution and the vote, and whether forcing the other side on people backfires. A four-arm experiment on 4,524 German adults, using two real narratives about the 2022–23 inflation surge, answers both.

Choice is confirmatory, and choice matters more than content. Right-leaning respondents pick the growth article and left-leaning respondents the redistribution article; the same article moves preferences three times as much when chosen as when imposed. The movement is specific to what the article argues, the left narrative moves taxes on the rich, the right narrative moves basic income, and the two narratives are not symmetric. The left article has the larger effect on policy preferences and none on party preference; the right article moves both, raising support for a far-right platform, because its fiscal and migration positions sit in one bundle while the left bundle asks a persuaded voter to accept a migration stance she may not hold.

Backfire is real, narrow, and not what the literature assumed. Forced exposure to the redistribution narrative pushes the most right-leaning fifth of respondents away from universal basic income and the most left-leaning fifth toward it, and touches nothing else. The right-leaning respondents who move against the article believe its central fact as fully as everyone else. Dissonance does not prevent them from learning that inflation hit the poor harder; it prevents them from drawing the policy conclusion. Information campaigns that aim to move contested policy positions by supplying facts should expect the facts to land and the positions to stay.

Three limitations bound these conclusions. The sample is a quota sample that over-represents AfD voters; results are unchanged under reweighting, except that the right article's aggregate effect weakens when the sample is weighted to the 2025 election. The moderator used to test

backfire is built from observables measured after treatment, balanced across arms but not elicited before it. And the belief estimates are group-level medians from knowledge questions that do not separate recall from acceptance. A second wave, pre-registered, elicits incentivized priors with confidence intervals before treatment, willingness to pay for congenial and uncongenial content, beliefs and preferences on several policies, and a platform menu that varies fiscal and migration positions independently. It will identify the dissonance cost at the individual level and test directly whether it attaches to the fact or to the policy.

The broader point is about the architecture of information rather than its content. In a fragmented environment the first decision a citizen makes is which reality to enter, and that decision is where most of the polarizing work is done. Correcting the facts inside each reality changes beliefs; it does not, on our evidence, change the positions that define the sides.

References

Akerlof, G. A. and Dickens, W. T. (1982). The economic consequences of cognitive dissonance. *American Economic Review*, 72(3):307–319.

Alesina, A. and Glaeser, E. L. (2004). *Fighting Poverty in the US and Europe: A World of Difference*. Oxford University Press, Oxford.

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

Andre, P., Haaland, I., Roth, C., Wiederholt, M., and Wohlfart, J. (2026). Narratives about the macroeconomy. *Review of Economic Studies*. Advance article, published online 18 February 2026.

Bail, C. A., Argyle, L. P., Brown, T. W., Bumpus, J. P., Chen, H., Hunzaker, M. B. F., Lee, J., Mann, M., Merhout, F., and Volfovsky, A. (2018). Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37):9216–9221.

Bakshy, E., Messing, S., and Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239):1130–1132.

Bénabou, R. (2015). The economics of motivated beliefs. *Revue d'économie politique*, 125(5):665–685.

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

Bénabou, R. and Tirole, J. (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–164.

Binetti, A., Nuzzi, F., and Stantcheva, S. (2024). People's understanding of inflation. *Journal of Monetary Economics*, 148:103652. Article 103652.

Cameron, A. C., Gelbach, J. B., and Miller, D. L. (2008). Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427.

Campante, F. R. and Hojman, D. A. (2013). Media and polarization: Evidence from the introduction of broadcast TV in the United States. *Journal of Public Economics*, 100:79–92.

Chopra, F., Haaland, I., and Roth, C. (2024). The demand for news: Accuracy concerns versus belief confirmation motives. *Economic Journal*, 134(661):1806–1834.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Stanford University Press, Stanford, CA.

Fryer, R. G., Harms, P., and Jackson, M. O. (2019). Updating beliefs when evidence is open to interpretation: Implications for bias and polarization. *Journal of the European Economic Association*, 17(5):1470–1501.

Gentzkow, M. and Shapiro, J. M. (2011). Ideological segregation online and offline. *Quarterly Journal of Economics*, 126(4):1799–1839.

Gerber, A. S., Karlan, D., and Bergan, D. (2009). Does the media matter? a field experiment measuring the effect of newspapers on voting behavior and political opinions. *American Economic Journal: Applied Economics*, 1(2):35–52.

Grossmann, M. and Hopkins, D. A. (2016). *Asymmetric Politics: Ideological Republicans and Group Interest Democrats*. Oxford University Press, New York.

Guess, A. and Coppock, A. (2020). Does counter-attitudinal information cause backlash? results from three large survey experiments. *British Journal of Political Science*, 50(4):1497–1515.

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

Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, 8(4):407–424.

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.

Levy, R. (2021). Social media, news consumption, and polarization: Evidence from a field experiment. *American Economic Review*, 111(3):831–870.

Little, A. T. (2019). The distortion of related beliefs. *American Journal of Political Science*, 63(3):675–689.

Lord, C. G., Ross, L., and Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11):2098–2109.

Martin, G. J. and Yurukoglu, A. (2017). Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–2599.

Nyhan, B. and Reifler, J. (2010). When corrections fail: The persistence of political misperceptions. *Political Behavior*, 32(2):303–330.

Rabin, M. and Schrag, J. L. (1999). First impressions matter: A model of confirmatory bias. *Quarterly Journal of Economics*, 114(1):37–82.

Roemer, J. E. (1998). Why the poor do not expropriate the rich: An old argument in new garb. *Journal of Public Economics*, 70(3):399–424.

Romano, J. P. and Wolf, M. (2005). Stepwise multiple testing as formalized data snooping. *Econometrica*, 73(4):1237–1282.

Sen, A. (2009). *The Idea of Justice*. Belknap Press of Harvard University Press, Cambridge, MA.

Suhay, E. (2025). *Debating the American Dream: How Explanations for Inequality Polarize Politics*. Russell Sage Foundation, New York.

Taber, C. S. and Lodge, M. (2006). Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3):755–769.

Wood, T. and Porter, E. (2019). The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior*, 41(1):135–163.

YouGov (2025). Public reaction to Elon Musk and the rows over child grooming gangs. https://yougov.co.uk/politics/articles/51298-public-reaction-to-elon-musk-and-the-rows-over-child-grooming-gangs.

Online Appendix

Biased Information, Biased Preferences

A Treatment Materials

Both articles were assembled from published German reporting on the 2022–23 inflation surge, with hyperlinks to the original sources retained in the text so that respondents could see the political alignment of the outlets each drew on. Both open with the same official statistic,overall inflation of 7.9% in 2022,and diverge thereafter. The estimated reading time of two minutes was displayed to respondents.

Exposure differed by arm. Respondents in T1 were shown the full text of the left article. Respondents in T2 were shown its headline and a one-sentence description and chose whether to open it. Respondents in T3 were shown the headline and description of both articles and chose one. The texts below are English translations; the survey was administered in German.

A.1 Left-Leaning Article (T1, T2, and T3)

Social gap in inflation at its highest level: Poorer people have an inflation rate almost three percentage points higher than the wealthy

Estimated reading time: 2 minutes

Germany's inflation crisis: Record high of 7.9% in 2022! According to the Federal Statistical Office, inflation in Germany was 7.9 percent overall in 2022. In September, October and November, inflation reached double-digits, the highest level in the country's post-war history. But it hits the poorest the hardest. Inflation hits poorer families much harder than everyone else. According to the Hans Böckler Foundation, the inflation rate for wealthy people in November 2022 was 7.8 percent on average. In contrast, the inflation rate for single people with low incomes was 10.7 percent, almost three percentage points higher. This means that the rising costs of everyday goods such as food and energy are putting a greater burden on poorer households. This makes it even more difficult for them to make ends meet.

