ECONtribute
MARKETS & PUBLIC POLICY

ECONtribute
Discussion Paper No. 445

Curating Political Information: Who Shares What with Whom?

Yero Samuel Ndiaye

November 2026

www.econtribute.de

UNIVERSITÄT BONN

UNIVERSITÄT ZU KÖLN

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/2 – 390838866 – is gratefully acknowledged.

Curating Political Information: Who Shares What with Whom?

Yero Samuel Ndiaye*

November 5, 2026

Click here for the most recent version

Abstract

Why do people supply information slanted toward their own positions? In four experiments, I investigate whether this slant reflects deliberate selection, which motives drive it, and how it affects those receiving information. People *curate the truth*, choosing which facts to pass on and to whom. Even with the usual explanations ruled out, they value sharing facts that support their own political position beyond anything they expect those facts to do to others. They also want to correct and persuade others. However, they share less with those who hold different views on policy, even though their own expectations give them strategic reason to share more with them. When receivers make real political donations, people select more on expected persuasion. Professional journalists select in a similar way, even without market or newsroom pressures. Individuals from whom facts are withheld end up with less accurate beliefs and barely notice the selection. Selection among true information thus distorts what people believe *without any falsehood*, at a margin current policies such as fact-checking do not target.

JEL classification: D83, D82, C91, D72, L82

Key words: Information Provision, Selective Disclosure, Political Information, Media Bias

*University of Cologne and Max Planck Institute for Behavioral Economics; ndiaye@wiso.uni-koeln.de. I am grateful for valuable comments and suggestions from Lara Marie Berger, Markus Dertwinkel-Kalt, Max Rainer Pascal Grossmann, Oliver Gürtler, Johannes Münster, Shakked Noy, Axel Ockenfels, Christopher Roth, Nicolas Röver, Carlo Schwarz, Christoph Schottmüller, Sebastian Siegloch, Till Stange, Johannes Wohlfahrt, and Florian Zimmermann, as well as by participants of seminars and conferences at the C-SEB Early Ideas Workshop, IMPRS BeSmart Thesis Workshop, RWI Workshop, DICE Brownbag Workshop, the University of Münster Research Seminar, BCFM PhD Conference, MPI Summer Summit, GfeW Annual Meeting. I thank Jan-Luca Bauss, Frauke Nora Gerster, Linda Plehn, Lena Schmelig, and Moritz Lakenbrink for their excellent research assistance. Support from the Center for Social and Economic Behavior (C-SEB) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2126/2 390838866 is gratefully acknowledged.

1 Introduction

People decide which information to share everywhere, all the time. Journalists and editors select findings for readers, fact-checkers select statements to examine, and people choose what to raise in conversation or post to followers. Observationally, people's and organizations' supply is slanted toward their own positions (e.g., Louis-Sidois 2025, Freitag et al. 2021, Shin and Thorson 2017).

As the information supply affects how well informed the public can be and thereby the functioning of democracies, legislation such as the EU Digital Services Act, fact-checks and content labels address distortions, largely by targeting false content.[1] Yet while false content spreads faster and further than true content (Vosoughi et al. 2018), it makes up only about 1 percent of Americans' news consumption, and sharing it is rare (Allen et al. 2020, Guess et al. 2019).[2] Selection among true information alone can distort what others believe without any false statement (Bowen et al. 2023). The slanted supply of true information is therefore an understudied and potentially under-regulated distortion.

Observed slant does not reveal motivations. Consider two facts about immigration in Germany: immigrants are overrepresented in jobs facing labor shortages, and foreigners are overrepresented among violent-crime suspects. Suppose a pro-immigration person posts only the first fact on social media. She may do so only because she never came across the second fact, may doubt its veracity, or may simply want to signal to her followers that she is pro-immigration. She may also do so because she believes her followers do not know the first fact, wants to win them over to her position, or values sharing a fact that shows her side is right independent of who sees it. Her reasons could also differ: she may have other goals when sharing with a like-minded friend rather than a skeptical co-worker. While the observed decision is consistent with all of these reasons, whether and how any intervention can affect her choice depends on which motivation drives it. Meanwhile, regardless of her motivation, some of her followers may never learn that the second fact exists unless she shares it, and may come away with inaccurate beliefs about immigration and crime despite having read nothing false.

In this paper, I use four experiments and a conceptual framework to show that *people deliberately curate the truth*, choosing which facts to pass on and to whom. They do so for many reasons and are less strategic than standard models of disclosure assume. Professional journalists, deciding without any market pressure, select the same way. People receiving information are adversely affected even without falsehoods.

[1] Other instruments address online safety (the UK's Online Safety Act), electoral integrity (the Canada Elections Act), media independence and pluralism (the European Media Freedom Act), and content provision (Canada's Broadcasting Act).

[2] Exposure is concentrated among a small minority of heavy consumers (Budak et al. 2024).

I begin by asking whether people select among facts when some usual explanations are muted. In a preregistered experiment, 1,200 quota-sampled U.S. adults first report what they themselves believe and then each decide whether to pass on four facts about immigration, two favoring each side, each to a different receiver who does not hold an accurate belief about it.[3] Every sender sees all four facts, so exposure cannot explain what they pass on, and passing on one fact never comes at the expense of another. The receiver never learns that a decision was made, so sharing cannot be used by senders to signal where they stand. As senders rate 97 percent of the facts credible, credibility concerns are minor. Still, senders pass on facts that support their own position 6.7 percentage points more often than facts that do not. When the receiver is randomly assigned to make a real political donation, the gap rises to 13.3 points. In this setting, the facts that support a sender's side are mostly those the sender expects to move the receiver toward it, so wanting to confirm one's position and wanting to persuade cannot be told apart.

Experiment 2 then disentangles the motivations. Here, 1,548 quota-sampled German adults each decide whether to pass on five facts about immigration or climate policy, each to a different receiver at one of five political positions, whose position they know and who learns theirs regardless. Senders also make incentivized predictions of how each fact would change each receiver's factual beliefs and political position. The results offer three main insights into senders' motives. First, they favor facts that support their own side beyond anything they expect those facts to do to the receiver. Second, they pass on facts they expect to correct the receiver or win the receiver over, and they share less with people who disagree with them: like-minded receivers get 78 percent of facts, receivers with completely opposite positions get only about half, although senders expect them to learn just as much and be more persuadable. Third, when the receiver is randomly assigned to make a real political donation, senders appear to select more on what they expect to persuade.

As journalists select information frequently and for large audiences, I next study the same decision among them. Slant in the information supply of media outlets is usually traced to the market environment, including competition, ownership and audience demand (e.g., Gentzkow and Shapiro 2006, Martin and McCrain 2019, Gentzkow and Shapiro 2010). Here, 309 professional German journalists make the same decisions as senders in Experiment 2, for receivers who make a real political donation, without market pressures. Journalists pass on facts that support their own position 20.4 percentage points more often than facts that do not, and their choices respond to the same returns as those of the general population. Their sharing is even more slanted than the general population's partly because they hold stronger political views. Among senders with

[3] Besides being prevalent, true statements limit the effect of lie aversion (Abeler et al. 2019) and reputational concerns.

equally strong views, journalists do pass on 11.4 percentage points more facts overall, and more of the facts that challenge the receiver's view, which they expect to be more persuasive. While the experiment captures individual choices rather than what outlets publish, it shows that some slant could originate in journalists' own choices before any market pressure applies.

Finally, I examine how this selection affects those receiving or not receiving information. In a preregistered experiment, 1,043 quota-sampled German adults are each matched to a sender from Experiment 2 or 3 and face that sender's actual choice. Receivers from whom the sender withheld a fact end up with nearly twice the belief error of receivers who got it. They barely notice the selection: receivers underestimate how strongly sender–fact alignment predicts sharing, and a quarter perceive no selection at all. A fact that receivers see moves their factual beliefs but not their political views, although senders expected it to move both. Selection among facts thus leaves people less informed in a way they barely notice.

Taken together, the results are hard to reconcile with standard models and with explanations based on exposure, doubt, signaling or market pressure. Standard models of strategic communication assume that senders care only about how the receiver responds (Grossman 1981, Milgrom 1981). Such a sender would not favor facts beyond what they do to the receiver and, given what senders here expect, would share more with those who disagree, not less.

These results point to a possible source of political polarization on the supply side: people may intentionally stop information from crossing political lines, and people on opposite sides may come to disagree about facts without hearing anything false. Instruments that target falsehoods do not reach this distortion, and to the extent journalists' own choices shape what outlets publish, neither do remedies that act on media markets alone. Receivers perceive little of the selection they face, so making selection visible may help, though evaluating such an instrument goes beyond this paper.

Related literature. This paper relates to five strands of literature on distortions in information markets.

First, a literature studies why people communicate about politics. Expressive acts may provide identity-confirming utility (Hillman 2010), while observable actions may be chosen for what others infer about one's type (Bénabou and Tirole 2006). Experimentally, providing anonymity increases willingness to voice stigmatized positions and changes audience inferences and sanctions (Bursztyn et al. 2023a), and making political-attitude reports public shifts them toward socially acceptable positions and makes them less informative (Braghieri 2024). In sharing, partisan polarization predicts the sharing of real news and low-credibility news (Osmundsen et al. 2021), moralization and attitude extremity amplify aligned sharing (Marie et al. 2023), paying

senders to be rated truthful can lead them to supply false messages that confirm receivers' motivated beliefs (Thaler 2025),[4] and Guriev et al. (2026) find evidence for reputation concerns and partisan persuasion, while partisan signaling appears comparatively less important. People also share more news favoring their own side when they believe their audience shares their politics (Marie and Petersen 2025). I pair each sender with real, individual receivers, rule out signaling by design, and elicit what the sender expects each fact to do to each receiver, which separates favoring one's own side from persuading.

A second literature investigates incentives in the news market as an explanation for slant. Models of profit-maximizing outlets predict slant toward consumers' priors, disciplined by later verification of the state and by the number of independently owned competitors (Gentzkow and Shapiro 2006). Competition need not reduce slant when readers hold opposed priors (Mullainathan and Shleifer 2005), and it can reduce the number of facts reported (Zhu and Dukes 2015). Empirically, consumer demand predicts measured slant (Gentzkow and Shapiro 2010), though political control (Durante and Knight 2012) and ownership (Martin and McCrain 2019, Cagé et al. 2022) also move content, and capture by incumbents weakens government accountability (Besley and Prat 2006, Prat and Strömberg 2013). Experimentally, attention incentives lead to more emotional and negative headlines (Berger 2026). Journalists' own ideology travels with them, though outlet policy views dominate (Boxell and Conway 2022), and most article-level slant variance lies within rather than between outlets (Braghieri et al. 2025). I remove market and newsroom pressures and investigate whether journalists' individual choices still produce slant.

A third literature examines what information receivers want and how they respond to it. Individuals seek and avoid information for reasons beyond its instrumental value (Golman et al. 2022), and in politics they disproportionately seek agreeing and avoid disconfirming information (Peterson and Iyengar 2021). Readers punish an outlet for selectively reporting one of two true statistics, but only when the selection runs against their own side (Chopra et al. 2024), strong ideologues reduce their demand for an outlet that shares their views when it is fact-checked (Chopra et al. 2022), and viewers choose opinion programming to inform factual estimates (Bursztyn et al. 2023b). Encountered information is then processed unevenly, with belief-threatening signals given less weight (Eil and Rao 2011, Möbius et al. 2021) or selectively forgotten (Zimmermann 2020). In politics the trust placed in a message tracks its agreement with the reader's views (Thaler 2024). That work investigates selectivity and valuation within receivers, while I focus on senders and on what their selection does to receivers.

[4] Direct sender–receiver communication is typically studied with direct financial incentives attached to the message, as in Thaler (2025). That setting is important but uncommon for non-professional information suppliers.

A fourth literature studies the disclosure of true information and its effect on receivers. In canonical models, senders disclose both favorable and unfavorable information because silence is informative to receivers (Grossman 1981, Milgrom 1981). Unraveling attenuates when the receiver is unsure the sender is informed (Dye 1985, Jung and Kwon 1988) or when disclosure is costly (Verrecchia 1983). In the laboratory, disclosure falls short of the benchmark (Jin et al. 2021, Benndorf et al. 2025), and related designs find that senders reliably select the evidence favoring their preferred conclusion (Degan et al. 2023, Burdea et al. 2023). Farina et al. (2026) separate withholding from selection and find that receivers do not fully correct for the selection they face, though their sender faces a single randomly matched receiver, so the object selected is which signal to disclose, never whom to address.[5] When senders pass on only aligned true news, receivers read too little into the resulting silence, and, when information is imprecise, beliefs polarize although no signal is false, in a model that takes sharing behavior as given (Bowen et al. 2023). A similar misperception is formalized in models of receivers who do not fully account for what a sender knows (Eyster and Rabin 2005, Fong et al. 2025). I measure receivers' anticipation of selectivity when they know the sender's policy position.

A fifth literature studies remedies against distortion in information supply. Fact-checking reduces engagement with claims rated outright false (Cagé et al. 2025) and lowers the rate at which they are shared (Henry et al. 2022), accuracy nudges improve discernment between true and false headlines (Pennycook et al. 2021), and media literacy produces broader and more persistent reductions in the perceived credibility of false posts (Berger et al. 2025).[6] Fact-checking nonetheless leaves intact the persuasive effect of the falsehoods it corrects (Barrera et al. 2020). Beliefs keep moving after a false signal is rendered uninformative, more for narratives than for statistics (Musolff et al. 2026). Comparing remedies within one design, Guriev et al. (2026) find that a content-neutral accuracy prime reduces the sharing of false posts while leaving true posts untouched, whereas offering a fact-check suppresses both. Removing reshared content wholesale cuts exposure to political news and untrustworthy sources, and lowers news knowledge within the experimental sample (Guess et al. 2023). I study selection among truths, a distortion none of these measures targets.

The rest of the paper proceeds as follows. Section 2 presents Experiment 1, which shows that people favor facts supporting their own position, more so when the receiver makes a real political donation. Section 3 develops a simple framework of the motives behind sharing. Section 4 presents Experiment 2, which disentangles these motives

[5] Cheap talk (Crawford and Sobel 1982) and Bayesian persuasion (Kamenica and Gentzkow 2011) vary the message technology and the sender's commitment power. Manili (2024) adds motivated belief formation on the receiver's side, and Suvorov et al. (2024) join selective supply to its measured cost for ego-relevant information, under motives that are posited rather than elicited.

[6] Lazer et al. (2018) review the evidence and the interventions.

and shows that senders select whom to inform as well as what to pass on. Section 5 presents Experiment 3, in which professional journalists make the same decision. Section 6 presents Experiment 4, which measures how withholding a fact affects receivers' beliefs. Section 7 concludes.

2 Experiment 1: Selective Provision of Facts

Experiment 1 provides evidence that senders selectively provide true information, and that downstream political action increases selection on alignment. The design makes several accounts of a slant in what senders provide other than selection on alignment—differential exposure, differential prior knowledge, disbelief, and signaling to the receiver—poor explanations of the pattern.

2.1 Design and Hypotheses

Experimental timeline. The experiment proceeds in three steps (Figure 1). I first elicit senders' prior knowledge about four immigration facts and measure their policy views through a donation decision. I then elicit their beliefs about the persuasiveness of information. Finally, senders make one provision decision per fact, each directed at a specific receiver, and a closing questionnaire records stated motivations, credibility assessments, and demographics.

Facts about immigration. Table 1 presents the four facts, selected from authoritative sources and balanced across policy directions. I validated them in a pre-registered survey ($N = 100$) that tested a larger candidate pool against criteria of veracity, perceived credibility, policy relevance, and directional balance.[7] Each is formulated as an invertible binary comparison, so beliefs are correct or incorrect.

Prior knowledge. I first elicit senders' prior knowledge. They complete each comparative statement, in random order, as "higher," "lower," or "don't know." Responses are incentivized: correct completions receive a bonus, incorrect completions do not, and "don't know" responses receive the bonus with 25% probability. Senders are then told the correct answer for each.

Policy views. I measure senders' policy views in an incentivized donation task: each sender allocates $500 between the *American Immigration Council* (AIC), which advocates

[7] Appendix A.1 describes the validation survey and the final selection.

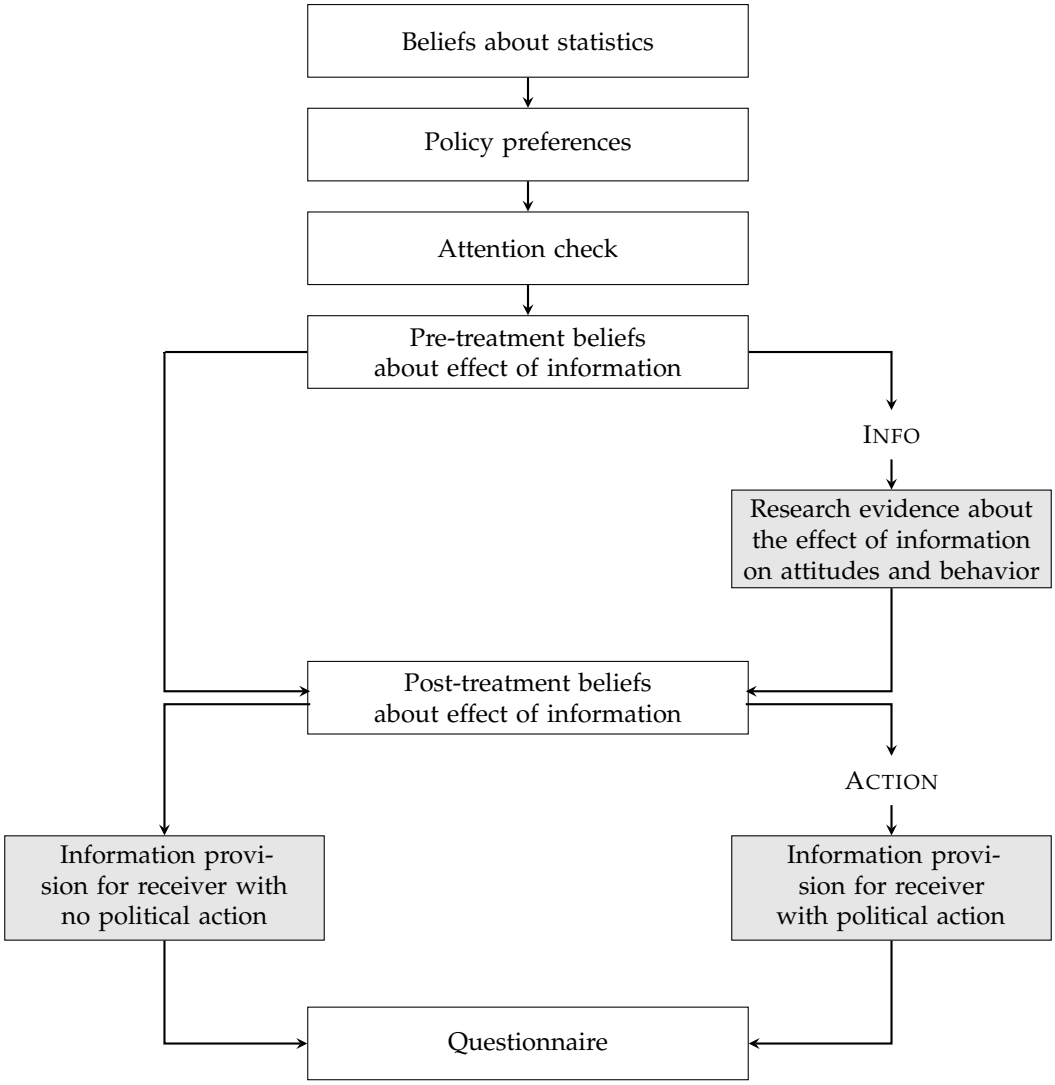


Figure 1: Timeline of Experiment 1

Table 1: Facts Shown to Senders

Direction	Fact
Pro	The crime rate of immigrants is *lower* than that of U.S.-born Americans.
Pro	The labor force participation of immigrants is *higher* than that of U.S.-born Americans
Anti	The average number of unauthorized border crossings is *higher* in the last five years compared to historical averages.
Anti	When comparing what first-generation immigrants use in government services versus what they pay in taxes, the net cost is *higher* than the same calculation for U.S.-born Americans at the state and local levels.

Notes: The facts are those used in Experiment 1, as verbatim claims shown to senders. Direction (Pro) = supports more-permissive immigration policy (Anti) = supports more-restrictive policy. Senders face all facts in random order. Sources and pre-validation details in Appendix A.1. The evidence shown for the crime statement is Texas data, and the evidence shown for the fiscal statement is an aggregate for first-generation households. Appendix A.1 documents the evidence for each comparison.

for increased immigration, and *NumbersUSA*, which advocates for decreased immigration. Senders are shown each organization's position in the organization's own words before deciding.[8]

Receiver belief states. For each sender, each fact is assigned to exactly one future receiver who will have either completed the comparative statement incorrectly or reported not knowing. The sender observes which case applies before deciding. Senders know they are deciding for receivers who are unaware that a disclosure decision was made. If information is provided, it appears as part of the survey flow. If information is withheld, the receiver does not learn they could have received information. Nondisclosure therefore carries no inferential content about the sender.

Treatments. Table 2 summarizes the four design dimensions. Two vary within sender and two between.

Table 2: Experimental Design and Identifying Variation

Dimension	Variation	Levels
Alignment with statistic	Within-subject	*Aligned* and *Unaligned*
Political action (ACTION)	Between-subject	ACTION = 1 and ACTION = 0
Receiver belief state	Within-subject	*False belief* and *Don't know*
Effect of information (INFO)	Between-subject	INFO = 1 and INFO = 0

Notes: The sample is Experiment 1 (N = 1,200 U.S. adults, Prolific, quota-sampled). Within-subject dimensions: each sender decides on all four statistics, yielding 4 decisions per sender. Each statistic's receiver belief state is drawn independently with probability one half. Between-subject dimensions are randomly assigned.

The primary within-sender dimension is alignment. A fact is *aligned* for a sender if the policy side it favors matches their own policy view, and *unaligned* otherwise. As each side has two facts favoring it, a non-indifferent sender faces two of each. Receiver belief state is drawn independently for each fact: each fact is randomly assigned to a receiver who either holds a false belief or reports not knowing. Between senders, the primary dimension is ACTION, under which receivers make the same incentivized donation decision as senders and otherwise complete the experiment without further decisions.[9] A second between-sender treatment, INFO, shows half of the senders the results of an information experiment that shifted attitudes toward immigrants (Haaland and Roth 2020). It is designed to shift their beliefs about how persuasive the facts are. Crossing it with ACTION gives the four cells CONTROL, ACTION, INFO and ACTION × INFO. Where a comparison names a single cell, ACTION and INFO denote the cell with only that treatment. Elsewhere they denote the treatment.

[8] Appendix A.2 describes the incentivized prior-knowledge elicitation and the donation task.

[9] Through randomization, any difference in alignment-based selectivity is attributable to the downstream political consequence, assuming that ACTION operates only through that consequence. ACTION may also shift the salience of the decision and the sender's construal of the task.

Hypotheses. I pre-registered four hypotheses about how senders provide information.

Hypothesis 1.0 (Baseline provision)**.** *Provision rates for both aligned and unaligned information exceed the indifference benchmark of* 50% *in* CONTROL*.*

Hypothesis 1.1 (Alignment-based selectivity)**.** *Senders provide aligned information at strictly higher rates than unaligned information in* CONTROL*.*

Hypothesis 1.2 (ACTION increase)**.** *The alignment gap in provision is larger in* ACTION *than* CONTROL*.*

Hypothesis 1.3 (INFO increase)**.** *The alignment gap in provision is larger in* INFO *than* CONTROL*.*

2.2 Results

2.2.1 Sample, Incentives, and Protocol

The sender experiment was run on Prolific from July 16 to August 12, 2025, quota-sampled on age, gender, ethnicity, and party identification.[10] This study was approved by the Ethics Council of the Max Planck Society under the Generalized Approval of Experiments Following the Protocol that is Standard in Experimental Economics (approval no. 2018_3 Renewal 2024_28) and was pre-registered on AsPredicted (No. 238343). Of 1,257 starters, 57 were excluded for failing the attention check, switching tabs during the fact questions, or timing out, leaving $N = 1{,}200$.[11] Senders took a median of 17.26 minutes and earned $3.32 on average.[12] Receiver data were collected on Prolific between September 1 and September 10, 2025.[13] 96 receivers took a median of 3 minutes and earned $0.80 on average. For each of the 16 combinations of fact, receiver answer and whether the receiver donates, one sender was drawn, and receivers were recruited until that sender's decision had reached one. The experiment was programmed in oTree (Chen et al. 2016).

[10] Appendix Table B.1 reports the sample composition.

[11] Twelve of these participants had already completed the study in an earlier session. Dropping their second completions, all in the INFO treatment, leaves the CONTROL and ACTION results unchanged and raises the alignment bias under INFO from 7.5 to 8.1 pp.

[12] They receive $2.50 for participating and can earn a bonus of $1.50, knowing that the task determining it is drawn at random: the prior-knowledge questions with probability 45%, the persuasiveness beliefs with probability 45%, and the provision task with probability 10%. The provision task pays the bonus whatever the sender decides.

[13] Receivers first answered the question. A sender's decision was implemented only for receivers who answered it falsely or said they did not know, so the description senders read holds for every receiver their decision reached.

2.2.2 Prior Knowledge, Policy Views, and Expected Attitude Change

Prior knowledge. I find substantial heterogeneity in senders' factual knowledge.[14] A majority correctly identify the facts on immigrants' crime rates (70.1%), labor force participation (79.4%), and recent border crossings (61.3%), but only 30.5% correctly assess the net fiscal impact of first-generation immigrants. Democrats and immigration supporters are more accurate on the pro-immigration facts, Republicans and immigration opponents on the anti-immigration facts (all differences $p < 0.001$).

Policy views. 30.2% of senders allocate the entire endowment to the pro-immigration organization, 14.1% to the anti-immigration one, and 7.2% split evenly (Figure B.2). The split correlates with stated policy views ($r = 0.562$) and with political affiliation (both $p < 0.001$), elicited at the end of the experiment.

Expected attitude change. For each sender i and fact j, I define the expected attitude change $\Delta\text{Attitude}_{i,j} \in [-1, 1]$ as the net share of ten hypothetical receivers whom the sender expects to become more, rather than less, supportive of immigration after receiving fact j. These beliefs are elicited unincentivized before the provision decisions. On average senders expect each fact to move attitudes in the direction of the policy it is hypothesized to support. Among senders not shown the study results, the expected shift is 25.5 percentage points (pp) toward pro-immigration for the crime fact and 20.7 pp for the labor fact, and 29.5 pp toward anti-immigration for the border fact and 26.7 pp for the fiscal fact (all $p < 0.001$).

2.2.3 Selective Provision and the ACTION Increase

Let $y_{i,j}$ equal 1 if sender i provides fact j, and let $AIC_i \in [0, 500]$ denote i's donation to the pro-immigration group (AIC). I consider sender i pro-immigration if $AIC_i > 250$ and anti-immigration if $AIC_i < 250$. A sender's *alignment bias* is the share of their aligned facts provided minus the share of their unaligned facts provided. It is positive when provision favors the sender's own policy position and zero when it does not.

Senders share accurate information at high rates in all four treatment conditions, including information that cuts against their own position. Unaligned provision rates range from 81.9% (ACTION) to 86.5% (CONTROL), all significantly above 50% ($p < 0.001$, one-sample t-tests).

Result 1.0 (Provision is high, including unaligned information)**.** *Senders share most unaligned facts:* 86.5 *percent in* CONTROL*, above the* 50 *percent benchmark in every treatment, confirming Hypothesis 1.0.*

[14] See Appendix Table B.2. The fiscal fact, where prior knowledge is lowest, is rated ex post as credible as the other three (Appendix Table B.3) and provided at comparable rates.

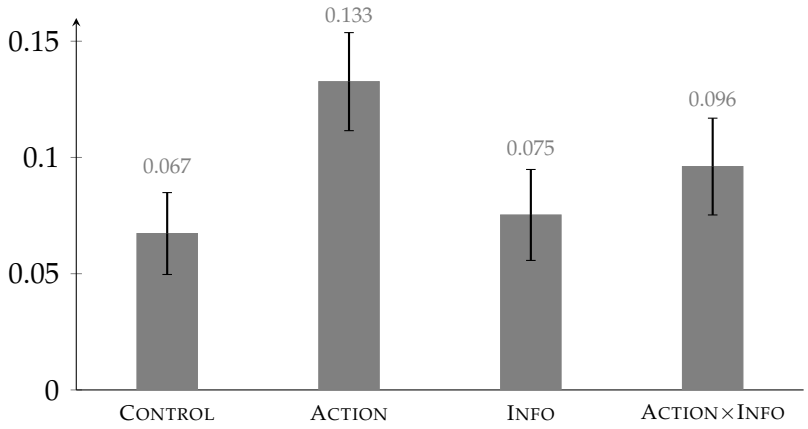


Figure 2: Alignment Bias by Treatment

Notes: The sample is Experiment 1's U.S. adults (N = 1,200, Prolific, quota-sampled). Alignment bias is the within-sender difference between the provision rate for aligned statistics and the provision rate for unaligned statistics. The figure is restricted to senders with revealed directional preference (N = 1,114, 93%; donation not an even split). Error bars are standard errors.