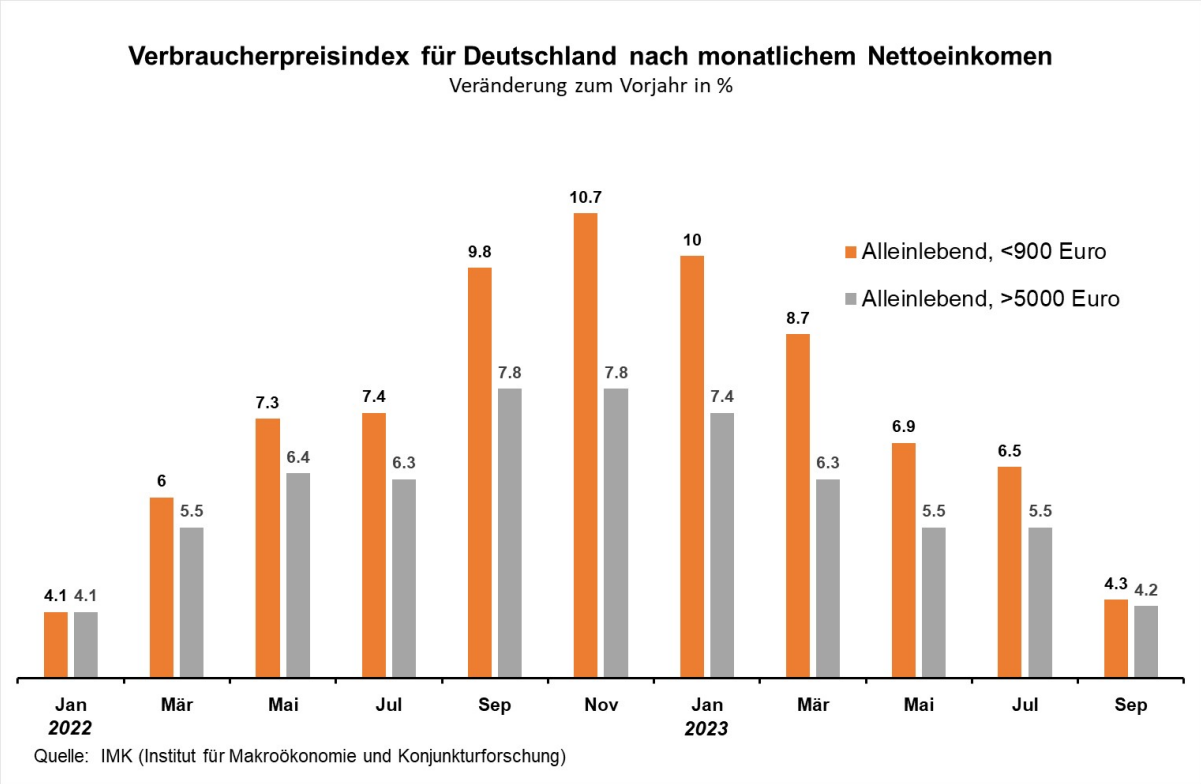

Figure A.1. Chart displayed in the left-leaning article.

Desperate decisions: families are starving while the number of personal bankruptcies is rising For Mrs. M., who wishes to remain anonymous, doing without is the reaction to the crisis: she receives a citizen's allowance, lives with her family near Bremen and all four have to get by on around 2,000 euros a month. Everything has to be paid for with that. And how can that be done when everything has become more expensive? Ms. M. has only one answer: starving. "I only eat twice a day to save money." But working people are also struggling. Stefanie Kaiser from the German Red Cross debt counseling service, Bremerhaven district association, reports: 'The number of employed people who come to us has increased by 50 percent in the last two years.' The vast majority of them have to file for bankruptcy. 'It's a catastrophe.' (Bremen Zwei).

The silent grip of inflation: prices remain high even when the storm has passed The time of galloping inflation seems to be over in 2024. But is everything OK again now? No. Now prices are stable, which slows inflation down. However, the price level remains high, even if the cost of individual products has fallen somewhat. Many people feel the inflation almost every day when they pay at the supermarket checkout. For the same money, you get significantly less goods (TAZ).

The inflation trap: How greed squeezes the poor Inflation hits low- to middle-income households particularly hard, as they struggle with rising prices for essential, almost irreplaceable goods. The minimum wage is 12.41 euros per hour and will rise to 12.82 euros in 2025. However, according to the German Trade Union Confederation (DGB), to compensate for inflation,

it would have to be 14 euros. The minimum wage commission has failed to ensure adequate inflation compensation,and the federal government allows it. The strategy of keeping wages low to fight inflation is only exacerbating the crisis, while the government is silent on the issue of "greedflation": excessive price increases by which companies use inflation as an excuse to boost their profits. Urgent action is needed to address these growing economic inequalities (TAZ).

What can be done? To combat the negative impact of inflation on poorer households, the German government must take decisive action, for example by capping electricity and gas prices. This is already providing them with some relief.

Increase corporate tax: The government should impose a tax on corporations that make extra profits during the energy crisis. Such excess profits should be taxed at a minimum of 50%, according to Oxfam, a global organization dedicated to fighting poverty, inequality, and injustice.

Promote a wealth tax: A re-introduction of the wealth tax on the rich or a one-off levy on large fortunes could be considered. Provide public services: This money should be used to improve education, healthcare, and social services to ensure that everyone can afford the basic necessities and to reduce the gap between rich and poor.

Raise the minimum wage: With real commitment, the government can help families struggling with rising costs, including by raising the minimum wage (Süddeutsche Zeitung). Remember, we are all in this stormy economic sea together. Let's make sure everyone has a sturdy ship.

A.2 Right-Leaning Article (T3 only)

Handouts won't save Germany: Why economic growth, not ever-increasing social benefits, is the real path to recovery

Estimated reading time: 2 minutes

According to the Federal Statistical Office, inflation in Germany was 7.9 percent overall in 2022. In September, October, and November, inflation reached double digits, the highest level in the country's post-war history.

The answer, however, is growth, not social benefits Rising prices in Germany are putting a strain on families and businesses. While immediate help is essential, experts are urging the government to focus on something even more important: economic growth. According to the European Commission, Germany's economic recovery has been sluggish, and the country's real GDP fell 0.3% in 2023, making it the only major economy in the developed world to contract in that year. The outlook for 2024 is not much better: according to DW, a 0.2% decline in GDP is forecast instead of the growth previously expected.

Instead of relying on inflation aid, Germany could focus on boosting its industry, attracting more investment and creating jobs. By prioritizing growth, the country can strengthen its resilience and ensure lasting stability, rather than merely alleviating the symptoms of the problem with temporary aid.

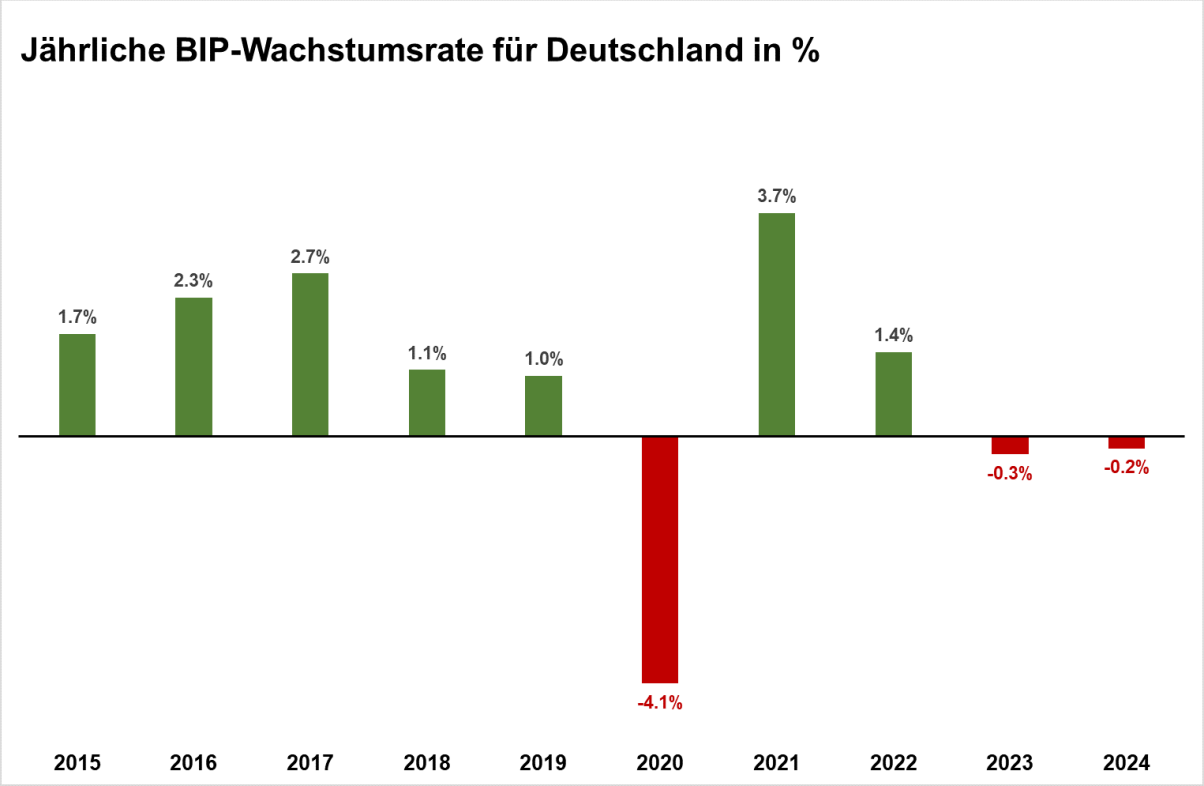

Figure A.2. Chart displayed in the right-leaning article.

We need more growth in Germany, not a "four-day week" Germany is at a critical point. It is facing an economic bottleneck, and boosting growth is more urgent than ever. As Christian Sewing, CEO of Deutsche Bank, emphasized at the WELT Economic Summit, focusing on growth and hard work, not a four-day week, is key to securing a stable economy for the future (Welt). However, the government's current emphasis on social benefits, such as the Citizen's Allowance, has raised concerns. The money, which is intended to support those in need, is likely to reduce the incentive to work (Welt). Many recipients prefer social benefits over employment. Especially since 40 percent of recipients are foreigners who do not have sufficient German language skills and therefore find it difficult to integrate into the labor market (FOCUS online). The CDU is bucking this trend. Its party leadership is calling for stricter penalties for those who turn down reasonable job offers. CDU leader Friedrich Merz noted that many of the 1.7 million employable recipients simply assume that it is more profitable to use citizen's allowance than to enter the workforce (Welt). The draft budget for 2025 reflects this: the proposed measures include the creation of stronger work incentives and stricter rules, which are expected to lead to budget savings of up to 5 billion euros in citizen's allowance. These savings would enable higher spending on pensions, investments in rail and road transport, child benefits and the military (Institute for Macrofinance).

Tax cuts to keep German jobs at home In September 2023, the German government announced a new corporate tax cut of €7 billion per year,finally a step towards reviving the economy. After years of watching jobs and companies move to other countries, this relief gives small

and medium-sized enterprises,the heart of local communities,the chance to thrive, invest, and create jobs here at home (DW). Joachim Nagel, President of the Deutsche Bundesbank, has expressed this clearly: 'In order to create an employment and investment-friendly environment, it is important to keep an eye on the tax burden on labor and capital.' Otherwise, Germany risks falling behind while global competitors with more business-friendly policies pull ahead (Financial Times). This tax break, totaling 32 billion euro over four years, is more than just a cut. It's about protecting Germany's economic strength and ensuring future generations are not left behind. When businesses thrive here, families feel the impact. Every job saved in Germany means stability for a family, better opportunities for young people and a stronger foundation for the country's future. By cutting taxes, Germany is finally giving its own citizens and businesses a real chance to succeed at home (Financial Times).

Reclaiming Germany's future: Fighting inflation with growth, not handouts To secure a future of prosperity, Germany must commit to policies that encourage economic growth, reduce government paternalism and reward hard work. While social support has its place, excessive spending risks fueling inflation and undermining productivity. With the right balance of incentives, tax relief and a focus on creating opportunities, Germany can enable its people to live independently and attract investment to sustain a strong economy.

Let us make growth, jobs and productivity a priority so that Germany becomes a country where prosperity is accessible to all,now and for generations to come.

A.3 Placebo Text (Control Group)

Half of the control group read the text below; the other half read nothing. The text is of comparable length to the treatment articles but unrelated to inflation, redistribution or politics. Table C.3 shows that outcomes do not differ between the two control subgroups, which is why they are pooled throughout.

> **The sweet journey of chocolate in Germany**
>
> Chocolate is a favourite treat all over the world, and Germany has played a significant part in its history and production. But have you ever wondered how this delicious delicacy is made? Let us take a journey through the chocolate-making process, paying particular attention to Germany's contribution to this sweet industry.
>
> It all begins with the humble cocoa bean, which grows in tropical regions far from Germany, such as West Africa and South America. The beans are harvested, fermented and dried before finding their way to chocolate factories in countries such as Germany, where the magic happens.
>
> In Germany the beans are carefully roasted to bring out their rich aroma. After roasting they are cracked open to release the cocoa nibs inside, which are then ground into a thick brown paste known as cocoa mass. This is the basis of every chocolate product.

Germany is home to some of the world's finest chocolate manufacturers, known for their precision and quality. They blend the cocoa mass with sugar, milk and sometimes other ingredients, depending on whether they are making dark, milk or white chocolate. The mixture is then refined and conched—a process that smooths the texture and improves the flavour.

Finally the chocolate is tempered, a technique that gives it its glossy surface and the satisfying snap when a piece is broken. It is then moulded into bars, truffles or other delicious shapes, ready to be enjoyed.

German chocolate brands such as Ritter Sport and Milka are known worldwide for their high-quality products, each with its own distinctive taste and history. The next time you enjoy a piece of German chocolate, you will taste the careful craftsmanship that goes into every bite.

B Survey Instrument

The questionnaire was administered in German; the text below is an English translation. Blocks are listed in the order in which respondents saw them: consent, stratification, treatment (Section A), outcomes, covariates, and the closing knowledge and demand-effect questions.