Figure 2 shows that provision is selective. On average across treatments, senders provide aligned facts 9.3 pp more often than their unaligned facts, and the bias is positive in all four conditions ($p < 0.001$).[15] The bias is defined for the 1,114 senders whose donation leaves them on one side of the issue.

I combine the two into a continuous alignment measure:

$$\text{Align}_{i,j} = \begin{cases} \frac{AIC_i - 250}{250} & \text{if } j \in \{\text{crime, labor}\} \\ \frac{AIC_i - 250}{250} \cdot (-1) & \text{if } j \in \{\text{border, fiscal}\} \end{cases}$$

This measure is bounded in $[-1, 1]$.[16] Its sign records whether fact j's direction matches sender i's policy view, and its magnitude is the intensity of that policy view. Only the sign, not the magnitude, varies within a sender. Thus, the within-sender comparison is between aligned and unaligned facts and the magnitude enters as a between-sender weight. This imposes, rather than tests, the assumption that a fact is more aligned, and the alignment gap correspondingly larger, the more one-sided the sender's position. Appendix Table B.7 shows that the results are robust to dropping this assumption.

I estimate the pre-registered linear probability model:

$$\begin{aligned} y_{i,j} = \beta_0 &+ \beta_1 \,\text{Align}_{i,j} + \beta_2 \,\text{ACTION}_i + \beta_3 \,\text{INFO}_i + \beta_4 \left(\text{Align}_{i,j} \times \text{ACTION}_i\right) \\ &+ \beta_5 \left(\text{Align}_{i,j} \times \text{INFO}_i\right) + \beta_6 \left(\text{ACTION}_i \times \text{INFO}_i\right) \\ &+ \beta_7 \left(\text{Align}_{i,j} \times \text{ACTION}_i \times \text{INFO}_i\right) + \gamma_1 \,\text{PriorKnowledge}_{i,j} + \gamma_2 \mathbf{S}_j + \gamma_3 \mathbf{X}_i + \eta_{i,j}, \end{aligned} \tag{1}$$

[15] Controlling for distribution of policy views the difference is +7.6 pp (Appendix B.3).

[16] The pre-registration defined alignment on $[0, 1]$. Re-centering is a one-to-one transformation that leaves all hypothesis tests invariant.

with standard errors clustered at the sender level. PriorKnowledge$_{i,j}$ denotes sender i's prior knowledge of fact j (1 if correct, 0.5 if "don't know," and 0 if incorrect, as pre-registered). This control adjusts for prior accuracy under the extreme assumption of no updating. Column (1) controls only for prior knowledge. Column (2) adds fact fixed effects $\mathbf{S}_j$, and column (3) adds demographic and political controls $\mathbf{X}_i$. β_1 captures baseline alignment-based selectivity (Hypothesis 1.1), β_4 whether alignment matters more when receivers act politically (Hypothesis 1.2), and β_5 whether information about persuasiveness increases selectivity (Hypothesis 1.3). The control γ_1 estimates whether senders who had more accurate priors about the fact provide it at higher rates.

Table 3 reports the results. Across all specifications, provision increases significantly with alignment ($p \leq 0.003$). Senders at half the maximum policy view intensity are 3.6 to 4.1 pp more likely to share aligned facts while senders at maximum policy view intensity are 7.2 to 8.2 pp more likely to share aligned facts.

Result 1.1 (Senders favor aligned information)**.** *Senders favor aligned information: the alignment bias is* 6.7 *pp in* CONTROL*, confirming Hypothesis 1.1, and positive in every condition.*

Political consequences increase alignment-based selectivity, confirming Hypothesis 1.2. Comparing CONTROL with ACTION, the sender-level alignment bias rises from 6.7 to 13.3 pp, an increase of 6.5 pp (one-sided Holm-adjusted $p = 0.045$). The decision-level interaction Align$_{i,j} \times$ ACTION$_i$ is $+0.041$ in every specification ($p \leq 0.025$). Averaging over the independently randomized INFO intervention, the bias rises from 7.1 to 11.4 pp, an increase of 4.3 pp (two-sided $p = 0.037$).

Result 1.2 (ACTION increases alignment-based selectivity)**.** *Political action increases alignment-based selectivity: the alignment bias rises from* 6.7 *to* 13.3 *pp, confirming Hypothesis 1.2.*

Political action also raises the slant of provision, the mean positional alignment of the facts senders provide, which measures how far they favor the sender's own side: averaging over INFO, it rises from 0.031 in CONTROL to 0.053 in ACTION, against a mean of zero among the facts offered (two-sided $p = 0.010$).

Senders provide facts about which they had accurate priors at higher rates (3.3 to 4.3 pp, $p \leq 0.006$). Controlling for prior knowledge does not attenuate the Align $\times$ ACTION interaction.

Hypothesis 1.3 is not confirmed. The INFO treatment raised the expected pro-immigration shift by 5.5 pp for the crime fact and 3.5 pp for the labor fact and did not detectably move it for the two anti-immigration facts, too little to identify effects.[17] Results here are only weakly informative. The pre-registered comparison of CONTROL with INFO gives $+0.8$ pp (one-sided $p = 0.385$), and neither the main effect of INFO

[17]The first stage is reported in detail by fact in Appendix A.3.

Table 3: Treatment Effects on Information Provision

	Probability of Statistic Provision		
	(1)	(2)	(3)
Align	0.041***	0.036***	0.037***
	(0.012)	(0.012)	(0.012)
INFO	−0.013	−0.012	−0.012
	(0.017)	(0.017)	(0.017)
ACTION	−0.002	−0.002	−0.003
	(0.016)	(0.016)	(0.016)
Prior Knowledge	0.043***	0.034***	0.033***
	(0.011)	(0.012)	(0.012)
Align × INFO	0.006	0.004	0.004
	(0.017)	(0.017)	(0.017)
Align × ACTION	0.041**	0.041**	0.041**
	(0.018)	(0.018)	(0.018)
INFO × ACTION	−0.007	−0.007	−0.008
	(0.024)	(0.024)	(0.024)
Align × INFO × ACTION	−0.028	−0.027	−0.027
	(0.026)	(0.026)	(0.026)
Clustered SE	Subject	Subject	Subject
Statistics FE	No	Yes	Yes
Controls	No	No	Yes
No. of Senders	1200	1200	1200
Num. obs.	4800	4800	4800
R^2	0.030	0.033	0.040
Adj. R^2	0.028	0.031	0.034

Notes: The sample is $N = 1{,}200$ Experiment 1 U.S. adults recruited on Prolific (quota-sampled), with 4,800 sender-statistic observations (4 statistics per sender). The outcome is the binary provision indicator. Align is $\text{Align}_{ij} \in [-1, 1]$, the sender's donation-based policy view signed by the direction of statistic j, as defined in the text. Prior Knowledge is the incentivized prior belief correctness, elicited before senders knew they would make a provision decision. INFO is a between-subject treatment in which the sender receives personalized feedback comparing their prediction to the actual results of Haaland and Roth (2020) (see Appendix A.3). ACTION is a between-subject treatment in which the receiver makes a downstream political donation. Controls are age, gender, education, ethnicity, income, party affiliation, partisanship. Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

nor its interaction with alignment is distinguishable from zero, including when INFO is allowed to act asymmetrically by fact direction (Appendix Table B.8).

Result 1.3 (INFO increase not confirmed)**.** *Information about persuasiveness does not detectably change the alignment bias; Hypothesis 1.3 is not confirmed.*

The estimates are unchanged under alternative specifications including sender fixed effects, presentation-position fixed effects, and the unrestricted starter sample. No single fact drives the result (Appendix B.3).

2.2.4 Heterogeneity and Alternative Explanations

Heterogeneity. Appendix B.4 reports pre-registered heterogeneity analyses by partisanship, political affiliation, policy view intensity and receiver belief state. Partisanship and affiliation with a political party increase selectivity in the baseline ($p = 0.001$ and $p = 0.032$) while policy view intensity points the same way ($p = 0.053$). Independents respond more strongly to political consequences, but estimates are imprecise ($p = 0.054$). The alignment-based selectivity does not detectably differ by receiver belief state.

Disbelief and prior knowledge. Selective provision could reflect selective belief: senders might withhold unaligned facts because they doubt them. Prior knowledge, incentivized and elicited before any provision decision, is controlled in every specification, and the alignment effect does not vary with it: the Align$\times$ Prior Knowledge interaction is $+0.009$ ($p = 0.579$), and the alignment coefficient is $+0.031$ for facts the sender had answered incorrectly and $+0.040$ for facts answered correctly. Perceived credibility is elicited unincentivized after the provision task. Although senders could rationalize withholding a fact by declaring it not credible, they rate 97.2 percent of facts credible. Credibility is slightly higher for aligned facts (98.3 and 96.0 percent). Even if every differentially non-credible unaligned fact were withheld purely from disbelief, that channel is bounded at 2.4 pp, against alignment biases of at least 6.7 pp in every treatment.[18]

Stated motives. Senders' own accounts provide descriptive evidence on the motives for their decisions. A closing question asked what considerations drove the decisions. Two in five senders gave a reason referring to the sharing decision. Those senders hold somewhat stronger issue policy views and are somewhat more partisan, yet are 4.0 pp *less* alignment-biased than the rest ($p = 0.043$). Two large language models classify each response under a fixed coding scheme, described in Appendix F.1.

Among senders giving a reason for sharing, correcting the receiver's beliefs is the most common consideration (84.9 percent), followed by influencing the receiver's opinion (7.9 percent) and the fact's fit with one's own position (2.6 percent). Appendix Table B.15 regresses each sender's alignment bias on these stated reasons. Senders citing belief correction are 9.9 pp less alignment-biased ($p < 0.001$), and senders citing persuasion or expression are more biased ($p < 0.001$). However, only persuasion is action-conditional: its association with the alignment bias rises when the receiver takes a political action (interaction $+0.417$, $p = 0.008$), while the expression coefficient is large

[18]The bound covers facts rated outright not credible. Additional tests are in Appendix B.5.

in both treatments and its interaction with ACTION is negative and imprecise.[19] The belief-correction coefficient likewise does not move with political action. Notably, ACTION changes what senders do without changing what they report considering. The share giving a codable reason and the shares citing each consideration are indistinguishable across treatments, individually and jointly ($F = 0.443$, $p = 0.722$).

Norms. Senders who have a larger alignment bias subsequently express weaker pro-accuracy norms and policy intervention and state facts are less relevant in forming opinions. There is no difference in perceived source of disagreement: values or facts (Appendix Table B.16).

2.2.5 What Experiment 1 Does Not Identify

Four accounts of the slant in what senders provide fit the observed pattern poorly. Differential exposure cannot generate it: every sender decides over the same four facts, two favoring each side, so no fact is withheld because a sender never encountered it. The alignment effect survives among senders who knew facts on both sides, a restriction that supplies two-sided support rather than balance. A costless disbelief rationalization was available and rarely used. Nor can provision signal the sender's position, because the receiver never learns that a decision was made.

The alignment gap however does not separate expression from persuasion. A sender may value passing on facts that favor their own side, or may expect those facts to move the receiver toward it. Here, senders expect net movement toward their own side for 78.0 percent of aligned facts and 17.2 percent of unaligned ones, a gap of 60.8 pp ($p < 0.001$). A potential return to correcting beliefs is also not fully identified. Every fact goes to a receiver who holds a false belief or reports not knowing, so there is a corrective rationale by construction. Provision does not detectably vary with the receiver's factual need: receivers who have false beliefs versus those who do not know. However, the experiment never elicits the correction a sender expects provision to achieve.

3 A Conceptual Framework for Selective Sharing

Alignment-based selectivity can reflect different motives. I model what a sender gains from passing on a fact, to show *when* observed disclosure and expected returns can tell these motives apart. The motives differ in whether she values what providing does for herself, what receipt does to the receiver, or who the receiver is.

[19] Notably, only twelve senders state an expressive reason, so this coefficient should be interpreted cautiously.

3.1 Motives, Returns, and Utility

Environment. Nature draws a binary state of the world $\omega \in \{0, 1\}$, the sender's and the receiver's policy positions $\pi^s, \pi^r \in \mathbb{R}$, and an indicator $A \in \{0, 1\}$ for whether the receiver goes on to take a political action. The fact is a statement whose content is ω.

Timeline. The sender observes ω, π^s, π^r and A, and then chooses whether to disclose the fact, $d = 1$, or withhold it, $d = 0$. A disclosed fact always reaches the receiver with a fixed arrival probability $\tau \in (0, 1]$.[20] Once received, the receiver's actual belief response is left unrestricted. The sender knows that the receiver observes π^s regardless of her choice. If $A = 1$, the receiver then chooses a political action informed by his beliefs about the world and his policy position.

Motives and returns to disclosure. I divide the motives by whether they depend on the receiver. The *expressive motive* with weight θ_E is *receiver-invariant*. It is the value of providing a fact that validates the sender's own position, whatever it does to the receiver. Let $\alpha \in \mathbb{R}$ denote the *alignment* of the fact to her political position, the extent to which she perceives it as validating that position. Larger values indicate larger validation, and negative values that she regards the fact as conflicting with her position. For a given sender, alignment, the expressive return, depends only on the fact, not on who receives it or on what he goes on to do.

The other motives are *receiver-dependent*: the first two depend on the receiver's response and the third on who he is. First, the *belief-correction motive* with weight θ_B is the value of moving the receiver's beliefs toward the truth, whichever side the truth favors. Let $\tilde{b}^0$ and $\tilde{b}^1$ denote her subjective forecasts of the receiver's belief, i.e., the probability he assigns to $\omega = 1$, under non-receipt and under receipt, with $\tilde{b}^0, \tilde{b}^1 \in [0, 1]$. Neither need equal the belief the receiver will in fact hold. The *expected gain in belief accuracy* from receipt is

$$\Delta b = |\omega - \tilde{b}^0| - |\omega - \tilde{b}^1| \in [-1, 1],$$

how much closer to the truth she expects him to be once he has the fact. It is positive whenever she expects any movement toward the truth, and negative when she expects receipt to move him away from it.

Second, the *persuasion motive* with weight θ_P is the value of moving him toward the sender's own position. Let $\tilde{\pi}^{r,0}$ and $\tilde{\pi}^{r,1}$ denote her forecasts of the receiver's policy position under non-receipt and under receipt. Neither need equal the position π^r she

[20] The arrival probability captures everything that stands between sending and being consumed, e.g., an unopened message, a skimmed article, or attention in conversation. Withholding and non-arrival both leave the receiver without the fact.

initially observes. The *expected persuasion return* is

$$\Delta\pi = |\pi^s - \tilde{\pi}^{r,0}| - |\pi^s - \tilde{\pi}^{r,1}| \in \mathbb{R},$$

the reduction she expects receipt to produce in the distance between the receiver's position and her own. Thus $\Delta\pi > 0$ means she expects receipt to move the receiver toward her, and $\Delta\pi < 0$ that she expects it to move the receiver away. Correction evaluates beliefs relative to the truth. Persuasion evaluates movement relative to the sender's own position.[21]

Third, the *relational motive* with weight θ_R is the additional value of providing any true information to a more like-minded receiver rather than a more distant one, whatever the fact. Let $\rho = |\pi^s - \pi^r|$ denote the initial political distance between the sender and the receiver.

The three returns and political distance need not move together: alignment is fixed by the fact and her position, while the two forecasts concern a particular receiver. She may expect a fact that validates her position to correct one receiver's belief a great deal without bringing him any closer to her.

Sender utility. Normalizing the utility of withholding to zero, $U(d = 0) \equiv 0$, the sender's net utility from disclosure is

$$U(d = 1) = \tau\,\theta_B(A)\,\Delta b + \theta_E(A)\,\alpha + \tau\,\theta_P(A)\,\Delta\pi - \theta_R(A)\,\rho + \kappa\,A, \tag{2}$$

where the four weights may depend on A. The returns describe what the sender expects provision to cause, and the weights describe how much she values those accomplishments. Finally, $\kappa \in \mathbb{R}$ shifts the value of disclosing anything when the receiver takes a political action, without changing what any return is worth at the margin. I use this linear benchmark for simplicity: the sender trades the returns at fixed rates. Every sign result below also holds when each determinant enters through its own strictly increasing transformation (Appendix I).

The arrival probability multiplies only the two returns that require the fact to arrive. Alignment and political distance arise from the act of providing and from who is at the other end, not from the receiver's response to what is sent.

Excluded motives. The receiver's response could affect the sender's payoff through his beliefs about the world and his position, or through what he comes to *believe about*

[21] The forecast under non-receipt is a primitive of the sender's problem and may incorporate whatever inference she expects the receiver to draw from not receiving the fact. When $\tau < 1$, non-receipt pools deliberate withholding with non-arrival and need not reveal the disclosure choice. When $\tau = 1$, non-receipt reveals withholding.

the sender. Equation (2) includes only the former. Because he observes her position in any case, disclosure earns no return from revealing that she holds one. Any discounting of the fact enters through her forecasts of his response, and I assign no separate return to its perceived credibility or to signaling an ability to tell true from false. Returns to how the sender is judged lie outside this analysis (see Braghieri 2024, Bursztyn et al. 2023a, Guriev et al. 2026).[22]

3.2 Disclosure and Selection

Disclosure decision. The sender discloses if and only if this utility exceeds an idiosyncratic shock capturing unmodeled determinants of disclosure, $U(d = 1) > \varepsilon$.

Assumption 1 (Conditional disclosure shocks)**.** *Conditional on A, the disclosure shock ε is distributed with an action-state-specific CDF F_A that does not vary with $(\Delta b, \alpha, \Delta\pi, \rho)$, has full support on $\mathbb{R}$, and has strictly positive density $f_A = F_A'$.*

Write $x = (\Delta b, \alpha, \Delta\pi, \rho)$ for the four determinants and $U_A(x)$ for the utility of disclosure in Equation (2) under A. The probability of disclosure is therefore

$$\Pr(d = 1 \mid x, A) = F_A\big(U_A(x)\big). \tag{3}$$

For a given A, disclosure is more likely the higher the expected utility to the sender.

Selectivity.

Proposition 1 (Selection)**.** *Fix A. Under Assumption 1, holding the other determinants fixed, the probability of disclosure*

(i) rises with the expected accuracy gain Δb if $\theta_B(A) > 0$,

(ii) rises with alignment α if $\theta_E(A) > 0$,

(iii) rises with the expected persuasion return $\Delta\pi$ if $\theta_P(A) > 0$, and

(iv) falls with political distance ρ if $\theta_R(A) > 0$.

In each case the probability response to a determinant x_j, $f_A(U_A)\,\partial U_A/\partial x_j$, has the sign of its marginal utility $\partial U_A/\partial x_j$.

The relational motive does not depend on the fact. When $\theta_R(A) > 0$, it makes her more willing to disclose to a like-minded receiver. A decline in disclosure with political

[22] Neither the observed position nor the truth of the fact removes every such inference. A receiver may still read how informed or how selective she is, and which fact she picks may convey how strongly she holds her position. Those inferences are assumed away.

distance reflects this motive only if distance does not stand in for receiver responses the returns omit. Assumption 1 maintains this.[23]

A higher disclosure rate for aligned facts is consistent with expression, persuasion, or correction. When the facts that validate a sender's position are also those she expects to persuade, as in Section 2, or to correct, she discloses aligned facts more often even if she places no value on alignment itself. Telling the motives apart requires observing each return, and that none perfectly move together.[24]

Assumption 2 (Separation)**.** *Across disclosure opportunities, conditional on A, none of Δb, α, $\Delta\pi$ and ρ is an exact affine function of the other three.*

3.3 The Effect of Receiver Political Action

When the receiver goes on to take a political action, what he believes and where he stands matter beyond the exchange itself. The model lets political action change how much the sender values the returns that run through the receiver's response.

Restrictions. Comparing disclosure of a given fact to a given receiver across A reveals a change in how much the sender values a return only if her forecasts do not change. I therefore assume that she expects the same fact to do the same to the receiver's beliefs and position whether or not he later acts. Otherwise, more frequent disclosure under political action could equally come from her expecting the fact to persuade more. Comparing how disclosure responds to a return, at given values of the returns, does not need this assumption. Alignment needs no such assumption, since it depends only on the fact and her position. I also assume that political action leaves the value of alignment and of political distance unchanged: action matters through what the receiver's beliefs and position lead to.

Assumption 3 (Invariance)**.** *For the sender, (i) for a given fact and receiver, Δb and $\Delta\pi$ take the same values whether or not the receiver goes on to take a political action; (ii) $\theta_E(A) = \theta_E$; (iii) $\theta_R(A) = \theta_R$.*

Three cases. Political action may leave the sender's valuations unchanged, $\theta_B(1) = \theta_B(0)$ and $\theta_P(1) = \theta_P(0)$. Under *consequential correction* it raises the value of correction, $\theta_B(1) > \theta_B(0)$. A sender who values the receiver holding accurate beliefs may value them more when he acts on them, because he can then make a more informed choice.

[23] Canonical disclosure models let the sender value disclosure only through the receiver's response (Grossman 1981, Milgrom 1981, Dye 1985). In Equation (2), objectives of this kind are the case $\theta_E(A) = \theta_R(A) = 0$.

[24] An alternative is to shift one return experimentally while holding the others fixed. The INFO treatment of Section 2 attempted this and moved expected persuasiveness too little to be informative.

This value does not depend on whether the informed choice moves toward or away from her position.

Under *consequential persuasion* it raises the value of persuasion, $\theta_P(1) > \theta_P(0)$. When a political action follows, a sender who prefers actions closer to her own position values the same expected movement more, provided a receiver closer to her position tends to choose actions she prefers. $\theta_P(1)$ then values his position both in itself and for the choice it leads to.

Proposition 2 (Consequentiality)**.** *Under Assumption 3(ii)–(iii), political action leaves the marginal utilities of alignment and political distance unchanged. It changes the marginal utility of the expected accuracy gain by $\tau[\theta_B(1) - \theta_B(0)]$ and that of the expected persuasion return by $\tau[\theta_P(1) - \theta_P(0)]$. Hence (i) if valuations are unchanged, neither moves; (ii) under consequential correction the first rises; (iii) under consequential persuasion the second rises.*

3.4 What Disclosure Reveals

Proposition 2 compares marginal utilities, but only disclosure is observed. Within an action-state, Proposition 1 links the two: each probability response has the sign of its marginal utility. Across action-states, the response to a determinant x_j, at the same x, changes by

$$f_1(U_1)\,\frac{\partial U_1}{\partial x_j} - f_0(U_0)\,\frac{\partial U_0}{\partial x_j}. \tag{4}$$

If the two densities are equal, this has the sign of the change in marginal utility, but they need not be. A change in disclosure probabilities across action-states therefore does not necessarily reveal a change in what the sender values. Two things can move a probability response without such a change. Political action may make the sender disclose more or less of everything. The same change in utility then moves the disclosure probability by a different amount, because the shock density differs at the new disclosure rate. Political action may also make her decisions more or less noisy, for instance more deliberate when the receiver's choice has consequences, which changes the responses to all determinants by a common factor. The first acts as a shift in the location of the shock, the second is a change in its scale.

Assumption 4 (Location–scale disclosure shocks)**.** *(i) For any A the shock has a known, strictly increasing and differentiable shape G up to a location μ_A and a scale $s_A > 0$: $F_A(u) = G\big((u - \mu_A)/s_A\big)$. (ii) Where index slopes are compared across A, $s_1 = s_0$.*

Part (i) leaves the shape general but requires it to be the same in both action-states. Because the shape is known, it can be undone: transforming the disclosure probability by G^{-1} gives an index, $(U_A(x) - \mu_A)/s_A$, that is linear in the sender's utility. Write

$\beta_j(A) = \partial\, G^{-1}\big(\Pr(d = 1 \mid x, A)\big)/\partial x_j$ for the slope of this index on determinant x_j, and $\beta_B(A)$, $\beta_E(A)$, $\beta_P(A)$ and $\beta_\rho(A)$ for the slopes on Δb, α, $\Delta\pi$ and ρ.

Corollary 1 (Index slopes and ratios)**.** *Under Assumptions 1 and 4(i), each index slope is the marginal utility of its determinant divided by the scale, $\beta_j(A) = s_A^{-1}\,\partial U_A/\partial x_j$, and the location drops out.*

(i) Under Assumptions 3(ii)–(iii) and 4(ii), with $s_1 = s_0 = s$, Proposition 2 carries over to the index up to the positive constant s^{-1}.

(ii) If $\theta_E(A) \neq 0$, the ratio $\beta_j(A)/\beta_E(A)$ does not depend on the scale. It is the sender's marginal rate of substitution between x_j and alignment: how much alignment she would trade for one more unit of x_j. For political distance, which lowers utility, the ratio is negative when $\theta_E(A) > 0$. If, in addition, $\theta_E > 0$ and Assumption 3(ii)–(iii) holds, then at given values of the determinants the persuasion ratio is higher under political action if and only if $\theta_P(1) > \theta_P(0)$, the correction ratio is higher if and only if $\theta_B(1) > \theta_B(0)$, and the distance ratio does not change with A.

A change in noise rescales every index slope by the same factor. If political action makes the sender more deliberate, all her index slopes steepen together, so noise cannot steepen her response to persuasion alone. A change in one slope relative to the others reflects a change in what she values. The corollary gives two ways to read this. The first compares each slope across action-states and assumes a common scale. The second compares each slope with the expression slope within an action-state, where the scale cancels, and assumes instead that political action leaves the value of expression unchanged. This restriction is natural because expression does not run through the receiver's response. Without it, a higher persuasion ratio still shows that persuasion gained in value relative to expression, given a positive expressive weight. The first comparison fails if the scale changes and the second if the value of expression does, so agreement between them is informative.[25]

4 Experiment 2: The Drivers of Selectivity

Experiment 2 asks what senders select on, and whether that selection changes when the receiver goes on to take a political action. Provision is predicted by how strongly a fact validates the sender's position, by the correction and the persuasion they expect it to achieve, and by how close the receiver's position is to their own. When the receiver

[25] Under the linear benchmark, $\beta_m(A)/\beta_E(A) = \tau\,\theta_m(A)/\theta_E(A)$ for $m \in \{B, P\}$ and $\beta_\rho(A)/\beta_E(A) = -\theta_R(A)/\theta_E(A)$. With θ_E fixed, the change in the persuasion ratio equals $\tau[\theta_P(1) - \theta_P(0)]/\theta_E$, so its size, as well as its sign, measures the change in the persuasion weight relative to the expressive weight (Appendix I).

takes a political action, the evidence suggests that senders select more on the expected persuasion return and that provision falls less steeply with distance. Overall the slant of provision does not detectably differ between the treatments.

4.1 Design and Hypotheses

Experimental timeline. Senders first report demographics and two policy view measures (Figure 3). They then make the provision decisions: senders decide whether to share each of five assigned facts, one decision at a time, directed at a receiver whose policy stance is visible. Finally, they give incentivized forecasts of how each fact affects the receiver's belief accuracy and policy position. A post-questionnaire, which also collects the ratings used to measure alignment, closes the study.