B.1 Consent

Respondents were shown the following before entering the survey and could proceed only by selecting "Yes".

This questionnaire is part of a study on the economy in Germany. Before you decide to answer the questionnaire, it is important that you understand why this study is being conducted and what your role is in it.

1. This survey aims to analyse the views of the German population on various economic issues. It is important that you answer honestly.
2. You will be asked to answer questions in a short online survey. There are no right or wrong answers. If you do not know the answer, simply give your best guess. The questionnaire takes about 10 minutes to complete.
3. There are no foreseeable risks associated with participating. Participation is completely voluntary, and you can withdraw at any time without any disadvantages.
4. You will be rewarded for your participation.
5. All responses will be treated strictly confidentially. No information that could be used to identify you will be published. Only the researchers involved in this study will have access to the information you provide. The data will not be passed on to third parties and will not be used for any purpose other than the study described above. In

accordance with European regulations on the protection of personal data, you have the right to access, rectify, and delete the information you have provided.

By participating, you consent to your data being stored and processed for the stated purposes. Would you like to continue? Yes / No

B.2 Stratification Block

Asked before treatment: year of birth, gender, and highest completed education. These three variables define the quota targets and the randomization blocks. **[TODO: Reproduce the exact wording and response options of the three stratification questions.]**

Q2 How old are you?

- 18–24
- 25–34
- 35–44
- 45–54
- 55–64
- 65–74
- 75–84
- 85 or older

Q3 What is your gender?

- Male
- Female
- Non-binary / other
- Prefer not to say

Q86 What school, university or vocational qualification do you have? Please think of your highest qualification.

- Secondary education (intermediate school certificate / Realschulabschluss, general higher education entrance qualification / Abitur, Hauptschulabschluss, university of applied sciences entrance qualification, Fachoberschule certificate)
- Vocational qualification (in-company or school-based vocational training, training at a technical school, Technikerschule, vocational or technical academy)
- Bachelor's degree or higher (Bachelor's, university of applied sciences degree, Diplom or Master's degree, completed teacher training, doctorate / habilitation)

B.3 Redistributive-Policy Outcomes

Three items use a 0 to 10 agreement scale, one asks for a tax rate in percent, and two use a slider running from "reduce" to "increase" around the status quo, recorded on a 0 to 10 scale with 5 denoting no change.

Wealth tax (0–10). Please answer the following question on a scale of 0 to 10. To what extent do you agree or disagree with the introduction of a wealth tax in Germany that requires high net worth individuals to pay a percentage of their total wealth as tax? *0 = completely disagree, 10 = completely agree.*

Unconditional basic income (0–10). In the following, we would like to ask you about your views on an unconditional basic income. The state provides all citizens with a monthly income that covers the basic costs of living. This income is intended to replace existing social benefits and is financed by tax revenues. Everyone receives the same amount, regardless of whether they are employed or not. *0 = completely disagree, 10 = completely agree.*

Minimum wage (0–10). To what extent do you support the idea of increasing the minimum wage in Germany? *0 = completely disagree, 10 = completely agree.*

Top income-tax rate (0 to 100%). In Germany, the average tax burden of a single person earning the minimum wage is about 27% of total gross income (including social security contributions). In your opinion, what percentage of the total income should a single person earning over 280,000 euros gross per year pay in taxes (including social security contributions)? *Slider from 0 to 100%.*

Inheritance tax (0–10). If a person over the age of 27 in Germany inherits or receives €1 million in cash from their parents, they have to pay around 9% of the sum as inheritance tax. In your opinion, should the inheritance tax decrease, stay the same, or increase? Move the slider to the left to lower the rate or to the right to increase it; if you think the rate should stay the same, leave the slider in the middle. *Recorded 0 = reduce, 5 = keep, 10 = increase.*

Corporate profit tax (0 to 10). The current corporate profit tax in Germany is about 30%. In your opinion, should the corporate profit tax decrease, stay the same, or increase? Slider as above. *Recorded 0 = reduce, 5 = keep, 10 = increase.*

All six are standardized to the control group before estimation (Section 5.2); Figure 2 shows their untransformed distributions.

B.4 Political Platforms

Respondents ranked four fictitious platforms from 1 (most preferred) to 4, with the order of presentation randomized and a blank-vote option available. Each platform pairs a fiscal with a

migration position.

Table B.1. The four political platforms ranked by respondents

Platform	Migration position	Fiscal position
Left (π)	Open borders with equal rights and access to social services for all migrants; no detention or deportation	Radical redistribution of wealth; high taxes on the wealthy and corporations to fund comprehensive public services; no taxes on low incomes
Center-Left (θ)	Inclusive immigration policy with accessible legal pathways, integration programs and a simplified path to citizenship	Progressive tax system; tax the wealthy and corporations to provide public services and relief to middle- and low-income households
Center-Right (ω)	Selective legal immigration focused on skilled workers; strict measures against illegal immigration; emphasis on cultural integration	Moderate tax cuts for middle incomes and small businesses; simplified tax code; reduction of certain social programs
Far-Right (σ)	Strict immigration controls; limited merit-based system for highly skilled workers; strong border security	Single (flat) tax to reduce government spending; significant corporate tax cuts to stimulate the private sector

Notes: Outcome variables are indicators for ranking each platform first; the restrictive-migration bundle is the union of ω and σ. Fiscal and migration positions are aligned across platforms, which is why Section 6.6 does not attribute the estimated effects to migration attitudes specifically.

B.5 Cognitive Dissonance

Three items, each on a 1–7 scale. The second and third are reverse-scaled relative to the first; they are used as recorded.

Discomfort. How much discomfort or inner tension do you feel when you encounter a political view that differs from your own? *1 = no discomfort at all, 7 = very severe discomfort.*

Belief revision. To what extent do you agree with the following statement? "Disagreement makes me question my current beliefs." *1 = totally agree, 7 = totally disagree.*

Reconciliation. To what extent do you agree with the following statement? "I find it difficult to reconcile opposing arguments with my own beliefs." *1 = totally agree, 7 = totally disagree.*

B.6 Distributive Justice: The Flute Question

You are asked to give a flute as a gift to one of three children. Child A is exceptionally talented at playing the flute and has shown remarkable dedication and effort in developing musical skills. Child B comes from a socially deprived background and has no access to toys or leisure activities. Child C inherited the flute from a family member who was a talented flautist; the flute has emotional value and symbolises a family tradition.

1. Give the flute to Child A because of its exceptional musical talent and achievements. *(talent)*
2. Give the flute to Child B to fulfil its need for entertainment and leisure. *(need)*
3. Give the flute to Child C as an inheritance, honouring family tradition. *(inheritance)*
4. I don't know.

The question follows Sen (2009). Responses are balanced across arms (Table C.1).

B.7 Political Self-Placement

On a scale from 0 (far left) to 10 (far right), where would you place yourself politically? Asked at the end of the outcome block rather than before treatment, to avoid priming; it is therefore post-treatment and used only descriptively and in Section D.

B.8 Media Consumption

Media consumption is elicited in three stages. Respondents first name the category that is their primary news source. They are then asked whether they additionally use each of the other two categories, and are shown the outlet list for a category only if it is their primary source or they report also using it. Outlet questions allow multiple responses. Excluding the open *Other* option in each list, the three lists contain 46 named outlets; these are the media covariates entering the propensity model in Section D.