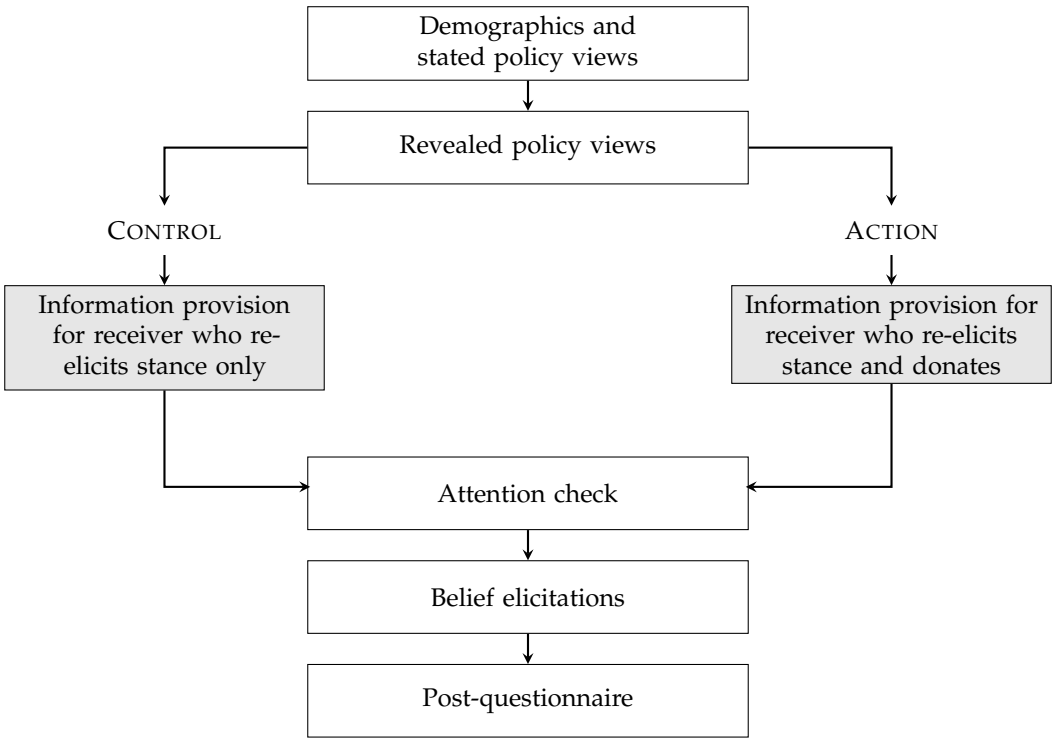


Figure 3: Timeline of Experiment 2

Facts. Table 4 lists the twelve facts, six in each of the two policy domains, immigration and climate. Each is presented as a binary, invertible comparison and drawn from official statistics and published research. Senders and receivers are told at the outset that every fact is true and sourced. Each sender is randomly assigned to one domain and sees five of its six facts with two for each policy direction plus one drawn at random. I pre-validated the facts in a pre-registered survey ($N = 319$) and balanced them

for credibility and relevance.[26]

Table 4: Facts Used in Experiments 2 and 3

Direction	Fact
Immigration policy	
Pro	The share of people with a migration background in shortage-occupation fields is *higher* than their population share.
Pro	The absolute increase in employed foreign nationals between 2013 and 2023 was *larger* than for German nationals.
Pro	*More* than half of the persons attending a general integration course for the first time complete it without interruption through the final module.
Anti	Among persons suspected of violent crimes, the share of non-German nationals is *higher* than their population share.
Anti	The share of persons receiving basic income support (*Bürgergeld*) is *higher* among foreign nationals than among German nationals.
Anti	Germany records *more* persons without a valid residence permit per capita than the EU average.
Climate policy	
Pro	In German electricity generation in 2023 and 2024, the share of renewable energies was *higher* than the share of fossil fuels.
Pro	The cost of generating one kilowatt-hour of electricity from new, ground-mounted solar installations is *lower* than from new gas-fired power plants.
Pro	The percentage increase in renewable energy jobs in Germany between 2022 and 2023 was *higher* than the global average.
Anti	Grid fees for household customers are expected to be *more* than double their current level by 2045 due to grid expansion.
Anti	There are *fewer* companies that view the effects of the energy transition positively than companies that view them negatively.
Anti	At higher shares of solar power, *more* price fluctuations are observed in European electricity markets.

Notes: Claims are verbatim as shown to senders in English translation (correct answer in *italics*). Direction (Pro) = supports more-permissive immigration or more-ambitious climate policy; (Anti) = supports more-restrictive or more-cautious policy. Pool of six facts per topic. Each sender draws five facts: two pro-policy, two anti-policy, and one random tiebreaker. Topics assigned between subjects. Sources and pre-validation details in Appendix A.4.

Policy views. I measure policy views two ways. Senders place themselves on a five-point policy scale on each of the two issues, and they split €250 between the pro-policy and the anti-policy advocacy organization for their assigned domain in a real donation decision. One donation decision from the study is drawn and implemented with real money. On immigration the organizations are *ProAsyl*, advocating for more permissive immigration policy, and *Republik21*, advocating for more restrictive immigration policy. On climate, *BUND*, supporting more national and less market-based climate policies,

[26]The fielded set departs from the pre-registered selection rule in its directional threshold and in retaining the climate domain (Appendix A.4).

and *Stiftung Marktwirtschaft*, supporting less national and more market-based climate policy. Senders see each organization's position in its own words before deciding.[27]

Receivers. Each sender faces five receivers sequentially, one at each position on the five-point scale. Before every decision the sender sees the receiver's stance stated in full, and knows the receiver will be told the sender's position in turn. Receiver positions are presented in randomized order, and facts are randomly assigned to receiver slots. Thus, there is within-sender variation in political distance that is unrelated to the fact by design. Senders know they are matched to a real receiver in at least four of ten cases. Conditional on a match, a decision to share is implemented with probability one half and a decision not to share is always implemented. Senders know a receiver who sees nothing therefore cannot tell withholding from a failed transmission.

Elicited returns. The final stage elicits the two returns that vary with the receiver. For each receiver–fact pair, senders predict the percentage of receivers holding that stance who would answer the fact's question correctly with and without receipt, with the difference the expected accuracy gain $\widehat{\Delta b}$. Senders also predict these receivers' average policy view on the five-point scale with and without the fact. $\widehat{\Delta \pi}$ is the distance between the sender's own position and the receivers' predicted mean position without the fact, minus that distance with the fact, so it is positive when the sender expects the fact to bring receivers closer. Both predictions are scored under a truncated quadratic scoring rule against matched receivers' realized outcomes.[28] In the post-questionnaire, senders also rate two facts from the domain they were *not* assigned, one hypothesized pro-policy and one hypothesized anti-policy: whether the information *"shows that your opinion is correct,"* and, separately, whether they are convinced it is true, which measures perceived credibility.[29] Because senders hold a stated stance on both issues, these ratings can be grouped by the rater's stance on the domain being rated and supply the stance–fact alignment index $\widehat{\alpha}$ for the senders who provide facts in that domain without rating it themselves.

Treatments and hypotheses. Senders are randomly assigned between subjects to CONTROL or ACTION, as in Experiment 1. They know in both treatments that receivers will restate their policy position after receiving or not receiving a fact. In the ACTION

[27] Appendix A.5 displays the organization descriptions shown to senders.

[28] Each forecast concerns all matched receivers who hold that stance rather than a single receiver, so it can be scored against their realized share answering correctly and their average policy view. Score is zero beyond an error of 20 percentage points or two scale points. Therefore, reporting the expectation is guaranteed to be optimal only for a sender sure of the outcome within that large range.

[29] Importantly, the instructions ask the rater to judge whether the fact validates their position, not whether their position is correct in the abstract, and not whether the fact would convince anyone else.

treatment receivers also make a real €250 donation decision between the same two organizations.

I pre-registered three hypotheses, mirroring Experiment 1.

Hypothesis 2.0 (Baseline provision)**.** *Provision rates for unaligned facts exceed the indifference benchmark of* 50% *in each treatment-topic cell.*

Hypothesis 2.1 (Selective provision)**.** *Senders are more likely to provide a fact aligned with their own policy view than one that is not, holding fact and topic fixed.*

Hypothesis 2.2 (ACTION increase)**.** *The alignment gap in provision is larger when the receiver makes a real political donation (*ACTION*) than when the receiver only restates their opinion (*CONTROL*).*

I also pre-registered a secondary analysis relating provision to the expressive return, the expected accuracy gain and the expected persuasion return, and to their interactions with ACTION. Analyses in Section 4.2.4 build on it.

4.2 Results

4.2.1 Sample, Incentives, and Protocol

Data collection ran from May 27 to June 4, 2026 on Bilendi, with quotas that match the German adult population on age, gender, and state. This study was approved by the Ethics Council of the Max Planck Society under the Generalized Approval of Experiments Following the Protocol that is Standard in Experimental Economics (approval no. 2018_3 Renewal 2024_28) and was pre-registered on AsPredicted (No. 293252). Of 1,694 completers, $N = 1{,}548$ senders providing 7,740 decisions form the main sample. Forty were excluded as speeders and 106 because they started before the study was shortened, as Appendix A.6 describes. Senders took a median of 17 minutes and earned €3.06 on average.[30] The experiment was programmed in oTree (Chen et al. 2016).

4.2.2 Policy Views and Measured Returns

Policy views. Senders' stated and revealed policy views correlate at $r = 0.627$ on immigration and $r = 0.503$ on climate (Appendix Figure C.1). As pre-registered, each measure is normalized to $[-1, 1]$ and averaged into a policy view index, which spans its full range in both domains.

[30] Senders receive €3 for participating. Twenty-five senders drawn at random can earn a bonus of €10, paid with a probability given by their score on one randomly chosen task, the belief elicitations or the comprehension questions.

Alignment index. Each climate fact is rated by senders who made their provision decisions on immigration, and each immigration fact by senders who decided on climate, according to whether the fact validates the rater's own stated position. A sender's own ratings never enter the index attached to their own decisions. For each fact and each of the five positions, twice the share of raters at that position who say the fact validates their position, minus one, gives the stance–fact alignment index $\hat{\alpha}_{ij} \in [-1, +1]$, and each sender's decisions are assigned the value for their own stated position. This treats the share at a sender's stated position as a measurement of that sender's alignment. The alignment index spans -0.429 to 0.830 and orders facts by their hypothesized direction: at the pro-policy pole, pro-policy facts score $+0.535$ against $+0.304$ for anti-policy facts, and at the anti-policy pole the ordering reverses, $+0.786$ against $+0.272$.

Expected accuracy gain. Senders expect receipt of a fact to improve factual knowledge, but only partly. They put the receiver's chance of answering correctly at 49.3% without the fact and 66.5% with it, an expected accuracy gain averaging 0.172. The elicited object is the share of receivers at the matched receiver's stance who answer correctly, which I use as the sender's forecast of a a single receiver's belief. Expected accuracy gains for facts that favor the sender's side and for facts that cut against it do not detectably differ ($p = 0.245$, Panel A of Appendix Figure C.2).

Expected persuasion. The average expected persuasion for receiver–fact pairs divide almost evenly between expected movement toward the sender (41.7%) and away from them (41.0%), and expected movement toward the sender grows with the distance between them (Appendix C.5). Whether a fact is expected to move the receiver toward the sender therefore depends on the receiver, not only on the fact. Senders expect movement toward their position for 45.6 percent of facts that favor their side and 37.9 percent of facts that cut against it, a gap below 8 pp in each treatment (7.7 in CONTROL, 7.6 in ACTION).[31]

Separability and balance. The returns and political distance are nearly uncorrelated. After fact fixed effects, no pair among the alignment index, the expected accuracy gain, the expected persuasion return and political distance correlates above 0.12 in absolute value—the largest is the persuasion return and distance, at 0.117—and no variance inflation factor exceeds 1.030, so the four are far from the exact affine dependence

[31] In Experiment 1, where every receiver held a false belief or did not know the answer and the facts were validated on which way learning them moves policy views, the gap is 60.8 pp. There, senders did not see the receiver's position, so expected persuasion is measured as movement toward the sender's side.

Assumption 2 rules out. The two elicited forecasts have similar means across treatments (sender-level mean-difference $p = 0.783$ and $p = 0.577$) and standard deviations within 0.015 of each other, as random assignment implies under Assumption 3(i).[32] Within fact, positional alignment correlates with the stance–fact alignment index at $r = 0.596$, but with the persuasion return only at $r = 0.084$ and with the expected accuracy gain at $r = 0.010$. These properties allow the returns and political distance to enter one specification.

4.2.3 Selective Provision and the ACTION Interaction

Senders provide unaligned facts well above the indifference benchmark: across the four treatment–topic cells, provision of unaligned facts ranges from 55.5% to 61.2% (one-sample t-tests, all $p < 0.001$). Overall provision is 63.9%.

Result 2.0 (Baseline provision)**.** *Senders share most unaligned facts, above the* 50 *percent benchmark in every treatment–topic cell, confirming Hypothesis 2.0.*

Table 5 reports the pre-registered tests. I estimate the linear probability model

$$y_{i,j} = \beta_0 + \beta_1 \,\text{Align}_{i,j} + \beta_2 \,\text{ACTION}_i + \beta_3 \left(\text{Align}_{i,j} \times \text{ACTION}_i\right) + \mathbf{S}_j + \gamma\, \mathbf{X}_i + \varepsilon_{i,j}, \quad (5)$$

where $\text{Align}_{i,j} = \text{Pref}_i \times \text{Dir}_j \in [-1, 1]$ is positional alignment, with Pref_i the sender's policy view index and $\text{Dir}_j = \pm 1$ the fact's pro- or anti-policy direction and $\text{ACTION}_i \in \{0, 1\}$ is an indicator equal to one if sender i is assigned to the ACTION treatment. The unit of observation is a sender–fact pair, and standard errors are clustered at the sender level throughout. β_1 captures alignment-based selectivity in CONTROL (Hypothesis 2.1) and β_3 whether the positional alignment slope is larger in the ACTION treatment (Hypothesis 2.2). Column (1) includes no controls, column (2) adds the fact fixed effects $\mathbf{S}_j$, and column (3) adds the demographic and political controls $\mathbf{X}_i$.

A fact is aligned for a sender if $\text{Align}_{i,j} > 0$ and unaligned if $\text{Align}_{i,j} < 0$, which defines both for the 1,496 senders with a directional policy view. Senders share aligned facts at higher rates in both treatments (Figure 4, Panel (A)): provision runs 68.2% on aligned facts against 56.6% on unaligned facts in CONTROL, and 71.3% against 58.5% in ACTION. The sender-level alignment bias averages 11.6 pp in CONTROL and 12.7 pp in ACTION, and the pre-registered coefficient is $+0.128$ per unit of positional alignment (one-sided $p < 0.001$, column (3)).

Result 2.1 (Selective provision)**.** *Senders favor aligned facts in both treatments, confirming Hypothesis 2.1.*

[32] Each sender is observed in one treatment only, so this is not a direct test of the assumption.

As in Section 2, the coefficient equals the difference in sharing between aligned and unaligned facts for a sender at half the maximum policy view intensity, because the two differ in positional alignment by twice that intensity. In CONTROL, the estimates imply that the median sender with a directional policy view, whose intensity is 0.52, is 13.3 pp more likely to share an aligned than an unaligned fact, and a sender at the maximum 25.6 pp.[33]

Hypothesis 2.2 is not confirmed. The Align × ACTION interaction is +0.014 (one-sided $p = 0.215$, column (3)).

Result 2.2 (ACTION increase not confirmed)**.** *The data do not show a larger alignment gap when the receiver takes a downstream political action.*

The pre-registered test asks whether the alignment gap widens under political action. Positional alignment depends only on the fact and the sender, whereas the returns that political action could make more valuable, correction and persuasion, also depend on the receiver. The test therefore bears on those returns only to the extent that they move with positional alignment.

4.2.4 The Four Drivers of Provision

Provision moves with all four determinants of Section 3, each in the direction the framework predicts. Within facts, it correlates positively with the stance–fact alignment index ($r = 0.094$, $p < 0.001$), the expected accuracy gain ($r = 0.026$, $p = 0.039$) and the expected persuasion return ($r = 0.026$, $p = 0.034$), and negatively with political distance ($r = -0.218$, $p < 0.001$). Provision varies with what a fact does for the sender and for the receiver, and most strongly with who the receiver is.

The specification. To estimate these responses jointly and on the scale of the framework, I take the shape G in Assumption 4 to be the standard normal, so that the disclosure probability is a probit in the four determinants:

$$\begin{aligned}\Pr(y_{ij} = 1) = \Phi\big(&\beta_0 + \beta_E\, \alpha_{ij} + \beta_B\, \Delta b_{ijk} + \beta_P\, \Delta\pi_{ijk} + \beta_\rho\, \rho_{ijk} + \gamma_A\, \text{ACTION}_i \\ &+ \delta_B\, (\Delta b_{ijk} \times \text{ACTION}_i) + \delta_P\, (\Delta\pi_{ijk} \times \text{ACTION}_i) + \mathbf{S}_j\big),\end{aligned} \tag{6}$$

where $\rho_{ijk} = |\text{Stance}_i - \text{Stance}_k|$ is the political distance between sender i and the receiver k of fact j, and $\mathbf{S}_j$ are fact fixed effects, which absorb each fact's direction, topic and other fact-level appeal. I refer to this as the four-determinant specification. By

[33] The construction $\text{Pref}_i \times \text{Dir}_j$ imposes this proportionality to intensity. With a binary aligned indicator in its place, an aligned fact is 12.2 pp more likely to be shared than an unaligned one in CONTROL ($p < 0.001$). Logistic and sender fixed effects specifications give consistent results (Appendix Table C.5). Leave-one-fact and sample-variant ranges are in Appendix C.3.

Table 5: Alignment-Based Selective Provision

	Probability of Fact Provision		
	(1)	(2)	(3)
Positional alignment	0.113***	0.127***	0.128***
	(0.013)	(0.013)	(0.013)
ACTION	0.026**	0.026**	0.023*
	(0.013)	(0.013)	(0.012)
Positional alignment × ACTION	0.011	0.014	0.014
	(0.018)	(0.018)	(0.018)
Clustered SE	Sender	Sender	Sender
Fact FE	No	Yes	Yes
Controls	No	No	Yes
Num. obs.	7740	7740	7715
R^2	0.025	0.032	0.045
Adj. R^2	0.025	0.030	0.039

Notes: The sample is the $N = 1{,}548$ Experiment 2 general-population senders (Bilendi panel, Germany, 2026), 7,740 sender-fact observations. Each sender makes five provision decisions, one per fact. Column (3) restricts to non-missing demographic controls (7,715 observations). The outcome is a binary provision indicator ($y_{ij} \in \{0, 1\}$). Positional alignment is $\text{Align}_{ij} = \text{Pref}_i \times \text{Dir}_j \in [-1, 1]$. Fact FE is an indicator for each of the 12 facts. Controls are age, gender, education, Bundesland, voting intention, left-right self-placement. Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Corollary 1, each slope is a marginal utility divided by the scale, so its sign is the sign of the corresponding weight. The interactions δ_B and δ_P measure how the correction and persuasion slopes change under political action. With a common scale across treatments, a positive δ_P is the observable counterpart of consequential persuasion, $\theta_P(1) > \theta_P(0)$ (Proposition 2). Columns (1)–(2) pool the treatments, and column (3) adds the interactions, so its main effects are CONTROL-treatment slopes. Standard errors are clustered by sender.[34]

What senders select on. Provision rises with what a fact does for the sender and for the receiver. Pooling the treatments (column (2)), moving the alignment index from its fifth to its ninety-fifth percentile raises predicted provision by 17.1 pp. The same move raises it by 6.6 pp for the expected persuasion return ($p < 0.001$) and by 3.6 pp for the expected accuracy gain (one-sided $p = 0.039$). Because the receiver learns the sender's position whichever fact arrives, the alignment-index association does not re-

[34]The interactions of ACTION with the alignment index and with distance are set to zero in line with Assumption 3(ii)–(iii) together with a common scale. Loosening this restrictions for the alignment-index interaction is reported with the robustness checks. The treatment indicator shifts the propensity to disclose anything and carries no information about the weights with a common scale, $\gamma_A = (\kappa - \mu_1 + \mu_0)/s$. Reading signs as weights treats $\hat{\alpha}$ and the elicited forecasts as measures of the sender's alignment and forecasts. One-sided p-values are reported where the framework signs a case.

Table 6: Determinants of Provision: Index Coefficients and Average Marginal Effects

	Provision of the Fact		
	(1)	(2)	(3)
Panel A: index coefficients			
Stance–fact alignment index	0.489^{***}	0.490^{***}	0.490^{***}
	(0.053)	(0.053)	(0.053)
Expected accuracy gain	0.086^{*}	0.085^{*}	0.131^{*}
	(0.048)	(0.048)	(0.069)
Persuasion	0.061^{***}	0.061^{***}	0.030
	(0.018)	(0.018)	(0.024)
Political distance	-0.250^{***}	-0.250^{***}	-0.250^{***}
	(0.013)	(0.013)	(0.013)
ACTION		0.070^{*}	0.081^{**}
		(0.036)	(0.040)
Expected accuracy gain × ACTION			−0.086
			(0.096)
Persuasion × ACTION			0.061^{*}
			(0.035)
Fact fixed effects	Yes	Yes	Yes
Senders	1,548	1,548	1,548
Observations	7,740	7,740	7,740
Panel B: average marginal effects by treatment, from column (3)			
	CONTROL	ACTION	Difference
Stance–fact alignment index	0.175	0.171	−0.004
Expected accuracy gain	0.047	0.016	−0.031
Persuasion	0.011	0.032	0.021
Political distance	−0.089	−0.087	0.002

Notes: The sample is the $N = 1{,}548$ Experiment 2 general-population senders (Bilendi panel, Germany, 2026), with 7,740 sender–fact observations in every column. The outcome is an indicator equal to one if the sender provides the fact. All columns are probit regressions with fact fixed effects and enter the stance–fact alignment index, the expected accuracy gain, the expected persuasion return and the political distance between sender and receiver. Column (2) adds the ACTION indicator, and column (3) interacts the expected accuracy gain and the expected persuasion return with ACTION, so its main effects are CONTROL-treatment slopes. The interactions rest on the randomization of ACTION. Panel B reports, for each treatment, the average marginal effect of each determinant from column (3). Inference is read from Panel A. Standard errors, clustered by sender, are in parentheses. Stars are two-sided. The text reports one-sided p-values where the framework signs a case. After Holm adjustment for the two pre-registered interactions with the elicited returns, the persuasion interaction has one-sided $p = 0.078$. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

flect signaling which side the sender is on. Canonical models of voluntary disclosure give the sender an objective that runs entirely through the receiver's beliefs, the special case $\theta_E = \theta_R = 0$ of Equation (2). Under Assumption 1, the estimates do not fit that case.

Whom senders inform. The largest difference in provisions corresponds to the sender–receiver distance. Holding the elicited returns fixed, moving the receiver from the

sender's own position to four points away lowers predicted provision by 36.4 pp, a distance only senders at the ends of the scale face. Each sender faces a receiver at every position, and facts are assigned to receivers at random, so within senders the decline of provision with distance is causal. The slant of what arrives also changes with distance. Because a provided fact arrives with probability one half, the slant of arrival equals the slant of provision in expectation. The slant of provision is 0.070 for receivers at the sender's own position and 0.156 for receivers four points away, rising by 0.027 per point (sender-clustered $p = 0.003$). Distant receivers get fewer facts, and more slanted ones. Among receivers of the same sender, the slant rises by only 0.015 per point ($p = 0.146$).[35]

Result 2.3 (What senders select on)**.** *Pooling the treatments, provision rises with the alignment index, the expected persuasion return and, more weakly, the expected accuracy gain, and falls with political distance.*

Political action. The evidence suggests that political action shifts selection toward persuasion and toward distant receivers, while the slant of provision does not detectably change. Provision is 65.2% under ACTION and 62.6% in CONTROL ($p = 0.038$). The correlation of provision with the expected persuasion return rises from 0.007 to 0.051 ($p = 0.064$), while those with the alignment index, the expected accuracy gain and positional alignment do not differ significantly (all $p > 0.3$). The slant of provision is 0.071 in both treatments ($p = 0.999$). Provision falls less steeply with political distance under ACTION ($p = 0.043$).

On the index, the two comparisons of Section 3.4 point the same way. With a common scale, the persuasion slope is higher under ACTION, $\delta_P = +0.061$ (one-sided $p = 0.039$), raising the probability response to the expected persuasion return from $+1.1$ to $+3.2$ pp per unit.[36] Appendix Table C.9 reports the ratios. With the expressive weight held fixed, the persuasion slope relative to the alignment-index slope, estimated within treatment, is higher under ACTION by 0.095. The correction interaction is not distinguishable from zero, $\delta_B = -0.086$ ($p = 0.371$). On a common scale, per standard deviation of each return, the persuasion interaction exceeds the correction interaction by 0.084 ($p = 0.072$). In the linear probability model, adding sender fixed effects lowers the persuasion interaction from $+0.023$ to $+0.013$ (one-sided $p = 0.170$). Part of the estimate may therefore rest on differences between senders.

[35] Reading the decline as the relational motive, a preference over whom one informs, requires Assumption 1: the design elicits forecasts of only the receiver's belief and position, so a sender who anticipates other responses from distant receivers, such as hostility, judgment of the sender or futility, could produce the same decline.

[36] Inference on this interaction is marginal. Appendix C.3 reports it under alternative clustering, specifications and sample definitions.

Result 2.4 (Motives and consequentiality)**.** *Under political action, the evidence suggests that senders select more on the expected persuasion return, but not detectably more on the expected accuracy gain, and, on the probability scale, provide relatively more to distant receivers. The slant of provision does not detectably differ between the treatments.*

This suggests the Align $\times$ ACTION interaction was unlikely to detect the change. Within facts, the expected persuasion return rises by only 0.129 per unit of positional alignment ($r = 0.084$), so the linear-probability persuasion interaction ($+0.023$) implies a very small Align $\times$ ACTION coefficient.[37] The pooled test also averages over the two topics: for immigration, the topic of Experiment 1, the interaction is $+0.050$ (one-sided $p = 0.011$), and for climate it is -0.042 ($p = 0.147$).

4.2.5 Heterogeneity and Alternative Explanations

Heterogeneity. Appendix C.4 reports the pre-registered heterogeneity analyses by the political distance between sender and receiver, political polarization, issue extremity and voting for a party of the political center. None of them detectably changes how strongly senders favor aligned facts. The one difference that survives correction for multiple testing is that the response to the stance–fact alignment index rises with how extreme the sender's position is (Appendix Table C.15, Benjamini–Hochberg $q = 0.008$).

Disbelief. As in Experiment 1, disbelief does not explain the selection on positional alignment. Holding fixed how credible other senders at the same stance find each fact, the positional alignment coefficient falls only from 0.365 to 0.331 ($p < 0.001$), as Appendix Table C.8 shows. With the credibility rating as an additional regressor, the stance–fact alignment index still predicts provision (0.253, $p < 0.001$).

5 Experiment 3: Selectivity Among Journalists

Experiment 3 asks whether journalists select on alignment absent the market incentives often used to explain observed slant. Journalists' provision rises with how strongly a fact validates their position and with the correction and the persuasion they expect it to achieve, and falls with the receiver's distance from their own position. Journalists provide more than the general population.

[37] The implied coefficient is $0.023 \times 0.129 = 0.003$, six percent of that test's minimum detectable effect of 0.049.

5.1 Design and Hypotheses

Protocol. Experiment 3 runs the ACTION treatment of Experiment 2 with professional journalists, using its facts, provision task and incentivized forecasts (Figure 3). There is no attention check. Every journalist is matched to a real receiver. The post-questionnaire here elicits outlet type, reach, role and editorial autonomy. Journalists rate no facts from the other domain, so the general-population stance–fact alignment index is assigned.

Hypotheses. I pre-registered the two hypotheses (AsPredicted No. 292902). The first mirrors Hypothesis 2.0: baseline provision above 50%. The second mirrors Hypothesis 2.1: alignment-based selective provision.

5.2 Results

5.2.1 Sample, Incentives, and Protocol

Journalists were recruited through the mailing lists of the Deutscher Journalisten-Verband (German Federation of Journalists) and the Deutsche Journalistinnen- und Journalisten-Union (German Union of Journalists). 309 journalists participated between May 25 and August 11, 2026.[38] The study was approved by the Ethics Council of the Max Planck Society under the same Generalized Approval. Journalists took a median of 22 minutes and earned €15.39 on average.[39] The sample is younger and more educated, has a higher share of females, and online outlets are overrepresented relative to the general German journalist population (Loosen et al. 2023).

5.2.2 Policy Views and Elicited Returns

Policy views. Journalists' policy views are more stronger than that of the general population, with mean of 0.689 compared with 0.541. In particular, the sample leans left.[40] On immigration, 58.6% of journalists favor more permissive policy, while 66.1% of the general population favors more restrictive policy. On climate, 83.7% favor expanded policy, compared with 52.8%.

Elicited returns. Journalists expect facts to correct receivers more than the general population does. They expect 44.7% of receivers without the fact to answer correctly

[38] One journalist completed the study seven times, identified by the payout details they submitted, and another twice, identified by the name they submitted. Only each first participation is retained.

[39] Journalists receive €15 for participating. Ten journalists drawn at random can earn a bonus of €30, determined as in Experiment 2.

[40] Mean left–right (0–10) self-placement is 3.2 against general population's 5.0. This is consistent with earlier survey evidence (Steindl et al. 2017).

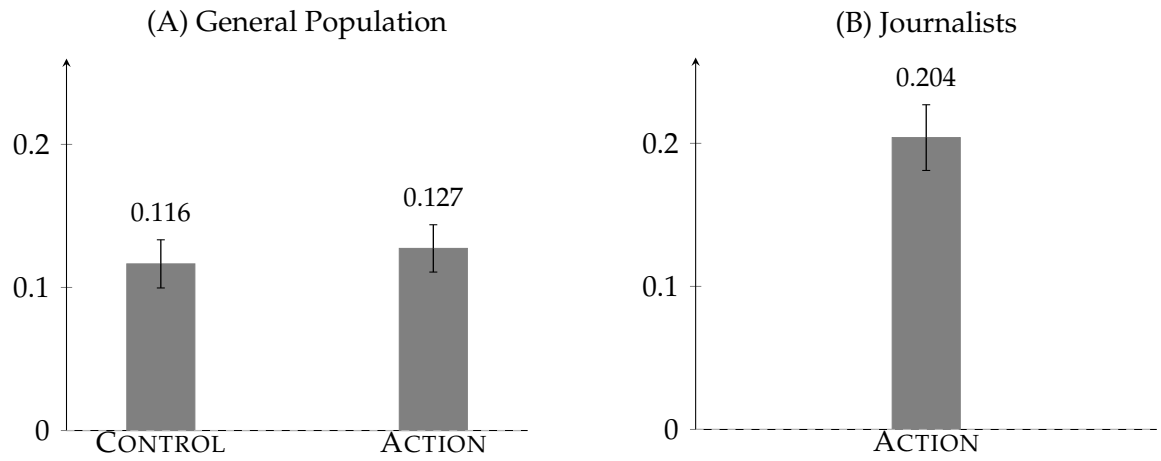


Figure 4: Alignment Bias in the General Population and Among Journalists

Notes: The sample is Experiments 2 and 3. The bias is the within-sender difference between provision rates for aligned and unaligned facts, averaged across senders. Panel (A) shows the general population ($N = 1{,}496$ senders; Bilendi panel, Germany, 2026). Panel (B) shows professional journalists ($N = 299$), who face the ACTION treatment only. The sample is senders with a directional policy view ($\text{Pref}_i \neq 0$), so that positional alignment is defined; 52 general-population and 10 journalist senders are excluded. Positional alignment is $\text{Align}_{ij} = \text{Pref}_i \times \text{Dir}_j \in [-1, +1]$. Error bars are standard errors of the sender-level mean.

and 76.9% of receiver who receive the fact to answer correctly, an expected accuracy gain of 0.322, against 0.172. For journalists, the facts that favor their side are also more often the ones they expect to move the receiver toward them: they expect such movement for 51.2% of facts that favor their side and 24.0% of facts that cut against it. The gap of 27.2 pp lies above the general population's 7.6 pp in the ACTION treatment. They also believe that distant receivers are more persuadable (0.100 against 0.015 per scale point).

Separability. The four determinants remain weakly correlated among journalists. After fact fixed effects, the largest correlation is again between the persuasion return and political distance, at 0.230, and no variance inflation factor exceeds 1.068. Within facts, positional alignment correlates with the persuasion return at $r = 0.192$, against 0.083 in the general population's ACTION treatment.

5.2.3 Selective Provision

Journalists share most unaligned facts: 64.4% on climate policy and 67.9% on immigration policy, above the 50 percent benchmark (one-sample t-tests, both $p < 0.001$). Overall provision is 12.5 pp higher than in the general population.

Result 3.0 (Baseline provision)**.** *Journalists share most unaligned facts.*

Table 7, columns (1)–(3), reports the pre-registered tests, Equation (5) without the ACTION terms. Journalists provide 86.6% of aligned facts against 66.2% of unaligned ones, a sender-level alignment bias of 20.4 pp ($p < 0.001$), against 12.7 pp in the general population's ACTION treatment.

Table 7: Selective Provision: Journalists and the Pooled Comparison with the General Population

	Journalists			Pooled: general population and journalists		
	(1)	(2)	(3)	(4)	(5)	(6)
Positional alignment	0.142***	0.126***	0.129***	0.126***	0.123***	0.124***
	(0.015)	(0.020)	(0.020)	(0.013)	(0.013)	(0.013)
ACTION				0.026**	0.023*	0.023*
				(0.013)	(0.012)	(0.012)
Positional alignment × ACTION				0.005	0.008	0.017
				(0.016)	(0.016)	(0.017)
Journalist				0.114***	0.097***	0.096***
				(0.016)	(0.019)	(0.019)
Positional alignment × Journalist						−0.028
						(0.020)
Clustered SE	Journalist	Journalist	Journalist	Sender	Sender	Sender
Fact FE	No	Yes	Yes	Yes	Yes	Yes
Controls	No	No	Yes	No	Yes	Yes
Sample	Journalists	Journalists	Journalists	Pooled	Pooled	Pooled
Num. obs.	1545	1545	1540	9285	9255	9255
R^2	0.062	0.077	0.137	0.045	0.099	0.099
Adj. R^2	0.062	0.069	0.098	0.044	0.093	0.093

Notes: The sample in cols. (1)–(3) is the $N = 309$ Experiment 3 journalists, 1,545 journalist-fact observations, all in the ACTION treatment. Col. (3) has 1,540 observations from 308 journalists. The sample in cols. (4)–(6) pools Experiment 2 general-population senders ($N = 1{,}548$; both treatments; 7,740 observations) and Experiment 3 journalists ($N = 309$; ACTION treatment only; 1,545 observations); col. (4) $N = 9{,}285$, cols. (5)–(6) $N = 9{,}255$ after covariate availability. The outcome is a binary provision indicator ($y_{ij} \in \{0, 1\}$), linear probability model, sender–fact observations; no sampling weights. Positional alignment is $\text{Align}_{ij} = \text{Pref}_i \times \text{Dir}_j \in [-1, 1]$. ACTION is an indicator for the ACTION treatment. Controls in col. (3) are the pre-registered demographic, political and journalist-background controls: age, gender, education, Bundesland, left-right self-placement, voting intention, medium type, reach, main role and editorial autonomy. Controls in cols. (5)–(6) are sender-receiver ideological distance, age, gender, education, Bundesland, left-right self-placement, and closest-party affiliation—the demographic variables measured in both samples. All specifications except col. (1) include fact fixed effects. Standard errors clustered at journalist or sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Result 3.1 (Selective provision among journalists)**.** *Journalists favor aligned facts.*

The journalist bias is larger, while the slope does not detectably differ: in the pooled specification the journalist difference is -0.028 ($p = 0.168$; column (6)). The bias is larger partly because journalists hold stronger policy views.

5.2.4 The Four Determinants of Provision

Table 8 estimates Equation (6) without the ACTION terms, on journalists in column (1) and on the general population's ACTION treatment in column (2).

What journalists select on. Provision responds to the fact and the receiver. Moving the alignment index from its fifth to its ninety-fifth percentile raises predicted provision by 12.7 pp ($p = 0.002$). The same move raises it by 11.4 pp for the expected persuasion return (one-sided $p = 0.003$) and by 10.0 pp for the expected accuracy gain ($p = 0.027$).[41] Selection on the alignment index thus persists with competition, ownership and audience demand removed.

Whom journalists inform. The largest contrast again concerns who the receiver is. Holding the elicited returns fixed, moving the receiver from the journalist's own position to four points away lowers predicted provision by 25.6 pp ($p < 0.001$), a distance

[41] Percentiles are those of the journalists and the general population's ACTION treatment combined.

Table 8: Determinants of Provision: Journalists and the General Population (ACTION Treatment)

	Provision of the Fact		
	Journalists	General population (ACTION treatment)	Difference
	(1)	(2)	(3)
Panel A: index coefficients			
Stance–fact alignment index	0.466***	0.517***	
	(0.151)	(0.074)	
Expected accuracy gain	0.299**	0.049	
	(0.135)	(0.067)	
Persuasion	0.137***	0.088***	
	(0.050)	(0.025)	
Political distance	−0.218***	−0.228***	
	(0.032)	(0.019)	
Fact fixed effects	Yes	Yes	
Senders	309	793	
Observations	1,545	3,965	
Mean provision	0.764	0.651	
Panel B: average marginal effects			
Stance–fact alignment index	0.130	0.182	−0.051
			(0.046)
Expected accuracy gain	0.084	0.017	0.066
			(0.044)
Persuasion	0.038	0.031	0.007
			(0.016)
Political distance	−0.061	−0.080	0.019*
			(0.011)

Notes: The samples are Experiment 3 journalists in column (1) and Experiment 2 general-population senders in the ACTION treatment in column (2). The estimator is a probit with fact fixed effects. The outcome is a binary provision indicator. The regressors are the stance–fact alignment index $\hat{\alpha}_{ij}$, assigned to journalists from general-population raters at the same stance. The incentivized expected accuracy gain $\widehat{\Delta b}_{ijk}$ and expected persuasion return $\widehat{\Delta \pi}_{ijk}$ and political distance ρ_{ijk}. Panel A reports index coefficients. Panel B reports average marginal effects. Column (3) is their difference, with a sender-cluster bootstrap standard error (1,000 replications, senders resampled within each sample). $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

only journalists at the ends of the scale face. Each journalist faces a receiver at every position, and facts are assigned to receivers at random, so within journalists the decline of provision with distance is causal.

Result 3.2 (What journalists select on)**.** *Journalists' provision rises with the alignment index, the expected persuasion return and the expected accuracy gain, and falls with political distance.*

Comparison with the general population. Journalists provide 11.4 pp more than the general population's ACTION treatment, and 9.7 pp more with political distance and demographic and political characteristics held constant (both $p < 0.001$; Table 7,

columns (4)–(6)). Their responses to the determinants do not detectably differ. Importantly, the two probit regressions have different latent scale. Therefore, Panel B of Table 8 compares the two in probability units: per unit of each determinant, the average marginal effects differ by -0.051 for the alignment index ($p = 0.262$), $+0.066$ for the expected accuracy gain ($p = 0.134$) and $+0.007$ for the expected persuasion return ($p = 0.647$), and provision falls somewhat less steeply with distance among journalists ($+0.019$, $p = 0.073$). The intervals are wide, so moderate differences cannot be ruled out. Since a provided fact is more likely to reach a journalist's receiver, the comparison of the responses to correction and persuasion may also reflect this difference in matching probability.

Result 3.3 (A level difference; no detected slope difference)**.** *Journalists provide more than the general population, and their responses to the determinants do not detectably differ.*

5.2.5 Heterogeneity

In the pre-registered heterogeneity analysis, by political distance, polarization, issue extremity and journalist background, one interaction survives correction for multiple testing: the positional alignment slope rises with the receiver's distance. So journalists favor aligned facts more strongly when the receiver is politically distant. Appendix Table D.2 shows that medium type, broadcaster type, reach, role and editorial autonomy do not detectably moderate alignment-based selectivity.

6 Experiment 4: The Epistemic Cost of Withholding

Experiment 4 asks how selective provision affects receivers. Each receiver is matched to a sender from Experiment 2 or 3. A receiver whose sender withheld a fact ends with nearly twice the squared belief error of one shown it. Receivers underestimate how often senders disclose and perceive little of how disclosure depends on sender–fact alignment.

6.1 Design and Outcomes

Experimental timeline. Receivers first report demographics and their positions on climate and immigration policy, and then their prior belief about a single fact, before senders are mentioned (Figure 5). They then receive the instructions and the comprehension checks, and either see the fact or learn that they will not receive it. Receivers next report their policy position again and make a donation if matched to an ACTION-treatment sender. They then report their posterior belief on the same scale as the prior,

and finally what they believe about how senders in their sender-fact-receiver pairing disclosed.

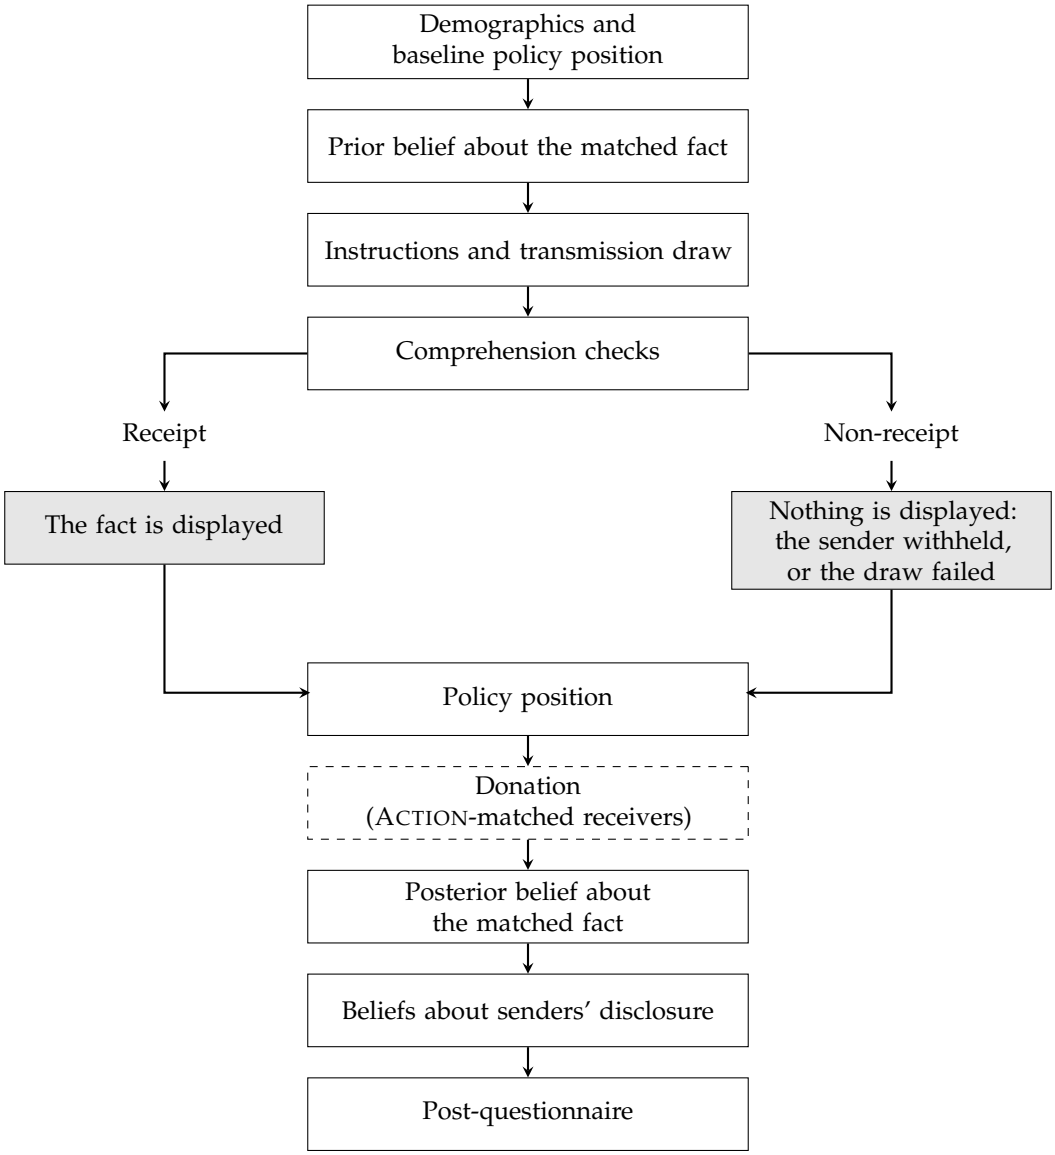


Figure 5: Timeline of Experiment 4

Reception. Each receiver is matched to a sender from Experiment 2 or 3 who decided whether to share the fact with a receiver at the receiver's policy stance. A shared fact reaches the receiver with probability one half. Each receiver therefore ends in one of three states: the sender shared the fact and it arrived (receipt, R), the sender shared it and it did not arrive (failed transmission, F), or the sender withheld it (withholding, W). The two causes of non-receipt look the same: a receiver who sees no fact cannot tell whether the sender withheld it or the random draw stopped it. Receivers know that both are possible. They are told the sender's stance, that the sender knew theirs, and that a coin decides whether the sender's decision is implemented.

Belief accuracy. Receivers report their belief about the fact twice, before senders are mentioned and after reception, as the probability they attach to each of two completions of the fact's sentence, one of which is correct. The second elicitation uses the same slider, with the earlier answer neither shown nor used to position it, and both are incentivized by a quadratic scoring rule.[42] I measure belief accuracy against ground truth by the squared loss $L = (1 - b)^2$, where b is the probability placed on the correct answer.

Policy views and donations. Receivers report their policy views on a five-point scale, the item senders answer, at the start for both topics and again after reception for the matched topic. Receivers matched to ACTION-treatment senders also split €250 between the two organizations of that topic. One decision across senders and receivers is implemented.

Beliefs about senders. Receivers estimate how many of up to ten real senders passed a fact on, once supposing for each potential correct completion. These rates are the perceived disclosure rates for the (in)correct completions. Their difference $\hat{\sigma}_1 - \hat{\sigma}_0$ is perceived selectivity: how far receivers believe disclosure depends on sender–fact alignment. Estimates are scored against real sender decision by a quadratic rule. Receivers do not know which completion is correct, so both estimates can be read as incentivized.[43]

6.2 Results

6.2.1 Sample, Incentives, and Protocol

Data collection ran from August 17 to 19, 2026 on Bilendi, with quotas that match the German adult population on age, gender, and state (Appendix Table E.1). This study was approved by the Ethics Council of the Max Planck Society under the Generalized Approval of Experiments Following the Protocol that is Standard in Experimental Economics (approval no. 2018_3 – Renewal 2024_28) and was pre-registered on AsPredicted (No. 306857). 1,081 receivers completed the study and I exclude 38 speeders leaving $N = 1{,}043$ receivers for the main analysis: 358 in receipt, 368 in failed transmission and 317 in withholding. Receivers took a median of 8.1 minutes and earned €2.21 on average.[44] The experiment was programmed in oTree (Chen et al. 2016).

[42] The slider layout is randomized across participants (Appendix E).

[43] Receivers shown the fact are asked about a different fact of the same topic at the same stances. The actual disclosure rate for the correct completion is the share of senders in the sender-fact-receiver pairing who passed the fact on (median of 36 decisions).

[44] Receivers are paid €2.15 for participating. Ten receivers drawn at random can earn a bonus of €10, paid with a probability given by their score on one randomly chosen task: the prior belief, the

6.2.2 The Cost of Withholding

Table 9 reports results of the primary specification,

$$L_r^{\text{post}} = \beta_0 + \beta_W W_r + \beta_F F_r + \gamma L_r^{\text{prior}} + \mathbf{S}_{j(r)} + \varepsilon_r, \tag{7}$$

where receipt is the omitted state, $\mathbf{S}_{j(r)}$ are fact fixed effects, and standard errors are HC3-robust. Mean posterior loss is 0.307 under withholding against 0.166 after receipt (Figure 6). Holding prior loss and the fact fixed, a receiver whose sender withheld the fact is left with $\beta_W = 0.128$ more posterior squared loss than a receiver shown it ($p < 0.001$, column (1)).

Table 9: Reception State, Belief Accuracy, and Downstream Outcomes

	Belief accuracy		Downstream outcomes	
	Squared loss (1)	Absolute error (2)	Policy view (3)	Donation (4)
Withholding (W)	0.128***	0.158***	-0.065	-5.55
	(0.018)	(0.020)	(0.053)	(9.14)
Failed transmission (F)	0.126***	0.152***	-0.058	-9.40
	(0.018)	(0.019)	(0.045)	(8.13)
$\beta_W - \beta_F$	0.001	0.006	-0.007	3.85
	(0.019)	(0.019)	(0.052)	(8.68)
Baseline outcome	0.478***	0.454***	0.852***	
	(0.035)	(0.034)	(0.020)	
Fact fixed effects	Yes	Yes	Yes	Yes
Demographic controls	No	No	Yes	Yes
Observations	1043	1043	1043	664

Notes: The sample is the N = 1,043 Experiment 4 receivers (317 withholding, 368 failed transmission, 358 receipt); column (4) restricts to the 664 receivers matched to ACTION-treatment senders. The estimates are OLS estimates of equation (7). Receipt is the omitted state. The outcomes are (1) posterior squared loss (2) posterior absolute error (3) policy view on the matched topic, signed by the fact's direction, controlling for the initial view (4) signed split of €250. Columns (3)–(4) add demographic controls. HC3 standard errors in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Result 4.0 (Less accurate beliefs under withholding)**.** *Receivers whose sender withheld a fact hold less accurate beliefs than receivers shown it: their posterior squared loss against ground truth is nearly twice as large.*

Non-receipt is costly. After failed transmission, receivers are left with $\beta_F = 0.126$ more posterior squared loss than receivers shown it ($p < 0.001$). This cause of non-receipt does not detectably differ from withholding in belief accuracy.[45]

comprehension questions, the posterior belief or the estimates about senders.

[45]Trimming for differential completion leaves this unchanged (Appendix Table E.2).

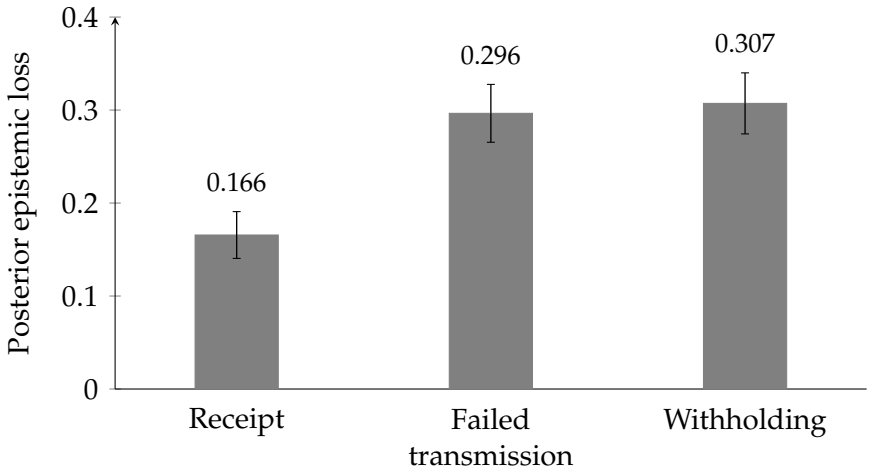


Figure 6: Posterior Epistemic Loss by Reception State

Notes: The sample is the canonical Experiment 4 sample, $N = 1{,}043$: 358 in receipt, 368 in failed transmission, 317 in withholding.

Withheld facts score 0.360 lower in positional alignment than facts whose transmission failed ($p < 0.001$), so a receiver who knew how senders disclose could learn about the truth from non-receipt. However, given the difference in rates at which senders actually disclose, inference is limited: Bayes' rule would lower mean loss after non-receipt by 0.002 relative to the prior, less than two percent of the cost of withholding.[46] Selectivity thus raises belief error through the withholding it causes, if receipt would also have helped the receivers of withheld facts.

Although receiving the fact fully reveals the correct answer, receivers shown it close about a third of the distance between their prior and certainty (32.9 percent, against 0.7 percent after failed transmission). Policy views and donations do not detectably differ between receipt and non-receipt (Table 9, columns (3)–(4)).

Senders anticipated the effect on knowledge but overestimated the effect on policy views. General-population senders expected receipt to raise the share of receivers who answer correctly by 15.8 pp, and it was 21.7 pp higher after receipt than after non-receipt. They expected receipt to move policy views 0.128 points toward the fact's side on the five-point scale, while policy views moved no differently after receipt than after non-receipt (0.001).[47]

6.2.3 Perceived Selectivity and Responses to Non-Receipt

Receivers underestimate how often senders disclose. On average, they believe 57.1 percent of the senders in their sender-fact-receiver combination passed the fact on, against an actual 67.7 percent ($p < 0.001$). Their estimates barely respond to how the actual rate depends on sender-fact-receiver combination: a ten-point higher actual rate raises

[46] Appendix I.2 derives the benchmark.

[47] Journalists expected 29.4 pp, against 28.9, and 0.232 points, against 0.235.

the perceived rate by less than one point.

Receivers also perceive little of the selection. Perceived selectivity, the difference between the disclosure rates receivers attribute to the two answers, barely moves with actual selectivity in their sender-fact-receiver pairing. The slope is 0.066. Correcting for sampling error in the sender-fact-receiver pairing rates raises it only to 0.085.[48] Figure 7 illustrates why. A quarter of receivers expect no alignment-based selection by senders at all, and those who expect any selection are almost as likely to get its direction wrong (35.2 percent of receivers) as right (39.7 percent).

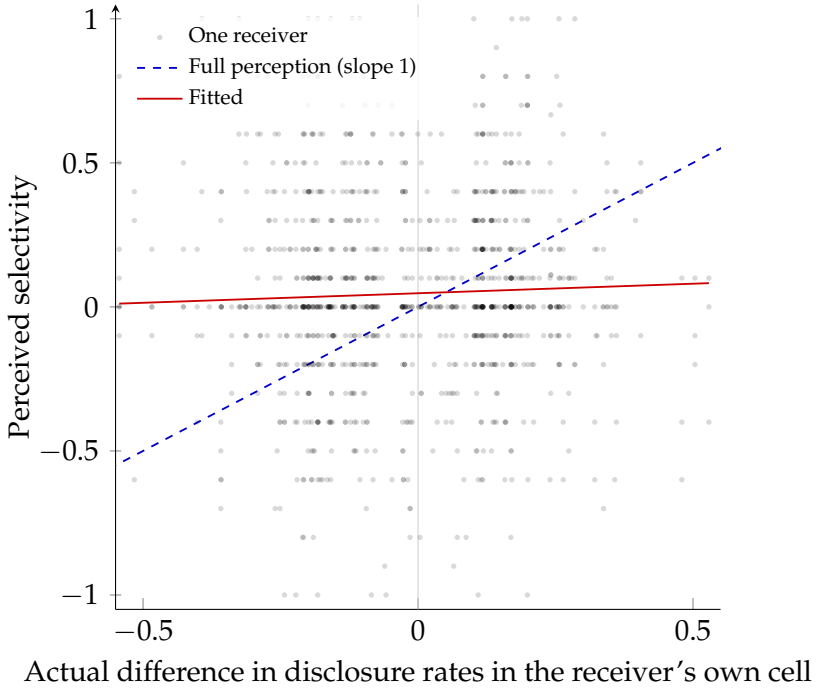


Figure 7: Perceived and Actual Selectivity in Sender Disclosure

Notes: The sample is the $N = 1{,}043$ Experiment 4 receivers. One point per receiver. Each receiver estimates how many senders in their own cell of sender stance, receiver stance, topic and fact passed the fact on under each of the two possible correct answers. The difference between those two estimates is their perceived selectivity. The actual difference on the horizontal axis is the corresponding difference among the senders who faced that cell, with the counterfactual state constructed from the opposite-direction facts of the same topic at the same stances. The dashed line is the 45-degree benchmark a receiver who perceived selectivity in full would trace, and the solid line is the fitted relation.

Given what they perceive, receivers should move little after non-receipt: at their own stated rates, Bayes' rule prescribes an average change of 2.3 points on the 0–100 scale. Their belief changes do not track even this. Regressed on the prescribed change, belief change has a slope of -0.180 ($p = 0.43$ against zero). Thus, I reject updating consistent with perceived selectivity, which implies a slope of one ($p < 0.001$).[49]

Result 4.1 (Receivers perceive little of the selection)**.** *Receivers underestimate how often senders disclose and perceive little of the selection. Their belief changes after non-receipt do not*

[48] Appendix E describes the correction.

[49] Notably, rates are elicited after the posterior belief, and the benchmark treats them as the beliefs held when updating. The pre-registered log-odds version implies the same, and none of these results depends on whether the fact supports the receiver's view (Appendix Table E.3).

follow what their own perception of sender selectivity implies.

7 Conclusion

This paper asks why people supply information slanted toward their own positions. I show that people deliberately curate the truth, choosing which facts to pass on and to whom. Senders, including professional journalists deciding without any market pressure, pass on facts that support their own position more often than facts that do not, even when exposure, doubt and signaling are largely ruled out. They favor these facts beyond anything they expect them to do to the receiver. They also pass on facts they expect to correct others or win them over, and they share less with those who disagree with them. When the receiver makes a real political donation, senders select more on what they expect to persuade. Those from whom a fact is withheld end up with less accurate beliefs and barely notice the selection. Standard models of strategic communication, in which senders care only about how the receiver responds, explain neither why senders favor facts that support their own side nor why they share less with those who disagree. More broadly, political information can be distorted without any false statement, through the choice of which facts to pass on and to whom, a margin current policies leave unaddressed.

These results may help explain why political disagreement extends to facts. People start from different factual beliefs predicted by their policy views, pass on the facts that favor their own side, and pass fewer facts across political lines. People also barely notice which information they miss, so inaccurate beliefs can persist among the like-minded without anyone hearing anything false. Demand reinforces this supply, as people prefer to discuss politics with those who agree and find it harder to learn from those who do not (Braghieri et al. 2024). A fact that receivers read moves their factual beliefs but not their political views. Whether this holds when people face the same selection repeatedly is an open question. Like falsehoods, selected facts may do their damage gradually.