Primary source. What is your main source of information for news and current events? *TV, radio and print / Online news / Social media, messaging and video networks / None.*

Additional use. Do you also generally use [each remaining category] as a news source? *Yes / No.*

Social media, messaging and video networks (9 named outlets).

- Facebook
- YouTube
- WhatsApp
- X (formerly Twitter)
- Instagram
- TikTok
- Snapchat
- LinkedIn
- Telegram
- Other

TV, radio and print (19 named outlets).

- ARD (Tagesschau, Tagesthemen)
- ZDF (heute, heute-journal)
- RTL News
- Regional television
- Regional and local newspapers
- Regional radio

- n-tv
- Sat.1 Nachrichten
- Welt
- Bild / Bild am Sonntag
- Der Spiegel
- Focus
- ProSieben :newstime
- Die Zeit
- Frankfurter Allgemeine Zeitung
- Handelsblatt
- Süddeutsche Zeitung
- Tagesspiegel
- taz
- Other

Online news (18 named outlets).

- ARD News online
- t-online
- Bild.de
- n-tv.de
- Focus online
- Regional online newspapers
- Spiegel Online
- GMX
- ZDF News online
- Regional online television
- Zeit Online
- Yahoo! News
- FAZ.NET
- MSN News
- Google News, Apple News
- BBC, CNN, Reuters
- Financial Times
- DW
- Other

B.9 Closing Questions

Two knowledge questions close the survey. They elicit precisely the two statistics on which the left article rests, the overall 2022 inflation rate, stated in both articles, and the rate for low-income single households, stated only in the left article, and provide the post-treatment beliefs used to estimate the updating weight λ in Section 8.3. A final question asks whether respondents perceived the survey as politically slanted, and is used to test for experimenter-demand effects in Table 5, Panel D.

Q41 Knowledge. Please answer the following questions. *Note: at most one decimal place may be entered. Please enter a value in percent.*

1. How high was the annual inflation rate in Germany in 2022? *(True value 7.9%; stated in both articles.)*
2. How high was the inflation rate specifically for single people with low incomes in Germany in 2022? *(Value asserted by the left article: 10.7%; not mentioned in the right article.)*

Q47 Perceived slant. Do you think this survey was biased?

- Yes, I had the impression that the survey leaned politically to the left.
- Yes, I had the impression that the survey leaned politically to the right.
- No, it did not feel biased.

In the control group, 85.0% report that the survey felt unbiased, 10.4% that it leaned left and 4.7% that it leaned right.

C Additional Results

Reading and choice rates by arm are reported in Table C.2: willingness to open the article is high (87%) and roughly constant across political groups, ranging from 78% among non-voters to 93% among Linke voters, whereas the direction of choice in the choice arm is strongly sorted, from 48% choosing the right article among Greens voters to 67% among CDU/CSU voters. Whether to read is not an ideological decision; what to read is.

Table C.3 reports the placebo, order and reweighting checks discussed in Section 7.

Table C.5 distinguishes exposure to the left article from uptake of its content. Uptake is measured from the closing knowledge questions: a respondent has absorbed the article if her answer lies within one or two percentage points of the figure it asserts. Assignment to forced exposure raises uptake of the low-income figure from 11% to 43%, so about a third of forced readers demonstrably absorb it, a rate that is stable across definitions. Instrumenting uptake with assignment rescales the intent-to-treat effects to a per-absorber basis, roughly three times larger (0.33 SD on the index, 0.51 on the corporate tax); because the article also carries an anecdote and explicit recommendations, the exclusion restriction is strong and the per-absorber estimates should be read as an upper bound. Two features are informative beyond the rescaling. Adding assignment to T2 and T3 as instruments, the over-identification test rejects for the wealth tax ($p = 0.01$–0.02) but not for the index or the corporate and inheritance taxes: choosing the article moves the wealth tax through a channel beyond absorbing its facts. And right-leaning respondents absorb the low-income figure at a higher rate than left-leaning ones (48% against 40% in T1) while responding to it less on the index and negatively on UBI (-0.26 SD per absorber, $p < 0.10$), which restates the dissociation of Section 8.3 in per-absorber units.

Table C.1. Means and balancing tests across randomized groups

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Mean of variable by group					p-value, equality of means			Signif. under
	T1	T2	T3	Control	Germany	T1 vs C	T2 vs C	T3 vs C	multiple testing
Stratification variables									
Male	0.49	0.50	0.48	0.50	0.51	0.63	0.99	0.41	No
18–24	0.10	0.10	0.11	0.11	0.17	0.80	0.66	0.77	No
25–34	0.17	0.17	0.15	0.15	0.14	0.19	0.24	0.93	No
35–44	0.18	0.16	0.15	0.19	0.15	0.79	0.09	0.03	No
45–54	0.22	0.20	0.23	0.20	0.14	0.23	0.85	0.10	No
55–64	0.22	0.22	0.23	0.23	0.18	0.38	0.30	0.79	No
65+	0.10	0.15	0.13	0.12	0.21	0.18	0.04	0.72	No
Bachelor+	0.28	0.28	0.25	0.28	0.26	0.97	0.82	0.10	No
Vocational	0.57	0.57	0.58	0.57	0.54	0.99	0.92	0.48	No
Secondary and below	0.15	0.15	0.17	0.15	0.20	0.98	0.67	0.30	No
Socio-economic characteristics									
No religion	0.45	0.46	0.46	0.45	0.38	0.83	0.72	0.70	No
Catholic	0.23	0.22	0.21	0.25	0.26	0.37	0.16	0.05	No
Protestant	0.16	0.17	0.16	0.15	0.33	0.52	0.15	0.41	No
Other religion	0.13	0.13	0.14	0.13	0.03	0.93	0.67	0.39	No
Not saying	0.03	0.02	0.02	0.02	0.003	0.64	0.83	0.79	No
Married	0.45	0.45	0.47	0.47	0.44	0.32	0.47	0.88	No
Full time	0.54	0.52	0.55	0.53	0.37	0.84	0.58	0.41	No
Part time	0.14	0.13	0.14	0.15	0.23	0.84	0.41	0.46	No
Self-employed	0.06	0.06	0.04	0.05	0.05	0.48	0.59	0.18	No
Not employed	0.26	0.29	0.27	0.27	0.35	0.66	0.33	0.78	No
HH income < €1,750	0.18	0.19	0.21	0.20	0.26	0.20	0.79	0.52	No
€1,751–3,000	0.32	0.30	0.30	0.29	0.25	0.19	0.49	0.78	No
€3,001–6,000	0.41	0.42	0.40	0.40	0.40	0.79	0.33	0.81	No
> €6,000	0.10	0.08	0.10	0.11	0.09	0.46	0.02	0.38	No
Prior voting behavior, 2021 (%)									
SPD	23.42	24.33	24.13	23.54	25.7	0.92	0.73	0.91	No
CDU/CSU	25.16	23.41	22.03	23.77	24.1	0.51	0.82	0.30	No
AfD	15.07	16.66	16.36	16.56	10.4	0.37	0.98	0.80	No
Greens	12.22	14.22	12.91	11.82	14.7	0.81	0.13	0.56	No
FDP	9.07	7.47	11.65	9.47	11.4	0.75	0.11	0.16	No
The Left	6.10	5.52	5.66	6.67	4.9	0.61	0.29	0.33	No
Other parties	8.97	8.18	7.95	8.18	8.8	0.55	0.89	0.40	No
Did not vote / blank	14.98	15.10	16.04	14.76	23.3	0.88	0.82	0.40	No
Notion of justice (flute question)									
Talent (utilitarian)	0.28	0.28	0.31	0.28		0.76	0.99	0.10	No
Need (egalitarian)	0.43	0.45	0.42	0.42		0.76	0.13	0.92	No
Inheritance (libertarian)	0.23	0.20	0.20	0.23		0.91	0.15	0.06	No
Don't know	0.07	0.06	0.07	0.07		0.89	0.55	0.87	No

Notes: Columns (1)–(4) report means by randomized group; column (5) the population benchmark (SOEP 2022 for demographics; official 2021 federal election result for voting). Columns (6)–(8) report p-values for the test of equality of means against the control group with heteroskedasticity-robust standard errors; column (9) indicates whether any difference remains significant after correcting for multiple hypothesis testing (Romano and Wolf, 2005).

Table C.2. Reading and choice rates by arm and by 2021 vote

	T2 (optional)		T3 (choice)	
	N	% who read	*N*	% choosing right
All respondents	1,151	87.1	1,133	56.8
By 2021 federal-election vote				
Linke	54	92.6	54	51.9
Greens	139	86.3	123	48.0
SPD	238	89.9	230	49.1
FDP	73	90.4	111	61.3
CDU/CSU	229	90.0	209	66.5
AfD	162	89.5	156	64.7
Others	82	80.5	69	53.6
Non-voter	174	77.6	181	54.1
By order in which the headlines were shown (choice arm only)				
Left headline shown first	–	–	601	55.4
Right headline shown first	–	–	532	58.3

Notes: Share of respondents in the optional arm who opened the left article, and share of respondents in the choice arm who opened the right (growth) article rather than the left one. Willingness to read is high and roughly constant across political groups; the direction of choice is strongly sorted. Party groups exclude respondents who did not report a 2021 vote choice.

Table C.3. Additional robustness: placebo control, article order, and reweighting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Placebo text vs. pure control (coefficient on the placebo indicator, pure control omitted)							
	Index	Wealth	UBI	Min. wage	Income	Inherit.	Corporate
Placebo (history of chocolate)	0.034	-0.006	0.067	0.025	0.040	0.045	-0.018
	(0.058)	(0.085)	(0.052)	(0.036)	(0.057)	(0.055)	(0.057)
Panel B: Redistribution index under alternative samples and weights							
	Baseline	Pure control only	T3: left shown first	T3: right shown first	SOEP weights	AfD share = 10.4%	AfD share = 20.8%
T1: Forced (left article)	0.103**	0.119***	0.104**	0.102**	0.127**	0.113**	0.147**
	(0.043)	(0.045)	(0.043)	(0.042)	(0.050)	(0.044)	(0.061)
T2: Read (left article)	0.190***	0.206***	0.192***	0.192***	0.182***	0.186***	0.192***
	(0.026)	(0.036)	(0.026)	(0.025)	(0.031)	(0.032)	(0.029)
T2: Skipped	-0.093	-0.078	-0.096	-0.086	-0.070	-0.086	-0.084
	(0.082)	(0.096)	(0.080)	(0.081)	(0.074)	(0.083)	(0.093)
T3: Chose left	0.265***	0.279***	0.196***	0.349***	0.269***	0.278***	0.274***
	(0.039)	(0.065)	(0.059)	(0.056)	(0.044)	(0.043)	(0.051)
T3: Chose right	-0.093**	-0.077	-0.129*	-0.050	-0.091**	-0.106**	-0.046
	(0.041)	(0.057)	(0.072)	(0.038)	(0.041)	(0.047)	(0.052)
Observations	4,524	3,960	3,992	3,923	4,514	4,514	4,514
Effective sample size	–	–	–	–	3,930	3,720	3,654

Notes: Panel A adds an indicator for the placebo control group (which read a short text on the history of chocolate) to the decomposed specification, with the pure control group as the omitted category; a zero coefficient justifies pooling the two control groups. Panel B re-estimates the decomposed specification for the redistribution index: (2) excluding the placebo group so that the pure control is the only comparison group; (3)–(4) restricting the choice arm to respondents who saw the left (right) headline first, with all non-T3 respondents retained; (5)–(7) with raking weights that match the sample to population margins for gender, age group and education (SOEP 2022), and additionally to the 2021 or 2025 federal election result and turnout (respondents reporting a non-binary gender are dropped from the weighted specifications). Maximum weight 2.9. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

Table C.4. Backfire: heterogeneity of the forced-exposure effect by right-leaning propensity (appendix)

	Redist. index (1)	Wealth (2)	UBI (3)	Min. wage (4)	Income (5)	Inherit. (6)	Corporate (7)	Far-Right (pp) (8)
Panel A: Continuous interaction with right-leaning propensity $\hat{p}_i$								
T1: Forced	0.105**	0.053	-0.025	0.045*	0.006	0.092**	0.168***	2.583**
	(0.043)	(0.034)	(0.025)	(0.026)	(0.064)	(0.038)	(0.043)	(1.115)
$\hat{p}_i$ (centered)	0.332**	0.293*	-0.261	-0.466***	0.352***	0.583***	0.323**	22.721***
	(0.163)	(0.177)	(0.251)	(0.156)	(0.123)	(0.109)	(0.145)	(6.290)
T1 × $\hat{p}_i$	0.013	-0.004	-0.341*	-0.126	-0.057	0.251	0.144	-7.028
	(0.245)	(0.231)	(0.184)	(0.229)	(0.194)	(0.233)	(0.143)	(9.705)
Panel B: T1 effect by quintile of $\hat{p}_i$ (Q1 = most left-leaning)								
T1 × Q1	0.082	0.013	0.123**	0.112*	0.002	-0.006	0.101*	-0.982
	(0.053)	(0.052)	(0.048)	(0.067)	(0.027)	(0.091)	(0.057)	(2.345)
T1 × Q2	0.030	0.060	-0.025	0.118**	-0.089	-0.030	0.074	1.722
	(0.052)	(0.048)	(0.063)	(0.054)	(0.086)	(0.062)	(0.080)	(2.971)
T1 × Q3	0.190***	0.097***	0.003	0.008	0.110	0.133**	0.265***	1.557
	(0.063)	(0.031)	(0.071)	(0.040)	(0.093)	(0.054)	(0.058)	(2.197)
T1 × Q4	0.195*	0.084	-0.041	-0.011	0.063	0.254***	0.261***	8.950***
	(0.116)	(0.087)	(0.093)	(0.080)	(0.125)	(0.081)	(0.086)	(3.299)
T1 × Q5	0.028	0.000	-0.178***	0.002	-0.055	0.108*	0.141**	1.247
	(0.069)	(0.056)	(0.044)	(0.089)	(0.076)	(0.060)	(0.062)	(2.016)
Other arms included	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,524	4,524	4,524	4,524	4,524	4,524	4,524	4,524

Notes: $\hat{p}_i$ is the predicted probability of choosing the right article from a two-step post-LASSO logit: an ℓ_1-penalized logit on 77 baseline predictors, estimated on the choice arm, selects 16 predictors; an unpenalized logit on the selected predictors (cross-validated AUC 0.63) is then applied to all respondents (Table D.1). The UBI interaction is robust to the penalty rule used for selection: under the sklearn min-CV rule (41 predictors) it is -0.32 ($p = 0.015$), and under the one-SE rule (11 predictors) -0.52 ($p = 0.014$). Columns (1)–(7) in control-group SD units; column (8) in percentage points. Panel A interacts forced exposure with the centered propensity; a negative interaction on a pro-redistribution outcome indicates that right-leaning respondents move against the left article (backfire). Panel B replaces the linear interaction with quintile indicators. All models include T2 Read, T2 Skipped, T3 Left and T3 Right. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