The same choices may shape the institutions that supply political information. In newsrooms, slant may originate in journalists' own choices before any market pressure applies, in line with evidence that journalists' ideology travels with them across outlets (Boxell and Conway 2022) and that most variation in article-level slant lies within outlets (Braghieri et al. 2025).

The typical response to distorted political information checks and labels what was said, which does not address slant. Fact-checkers themselves select which statements to examine, and examine those on their own side less often (Louis-Sidois 2025). Appeals to accuracy reach only one of the reasons people share selectively. Senders also favor their own side and hold back from those who disagree, and when the stakes are

real, as before elections, they care more about persuading than about correcting others. Remedies that act on media markets miss slant that begins in individual choices.

No one observes the full set of facts a sender could have passed on, and every supplier must select, so selection itself is hard to address. What can be required is balance. The Fairness Doctrine, which bound U.S. broadcasters until 1987 (Federal Communications Commission 1987), acted on which facts are supplied, whatever the motive, though not on whom senders inform. Whether balance alone is a desirable standard for an undistorted information supply is, however, questionable.

These results concern one part of the information supply chain: which of a given set of facts people pass on, and to whom. The pattern holds across two countries, two policy domains, three samples and two disclosure regimes, and for journalists it captures individual choices rather than what outlets publish. Upstream, people find, verify and produce facts. Downstream, news organizations, recommendation algorithms and artificial intelligence increasingly decide which facts reach whom, and may amplify or counteract these individual tendencies. How these forces interact to shape what people believe is a question for future work.

References

Abeler, J., Nosenzo, D., and Raymond, C. (2019). Preferences for truth-telling. *Econometrica*, 87(4):1115–1153.

Abramitzky, R., Boustan, L., Jácome, E., Pérez, S., and Torres, J. D. (2023). Law-abiding immigrants: The incarceration gap between immigrants and the US-born, 1870–2020. Working Paper 31440, National Bureau of Economic Research.

Allen, J., Howland, B., Mobius, M., Rothschild, D., and Watts, D. J. (2020). Evaluating the fake news problem at the scale of the information ecosystem. *Science Advances*, 6(14):eaay3539.

Barrera, O., Guriev, S., Henry, E., and Zhuravskaya, E. (2020). Facts, alternative facts, and fact checking in times of post-truth politics. *Journal of public economics*, 182:104123.

Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American economic review*, 96(5):1652–1678.

Benndorf, V., Kübler, D., and Normann, H.-T. (2025). Information unraveling and limited depth of reasoning. *Games and Economic Behavior*.

Berger, L. M. (2026). How digital media markets amplify news sentiment. Working paper, University of Cologne.

Berger, L. M., Kerkhof, A., Mindl, F., and Münster, J. (2025). Debunking "fake news" on social media: Immediate and short-term effects of fact-checking and media literacy interventions. *Journal of Public Economics*, 245:105345. Open access, CC BY. Available online 5 April 2025. Pre-registered at AEA Registry AEARCTR-0008199.

Besley, T. and Prat, A. (2006). Handcuffs for the grabbing hand? media capture and government accountability. *American Economic Review*, 96(3):720–736.

Bowen, T. R., Dmitriev, D., and Galperti, S. (2023). Learning from shared news: When abundant information leads to belief polarization. *The Quarterly Journal of Economics*, 138(2):955–1000.

Boxell, L. and Conway, J. (2022). Journalist ideology and the production of news: Evidence from movers. Working paper.

Braghieri, L. (2024). Political correctness, social image, and information transmission. *American Economic Review*, 114(12):3877–3904.

Braghieri, L., Eichmeyer, S., Levy, R., Mobius, M. M., Steinhardt, J., and Zhong, R. (2025). Level slant and polarization of news consumption on social media. *SSRN Working Paper*.

Braghieri, L., Schwardmann, P., and Tripodi, E. (2024). Talking across the aisle.

Budak, C., Nyhan, B., Rothschild, D. M., Thorson, E., and Watts, D. J. (2024). Misunderstanding the harms of online misinformation. *Nature*, 630:45–53.

Burdea, V., Montero, M., and Sefton, M. (2023). Communication with partially verifiable information: An experiment. *Games and Economic Behavior*, 142:113–149.

Bursztyn, L., Egorov, G., Haaland, I., Rao, A., and Roth, C. (2023a). Justifying dissent. *The Quarterly Journal of Economics*, 138(3):1403–1451.

Bursztyn, L., Haaland, I. K., Röver, N., and Roth, C. (2025). The social desirability atlas. Technical report, National Bureau of Economic Research.

Bursztyn, L., Rao, A., Roth, C., and Yanagizawa-Drott, D. (2023b). Opinions as facts. *The Review of Economic Studies*, 90(4):1832–1864.

Cagé, J., Gallo, N., Hengel, M., Henry, E., and Huang, Y. (2025). Fact-checking and misinformation: Evidence from the market leader. CESifo Working Paper 12319, CESifo.

Cagé, J., Hengel, E., Hervé, N., and Urvoy, C. (2022). Hosting media bias: Evidence from the universe of French broadcasts 2002–2020. HAL preprint hal-03878119.

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

Chopra, F., Haaland, I., and Roth, C. (2022). Do people demand fact-checked news? evidence from us democrats. *Journal of Public Economics*, 205:104549.

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

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, 50(6):1431–1451.

de Quidt, J., Haushofer, J., and Roth, C. (2018). Measuring and bounding experimenter demand. *American Economic Review*, 108(11):3266–3302.

Degan, A., Li, M., and Xie, H. (2023). An experimental investigation of persuasion through selective disclosure of evidence. *Canadian Journal of Economics/Revue canadienne d'économique*, 56(4):1490–1516.

Durante, R. and Knight, B. (2012). Partisan control, media bias, and viewer responses: Evidence from berlusconis italy. *Journal of the European Economic Association*, 10(3):451–481.

Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of accounting research*, pages 123–145.

Eil, D. and Rao, J. M. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics*, 3(2):114–138.

Eyster, E. and Rabin, M. (2005). Cursed equilibrium. *Econometrica*, 73(5):1623–1672.

Farina, A., Fréchette, G. R., Ispano, A., Lizzeri, A., and Perego, J. (2026). The selective disclosure of evidence: An experiment. *The Review of Economic Studies*. Forthcoming.

Federal Communications Commission (1987). Syracuse Peace Council, Memorandum Opinion and Order. 2 FCC Rcd 5043; recon. denied, 3 FCC Rcd 2035 (1988); aff'd sub nom. Syracuse Peace Council v. FCC, 867 F.2d 654 (D.C. Cir. 1989), cert. denied, 493 U.S. 1019 (1990).

Fong, M.-J., Lin, P.-H., and Palfrey, T. R. (2025). Cursed sequential equilibrium. *American Economic Review*, 115(8):2616–2658.

Freitag, J., Kerkhof, A., and Münster, J. (2021). Selective sharing of news items and the political position of news outlets. *Information Economics and Policy*, 56:100926.

Gentzkow, M. and Shapiro, J. M. (2006). Media bias and reputation. *Journal of Political Economy*, 114(2):280–316.

Gentzkow, M. and Shapiro, J. M. (2010). What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71.

Golman, R., Loewenstein, G., Molnar, A., and Saccardo, S. (2022). The demand for, and avoidance of, information. *Management Science*, 68(9):6454–6476.

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of law and Economics*, 24(3):461–483.

Guess, A., Nagler, J., and Tucker, J. (2019). Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science advances*, 5(1):eaau4586.

Guess, A. M., Malhotra, N., Pan, J., Barberá, P., Allcott, H., Brown, T., Crespo-Tenorio, A., Dimmery, D., Freelon, D., Gentzkow, M., González-Bailón, S., Kennedy, E., Kim, Y. M., Lazer, D., Moehler, D., Nyhan, B., Rivera, C. V., Settle, J., Thomas, D. R., Thorson, E., Tromble, R., Wilkins, A., Wojcieszak, M., Xiong, B., Kiewiet de Jonge, C., Franco, A., Mason, W., Stroud, N. J., and Tucker, J. A. (2023). Reshares on social media amplify political news but do not detectably affect beliefs or opinions. *Science*, 381(6656):404–408.

Guriev, S., Henry, E., Marquis, T., and Zhuravskaya, E. (2026). Curtailing false news, amplifying truth. *Econometrica*. Forthcoming.

Haaland, I. and Roth, C. (2020). Labor market concerns and support for immigration. *Journal of Public Economics*, 191:104256.

Henry, E., Zhuravskaya, E., and Guriev, S. (2022). Checking and sharing alt-facts. *American Economic Journal: Economic Policy*, 14(3):55–86.

Hillman, A. L. (2010). Expressive behavior in economics and politics. *European Journal of political economy*, 26(4):403–418.

Jin, G. Z., Luca, M., and Martin, D. (2021). Is no news (perceived as) bad news? an experimental investigation of information disclosure. *American Economic Journal: Microeconomics*, 13(2):141–173.

Jung, W.-O. and Kwon, Y. K. (1988). Disclosure when the market is unsure of information endowment of managers. *Journal of Accounting research*, pages 146–153.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Lazer, D. M., Baum, M. A., Benkler, Y., Berinsky, A. J., Greenhill, K. M., Menczer, F., Metzger, M. J., Nyhan, B., Pennycook, G., Rothschild, D., et al. (2018). The science of fake news. *Science*, 359(6380):1094–1096.

Lee, D. S. (2009). Training, wages, and sample selection: Estimating sharp bounds on treatment effects. *Review of Economic Studies*, 76(3):1071–1102.

Light, M. T., He, J., and Robey, J. P. (2020). Comparing crime rates between undocumented immigrants, legal immigrants, and native-born us citizens in texas. *Proceedings of the National Academy of Sciences*, 117(51):32340–32347.

Loosen, W., von Garmissen, A., Bartelt, E., and van Olphen, T. (2023). Journalismus in Deutschland 2023: Aktuelle Befunde zu Situation und Wandel. Arbeitspapiere des Hans-Bredow-Instituts 68, Hans-Bredow-Institut, Hamburg.

Louis-Sidois, C. (2025). Both judge and party? investigating the political unbiasedness of fact-checkers. *Journal of the European Economic Association*, 23(6):2137–2164.

Manili, J. (2024). Disclosure to an heterogeneous audience: The role of motivated beliefs. Working paper, Northwestern University. SSRN 4696951.

Marie, A., Altay, S., and Strickland, B. (2023). Moralization and extremism robustly amplify myside sharing. *PNAS Nexus*, 2(4):pgad078.

Marie, A. and Petersen, M. B. (2025). Motivations to connect with like-minded audiences increase partisan sharing on social media. *PNAS nexus*, 4(7):pgaf197.

Martin, G. J. and McCrain, J. (2019). Local news and national politics. *American Political Science Review*, 113(2):372–384.

Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391.

Möbius, M. M., Niederle, M., Niehaus, P., and Rosenblat, T. S. (2021). Managing self-confidence: Theory and experimental evidence. Working paper, October 2021.

Mullainathan, S. and Shleifer, A. (2005). The market for news. *American Economic Review*, 95(4):1031–1053.

Musolff, R., Roth, C., and Zimmermann, F. (2026). Learning from false stories. ECONtribute Discussion Paper No. 412.

National Academies of Sciences, Engineering, and Medicine (2017). *The Economic and Fiscal Consequences of Immigration*. The National Academies Press, Washington, DC.

Osmundsen, M., Bor, A., Vahlstrup, P. B., Bechmann, A., and Petersen, M. B. (2021). Partisan polarization is the primary psychological motivation behind political fake news sharing on twitter. *American political science review*, 115(3):999–1015.

Pennycook, G., Epstein, Z., Mosleh, M., Arechar, A. A., Eckles, D., and Rand, D. G. (2021). Shifting attention to accuracy can reduce misinformation online. *Nature*, 592:590–595.

Peterson, E. and Iyengar, S. (2021). Partisan gaps in political information and information-seeking behavior: Motivated reasoning or cheerleading? *American Journal of Political Science*, 65(1):133–147.

Prat, A. and Strömberg, D. (2013). The political economy of mass media. *Advances in economics and econometrics*, 2:135.

Shin, J. and Thorson, K. (2017). Partisan selective sharing: The biased diffusion of fact-checking messages on social media. *Journal of Communication*, 67(2):233–255.

Steindl, N., Lauerer, C., and Hanitzsch, T. (2017). Journalismus in Deutschland: Aktuelle Befunde zu Kontinuität und Wandel im deutschen Journalismus. *Publizistik*, 62(4):401–423.

Suvorov, A., van de Ven, J., and Villeval, M. C. (2024). Selective information sharing and group delusion. Working paper, New Economic School, University of Amsterdam, and CNRS-GATE Lyon.

Thaler, M. (2024). The fake news effect: Experimentally identifying motivated reasoning using trust in news. *American Economic Journal: Microeconomics*, 16(2):1–38.

Thaler, M. (2025). The supply of motivated beliefs. Working paper.

Verrecchia, R. E. (1983). Discretionary disclosure. *Journal of Accounting and Economics*, 5:179–194.

Vosoughi, S., Roy, D., and Aral, S. (2018). The spread of true and false news online. *Science*, 359(6380):1146–1151.

Zhu, Y. and Dukes, A. (2015). Selective reporting of factual content by commercial media. *Journal of Marketing Research*, 52(1):56–76.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review*, 110(2):337–361.

A Experimental Design Details

A.1 Fact Selection and Validation Survey (Experiment 1)

I selected statistics for the main experiment from a larger pool of candidates using a pre-registered validation survey ($N = 100$, AsPredicted No. 232064) conducted on Prolific between June 10 and June 21, 2025. The survey evaluated eight candidate statistics (four hypothesized as pro-immigration and four as anti-immigration) against three survey-measured criteria (perceived credibility, policy relevance, and perceived political direction). Facts were selected by the following rule: within each direction, the two statistics with the highest policy relevance among those with adequate perceived credibility were selected, subject to the constraint that the final set contains two pro-immigration and two anti-immigration statistics. The pre-registration (AsPredicted No. 232064) governs fact selection. It also lists a belief-accuracy outcome, two order conditions, correlations of initial beliefs with the ratings, a pre–post attitude test, and five exclusion rules.

Candidate facts and selection criteria. Table A.1 reports the eight candidate statistics with their survey-measured scores. Four criteria were operationalized as follows:

Veracity was established through source verification. Each statistic is drawn from a peer-reviewed publication, a federal statistical agency, or a research institution. The sources shown to participants are listed below.

Policy relevance was measured by asking respondents: "How relevant do others consider the statistic when forming immigration policy positions?" (1 = completely irrelevant, 5 = extremely relevant).

Perceived credibility was measured by asking respondents: "How credible is the statistic?" (1 = not at all, 5 = entirely).

Perceived political direction was measured through two items: *usability* ("For whom would the statistic be most useful in a policy debate?"; 1 = exclusively supporting expansive immigration, 5 = exclusively supporting restrictive immigration) and *opinionability* ("How does learning about the statistic change someone's immigration policy views?"; 1 = convince to support expansion, 5 = convince to support limitation). These directional measures serve the classification of statistics as pro- or anti-immigration and the balance criterion.

Table A.1: Candidate Facts: Survey-Measured Indicators

	Policy	Perceived	Perceived Direction	
Statistic	Relevance	Credibility	Usability	Opinionability
Pro-immigration				
Crime Rate†	3.74 (1.18)	3.50 (0.97)	2.23 (1.06)	2.36 (1.11)
Labor†	3.64 (1.04)	3.81 (1.02)	2.40 (1.09)	2.57 (1.23)
Entrepreneurship	3.12 (1.08)	3.54 (1.00)	2.09 (1.05)	2.18 (1.15)
Healthcare	3.28 (1.03)	3.59 (1.07)	2.03 (1.05)	2.19 (1.15)
Anti-immigration				
Border Crossings†	3.85 (1.14)	3.97 (0.90)	3.94 (1.14)	4.08 (1.11)
Fiscal Impact†	3.79 (1.01)	3.55 (0.93)	3.68 (1.15)	3.73 (1.01)
Poverty Rate	3.41 (1.01)	3.93 (0.93)	3.40 (1.13)	3.52 (1.11)
Wage Inequality	3.35 (0.95)	3.53 (0.90)	3.41 (1.06)	3.43 (0.99)

Notes: The sample is the $N = 100$ respondents of the Experiment 1 validation survey. Cells report means with standard deviations in parentheses. Selected statistics are marked with †. Policy relevance: 1 = completely irrelevant, 5 = extremely relevant. Perceived credibility: 1 = not at all credible, 5 = entirely credible. Usability and opinionability: 1 = supports expansive immigration, 5 = supports restrictive immigration. One-sample t-tests against the midpoint of 3 reject the null for all eight statistics on both directional measures ($p < 0.001$ for all selected statistics).

Selection. The selection rule applied within each direction independently. Among the four pro-immigration candidates, Crime Rate (3.74) and Labor Force Participation (3.64) had the highest policy relevance scores, both with adequate perceived credibility (≥ 3.50). Among the four anti-immigration candidates, Border Crossings (3.85) and Fiscal Impact (3.79) had the highest policy relevance, again with adequate credibility. The one rejected statistic with notably high credibility—Poverty Rate (3.93)—scored lower on policy relevance (3.41) than both selected anti-immigration statistics. Perceived direction does not detectably differ across Democrats, Republicans and Independents for the four selected statistics (one-way ANOVA).

Sources. For each statistic I report the values and the sources shown to participants. All sources were available at the time of the experiment (July–August 2025).

Crime Rate. The screen stated that immigrants have lower crime rates than U.S.-born Americans across multiple studies and time periods, and illustrated this with Texas data, described as the most detailed available: 213 violent crimes per 100,000 residents for U.S.-born citizens, 185 for legal immigrants (13% lower) and 96 for undocumented immigrants (55% lower). These are violent felony arrest rates for 2012–2018 from Light et al. (2020), who link Texas Department of Public Safety criminal-history

records to Department of Homeland Security immigration status checks. The gap persists when they use convictions instead of arrests.[50] The screen listed Light et al. (2020), the Cato Institute,[51] and "Texas Department of Public Safety conviction records (2012–2018)" as sources, although the figures shown are arrest rates. National evidence points in the same direction (Abramitzky et al. 2023).

Labor Force Participation. Participants were shown 2023 annual average labor force participation rates: 66.6% for the foreign-born population and 61.8% for the native-born population (4.8 pp higher). The source shown was the Bureau of Labor Statistics.[52]

Border Crossings. Participants were shown a historical average of 890,435 Border Patrol apprehensions per year over fiscal years 1970–2019 and a recent average of 1,584,573 per year over fiscal years 2020–2024 (labeled 2019–2024 on the screen), 78% higher.[53] Both are U.S. Border Patrol counts. The statistic reflects data available at experiment time (July–August 2025).

Fiscal Impact. The statement compares first-generation immigrants with U.S.-born Americans. The screen stated that first-generation immigrants have higher net costs at the state and local levels and showed the aggregate: first-generation immigrant households generate $57.4 billion more in annual state and local costs than they contribute in taxes, averaged over 2011–2013. The comparison rests on the same report, which the screen did not show: per independent person, with dependents assigned to their parents, the first generation costs state and local budgets about $1,600 a year, while the U.S.-born generations contribute on net.[54]

A.2 Belief Elicitation and Donation Task (Experiment 1)

Prior-knowledge elicitation. Senders complete each comparative statement by choosing "higher," "lower," or "don't know." Correct completions receive a bonus, incorrect completions do not, and "don't know" responses receive the bonus with 25% probability. Because exactly one completion is correct, choosing the more likely one wins the bonus with probability at least one half, against one quarter for "don't know." "Don't know" accounts for 8.8 percent of answers, and replacing its 0.5 coding with separate indicators for correct and "don't know" answers leaves every alignment coefficient in

[50] Published in *Proceedings of the National Academy of Sciences*. Data and replication materials available at https://www.openicpsr.org/openicpsr/project/124923.

[51] Cato Institute Policy Analysis no. 977; https://www.cato.org/policy-analysis/illegal-immigrant-murderers-texas-2013-2022.

[52] U.S. Bureau of Labor Statistics, *Foreign-Born Workers: Labor Force Characteristics – 2023* (May 21, 2024), https://www.bls.gov/news.release/archives/forbrn_05212024.pdf.

[53] CBP historical data available at https://www.cbp.gov/newsroom/stats/cbp-enforcement-statistics.

[54] National Academies of Sciences, Engineering, and Medicine (2017), Chapter 9, Section 9.6 and Table 9-6, https://www.nationalacademies.org/read/23550/chapter/14.

Table 3 unchanged to four decimal places.

Donation task. Senders allocate $500 between the *American Immigration Council* (AIC), which advocates for increased immigration, and *NumbersUSA*, which advocates for decreased immigration. I briefly describe the stated goals of each organization and present their missions in their own words, with the display order randomized. Participants are informed that one donation decision from the experiment will be randomly selected and implemented, and they receive a verification link to confirm the donation was placed.

A description text and part of both mission statements are reproduced verbatim below. Senders saw the two organizations side by side as bulleted lists, with their left–right order randomized.

> *Description text.* American immigration organizations range from those advocating for more open policies to those pushing for greater restrictions. The *American Immigration Council* champions seeing immigrants as valuable contributors who deserve stronger protections and family-friendly policies. In contrast, *NumbersUSA* advocates for reduced immigration levels to prevent overcrowding and promotes stricter enforcement of limitations.
>
> Below, you will find a brief description of each organization in their own words.
>
> **American Immigration Council (AIC)**
>
> - The American Immigration Council envisions a nation where immigrants are embraced, communities are enriched, and justice prevails for all.
> - We strive to create a society that values immigrants as vital contributors to our nation [. . .].
> - Our immigration system and the way it treats immigrants often fail to live up to our highest values and ideals.
> - Not only have we failed, in many cases, to afford basic due process rights, but in countless ways, our outdated immigration laws and policies do not respect the family unit, discourage innovation, and turn away those in need of protection.
>
> **NumbersUSA**
>
> - NumbersUSA advocates for sensible immigration policies that ensure Americans' safety and promote economic fairness while protecting the environment and quality of life.

- The current system creates endless "chain migration," of distant relatives, with backlogs, frustration, and a level of immigration that is incompatible with conservation and economic justice.
- The current system displaces qualified Americans with exploitable foreign workers.
- The current system raffles away green cards and puts tourists who give birth in the U.S. on a path to citizenship.

A.3 Information Condition Implementation (Experiment 1)

The persuasiveness elicitation proceeds in three stages: study presentation and elicitation (all senders), information feedback (INFO = 1 only), and expected attitude change elicitation (all senders). I describe each stage and reproduce the text shown to participants.

Stage 1: Study presentation and elicitation (all senders). All senders read the following description of the Haaland and Roth (2020) study:

> A recent study published in the prestigious academic journal, *Journal of Public Economics*, examined how people form political opinions.
>
> The researchers surveyed more than 3,000 Americans that are representative of the US population.
>
> People were told about a 1980 event when 55,000 Cuban immigrants unexpectedly arrived in Miami, increasing the city's workforce by 8% almost overnight. They were asked how this affected the wages and employment of U.S.-born workers.
>
> Then, half of the participants learned that research found this wave of immigration had virtually no effect on the wages or unemployment for U.S.-born workers, because the new immigrants increased demand for goods and services, which created more jobs.
>
> The other half of the participants received no information.

Senders then make two incentivized predictions:

1. *Behavioral effect (continuous):* "By what percentage did learning about the research findings change people's petition signing rates?" Senders enter a percentage (positive or negative).

2. *Differential updating (categorical):* "After learning about the research findings, who do you think changed their attitudes toward immigration more?" Senders choose

among: (a) people who initially believed immigration harmed U.S.-born workers, (b) people who initially believed immigration did not harm U.S.-born workers, or (c) both groups changed about the same amount.

Each question is equally likely to be selected for the bonus payment. For the continuous estimate, the probability of receiving the bonus is 100% if the estimate equals the true value (69%) and falls to 0% at a deviation of 20 percentage points or more, as $(1 - 5|d|)^2$ for a deviation d expressed as a share. For the categorical estimate, the probability is 100% if correct and 0% otherwise. Senders are instructed not to search for the answers and tab-switching is monitored to enforce compliance. Senders who switched tabs during this stage were excluded as pre-registered.

Stage 2: Information feedback (INFO = 1 only). Senders in the INFO condition receive personalized feedback comparing their predictions to the actual results. Senders in the no-INFO condition proceed directly to Stage 3. The feedback screen displays the following:

> **Petition Signatures**
>
> The percentage of people who signed the White House petition actually increased by **69%**.
>
> You estimated that the percentage of people who signed the White House petition changed by ***[participant's estimate]%***.
>
> That means ***you [underestimated/overestimated] the change by [participant-specific gap] percentage points***.
>
> **Effect on Attitudes toward Immigration**
>
> Actual Result: People who initially believed immigration harmed U.S.-born workers changed their attitudes toward immigration **significantly more** than those who already held correct beliefs.
>
> Your Answer: ***[participant's answer]***.

The feedback is accompanied by a bar chart that displays the sender's estimate and the true value (69%) side by side, with bar heights scaled proportionally. The sender's categorical prediction is marked as correct or incorrect with a color-coded indicator. A summary reiterates the two key findings: (i) receiving information strongly changed behavior and (ii) the effect was strongest for participants who initially held false beliefs.

Stage 3: Expected attitude change elicitation (all senders). All senders then predict how receiving each of the four immigration statistics would shift the attitudes of hypothetical receivers. For each statistic j, senders allocate 10 hypothetical participants—described as holding the receiver belief state realized for that statistic (false belief or

"don't know")—among three categories: those who become *more supportive* of immigration, those who become *less supportive* of immigration, and those whose opinions *do not change*. The allocations must sum to 10. This procedure generates the expected attitude change measure $\Delta\text{Attitude}_{i,j}$ defined in Section 2.2.2, which provides within-sender variation in perceived persuasiveness across statistics, and between-sender variation across receiver belief states.

> You will estimate how ten participants who hold **false beliefs** or **do not know the answer** would change their attitudes about immigration if they **are provided with research findings**.
>
> In particular, you will estimate how many of them...
>
> - become less supportive of immigration.
> - don't change their opinion.
> - become more supportive of immigration.

The interface displays 10 person icons to make the allocation task concrete. Senders complete this task four times: once for each of the four statistics, with the hypothetical receivers described by the belief state of the receiver group matched to that statistic. The order of statistics is randomized across senders.

Senders' beliefs and the INFO first stage. Senders believe the information in Haaland and Roth (2020) moved policy views, with roughly 90 percent expecting it to raise petition signatures, though they underestimate the size of the effect it had (Figure B.1). Senders shown the Haaland and Roth (2020) results expect a larger pro-immigration shift from the crime fact (+5.5 pp, one-sided $p = 0.002$) and from the labor fact (+3.5 pp, one-sided $p = 0.033$). For the two anti-immigration facts the changes are not distinguishable from zero (Table B.4, Panel (A)).

A.4 Fact Selection and Pre-Validation Survey (Experiments 2 and 3)

I selected facts for Experiments 2 and 3 from a larger candidate pool using a pre-registered pre-validation survey ($N = 319$, AsPredicted #290989) conducted on the Bilendi panel in Germany in May 2026. The survey evaluated 36 candidate facts. Six were hypothesized as pro-policy and six as anti-policy in three domains: immigration, climate, redistribution. They were evaluated against four criteria, and additionally validated the domain-specific advocacy organizations used in the ACTION treatment.

Criteria and selection rule. Four criteria were applied to each candidate fact. *Prior knowledge* was elicited via an incentivized completion of the fact statement with a confidence slider from 50 to 100. The probabilistic prior knowledge score is PK $=$ confidence$/100$ for correct completions and $1 -$ confidence$/100$ for incorrect completions. *Directional consistency* was measured by a 5-point rating (1 = strongly pro-policy, 5 = strongly anti-policy). The pre-registration counts a fact as directional if at least 60 percent of all respondents rate it in its hypothesized direction. Because many respondents chose the neutral midpoint, four of the twelve fielded facts meet this threshold, and no pro-immigration or anti-climate candidate does, so the pre-registered rule would have left two topic–direction cells without a fact. I therefore applied a weaker criterion, a deviation from the pre-registration: a fact counts as directional if more respondents rate it in its hypothesized direction than in the opposite one. In the main experiment, where senders also rate two facts from the domain they were not assigned, the share who say a fact shows their opinion is correct rises with their stance in the fact's hypothesized direction for eleven of the twelve fielded facts, including the reformulated *Netzausbaukosten*, and is flat for *Beschäftigungswachstum*. Prior-knowledge balance was also checked on the pre-registered binary correct rate, on which the selected sets differ by 0.21 (immigration) and 0.23 (climate), within the bound of 0.25. *Policy relevance* and *perceived credibility* were each measured on a 1–5 scale.