Table C.5. Information uptake: intent-to-treat versus effect per absorber

Uptake	First stage	F	Index	Wealth	Corp.	Inherit.	UBI	Min. wage
Panel A: Intent-to-treat reference (T1 vs. control)								
Forced exposure (ITT)	–	–	0.106**	0.057	0.162***	0.088**	-0.019	0.048*
			(0.041)	(0.035)	(0.043)	(0.039)	(0.026)	(0.025)
Panel B: Effect per absorber — 2SLS, uptake instrumented by T1 assignment								
Low-income, ±1	0.321	279	0.330***	0.177	0.506***	0.273**	-0.060	0.151**
(11% → 43%)	(0.019)		(0.125)	(0.109)	(0.126)	(0.121)	(0.077)	(0.071)
Low-income, ±2	0.308	342	0.343***	0.185	0.526***	0.284**	-0.062	0.157**
(21% → 52%)	(0.017)		(0.130)	(0.113)	(0.134)	(0.123)	(0.080)	(0.076)
Both, ±1	0.341	486	0.310***	0.167*	0.475***	0.256**	-0.056	0.141**
(7% → 41%)	(0.015)		(0.113)	(0.100)	(0.119)	(0.106)	(0.072)	(0.070)
Both, ±2	0.345	598	0.307***	0.165*	0.470***	0.254**	-0.056	0.140**
(14% → 49%)	(0.014)		(0.114)	(0.098)	(0.117)	(0.107)	(0.071)	(0.069)
Panel C: Over-identified 2SLS, uptake instrumented by T1, T2 and T3 assignment								
Low-income, ±1	–	169	0.403***	0.289***	0.546***	0.312***	-0.052	0.180**
overid. p			[0.39]	[0.01]	[0.61]	[0.17]	[0.06]	[0.52]
Low-income, ±2	–	210	0.420***	0.304***	0.567***	0.329***	-0.059	0.185**
overid. p			[0.40]	[0.02]	[0.57]	[0.19]	[0.06]	[0.51]
Both, ±1	–	260	0.380***	0.273***	0.516***	0.294***	-0.048	0.170**
overid. p			[0.39]	[0.01]	[0.62]	[0.16]	[0.06]	[0.53]
Both, ±2	–	317	0.376***	0.276***	0.507***	0.300***	-0.060	0.163**
overid. p			[0.37]	[0.02]	[0.49]	[0.20]	[0.06]	[0.48]
Panel D: Effect per absorber by right-leaning propensity (T1 instrument; low-income, ±1)								
Left-leaning half	0.299	123	0.481***	0.312**	0.559***	0.176	0.139	0.379***
(9% → 40%)	(0.027)		(0.114)	(0.133)	(0.166)	(0.164)	(0.160)	(0.124)
Right-leaning half	0.340	153	0.274	0.127	0.507***	0.439***	-0.262*	-0.006
(13% → 48%)	(0.027)		(0.206)	(0.180)	(0.141)	(0.158)	(0.154)	(0.139)

Notes: Uptake definitions: *Low-income* = answer to the low-income question within the margin of 10.7%; *Both* = answers to both questions within the margin of 10.7% and 7.9%; parentheses under each definition give the uptake rate in the control group and in T1. Information uptake is an indicator that the respondent's answer to the closing knowledge question(s) lies within the stated margin of the figure asserted by the left article (10.7% for low-income singles; 7.9% overall). Panel A reports the intent-to-treat effect of forced exposure for reference. Panel B reports two-stage least squares estimates of the effect of uptake on each outcome, instrumenting uptake with assignment to T1 on the sample of T1 and control respondents; the first stage is the effect of assignment on uptake and F its clustered F-statistic. Because assignment is random and never reduces uptake, the estimate is the effect of the article on respondents who absorb its facts (compliers), under the assumption that the article does not affect respondents who do not absorb them; since the article also contains an anecdote and explicit policy recommendations, this assumption is strong, and the estimates are best read as the intent-to-treat effect scaled to a per-absorber basis. Panel C adds assignment to T2 and T3 as instruments on the full sample; the bracketed p-value is Wooldridge's robust score test of overidentification, which rejects when the three arms imply different per-absorber effects. Panel D splits the T1–control sample at the median of the right-leaning propensity of Table D.1. Outcomes in control-group SD units. Standard errors clustered at the federal-state level ($G = 16$) in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$. All specifications include stratification and demographic controls (nationality, age, gender, education, household income, marital status, religion, place of birth, employment, housing tenure, income source, and 2021 vote). Omitted category: no-article control group.

D The Right-Leaning Propensity Score

Prediction P3 requires a measure of how congenial the left article is to a given respondent. We construct one from behaviour in the choice arm, where the preferred article is observed, and extrapolate it to respondents who were never given a choice.

Construction. The procedure has two steps. First, an ℓ_1-penalized logit of choosing the right article on all candidate predictors is estimated on the choice arm, with the penalty selected by 5-fold cross-validation; it retains 16 of 77 candidate columns. Second, an unpenalized logit on the retained predictors is estimated on the same arm and applied to every respondent, giving $\hat{p}_i$. The second step matters: predictions from the penalized fit are shrunk toward the mean, which compresses $\hat{p}_i$ and attenuates the interactions that test Prediction P3. The candidate set comprises the stratification and socio-demographic variables, the 46 named media outlets and the three main-source and secondary-source indicators, the three cognitive-dissonance items, and the justice question.

What is retained. Table D.1 reports the post-selection logit. Choice is predicted by partisanship, media diet and a few demographics. Readers of *Welt* and *Bild*,the conservative broadsheet and the right-leaning tabloid,and of Yahoo News are more likely to choose the growth article, as are full-time employees and, more weakly, CDU/CSU voters; SPD voters, respondents placing themselves on the left, those aged 25–34, and users of Instagram, Snapchat and TikTok are less likely. Neither education nor household income is retained. Neither is any of the three cognitive-dissonance items, nor the justice question: once partisanship and media habits are accounted for, general psychological dispositions add nothing to the prediction of which article a respondent opens. Coefficients are post-selection and should be read descriptively rather than as parameters with valid standard errors.

Fit. The model has a McFadden pseudo-R^2 of 0.054, an in-sample AUC of 0.659, and a 5-fold cross-validated AUC of 0.635 against a benchmark of 0.5. It predicts choice, but coarsely: 62% of choices are classified correctly against 57% from always predicting the modal choice. Two consequences follow. Most of the variation in what people choose to read is not accounted for by observables, which is a substantive finding as much as a limitation. And because $\hat{p}_i$ is a noisy proxy, the interactions in Section 6.5 are attenuated, so the backfire estimates are conservative.

Distribution and balance. $\hat{p}_i$ ranges from 0.09 to 0.91 with a standard deviation of 0.135. Within the choice arm it separates the two groups as intended: mean 0.53 among those who chose the left article and 0.60 among those who chose the right one. Across the randomized arms it is balanced,differences from the control group of −0.009 (T1, $p = 0.11$), −0.006 (T2,

$p = 0.30$) and -0.002 (T3, $p = 0.70$),as it must be, since assignment is random and $\hat{p}_i$ is a function of respondent characteristics only.

Post-treatment predictors. Five candidate variables,the three dissonance items, the justice question and political self-placement,are elicited after the treatment block, because eliciting political attitudes beforehand would have primed the very responses the experiment measures. Three considerations bound the concern. They are balanced across the randomized arms at conventional levels (Table C.1). Only one of the five, self-placement, is retained by the LASSO. And re-estimating $\hat{p}_i$ with all five excluded leaves the results unchanged: the interaction of forced exposure with $\hat{p}_i$ on UBI is -0.54 ($p = 0.04$) against -0.35 ($p = 0.06$) in the specification used in the text. The concern is not eliminated,a measure taken before treatment would remove it by construction,and it is one of the motivations for the second wave described in Section 8.4.

Table D.1. Right-leaning propensity: retained predictors and post-selection logit

Predictor	Coefficient	(SE)
Yahoo News	+0.749**	(0.363)
Self-placement: left (2 of 5)	-0.534***	(0.155)
Other traditional media	-0.512*	(0.284)
Age 25–34	-0.501***	(0.180)
2021 vote: SPD	-0.443***	(0.164)
Other online sources	-0.441*	(0.235)
Reads Bild	+0.356*	(0.212)
Snapchat	-0.342	(0.254)
Employed full-time	+0.307**	(0.132)
Instagram	-0.243	(0.159)
Reads Welt	+0.204	(0.156)
2021 vote: CDU/CSU	+0.181	(0.172)
Married	+0.176	(0.128)
Main source: social media	-0.168	(0.173)
TikTok	-0.131	(0.185)
Female	-0.129	(0.132)
Candidate predictors (after dummy expansion)	77	
Predictors retained by LASSO	16	
McFadden pseudo-R^2	0.053	
AUC, in-sample	0.658	
AUC, 5-fold cross-validated	0.632	
Observations (choice arm)	1,133	

Notes: Two-step post-LASSO procedure. Step 1: an ℓ_1-penalized logit of choosing the right article on all candidate predictors, estimated on the choice arm (T3) with the penalty chosen by 5-fold cross-validation (Stata `lasso logit, selection(cv)`), selects the predictors listed. Step 2: an unpenalized logit on the selected predictors, estimated on T3, yields $\hat{p}_i$ for every respondent. Coefficients are log-odds of choosing the right article with heteroskedasticity-robust standard errors; inference is post-selection and should be read descriptively. AUC = area under the ROC curve (0.5 = no predictive power); the cross-validated AUC holds the selected set fixed and is therefore mildly optimistic. Candidate predictors: stratification and socio-demographic variables, media consumption (main and secondary source and outlet-level indicators), the three cognitive-dissonance items, and the justice question. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

E Theory Appendix

E.1 Derivation of the Updating Weight

The posterior minimizes $L(\theta) = \alpha_i(\theta - \hat{\theta}_s)^2 + \kappa_i(\theta - \mu_i)^2$. The first-order condition $2\alpha_i(\theta - \hat{\theta}_s) + 2\kappa_i(\theta - \mu_i) = 0$ gives $(\alpha_i + \kappa_i)\theta = \alpha_i\hat{\theta}_s + \kappa_i\mu_i$, hence

$$\mu_i^{\text{post}} = \lambda_i\hat{\theta}_s + (1 - \lambda_i)\mu_i = \mu_i + \lambda_i(\hat{\theta}_s - \mu_i), \qquad \lambda_i = \frac{\alpha_i}{\alpha_i + \kappa_i}, \tag{E.1}$$

with $L'' = 2(\alpha_i + \kappa_i) > 0$ and $\partial\lambda_i/\partial\kappa_i = -\alpha_i/(\alpha_i + \kappa_i)^2 < 0$. Because λ_i depends only on the ratio $\alpha_i : \kappa_i$, the normalization $\alpha_i + \kappa_i = 1$ is without loss of generality at the updating stage, and then $\lambda_i = \alpha_i$.

E.2 Backfire as Source Skepticism

If the reader discounts the message as coming from a biased persuader, the perceived signal is $\tilde{\theta} = \hat{\theta}_s - \rho(\hat{\theta}_s - \mu_i)$ with $\rho \geq 0$, and the effective weight is $\lambda_i^{\text{eff}} = (1 - \rho)\lambda_i$. For $\rho < 1$ the message is attenuated; for $\rho > 1$ it becomes evidence of the opposite and μ_i^{post} moves away from $\hat{\theta}_s$. For $\rho \in [0, 1]$ this is the confirmatory misperception of Rabin and Schrag (1999), in which a signal is read as more prior-consistent than it is, and the encoding-then-updating structure of Fryer et al. (2019). Values $\rho > 1$ require the reader to treat the message as evidence about the sender's intent rather than about the world, which is outside both models. Since ρ plausibly rises with the distance $|\hat{\theta}_s - \mu_i|$ and with exposure the reader did not choose, the model predicts backfire for incongruent, forced content,and, as Section 8.3 shows, at the policy tier rather than the factual one.

E.3 Two-Tier Extension: Fact and Policy

Let θ denote a factual quantity (the low-income inflation rate) and π a policy position (support for UBI), with the respondent's prior policy position π_i^0 and a mapping $g(\theta)$ from the fact to the policy position the article implies. Updating on the fact carries a cost κ_i^{fact} and yields $\theta_i^{\text{post}} = \theta_i^0 + \lambda_i^{\text{fact}}(\hat{\theta}_s - \theta_i^0)$ with $\lambda_i^{\text{fact}} = \alpha_i/(\alpha_i + \kappa_i^{\text{fact}})$. The policy position then minimizes $\alpha_i(\pi - g(\theta_i^{\text{post}}))^2 + \kappa_i^{\text{policy}}(\pi - \pi_i^0)^2$, giving $\pi_i^{\text{post}} = \pi_i^0 + \lambda_i^{\text{policy}}\,(g(\theta_i^{\text{post}}) - \pi_i^0)$ with $\lambda_i^{\text{policy}} = \alpha_i/(\alpha_i + \kappa_i^{\text{policy}})$. The evidence of Table 8 and Figure 9 corresponds to $\kappa_i^{\text{fact}} \approx 0$ for all i and κ_i^{policy} large for right-leaning respondents on UBI only; backfire (π_i^{post} moving away from $g(\theta_i^{\text{post}})$) requires, at the policy tier, the source-skepticism term ρ of the previous subsection applied to g rather than to $\hat{\theta}_s$.

E.4 Identification Note

Equation (3) has the algebraic form of a normal-normal posterior mean with α_i in the role of signal precision and κ_i in the role of prior precision. Updating data alone therefore cannot distinguish a dissonance-averse respondent from a confident Bayesian: both move little. The two differ at the selection stage. For a Bayesian the value of free information is weakly positive, so avoidance is at most indifference; for a respondent with $\kappa_i > 0$ whose beliefs enter utility directly, information about which she expects to be uncomfortable has strictly negative value and is avoided. The 13% who decline a free article in the optional arm and the confirmatory sorting in the choice arm are evidence of the second, though they do not measure κ_i. Separating the levels of α_i and κ_i, rather than their ratio, requires the willingness-to-pay module of Section 8.4, which prices dissonance in money, and a measure of prior confidence, which observes precision directly.