The selection rule chose the top three facts by policy relevance within each topic and direction cell among those passing directional consistency, subject to balance constraints across the pro and anti sets for credibility and policy relevance.

The ACTION treatment organizations were additionally validated against two criteria: (i) at least 50% of respondents correctly recognize each organization as representing the pro-policy or anti-policy position respectively, and (ii) a positive Spearman correlation between topic-specific attitude and hypothetical donation to the pro-organization, restricted to respondents with directional attitudes.

Domain selection. The redistribution domain failed on both grounds and is excluded. All six redistribution anti-direction candidate facts failed the directional consistency criterion, and the redistribution anti-organization had a recognition rate of 33.3%, well below the 50% threshold. Immigration passed both organization criteria (anti-organization recognition: 93.6%; Spearman $r = 0.571$). Climate passed the correlation criterion ($r = 0.566$), but its anti-organization fell below the recognition threshold (42.6% vs. 50%). The pre-registration both drops any topic whose organization pair fails and requires at least two topics to proceed, so with two of the three topics failing, one of the two rules had to be violated. I kept two topic because climate's deviation was confined to one organization, whereas redistribution failed on its facts as well.

Deviations from algorithmic fact selection. *Immigration.* The algorithm selected *Bildungsabschluss* (correct rate 0.78), which produced a pro–anti gap in the correct rate of 0.29, exceeding the 0.25 balance threshold. It was replaced by *Illegale Aufenthalte* (correct rate 0.55), closing the gap to 0.21 while satisfying all other balance constraints.

Climate. The algorithmic pro set (*Emissionsminderung*, *Klimaziele*, *Extremwetter*) failed the pre-registered prior-knowledge balance, with a gap of 0.255 in the correct rate against the bound of 0.25, and was revised to *Erneuerbare Energien*, *Stromgestehungskosten*, and *Grüne Beschäftigung*. *Netzausbaukosten* (anti) was reformulated between the pre-validation survey and the main experiment. Directional ratings in Table A.2 reflect the pre-reformulation version, which failed the consistency criterion (17.5% consistent vs. 42.1% inconsistent). The reformulated version is what was fielded in the main experiment. Excluding this fact leaves the estimates unchanged (positional alignment +0.130; persuasion interaction +0.071).

Candidate facts and pre-validation ratings. Table A.2 reports the pre-validation ratings of all 36 candidate facts and marks the 12 selected. Table A.3 reports the balance check.

Table A.2: Candidate Facts: Pre-Validation Ratings

					% Directional rating		
Fact	Dir.	Relevance	Credibility	PK	Consistent	Inconsistent	Sel.
Immigration							
Fachkräftemangel	Pro	3.79	3.65	0.52	45.8	31.2	✓
Beschäftigungswachstum	Pro	3.43	3.45	0.55	38.8	28.6	✓
Rente	Pro	3.31	3.55	0.65	14.3	42.9	
Wohnungsmarkt	Pro	3.16	3.34	0.32	8.1	41.9	
Integrationskurs	Pro	3.61	3.29	0.37	32.1	25.0	✓
Unternehmertum	Pro	2.94	3.44	0.56	27.8	30.6	
Kriminalität	Anti	4.04	3.76	0.70	90.0	4.0	✓
Bürgergeld‡	Anti	3.98	3.74	0.67	78.7	4.3	✓
Bildungsabschluss	Anti	3.74	3.74	0.69	60.0	6.0	
Armut	Anti	3.45	3.63	0.64	44.9	14.3	
Illegale Aufenthalte	Anti	3.69	3.45	0.51	76.5	7.8	✓
Duldung	Anti	3.68	3.51	0.64	52.8	22.6	
Climate policy							
Emissionsminderung	Pro	3.90	3.43	0.44	44.8	20.7	
Erneuerbare Energien	Pro	3.67	3.61	0.52	51.0	14.3	✓
Klimaziele	Pro	3.75	3.36	0.35	33.9	13.6	
Extremwetter	Pro	3.80	4.04	0.73	76.1	2.2	
Stromgestehungskosten	Pro	3.46	3.51	0.56	61.4	14.0	✓
Grüne Beschäftigung	Pro	3.28	3.46	0.57	47.8	8.7	✓
Strompreise	Anti	3.78	4.08	0.84	17.6	51.0	
CO_2-Intensität	Anti	3.16	3.44	0.52	23.6	36.4	
Netzausbaukosten†	Anti	3.84	3.53	0.72	17.5	42.1	✓
Energiekosten Industrie	Anti	4.09	3.60	0.80	35.8	34.0	✓
Preisvolatilität	Anti	3.73	3.36	0.66	33.9	25.0	✓
Energieimporte	Anti	3.98	3.74	0.73	27.9	44.2	
Redistribution (excluded)							
Sozialer Aufstieg	Pro	3.78	3.38	0.44	46.0	16.0	
Erbschaften	Pro	3.35	3.52	0.63	61.7	11.7	
Einkommensmobilität	Pro	3.30	3.42	0.54	38.0	14.0	
Umverteilungswirkung	Pro	3.86	3.59	0.50	32.8	26.6	
Armutsrisiko	Pro	3.80	3.46	0.55	50.8	13.1	
Nettovermögen	Pro	3.72	3.56	0.71	59.6	10.5	
Steuer-/Abgabenquote	Anti	3.96	3.51	0.68	25.5	45.1	
Beitragszahler	Anti	4.40	3.77	0.75	16.1	46.8	
Bürgergeld-Anreize	Anti	3.60	3.42	0.60	12.9	45.2	
Rentenausgaben	Anti	3.61	3.46	0.47	29.6	31.5	
Bürgergeld/Reallöhne	Anti	3.53	3.42	0.51	34.0	35.8	
Abgabenlast	Anti	3.92	3.63	0.75	30.0	41.7	

Notes: The sample is the $N = 319$ respondents of the pre-validation survey for Experiments 2 and 3; each fact was rated by 36–64 respondents. Policy relevance and perceived credibility: 1 = lowest, 5 = highest. PK is the confidence-weighted prior knowledge probability (see text). Consistent (Inconsistent) is the share of respondents whose directional rating matches (opposes) the assigned direction. The remainder chose the neutral midpoint. ✓ Selected. † Reformulated after pre-validation; ratings are pre-reformulation.

Table A.3: Balance Check: Selected Facts by Domain

Domain	ΔRelevance	ΔCredibility	ΔPK
Immigration	0.29	0.19	0.15
Climate	0.42	0.03	0.18

Notes: The table compares the facts selected for Experiments 2 and 3 by domain. Δ is the absolute difference between mean pro-direction and mean anti-direction facts. Thresholds: ΔRelevance $\leq$ 1.0, ΔCredibility $\leq$ 1.0, ΔPK $\leq$ 0.25. The ΔPK row uses the confidence-weighted score. On the binary correct rate the gaps are 0.21 and 0.23. Climate ΔPK uses pre-reformulation statistics for *Netzausbaukosten*.

Sources. Senders saw each fact with its correct answer and a one-line source beneath it. Table A.4 reproduces these lines.

Table A.4: Sources Shown to Senders

Fact	Source
Immigration	
Fachkräftemangel	Statistisches Bundesamt (2025): "Fachkräftemangel: Anteil ausländischer Beschäftigter in Engpassberufen gestiegen"
Beschäftigungswachstum	Bundesagentur für Arbeit (2024): "Ausländische Arbeitskräfte am deutschen Arbeitsmarkt"
Integrationskurs	Bundesamt für Migration und Flüchtlinge (2021): "Verläufe in Integrationskursen" (Kurzanalyse 7/2021)
Kriminalität	Bundeskriminalamt (2024): "Polizeiliche Kriminalstatistik 2024"; Statistisches Bundesamt (2024): "Bevölkerung nach Staatsangehörigkeit"
Bürgergeld	Bundesagentur für Arbeit (2025): "Analyse Arbeitsmarkt – Grundsicherung für Arbeitsuchende"; Statistisches Bundesamt (2025): "Bevölkerung nach Staatsangehörigkeit"
Illegale Aufenthalte	Eurostat (2025): "Migration and Asylum in Europe 2025"; Statistisches Bundesamt (2025): "Bevölkerung – Basistabelle"
Climate policy	
Erneuerbare Energien	Fraunhofer ISE (2025): "Öffentliche Stromerzeugung 2024"
Stromgestehungskosten	Fraunhofer ISE (2024): "Stromgestehungskosten erneuerbare Energien"
Grüne Beschäftigung	Internationale Energieagentur (2023): "World Energy Employment 2023"; Bundesministerium für Wirtschaft und Energie (2024): "Erneuerbare Energien in Zahlen"
Netzausbaukosten	EWI Köln (2023): "Netzausbaukosten und Netzentgelte bis 2045"
Energiekosten Industrie	Deutscher Industrie- und Handelskammertag (2025): "Energiewende-Barometer 2025"
Preisvolatilität	Schischke & Rathgeber (2026): "Unraveling European electricity price volatility: The impact of renewables," *Energy Strategy Reviews*

Notes: The facts are those of Table 4, labeled as in Table A.2. Source lines are reproduced as displayed to senders in Experiments 2 and 3, in German.

A.5 Donation Task Organizations (Experiments 2 and 3)

Senders allocate €250 between the two advocacy organizations for their assigned domain. For immigration these are *ProAsyl*, which advocates a more open migration and asylum policy, and *Republik21*, which advocates a stricter one; for climate, *BUND*, which advocates strict national climate targets for individual economic sectors, and

Stiftung Marktwirtschaft, which opposes them. Senders first see a description sentence naming both organizations, then each organization's positions in its own words, presented side by side as bulleted lists with the left–right order randomized within sender.

The task was administered in German and the text below is translated. The description text is reproduced for the immigration domain and for a sender who saw *ProAsyl* on the left. The alternative order reverses the two clauses and the two blocks without altering any wording.

> *Description text.* The following two organizations hold different positions on immigration. On one side is *ProAsyl*, which advocates a more open migration and asylum policy. On the other side is *Republik21*, which advocates a stricter migration and asylum policy.
>
> Below you will find some statements from the two organizations in their own words.
>
> **ProAsyl**
>
> - Refugee protection is a human right.
> - The sealing of borders, refusals of entry without individual case assessment, and political measures that in effect hollow out the right to asylum are to be critically examined and rejected.
> - Refugees should be given access to education, work, counseling, and participation in society.
>
> **Republik21**
>
> - Immigration policy should take into account not only the well-being of immigrants but also the interests of the native population.
> - The inflow of refugees is to be reduced in order to end the overburdening of institutions.
> - In cases of refusal to integrate, criminal conduct, extremist views, or holiday travel to the country of origin, refugees' residence permits should lapse.
>
> **BUND**
>
> - The barriers currently holding back the environmentally compatible expansion of renewable energy must fall.
> - Turn off the money tap for oil, gas, and coal by abolishing climate-damaging subsidies such as preferential tax treatment for diesel and kerosene and ending support for new gas infrastructure.

- Climate action requires binding commitments: the Federal Climate Action Act must firmly set annual climate targets for all economic sectors.

Stiftung Marktwirtschaft

- Germany's energy transition is uniquely expensive, the security of its energy supply is not assured over the longer term, the CO_2 savings achieved are modest, and the promised "green economic miracle" has failed to materialize.
- The domestic energy supply should be expanded in a technology-open manner in order to ensure a secure, clean, and affordable energy supply beyond solar, wind, and hydro(gen) power.

A.6 Revision of Experiment 2 During Fielding

Experiment 2 had to be shortened during fielding. Senders who started before the revision took a median of 29.4 minutes, against the fifteen agreed with the panel provider. One part was therefore removed. An instruction page and four pages on which senders estimated how people sharing their position would perceive facts from the other topic. After the revision the median was 17.1 minutes. The donation task, the five provision decisions, both incentivized belief elicitations, the cross-domain rating task and the post-questionnaire were unchanged. AsPredicted #293252, filed on 27 May 2026, supersedes #292898 and states the removal and its reason. The provision task, outcome, positional alignment measure, estimating equations and exclusion rules are identical in both. Every sender in the main sample began after it was filed.

The removed pages carried the perceived-believability control of the pre-registered column (2). The specification in #293252 retained this control from #292898 by error, so column (2) of Table 5 omits a control the pre-registration lists. The 106 senders are excluded from the main sample. Adding them back leaves every sign unchanged (Appendix Table C.12, column (3)).

B Experiment 1: Supplementary Results

B.1 Sample Characteristics

Table B.1: Sample Characteristics and Treatment Balance

	Population (1)	Sample (2)	CONTROL (3)	ACTION (4)	INFO (5)	ACTION×INFO (6)
Gender						
Female	0.509	0.502	0.507	0.490	0.503	0.510
Male	0.491	0.478	0.470	0.477	0.477	0.487
Ethnicity						
White	0.578	0.587	0.577	0.590	0.587	0.593
Black	0.121	0.129	0.137	0.120	0.140	0.120
Hispanic	0.187	0.182	0.180	0.183	0.177	0.187
Asian	0.059	0.058	0.063	0.063	0.057	0.050
Other	0.055	0.044	0.043	0.043	0.040	0.050
Political Affiliation						
Democrat	—	0.388	0.377	0.393	0.393	0.387
Independent	—	0.308	0.323	0.303	0.303	0.300
Republican	—	0.298	0.297	0.297	0.300	0.300

Notes: The sample is $N = 1{,}200$ U.S. adults recruited on Prolific (quota-sampled) for Experiment 1. Column (1) reports 2020 U.S. Census population proportions for gender and ethnicity. Column (2) reports full sample proportions. Columns (3)–(6) report proportions by treatment cell.

B.2 Additional Summary Statistics

Table B.2: Prior Knowledge of the Four Immigration Facts

	Share of Responses		
	Correct	False	Don't know
Crime	0.701	0.223	0.076
Labor	0.794	0.143	0.063
Border	0.613	0.317	0.070
Fiscal	0.305	0.553	0.143

Notes: The sample is the $N = 1{,}200$ Experiment 1 senders. The beliefs were elicited before senders were informed they would make provision decisions. Correct means the stated belief matches the verified fact direction; False means the stated belief contradicts the fact.

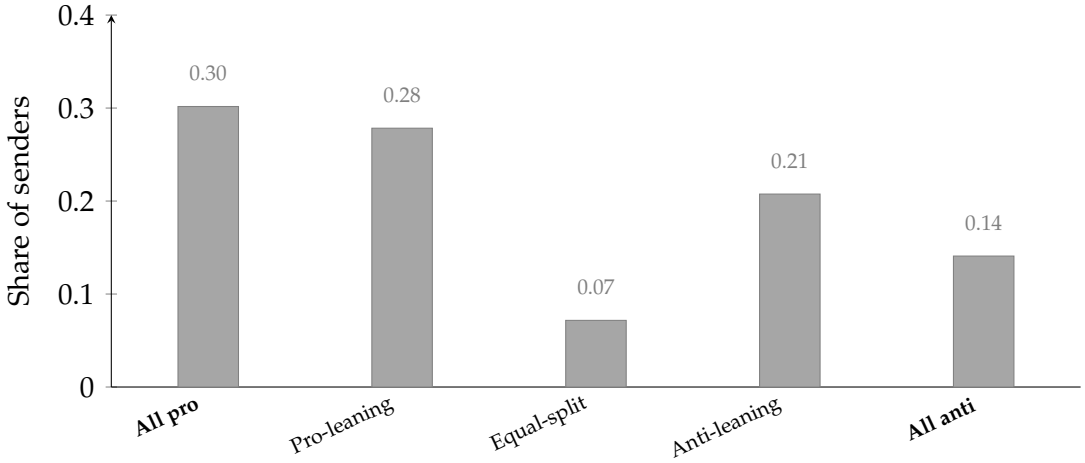


Figure B.2: Distribution of Senders' Donation Allocations to the Immigration Advocacy Groups

Notes: The sample is the $N = 1{,}200$ senders of Experiment 1. The horizontal axis shows the allocation to the pro-immigration group collapsed into five categories ("All pro = \$500, "Pro-leaning = \$251–499, "Equal-split = \$250, "Anti-leaning = \$1–249, "All anti = \$0) from a total \$500 stake.

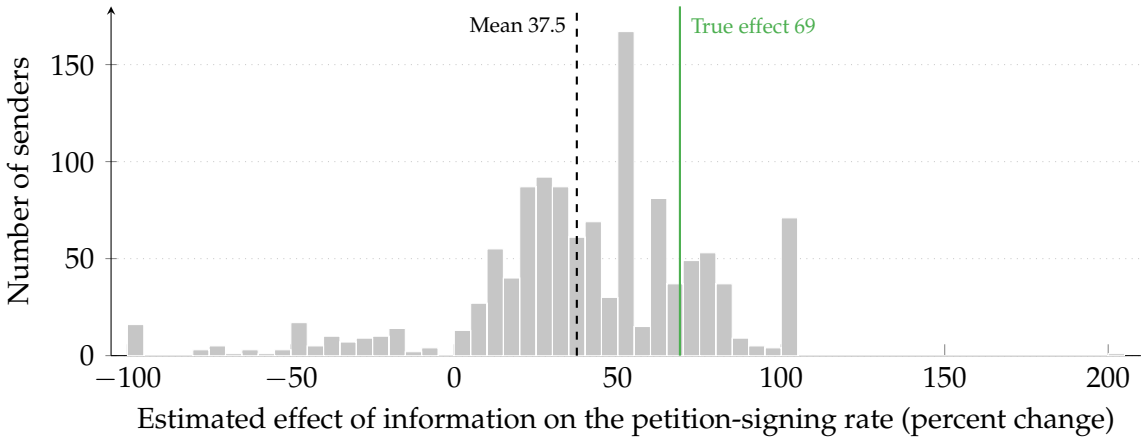


Figure B.1: Senders' Beliefs About the Persuasive Effect of Information

Notes: The sample is the $N = 1{,}200$ Experiment 1 senders. The figure shows the distribution of senders' estimates of the percent change in the petition-signing rate caused by providing information, elicited before the provision task (pre-treatment). The dashed line marks the sample mean (37.5 percent); the solid green line marks the true effect, a 69 percent increase, from Haaland and Roth (2020).

B.3 Robustness of Main Results

Demeaning by fact. The sender-level alignment bias is 9.3 pp, averaged over the four treatments. The facts differ in how often they are provided: the two pro-immigration facts by 6.6 pp more than the two anti-immigration ones. A sender's aligned facts are the pair matching their position, and 62.5 percent of senders are pro-immigration, so this difference raises the average bias. Subtracting each fact's mean provision leaves 7.6 pp. Averaged over INFO, the ACTION increase is 4.3 pp raw and 4.5 pp demeaned.

Table B.3: Perceived Credibility of Facts by Political Affiliation

	Percentage Perceiving Statistic as Credible			
	All	Democrats	Independents	Republicans
Crime	0.962	0.996	0.949	0.933
Labor	0.978	0.998	0.965	0.964
Border	0.978	0.985	0.976	0.972
Fiscal	0.968	0.981	0.962	0.958

Notes: The sample is the $N = 1{,}200$ Experiment 1 senders. Credibility is assessed post-provision. Values shown are the share of senders rating each statistic as credible (1 = credible, 0 = not credible).

Power for Hypothesis 1.3. INFO raises the expected movement toward the sender's side by $+0.034$ per unit of alignment ($p = 0.032$), and provision rises by $+0.079$ per unit of that forecast ($p < 0.001$, sender and fact fixed effects). Together they imply an Align $\times$ INFO interaction of $+0.0027$, 6.3 percent of the one-sided 80 percent-power minimum detectable effect of 0.043. The estimated interaction is $+0.004$ ($p = 0.831$). The calculation applies to this design under a linear mapping and treats the association of provision with the forecast as the response to a forecast change.

Alternative information provision specification. Table B.8 allows asymmetric INFO effects by fact direction to accommodate the asymmetric belief-shifting effect. The additional terms $\text{INFO}_i \times \text{PROSTAT}_j$ and $\text{Align}_{i,j} \times \text{INFO}_i \times \text{PROSTAT}_j$ are not statistically significant in the specifications with fact fixed effects. Wald tests on column (3) detect no effect of INFO on how provision responds to alignment, for pro-immigration facts ($p = 0.403$) or anti-immigration facts ($p = 0.606$), and none on overall provision for pro-immigration facts ($p = 0.517$).

Table B.10 replaces positional alignment with senders' own forecast of how far each fact moves receivers toward their position, elicited before any provision decision.

B.4 Heterogeneity

Table B.11 reports the pre-registered heterogeneity of selective provision by partisanship, party affiliation and policy view intensity, at the decision level with a binary aligned indicator and fact fixed effects. Table B.12 interacts positional alignment and both treatments with the receiver's belief state. No interaction with the belief state is distinguishable from zero.

Table B.4: Expected Effects of Information on Political Preferences with Controls

	Expect Effect of Information on Political Preferences			
Statistic	Crime (1)	Labor (2)	Border (3)	Fiscal (4)
Panel (A): INFO				
Constant	0.236***	0.033	−0.336***	−0.259***
	(0.067)	(0.071)	(0.073)	(0.080)
INFO	0.055***	0.035*	−0.034	0.028
	(0.019)	(0.019)	(0.021)	(0.023)
Estimated Effect Size	0.000	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Prior Knowledge	0.116***	0.133***	0.003	0.129***
	(0.023)	(0.027)	(0.024)	(0.026)
Panel (B): Receiver Belief State				
Constant	0.249***	0.043	−0.347***	−0.255***
	(0.068)	(0.071)	(0.074)	(0.080)
Receiver Belief False	0.014	0.006	−0.000	0.018
	(0.019)	(0.019)	(0.021)	(0.023)
Estimated Effect Size	0.000	0.001***	0.001***	0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Prior Knowledge	0.116***	0.134***	0.001	0.129***
	(0.023)	(0.027)	(0.024)	(0.026)
Clustered SE	No	No	No	No
Controls	Yes	Yes	Yes	Yes
Panel (A): R^2	0.049	0.054	0.047	0.076
Panel (B): R^2	0.042	0.051	0.044	0.075
Panel (A) Adj. R^2	0.031	0.036	0.029	0.058
Panel (B) Adj. R^2	0.025	0.033	0.027	0.058
Num. obs.	1200	1200	1200	1200

Notes: The sample is the $N = 1{,}200$ senders of Experiment 1. The outcome is the expected effect of each statistic on receivers' immigration policy preferences, $[-1, 1]$ (positive = pro-immigration shift). In Panel (A) the regressor is the INFO indicator. In Panel (B) the regressor is Receiver Belief False (= 1 if the matched receiver holds an incorrect prior; = 0 if the receiver reports not knowing). Estimated Effect Size is the sender's incentivized estimate of the Haaland & Roth (2020) treatment effect. Prior Knowledge is the incentivized prior knowledge (1 correct, 0.5 "don't know," 0 incorrect). Controls are demographic covariates. Standard errors are not clustered. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

B.5 Disbelief, Stated Motives, and Norms

Stated-motive cell sizes and the second classifier. Table B.15 regresses each sender's alignment bias on the considerations their closing open-text-response was classified as invoking. Twelve of 1,102 senders state an expressive reason, four in CONTROL and eight under ACTION, and the coefficient is significant at the one percent level within each treatment taken on its own. Any estimated difference between the treatments is

Table B.5: Treatment Effects on Alignment Bias

	Alignment Bias	
	(1)	(2)
Intercept	0.067^{***}	-0.026
	(0.021)	(0.071)
Treatment: ACTION	0.065^{**}	0.066^{**}
	(0.029)	(0.029)
Treatment: INFO	0.008	-0.000
	(0.029)	(0.029)
Treatment: ACTION×INFO	0.029	0.024
	(0.029)	(0.029)
Controls	No	Yes
R^2	0.005	0.048
Adj. R^2	0.003	0.029
Num. obs.	1114	1114

Notes: The sample is the $N = 1{,}114$ Experiment 1 senders with revealed directional policy preference ($AIC_i \neq 250$). The outcome is sender-level alignment bias (mean aligned minus mean unaligned provision rate). The reference is CONTROL (no-ACTION, no-INFO). Treatment ACTION is ACTION only; Treatment INFO is INFO only; Treatment ACTION×INFO is ACTION and INFO. Column (2) adds age, gender, education, ethnicity, and party identification. Standard errors are not clustered. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

thus imprecise. 23 senders cite persuasion in the ACTION treatment. Table B.17 reports the second classifier, which gives a smaller but same-signed estimate on the 30 such senders it identifies.

Table B.6: Treatment Effects on Information Provision: Logit

	Probability of Statistic Provision		
	(1)	(2)	(3)
Align	0.460^{***}	0.416^{***}	0.430^{***}
	(0.133)	(0.134)	(0.137)
INFO	-0.126	-0.126	-0.126
	(0.183)	(0.184)	(0.188)
ACTION	0.173	0.172	0.167
	(0.195)	(0.196)	(0.200)
Prior Knowledge	0.399^{***}	0.308^{***}	0.317^{***}
	(0.103)	(0.114)	(0.113)
Align × INFO	0.005	-0.009	-0.010
	(0.177)	(0.177)	(0.183)
Align × ACTION	0.535^{**}	0.538^{**}	0.553^{***}
	(0.210)	(0.209)	(0.214)
INFO × ACTION	-0.234	-0.233	-0.241
	(0.259)	(0.260)	(0.268)
Align × INFO × ACTION	-0.434	-0.425	-0.433
	(0.270)	(0.270)	(0.277)
Clustered SE	Subject	Subject	Subject
Statistics FE	No	Yes	Yes
Controls	No	No	Yes
No. of Senders	1200	1200	1200
AIC	3258.054	3247.312	3248.924
BIC	3316.341	3325.029	3449.691
Log Likelihood	-1620.027	-1611.656	-1593.462
Deviance	3240.054	3223.312	3186.924
Num. obs.	4800	4800	4800

Notes: The sample is the $N = 1{,}200$ senders of Experiment 1 (4,800 sender-statistic observations). The estimator is logistic regression. Column (1) controls for prior knowledge only; column (2) adds statistic fixed effects; column (3) adds statistic fixed effects and demographic controls. Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.7: Treatment Effects on Information Provision (Binary Alignment)

	Probability of Statistic Provision		
	(1)	(2)	(3)
Aligned	0.056***	0.050***	0.050***
	(0.019)	(0.019)	(0.019)
INFO	−0.017	−0.014	−0.014
	(0.026)	(0.026)	(0.025)
ACTION	−0.046*	−0.045*	−0.048*
	(0.027)	(0.027)	(0.027)
Prior Knowledge	0.047***	0.036***	0.035***
	(0.012)	(0.013)	(0.013)
Aligned × INFO	0.009	0.004	0.004
	(0.027)	(0.027)	(0.027)
Aligned × ACTION	0.065**	0.064**	0.064**
	(0.028)	(0.028)	(0.028)
INFO × ACTION	0.028	0.025	0.024
	(0.038)	(0.038)	(0.038)
Aligned × INFO × ACTION	−0.047	−0.042	−0.041
	(0.041)	(0.041)	(0.041)
Clustered SE	Subject	Subject	Subject
Statistics FE	No	Yes	Yes
Controls	No	No	Yes
No. of Senders	1114	1114	1114
Num. obs.	4456	4456	4456
R^2	0.028	0.032	0.041
Adj. R^2	0.026	0.030	0.034

Notes: The sample is the 1,114 Experiment 1 senders with a directional policy preference, 4,456 sender-statistic observations. The model is a linear probability model. *Aligned* is an indicator equal to one if the statistic's policy direction matches the sender's own position. Column (1) controls for prior knowledge only; column (2) adds statistic fixed effects; column (3) adds statistic fixed effects and demographic controls. Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.8: Treatment Effects on Information Provision: INFO Effects by Fact Direction

	Probability of Statistic Provision		
	(1)	(2)	(3)
Align	0.041***	0.037***	0.037***
	(0.012)	(0.012)	(0.013)
INFO	−0.034*	−0.018	−0.018
	(0.019)	(0.020)	(0.020)
ACTION	−0.002	−0.002	−0.003
	(0.016)	(0.016)	(0.016)
Prior Knowledge	0.038***	0.035***	0.034***
	(0.012)	(0.012)	(0.012)
Align × INFO	−0.015	−0.009	−0.011
	(0.021)	(0.021)	(0.022)
Align × ACTION	0.041**	0.041**	0.041**
	(0.018)	(0.018)	(0.018)
INFO × ACTION	−0.007	−0.007	−0.008
	(0.024)	(0.024)	(0.024)
Align × INFO × ACTION	−0.027	−0.027	−0.027
	(0.026)	(0.026)	(0.026)
INFO × PROSTAT	0.038***	0.005	0.005
	(0.015)	(0.020)	(0.020)
Align × INFO × PROSTAT	0.025	0.025	0.029
	(0.023)	(0.023)	(0.025)
Clustered SE	Subject	Subject	Subject
Statistics FE	No	Yes	Yes
Controls	No	No	Yes
No. of Senders	1200	1200	1200
Num. obs.	4800	4800	4800
R^2	0.032	0.034	0.041
Adj. R^2	0.030	0.031	0.034

Notes: The sample is the $N = 1{,}200$ Experiment 1 senders (4,800 sender-statistic observations). ProStat is a binary indicator for pro-immigration statistics (Crime and Labor). Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.9: Leave-One-Fact-Out Robustness of the Alignment Effect and the ACTION Increase

Coefficient	Sign-positive	Min β	Max β	Max p
Align	4/4	0.031	0.043	0.028
Align × ACTION	4/4	0.039	0.044	0.068
Statistics FE		Yes		
Controls		Yes		
Clustered SE		Sender		

Notes: The sample is Experiment 1 senders. Each coefficient is re-estimated four times in the main specification (Table 3, column 3; statistic fixed effects, full demographic and political controls, sender-clustered standard errors), each time dropping one of the four statistics (crime, labor, border, fiscal). Columns report the number of leave-one-out estimates with a positive coefficient, the minimum and maximum coefficient across the four runs, and the largest p-value observed. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.10: Alignment and Expected Persuasion in Experiment 1

	Probability of Statistic Provision		
	(1)	(2)	(3)
Align	0.038***		0.043***
	(0.009)		(0.011)
Align × ACTION	0.028**		0.024
	(0.013)		(0.016)
Expected persuasion		0.023	−0.013
		(0.016)	(0.021)
Expected persuasion × ACTION		0.042*	0.013
		(0.024)	(0.030)
ACTION	−0.006	−0.007	−0.006
	(0.012)	(0.012)	(0.012)
Prior knowledge	0.034***	0.050***	0.034***
	(0.012)	(0.012)	(0.012)
Fact fixed effects	Yes	Yes	Yes
INFO control	Yes	Yes	Yes
Senders	1,200	1,200	1,200
Observations	4,800	4,800	4,800

Notes: Linear probability estimates for the Experiment 1 senders, pooled across the INFO treatment. The outcome is the provision decision. Expected persuasion is the net share of ten hypothetical receivers whom the sender expects to move toward the sender's own side after receiving the fact, elicited without incentives before any provision decision. Standard errors clustered by sender in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.11: Heterogeneity in Selective Provision with Binary Alignment

	Probability of Statistic Provision		
	(1)	(2)	(3)
Aligned	−0.008	0.083***	0.004
	(0.024)	(0.024)	(0.033)
Aligned × ACTION	0.075**	0.021	−0.011
	(0.035)	(0.033)	(0.048)
Aligned × Partisanship	0.123***		
	(0.038)		
Aligned × ACTION × Partisanship	−0.056		
	(0.059)		
Aligned × Independent		−0.071**	
		(0.033)	
Aligned × Republican		−0.006	
		(0.035)	
Aligned × ACTION × Independent		0.096*	
		(0.050)	
Aligned × ACTION × Republican		−0.031	
		(0.050)	
Aligned × Preference intensity			0.082*
			(0.042)
Aligned × ACTION × Preference intensity			0.069
			(0.061)
Moderator	Partisanship	Political party	Pref. intensity
Clustered SE	Subject	Subject	Subject
Statistics FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
No. of Senders	1114	1114	1114
Num. obs.	4456	4456	4456
R^2	0.041	0.040	0.044
Adj. R^2	0.036	0.033	0.038

Notes: The sample is Experiment 1 senders with a directional policy preference. Partisanship is the distance of the sender's political orientation from the scale midpoint, rescaled to $[0, 1]$; policy view intensity is $|\text{Pref}_i|$. Each column interacts a binary aligned indicator, the ACTION treatment, and one moderator, with statistic fixed effects and demographic controls. Reference category for political party is Democrat. Alignment enters as a binary indicator rather than the continuous measure, and the specification includes statistic fixed effects. Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.12: Heterogeneous Treatment Effects of Receiver Beliefs

	Probability of Statistic Provision	
	(1)	(2)
Intercept	0.863***	0.894***
	(0.018)	(0.045)
Align	0.026*	0.025*
	(0.014)	(0.014)
INFO	−0.034	−0.034
	(0.021)	(0.022)
ACTION	−0.021	−0.022
	(0.020)	(0.021)
Receiver Belief False	−0.025	−0.024
	(0.018)	(0.018)
Prior Knowledge	0.035***	0.034***
	(0.012)	(0.012)
Align × INFO	0.025	0.025
	(0.023)	(0.023)
Align × ACTION	0.053**	0.054**
	(0.024)	(0.024)
INFO × ACTION	0.014	0.014
	(0.030)	(0.031)
Align × Receiver Belief False	0.019	0.020
	(0.022)	(0.022)
INFO × Receiver Belief False	0.043	0.043
	(0.027)	(0.027)
ACTION × Receiver Belief False	0.034	0.034
	(0.026)	(0.026)
Align × INFO × ACTION	−0.050	−0.051
	(0.035)	(0.035)
Align × INFO × Receiver Belief False	−0.042	−0.042
	(0.032)	(0.032)
Align × ACTION × Receiver Belief False	−0.020	−0.022
	(0.033)	(0.033)
INFO × ACTION × Receiver Belief False	−0.041	−0.043
	(0.038)	(0.038)
Align × INFO × ACTION × Receiver Belief False	0.047	0.048
	(0.045)	(0.045)
Clustered SE	Subject	Subject
Statistics FE	Yes	Yes
Controls	No	Yes
No. of Senders	1200	1200
R^2	0.034	0.041
Adj. R^2	0.030	0.033
Num. obs.	4800	4800

Notes: The sample is N = 1,200 Experiment 1 senders, 4,800 sender-statistic observations. The outcome is a binary provision indicator. The model is the full four-way interaction of Align, ACTION, INFO, and Receiver Belief False (= 1 if the matched receiver holds a factually incorrect prior belief; = 0 if the receiver reports not knowing). Column (2) adds demographic controls. Statistic fixed effects and sender-clustered standard errors throughout. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.13: The Alignment Effect by Sender Beliefs

	Probability of Statistic Provision		
	(1)	(2)	(3)
Align	0.031*	0.052	0.044
	(0.016)	(0.048)	(0.049)
Prior Knowledge	0.035***	0.024**	0.026**
	(0.012)	(0.012)	(0.012)
Credibility		0.253***	0.253***
		(0.063)	(0.063)
Align × Prior Knowledge	0.009		0.015
	(0.016)		(0.015)
Align × Credibility		−0.021	−0.025
		(0.047)	(0.047)
Align × ACTION	0.041**	0.044**	0.044**
	(0.018)	(0.018)	(0.018)
Align × INFO	0.004	0.007	0.007
	(0.017)	(0.017)	(0.017)
Align × INFO × ACTION	−0.027	−0.032	−0.032
	(0.026)	(0.025)	(0.025)
Clustered SE	Subject	Subject	Subject
Statistics FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes
No. of Senders	1200	1200	1200
Num. obs.	4800	4800	4800
R^2	0.040	0.058	0.058
Adj. R^2	0.034	0.052	0.052

Notes: The sample is the $N = 1{,}200$ senders of Experiment 1 (4,800 sender-statistic observations). Each column is specification (1) with statistic fixed effects and demographic controls, augmented by an interaction between alignment and one of the two belief measures. *Prior Knowledge* is the sender's incentivized prior about the statistic, elicited before any provision decision (1 correct, 0.5 "don't know," 0 incorrect). *Credibility* is the sender's unincentivized rating of the statistic, elicited after the provision task. Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.14: Selective Provision Among Facts the Sender Knew

	Probability of Statistic Provision		
	(1)	(2)	(3)
Align	0.025**	0.024**	
	(0.011)	(0.012)	
Align × ACTION	0.040**	0.037**	
	(0.018)	(0.018)	
Aligned (binary)			0.041**
			(0.019)
Aligned (binary) × ACTION			0.045
			(0.029)
Statistics FE	Yes	Yes	Yes
Sender FE	No	Yes	Yes
Controls	Yes	No	No
Clustered SE	Sender	Sender	Sender
No. of Senders	673	673	673
Num. obs.	1952	1952	1952

Notes: The sample is Experiment 1 decisions on statistics the sender answered correctly in the incentivized prior elicitation, restricted to senders who did so for at least one aligned and one unaligned statistic ($N = 673$ senders, 1,952 decisions). Every statistic in the sample is one the sender had answered correctly. Prior knowledge is not a regressor here because it is constant by construction. Column (1) is specification (1) restricted to alignment and its ACTION interaction, with statistic fixed effects and the full control set; column (2) replaces the controls with sender fixed effects; column (3) repeats column (2) with a binary aligned indicator, which imposes no proportional scaling in preference intensity. Standard errors clustered at the sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.15: Self-Reported Motivations and Alignment Bias

	Alignment Bias		
	(1)	(2)	(3)
Intercept	0.109***	0.091***	0.094***
	(0.014)	(0.017)	(0.019)
Belief correction	−0.099***	−0.099***	−0.084***
	(0.017)	(0.017)	(0.022)
Expression	0.636***	0.630***	0.753***
	(0.105)	(0.105)	(0.129)
Persuasion	0.276***	0.276***	0.056
	(0.084)	(0.083)	(0.118)
Other stated reason	−0.000	0.001	0.001
	(0.129)	(0.131)	(0.130)
ACTION		0.036*	0.029
		(0.020)	(0.028)
Persuasion × ACTION			0.417***
			(0.156)
Expression × ACTION			−0.224
			(0.191)
Belief correction × ACTION			−0.022
			(0.033)
Receiver Action	No	Yes	Yes
Num. obs.	1102	1102	1102
R^2	0.099	0.102	0.115
Adj. R^2	0.096	0.098	0.109

Sample: $N = 1{,}102$ senders with a revealed directional policy preference ($AIC_i \neq 250$), Experiment 1, counting once each of the 12 participants who completed the study twice. The outcome is sender-level alignment bias. Indicators for each motive the sender's open-ended decision reason was coded as citing, by `claude-sonnet-5`. The omitted reference is a sender whose response was coded as citing none of the listed motives. Categories are not mutually exclusive, so the indicators are not a partition. *Belief correction* refers to direction-blind truth-telling or the receiver's right to accurate information; *Expression* refers to reasons citing the fact's fit with the sender's own position; *Persuasion* refers to receiver-directed reasons. ACTION is the receiver political-action treatment. Column (3) interacts each coded motive with ACTION, so each channel carries its own treatment slope; *Other stated reason* is a residual category and is not interacted. Heteroskedasticity-robust (HC1) standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.16: Stated Accuracy Norms and Alignment Bias

	Correct false beliefs	Facts matter for opinions	Platforms should fact-check	Disagreement is about facts
	(1)	(2)	(3)	(4)
Alignment bias	-0.214^{***}	-0.138^{**}	-0.252^{**}	-0.090
	(0.066)	(0.064)	(0.101)	(0.090)
Controls	Yes	Yes	Yes	Yes
No. of Senders	1114	1114	1114	1114
R^2	0.031	0.021	0.089	0.023
Adj. R^2	0.020	0.010	0.079	0.013
Num. obs.	1114	1114	1114	1114

Notes: The sample is Experiment 1 senders with a directional policy preference. The dependent variable is agreement with one post-task statement about accurate information, on a five-point scale; column (4) is reverse-coded so that higher values indicate the pro-accuracy pole on every item. *Alignment bias* is the share of aligned statistics provided minus the share of unaligned statistics provided. The items were elicited after the provision task. All columns include demographic and political controls and treatment indicators. Heteroskedasticity-robust standard errors. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table B.17: Self-Reported Motivations and Alignment Bias: Secondary Classifier

	Alignment Bias		
	(1)	(2)	(3)
Intercept	0.112***	0.093***	0.097***
	(0.014)	(0.017)	(0.019)
Belief correction	−0.099***	−0.100***	−0.091***
	(0.017)	(0.017)	(0.023)
Expression	0.628***	0.620***	0.757***
	(0.092)	(0.093)	(0.113)
Persuasion	0.187***	0.187***	0.028
	(0.070)	(0.069)	(0.086)
Other stated reason	−0.045	−0.048	−0.046
	(0.053)	(0.053)	(0.053)
ACTION		0.037*	0.026
		(0.020)	(0.028)
Persuasion × ACTION			0.335**
			(0.132)
Expression × ACTION			−0.320*
			(0.189)
Belief correction × ACTION			−0.010
			(0.033)
Receiver Action	No	Yes	Yes
Num. obs.	1102	1102	1102
R^2	0.078	0.081	0.092
Adj. R^2	0.075	0.077	0.085

Sample: $N = 1{,}102$ senders with a revealed directional policy preference ($AIC_i \neq 250$), Experiment 1, counting once each of the 12 participants who completed the study twice. Outcome is sender-level alignment bias. Indicators for each motive the sender's open-ended decision reason was coded as citing, from the *secondary* classifier (`gpt-5.6-terra`). The omitted reference is a sender whose response was coded as citing none of the listed motives. Categories are not mutually exclusive. ACTION is the receiver political-action treatment. Heteroskedasticity-robust (HC1) standard errors in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

C Experiment 2: Supplementary Results

C.1 Policy Views of Senders

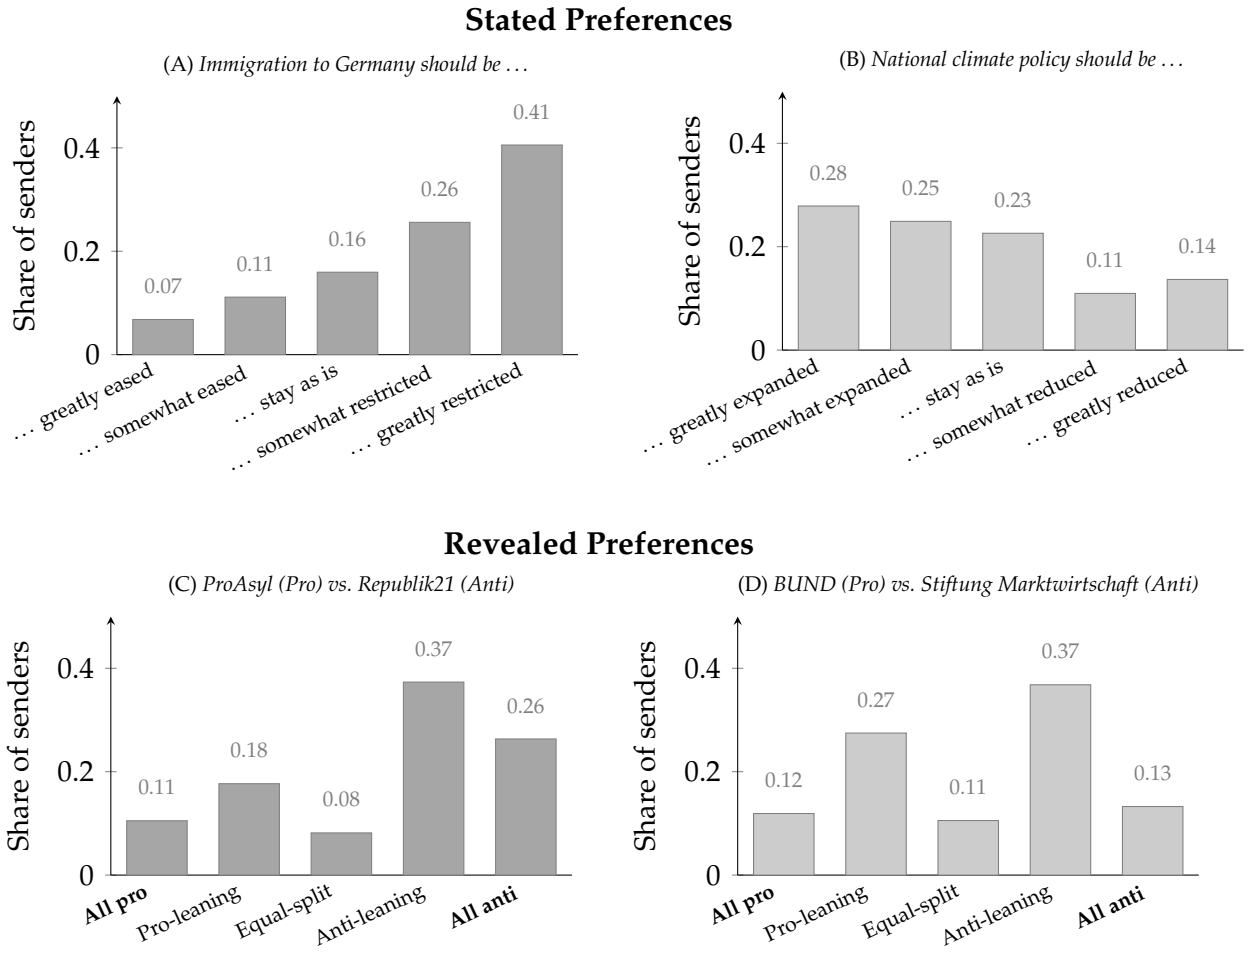


Figure C.1: Distribution of Policy Preferences Among Senders

Notes: The sample is the Experiment 2 senders ($N = 1{,}548$). The top row (Panels A–B) shows stated preferences from the 5-point Likert item. Labels complete the stance statement shown in the panel title (e.g., *Immigration to Germany should be ... greatly eased*). The bottom row (Panels C–D) shows revealed preferences from the real donation task, grouped into five buckets from All pro (€250 to the pro-policy organization) to All anti (€0), displayed pro-to-anti left to right. The left column shows immigration senders ($N = 809$) and the right column climate senders ($N = 739$).

C.2 Controls and Additional Provision Estimates

C.3 Robustness of the Motive-Channel Estimates

C.4 Heterogeneity

Tables C.14 and C.15 report the pre-registered heterogeneity analyses.

Table C.1: Provision Decisions with Demographic Controls

	Probability of Fact Provision	
	(1)	(2)
Positional alignment	0.110***	0.128***
	(0.013)	(0.013)
ACTION	0.021*	0.023*
	(0.012)	(0.012)
Positional alignment × ACTION	0.014	0.014
	(0.017)	(0.018)
Clustered SE	Sender	Sender
Fact FE	No	Yes
Controls	Yes	Yes
Num. obs.	7715	7715
R^2	0.085	0.045
Adj. R^2	0.080	0.039

Notes: The sample is the $N = 1{,}548$ Experiment 2 general-population senders, both treatments, with 7,715 observations with non-missing controls. The outcome is the binary provision indicator. Positional alignment is $\text{Align}_{ij} = \text{Pref}_i \times \text{Dir}_j \in [-1, +1]$. Column (1) includes sender–receiver ideological distance, age, gender, education, Bundesland, closest-party affiliation and left–right self-placement, without fact fixed effects. Column (2) corresponds to column (3) of Table 5: fact fixed effects and the pre-registered controls (age, gender, education, Bundesland, voting intention and left–right self-placement). Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

C.5 The Elicited Returns

Figure C.2 shows the distributions of the two elicited returns.

Table C.2: Provision and the Three Returns: Linear Probability Reduced Form

	Probability of Fact Provision		
	(1)	(2)	(3)
Stance–fact alignment index	0.148***	0.148***	0.138***
	(0.020)	(0.020)	(0.028)
Expected accuracy gain	0.034*	0.034*	0.053**
	(0.017)	(0.017)	(0.025)
Persuasion	0.008	0.009	−0.004
	(0.007)	(0.007)	(0.009)
ACTION		0.026**	0.024
		(0.013)	(0.020)
Stance–fact alignment index × ACTION			0.020
			(0.037)
Expected accuracy gain × ACTION			−0.036
			(0.034)
Persuasion × ACTION			0.025*
			(0.013)
Clustered SE	Sender	Sender	Sender
Fact FE	Yes	Yes	Yes
Num. obs.	7740	7740	7740
R^2	0.014	0.015	0.016
Adj. R^2	0.012	0.013	0.013

Notes: The sample is Experiment 2's general-population senders. The dependent variable is the binary provision decision. It is the three-return linear probability reduced form, with the interaction of the stance–fact alignment index and ACTION estimated freely. Coefficients are linear-probability coefficients. Panel B, reports the four-determinant average marginal effects. Column (1) regresses provision on the stance–fact alignment index and the two elicited returns. Column (2) adds the ACTION main effect. Column (3) interacts each return with ACTION. Main effects are pooled across treatments in columns (1) and (2) and are CONTROL-treatment slopes in column (3). All columns include fact fixed effects; standard errors are clustered at the sender level (in parentheses). $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.3: Three-Return Reduced Form with Demographic Controls

	Probability of Fact Provision	
	(1)	(2)
Stance–fact alignment index	0.138***	0.139***
	(0.028)	(0.028)
Expected accuracy gain	0.053**	0.066***
	(0.025)	(0.024)
Persuasion	−0.004	−0.003
	(0.009)	(0.009)
ACTION	0.024	0.021
	(0.020)	(0.020)
Stance–fact alignment index × ACTION	0.020	0.023
	(0.037)	(0.037)
Expected accuracy gain × ACTION	−0.036	−0.046
	(0.034)	(0.034)
Persuasion × ACTION	0.025*	0.023*
	(0.013)	(0.013)
Clustered SE	Sender	Sender
Fact FE	Yes	Yes
Controls	No	Yes
Num. obs.	7740	7715
R^2	0.016	0.029
Adj. R^2	0.013	0.022

Notes: The sample is the $N = 1{,}548$ Experiment 2 general-population senders, both treatments; column (1) has 7,740 observations and column (2) 7,715 with non-missing controls. The outcome is the binary provision indicator. The stance–fact alignment index is the cross-domain behavioral rating $\alpha_{ij} \in [-1, +1]$. Persuasion is the incentivized expected movement of the receiver toward the sender's position $\Delta\pi_{ijk}$. Expected accuracy gain is the incentivized expected correction of the receiver's factual belief $\Delta b_{ijk} \in [-1, 1]$. Column (1) corresponds to the three-return reduced form of Table C.2; column (2) adds demographic controls (age, gender, education, Bundesland, closest-party affiliation, left-right self-placement). Sender-receiver ideological distance is excluded as a heterogeneity variable. All specifications include fact fixed effects. Standard errors clustered at sender level. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.4: Motive Channels with Positional Alignment Substituted for the Stance–Fact Alignment Index

	Probability of Fact Provision			
	(1)	(2)	(3)	(4)
Stance–fact alignment index	0.490***		0.448***	
	(0.053)		(0.073)	
Positional alignment		0.365***		0.339***
		(0.027)		(0.036)
Expected accuracy gain	0.131*	0.129*	0.130*	0.128*
	(0.069)	(0.070)	(0.069)	(0.070)
Persuasion	0.030	0.025	0.032	0.027
	(0.024)	(0.024)	(0.024)	(0.024)
Political distance	−0.250***	−0.240***	−0.250***	−0.240***
	(0.013)	(0.013)	(0.013)	(0.013)
ACTION	0.081**	0.078**	0.052	0.080**
	(0.040)	(0.040)	(0.054)	(0.040)
Expected accuracy gain × ACTION	−0.086	−0.074	−0.085	−0.072
	(0.096)	(0.097)	(0.096)	(0.097)
Persuasion × ACTION	0.061*	0.062*	0.058*	0.059*
	(0.035)	(0.035)	(0.035)	(0.035)
Stance–fact alignment index × ACTION			0.082	
			(0.101)	
Positional alignment × ACTION				0.053
				(0.049)
Clustered SE	Sender	Sender	Sender	Sender
Fact FE	Yes	Yes	Yes	Yes
Distance form	Linear	Linear	Linear	Linear
Alignment measure	Index	Positional	Index	Positional
Alignment measure × ACTION	Restricted	Restricted	Freed	Freed
Num. obs.	7740	7740	7740	7740
Adj. pseudo R^2	0.049	0.060	0.049	0.059

Notes: The sample is the $N = 1{,}548$ Experiment 2 senders (Bilendi panel, Germany, 2026), 7,740 sender–fact observations. The estimator is a probit; coefficients are index coefficients, matching the four-determinant specification of Table 6. The outcome is a binary provision indicator. The four-determinant specification is estimated twice: once with the stance–fact alignment index α_{ij}, assigned from cross-domain raters' validation judgments, and once with positional alignment $\text{Align}_{ij} = \text{Pref}_i \times \text{Dir}_j \in [-1, +1]$, the pre-registered confirmatory measure, in its place. Columns (1)–(2) set the interactions of ACTION with the alignment measure of that column and with distance to zero. Columns (3)–(4) free the interaction of ACTION with the alignment measure. All specifications include linear political distance, fact fixed effects and sender-clustered standard errors. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.5: Specification Robustness of the Treatment Effects

	(1) Linear probability (pre-registered)	(2) Logit (scaled logit)	(3) Sender fixed effects
Positional alignment (Hyp. 2.1)	0.128***	0.125***	0.129***
	(0.013)	(0.013)	(0.013)
Positional alignment × ACTION (Hyp. 2.2)	0.014	0.018	0.005
	(0.018)	(0.018)	(0.018)
Fact fixed effects	Yes	Yes	Yes
Sender fixed effects	No	No	Yes
Demographic controls	Yes	Yes	Absorbed
Senders	1,543	1,543	1,543
Observations	7,715	7,715	7,715

Notes: The sample is Experiment 2's general-population senders. The dependent variable is the binary provision decision. All three columns estimate the pre-registered confirmatory specification—provision on positional alignment, ACTION, and their interaction—on the same estimation sample. Column (1) reproduces column (3) of Table 5. Column (2) is the pre-registered logistic specification, reported as coefficients scaled by the sample mean of $p(1-p)$, so that magnitudes are comparable with the linear columns. Column (3) adds sender fixed effects, so the positional alignment coefficient is identified only from each sender's own comparisons across facts; the ACTION main effect is absorbed, and the demographic controls with it. Standard errors clustered at the sender level (in parentheses). $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.6: Provision and Positional Alignment Under Alternative Preference Measures

	(1) Composite (registered)	(2) Revealed only	(3) Stated only	(4) Excluding zero index	(5) Excluding offsetting
Positional alignment	+0.128***	+0.100***	+0.102***	+0.128***	+0.129***
	(0.013)	(0.012)	(0.012)	(0.013)	(0.013)
Positional alignment × ACTION	+0.014	+0.009	+0.009	+0.014	+0.014
	(0.018)	(0.017)	(0.016)	(0.018)	(0.018)
Observations	7,715	7,715	7,715	7,455	7,650
Senders	1,548	1,548	1,548	1,496	1,535
Fact fixed effects	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes

Notes: The sample is Experiment 2's general-population senders. Linear probability models of provision on positional alignment, its interaction with ACTION, and the pre-specified controls (age, gender, education, state, vote intention, left–right placement), with fact fixed effects and standard errors clustered at the sender level. Column (1) is the pre-registered specification, in which the sender's policy preference is the average of the revealed measure from the €250 donation split and the stated measure from the five-point stance item. Columns (2) and (3) rebuild alignment from each component alone. Column (4) drops the 52 senders whose composite preference is exactly zero; column (5) drops only the 13 of those who state a position at one pole of the scale and donate against it. Stars denote two-sided significance: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table C.7: The Four-Determinant Specification Under Relaxations: Distance Functional Form and the Expressive Interaction

	Four-determinant	Flexible distance	Freed expressive interaction
Stance–fact alignment index	0.490***	0.481***	0.448***
	(0.053)	(0.053)	(0.073)
Expected accuracy gain	0.131*	0.128*	0.130*
	(0.069)	(0.069)	(0.069)
Persuasion	0.030	0.033	0.032
	(0.024)	(0.024)	(0.024)
Political distance	−0.250***		−0.250***
	(0.013)		(0.013)
ACTION	0.081**	0.081**	0.052
	(0.040)	(0.040)	(0.054)
Expected accuracy gain × ACTION	−0.086	−0.084	−0.085
	(0.096)	(0.096)	(0.096)
Persuasion × ACTION	0.061*	0.060*	0.058*
	(0.035)	(0.035)	(0.035)
Stance–fact alignment index × ACTION			0.082
			(0.101)
Distance form	Linear	Indicators	Linear
Fact fixed effects	Yes	Yes	Yes
Senders	1,548	1,548	1,548
Observations	7,740	7,740	7,740

Notes: The sample is Experiment 2's general-population senders. Column (1) is column (3) of Table 6, the four-determinant specification: linear distance, the interactions of ACTION with the expressive return and with distance restricted to zero. Column (2) replaces linear distance with distance indicators; a Wald test with sender-clustered standard errors favors the flexible form ($\chi^2(3) = 26.4$, $p < 0.001$). Column (3) frees the expressive interaction with ACTION; a Wald test does not reject the restriction ($\chi^2(1) = 0.656$, $p = 0.418$). The freed interaction tests equality of index coefficients across treatments. Same sample, estimator, fixed effects, and clustering as Table 6. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.8: Alignment Measures and the Credibility Rating

	General population (Experiment 2)				Journalists (Experiment 3)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Stance–fact alignment index	0.490***	0.253***			0.466***	0.192		
	(0.053)	(0.076)			(0.151)	(0.218)		
Positional alignment			0.365***	0.331***			0.421***	0.393***
			(0.027)	(0.037)			(0.070)	(0.078)
Credibility rating		0.898***		0.264		0.937*		0.236
		(0.191)		(0.184)		(0.478)		(0.374)
Expected accuracy gain	0.131*	0.126*	0.129*	0.128*	0.299**	0.300**	0.345***	0.340**
	(0.069)	(0.069)	(0.070)	(0.070)	(0.135)	(0.133)	(0.133)	(0.133)
Expected persuasion	0.030	0.031	0.025	0.025	0.137***	0.135***	0.092*	0.093*
	(0.024)	(0.024)	(0.024)	(0.024)	(0.050)	(0.050)	(0.050)	(0.050)
Political distance	−0.250***	−0.246***	−0.240***	−0.240***	−0.218***	−0.221***	−0.208***	−0.209***
	(0.013)	(0.013)	(0.013)	(0.013)	(0.032)	(0.032)	(0.032)	(0.032)
ACTION	0.081**	0.084**	0.078**	0.080**				
	(0.040)	(0.040)	(0.040)	(0.040)				
Expected accuracy gain × ACTION	−0.086	−0.084	−0.074	−0.074				
	(0.096)	(0.097)	(0.097)	(0.097)				
Expected persuasion × ACTION	0.061*	0.058*	0.062*	0.061*				
	(0.035)	(0.035)	(0.035)	(0.035)				
Fact fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Senders	1,548	1,548	1,548	1,548	309	309	309	309
Observations	7,740	7,740	7,740	7,740	1,545	1,545	1,545	1,545

Notes: Probit estimates of the four-determinant specification: column (1) corresponds to column (3) of Table 6 and column (5) corresponds to column (1) of Table 8. Columns (2) and (6) add the credibility rating, the share of senders at the same stance who rate the fact true, from the same cross-domain rating task and page as the stance–fact alignment index. Columns (3)–(4) and (7)–(8) measure expression by positional alignment, built from the sender's own donation and stated position. Standard errors clustered by sender in parentheses. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.9: Within-Treatment Ratios to the Stance–Fact Alignment Index Coefficient

	Ratio to the stance–fact alignment index coefficient		
	CONTROL (1)	ACTION (2)	Difference (3)
Persuasion	0.075	0.171	0.095
	[-0.032, 0.205]	[0.062, 0.306]	[-0.065, 0.258]
Expected accuracy gain	0.277	0.094	−0.183
	[0.009, 0.603]	[-0.154, 0.384]	[-0.577, 0.247]
Political distance	−0.592	−0.441	0.151
	[-0.883, -0.433]	[-0.608, -0.331]	[-0.083, 0.447]

Notes: The sample is Experiment 2's general-population senders. Each entry is the ratio of a determinant's probit index coefficient to the stance–fact alignment index coefficient, estimated separately in each treatment, so that each treatment's latent scale cancels (Corollary 1(ii)). Column (3) is the difference between treatments. Brackets are 95% percentile intervals from 1,000 sender-cluster bootstrap replications. The ratios of the two returns are in the units of the elicited measures, and the intervals do not reflect sampling error in the alignment index (Appendix C.3).

Table C.10: Inference Robustness: Clustering and Wild Bootstrap

	(1) Sender SE	(2) Two-way SE (Sender + Fact)
Stance–fact alignment index	0.173***	0.173***
	(0.019)	(0.016)
Expected accuracy gain	0.030*	0.030
	(0.017)	(0.019)
Persuasion	0.011	0.011
	(0.009)	(0.009)
Political distance	−0.090***	−0.090***
	(0.005)	(0.005)
Persuasion × ACTION	0.022*	0.022
	(0.012)	(0.014)
ACTION	0.024*	0.024*
	(0.013)	(0.012)
Clustered SE	Sender	Sender + Fact
Fact FE	Yes	Yes
Wild bootstrap p (fact, exact 2^{12})	—	0.118
Num. obs.	7,740	7,740

Notes: The sample is Experiment 2's general-population senders, with 7,740 sender-fact observations. Both columns are linear probability models. Both columns estimate the four-determinant specification (*Stance–fact alignment index*, *Expected accuracy gain*, *Persuasion*, *Political distance*, *Persuasion* × ACTION, ACTION) with fact fixed effects. Column (1): standard errors clustered at the sender level. Column (2): two-way clustering by sender and fact ($K = 12$ fact clusters, asymptotic standard errors indicative only). The wild bootstrap is an exact enumeration of all $2^{12} = 4{,}096$ sign patterns at the fact level, and reports the p-value for *Persuasion* × ACTION. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.11: Sample Characteristics and Population Benchmarks

	Population (1)	Sample (2)	CONTROL (3)	ACTION (4)
Gender				
Female	0.499	0.481	0.501	0.462
Male	0.501	0.517	0.497	0.536
Diverse	—	0.003	0.003	0.003
Age band				
18–29	0.179	0.197	0.199	0.195
30–39	0.178	0.176	0.177	0.175
40–49	0.170	0.160	0.150	0.169
50–59	0.193	0.200	0.188	0.212
60–75	0.279	0.267	0.286	0.248
Bundesland				
Baden-Württemberg	0.134	0.129	0.135	0.122
Bayern	0.159	0.156	0.144	0.166
Berlin	0.045	0.042	0.040	0.044
Brandenburg	0.030	0.032	0.036	0.029
Bremen	0.008	0.010	0.011	0.010
Hamburg	0.023	0.023	0.021	0.024
Hessen	0.075	0.072	0.068	0.076
Mecklenburg-Vorpommern	0.019	0.023	0.023	0.024
Niedersachsen	0.095	0.099	0.102	0.096
Nordrhein-Westfalen	0.216	0.215	0.204	0.226
Rheinland-Pfalz	0.050	0.050	0.046	0.054
Saarland	0.012	0.010	0.016	0.005
Sachsen	0.047	0.048	0.050	0.047
Sachsen-Anhalt	0.025	0.032	0.029	0.034
Schleswig-Holstein	0.035	0.031	0.041	0.021
Thüringen	0.025	0.028	0.034	0.021
School-leaving qualification				
Hauptschulabschluss or none	0.297	0.265	0.277	0.253
Mittlerer Abschluss (Realschule)	0.314	0.293	0.270	0.315
Fachhochschul-/Hochschulreife or higher	0.390	0.442	0.453	0.431

Notes: The Experiment 2 sample is $N = 1{,}548$ German adults recruited on Bilendi (quota-sampled). Column (1) reports German population proportions. For gender, age band, and Bundesland, the benchmark is the population aged 18–75 as of 31 December 2024 (Statistisches Bundesamt, Statistischer Bericht 12411). For school-leaving qualification, the benchmark is the Mikrozensus 2023 (Statistisches Bundesamt, Bildungsstand der Bevölkerung), persons aged 15 and over excluding those still in school. Column (2) reports full-sample proportions. Columns (3)–(4) report proportions by treatment.

Table C.12: Sample-Inclusion Robustness of the Estimates

	Sample additions				Quality-check exclusions		
	Canonical	+ Speeders	+ Early starts	+ Both	Attention failers	Any tab switching	Comprehension attempts > 2
Positional alignment (Hyp. 2.1)	0.128^{***}	0.131^{***}	0.123^{***}	0.125^{***}	0.135^{***}	0.132^{***}	0.147^{***}
	(0.013)	(0.013)	(0.013)	(0.012)	(0.015)	(0.014)	(0.018)
Positional alignment $\times$ ACTION (Hyp. 2.2)	0.014	0.012	0.015	0.013	0.007	0.006	0.007
	(0.018)	(0.017)	(0.017)	(0.017)	(0.020)	(0.019)	(0.023)
Stance–fact alignment index (α_{ij})	0.489^{***}	0.483^{***}	0.477^{***}	0.472^{***}	0.551^{***}	0.523^{***}	0.462^{***}
	(0.053)	(0.052)	(0.050)	(0.049)	(0.061)	(0.057)	(0.071)
Expected accuracy gain (Δb_{ijk})	0.086^{*}	0.076	0.087^{*}	0.077^{*}	0.136^{**}	0.082	0.058
	(0.048)	(0.048)	(0.047)	(0.046)	(0.057)	(0.054)	(0.064)
Persuasion ($\Delta\pi_{ijk}$)	0.061^{***}	0.061^{***}	0.055^{***}	0.056^{***}	0.081^{***}	0.053^{***}	0.072^{***}
	(0.018)	(0.017)	(0.017)	(0.017)	(0.021)	(0.020)	(0.025)
Persuasion $\times$ ACTION	0.061^{*}	0.055	0.046	0.041	0.073^{*}	0.033	0.097^{**}
	(0.035)	(0.034)	(0.034)	(0.033)	(0.041)	(0.038)	(0.048)
Senders	1548	1588	1654	1694	1156	1277	860
Observations (confirmatory)	7715	7910	8245	8440	5760	6365	4285
Observations (motive specifications)	7740	7940	8270	8470	5780	6385	4300

Notes: The sample is Experiment 2's general-population senders under alternative sample definitions, and the dependent variable is the binary provision decision. Canonical: the main analysis sample. + Speeders: adds completers faster than $0.25\times$ the median completion time. + Early starts: adds participants who started before the revised instrument went live on 27 May 2026, at 09:00 UTC (Appendix A.6). + Both: adds both groups. Columns (5)–(7) drop, in turn, the three groups the pre-registration retains in the primary analysis and names for robustness: senders failing the attention check, senders who switched a browser tab at any point, and senders needing more than two attempts at the comprehension checks. None of the three differs across treatments (all $p > 0.30$). The table reports two estimators, each matching the specification it belongs to. Rows (1)–(2) are the pre-registered confirmatory specification and are estimated by linear probability, which is the form the pre-registration names: provision on positional alignment, ACTION, their interaction, demographic controls (age, gender, education, Bundesland, voting intention, left–right), and fact fixed effects. Rows (3)–(6) are the motive specification of Table 6 and are *probit index coefficients*. Rows (3)–(5) enter the stance–fact alignment index, the two elicited returns and political distance with fact fixed effects and row (6) is column (3) of Table 6, which adds ACTION and its interactions with both returns. Rows (3)–(6) report index coefficients; rows (1)–(2) report linear-probability coefficients. The two observation rows differ because the confirmatory specification drops observations with missing demographic controls (25 observations from five senders in the canonical column). Standard errors clustered at the sender level (in parentheses). $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.13: Leave-One-Fact-Out Robustness of the Motive Channels

Determinant	Sign-positive	Min β	Max β	Max p
Stance–fact alignment index	12/12	0.471	0.511	0.000
Expected accuracy gain	12/12	0.057	0.108	0.262
Persuasion	12/12	0.052	0.070	0.005
Political distance	0/12	-0.257	-0.245	0.000
Fact FE		Yes		
Clustered SE		Sender		

Notes: The sample is Experiment 2's general-population senders. Each determinant is re-estimated twelve times in the four-determinant specification of column (1) of Table 6 (*Stance–fact alignment index*, *Expected accuracy gain*, *Persuasion*, *Political distance*), each time dropping one of the twelve facts; fact fixed effects, sender-clustered standard errors. Coefficients are *probit index* coefficients. Columns report the number of leave-one-out estimates with a positive coefficient, the minimum and maximum coefficient across the twelve runs, and the largest p-value observed.

Table C.14: Heterogeneity by Sender Profile: Positional Alignment and ACTION

	Probability of Fact Provision					
	(1)	(2)	(3)	(4)	(5)	(6)
Moderator	Distance	Polarization	Pol. group	Issue extremity		
Topic	Pooled	Pooled	Pooled	Pooled	Immigration	Climate
Positional alignment	0.119***	0.115***	0.124***	0.101**	0.119***	0.143***
	(0.019)	(0.019)	(0.021)	(0.041)	(0.017)	(0.020)
ACTION	−0.006	0.040**	0.004	0.008	0.026	0.026
	(0.017)	(0.018)	(0.018)	(0.024)	(0.018)	(0.017)
Positional alignment × ACTION	−0.018	0.002	0.008	−0.011	0.050**	−0.042
	(0.027)	(0.026)	(0.028)	(0.057)	(0.022)	(0.029)
Receiver distance	−0.094***					
	(0.007)					
Positional alignment × Distance	0.001					
	(0.008)					
ACTION × Distance	0.019**					
	(0.009)					
Positional alignment × ACTION × Distance	0.020*					
	(0.012)					
Sender polarization		0.010*				
		(0.006)				
Positional alignment × Polarization		0.007				
		(0.009)				
ACTION × Polarization		−0.010				
		(0.009)				
Positional alignment × ACTION × Polarization		0.007				
		(0.012)				
Extremes (ref: center)			−0.045**			
			(0.020)			
Others (ref: center)			−0.093***			
			(0.026)			
Positional alignment × Extremes			0.001			
			(0.028)			
Positional alignment × Others			0.024			
			(0.040)			
ACTION × Extremes			0.024			
			(0.027)			
ACTION × Others			0.067*			
			(0.037)			
Positional alignment × ACTION × Extremes			0.016			
			(0.038)			
Positional alignment × ACTION × Others			−0.007			
			(0.052)			
Issue extremity				−0.018		
				(0.012)		
Positional alignment × Issue extremity				0.016		
				(0.023)		
ACTION × Issue extremity				0.014		
				(0.016)		
Positional alignment × ACTION × Issue extremity				0.015		
				(0.032)		
Clustered SE	Sender	Sender	Sender	Sender	Sender	Sender
Fact FE	Yes	Yes	Yes	Yes	Yes	Yes
Moderator	Distance	Polarization	Pol. group	Issue extremity	—	—
Topic	Pooled	Pooled	Pooled	Pooled	Immigration	Climate
Spec	(1)	(2)	(3)	(4)	(5)	(6)
Num. obs.	7740	7740	7715	7740	4045	3695
R^2	0.079	0.033	0.035	0.033	0.035	0.030
Adj. R^2	0.077	0.030	0.032	0.030	0.033	0.028

Notes: The sample is the $N = 1{,}548$ Experiment 2 general-population senders (Bilendi panel, Germany, 2026). Columns (1), (2) and (4) use 7,740 sender-fact observations. Column (3) uses 7,715 (non-missing political group). Columns (5)–(6) are topic subsamples (immigration: $N = 809$ senders; climate: $N = 739$ senders). The outcome is a binary provision indicator. Positional alignment is $\text{Align}_{ij} = \text{Pref}_i \times \text{Dir}_j \in [-1, 1]$ (pre-registered confirmatory measure). Receiver distance is $|\text{sender_likert} - \text{recv_stance}| \in \{0, 1, 2, 3, 4\}$. Sender polarization is the absolute left-right distance from center. For political group, Extremes = AfD, Linke, BSW; Others = other/no vote; Center = CDU/CSU, SPD, Grüne, FDP (reference). Issue extremity is the distance of the sender's stated position from the scale midpoint, $|\text{Stance}_i - 3| \in \{0, 1, 2\}$. Columns (1)–(4) include the moderator and all interactions; columns (5)–(6) are subsample estimates with no moderator. All specifications include fact fixed effects and sender-clustered standard errors. For multiple testing, the heterogeneity family spans twenty-seven interaction terms across columns (1)–(4) of this table and Table C.15, with Benjamini–Hochberg adjustment. The topic subsamples in columns (5)–(6) are outside it. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table C.15: Heterogeneity by Sender Profile: Motive Channels

	Probability of Fact Provision			
	(1) Distance	(2) Polarization	(3) Political group	(4) Issue extremity
Stance–fact alignment index	0.525***	0.429***	0.465***	0.042
	(0.083)	(0.074)	(0.080)	(0.125)
Persuasion	0.045	−0.010	−0.018	0.054
	(0.045)	(0.033)	(0.037)	(0.064)
ACTION	−0.017	0.102**	0.002	0.025
	(0.055)	(0.050)	(0.054)	(0.068)
Persuasion × ACTION	−0.048	0.072	0.136***	−0.050
	(0.062)	(0.048)	(0.052)	(0.091)
Receiver distance	−0.265***	−0.251***	−0.249***	−0.256***
	(0.025)	(0.013)	(0.013)	(0.014)
Stance–fact alignment index × Distance	−0.023			
	(0.042)			
Persuasion × Distance	−0.005			
	(0.019)			
ACTION × Distance	0.046*			
	(0.027)			
Persuasion × ACTION × Distance	0.049*			
	(0.026)			
Sender polarization		0.023		
		(0.023)		
Stance–fact alignment index × Polarization		0.041		
		(0.036)		
Persuasion × Polarization		0.028*		
		(0.017)		
ACTION × Polarization		−0.025		
		(0.025)		
Persuasion × ACTION × Polarization		−0.010		
		(0.023)		
Extremes (ref: center)			−0.156**	
			(0.072)	
Others (ref: center)			−0.274***	
			(0.087)	
Stance–fact alignment index × Extremes			0.087	
			(0.117)	
Stance–fact alignment index × Others			0.030	
			(0.142)	
Persuasion × Extremes			0.063	
			(0.053)	
Persuasion × Others			0.111*	
			(0.067)	
ACTION × Extremes			0.076	
			(0.078)	
ACTION × Others			0.182*	
			(0.103)	
Persuasion × ACTION × Extremes			−0.104	
			(0.075)	
Persuasion × ACTION × Others			−0.179*	
			(0.103)	
Issue extremity				−0.120**
				(0.050)
Stance–fact alignment index × Issue extremity				0.357***
				(0.099)
Persuasion × Issue extremity				−0.018
				(0.038)
ACTION × Issue extremity				0.033
				(0.047)
Persuasion × ACTION × Issue extremity				0.073
				(0.055)
Clustered SE	Sender	Sender	Sender	Sender
Fact FE	Yes	Yes	Yes	Yes
Moderator	Distance	Polarization	Political group	Issue extremity
Spec	(1)	(2)	(3)	(4)
Num. obs.	7740	7740	7715	7740
Deviance	9580.042	9575.748	9531.625	9570.204
Log Likelihood	−4790.021	−4787.874	−4765.812	−4785.102
Adj. pseudo R^2	0.050	0.050	0.050	0.050

Notes: The sample is the $N = 1{,}548$ Experiment 2 general-population senders (Bilendi panel, Germany, 2026); 7,740 sender-fact observations in columns (1), (2) and (4), 7,715 in column (3) (non-missing political group). The estimator is probit; coefficients are index coefficients from the four-determinant specification of Table 6, extended with the moderator and its interactions. The outcome is a binary provision indicator. The stance–fact alignment index is the cross-domain behavioral rating α_{ij}, assigned from validation judgments elicited post-provision on a different topic. Persuasion is the incentivized expected preference shift $\Delta\pi_{ijk}$ toward the sender's position. Receiver distance is $|\text{sender_likert} - \text{recv_stance}| \in \{0, 1, 2, 3, 4\}$. Sender polarization is the absolute left-right distance from center. Political group has Extremes = AfD, Linke, BSW; Others = other/no vote; Center = CDU/CSU, SPD, Grüne, FDP (reference). Issue extremity is the distance of the sender's stated position from the scale midpoint, $|\text{Stance}_i - 3| \in \{0, 1, 2\}$. All specifications include fact fixed effects and sender-clustered standard errors. For multiple testing, the heterogeneity family spans twenty-seven interaction terms across this table and columns (1)–(4) of Table C.14, with Benjamini–Hochberg adjustment. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

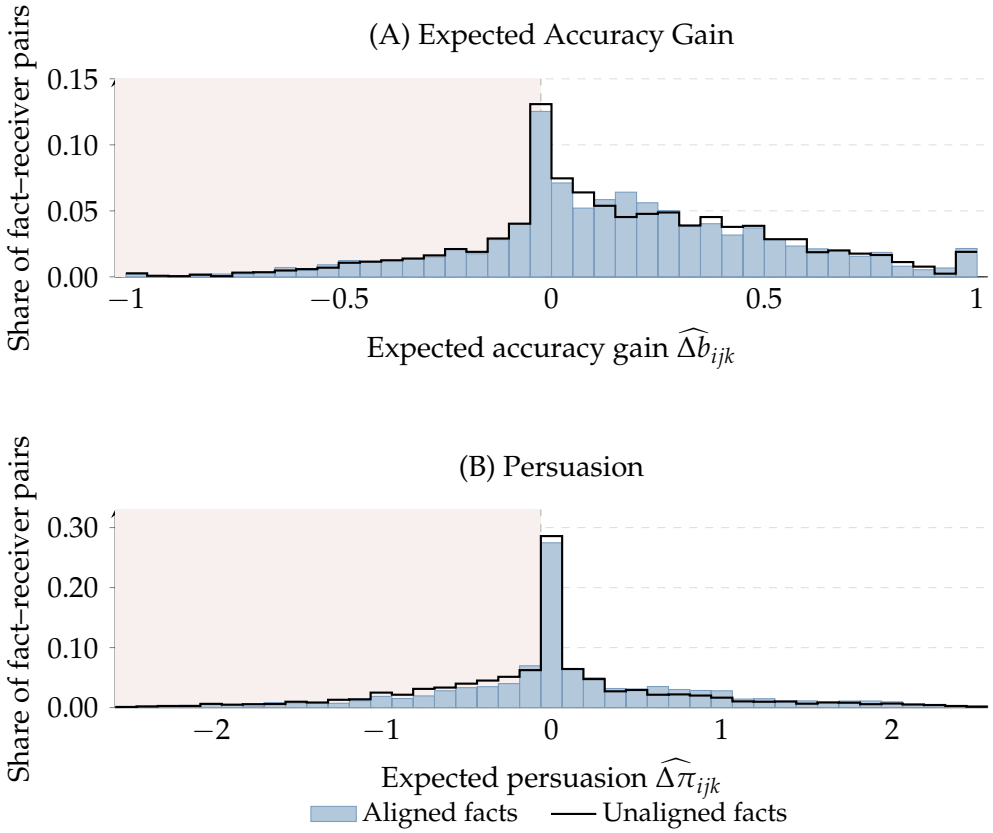


Figure C.2: Elicited Returns to Disclosure

Notes: The sample is Experiment 2's general-population senders. The figure shows distributions of the two incentivized elicitations for positionally aligned facts ($\text{Align}_{ij} > 0$; bars) and unaligned facts ($\text{Align}_{ij} < 0$; black outline), as shares within each group; observations of senders with $\text{Pref}_i = 0$ (260 of 7,740) are omitted from the split. Panel (A) shows expected accuracy gain $\widehat{\Delta b}_{ijk}$. Panel (B) shows expected persuasion $\widehat{\Delta\pi}_{ijk}$. Shaded regions mark negative values (disclosure expected to move the receiver's factual belief away from the truth, respectively the receiver away from the sender's position). Elicitations are incentivized via Quadratic Scoring Rule.

D Experiment 3: Supplementary Results

D.1 Sample Composition

Table D.1 compares the demographic and occupational characteristics of the Experiment 3 journalist sample to Loosen et al. (2023). The post-questionnaire collected gender, age, education, medium type, conditional broadcaster affiliation, geographic reach, role, and editorial autonomy.

Figure D.1 reports the stated and revealed policy view distributions of the journalist sample, by topic.

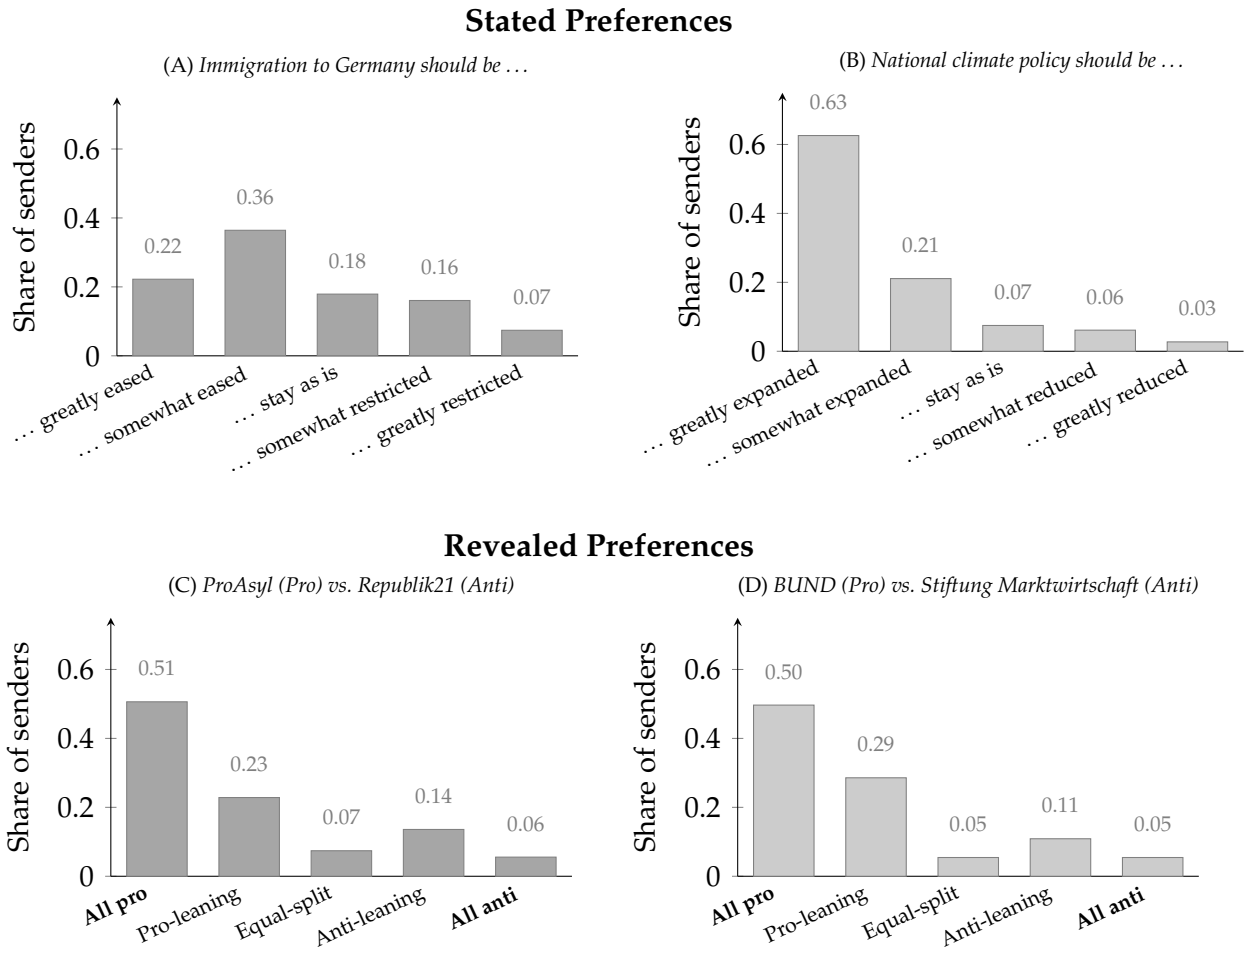


Figure D.1: Distribution of Policy Preferences Among Journalists

Notes: The sample is the Experiment 3 journalists ($N = 309$). The top row (Panels A–B) shows stated preferences from the 5-point Likert item; tick labels complete the stance statement shown in the panel title (e.g., *Immigration to Germany should be … greatly eased*). The bottom row (Panels C–D) shows revealed preferences from the real donation task, grouped into five buckets from All pro (€250 to the pro-policy organization) to All anti (€0), displayed pro-to-anti left to right. The left column shows immigration senders ($N = 162$, no non-response). The right column shows climate senders ($N = 147$, no non-response).

Table D.1: Journalist Sample Characteristics: Comparison to Population Benchmark

	Experiment 3 ($N = 309$)	Loosen et al. (2023) ($N = 1{,}221$)
Gender		
Female	53.1	44.0
Male	45.0	55.9
Age		
Under 50	76.7	62.4
Education		
University degree (B.A.+)	84.7	67.1
Medium type		
Print	28.2	57.0
Online	45.8	3.1
Television	7.1	17.3
Radio/podcast	10.7	17.2
News agency	3.6	3.5
Broadcaster type (TV/radio respondents only)		
Public broadcaster	80.0	61.4
Role (multiple selection; shares exceed 100 %)		
Editor	57.5	—
Editorial autonomy		
High ($\geq 4/5$)	76.3	—
Political affiliation		
Grüne or Linke	58.1	—
Mean left-right self-placement (0–10)	3.2	—

Notes: The Experiment 3 sample is $N = 309$ professional German journalists recruited via DJV and DJU mailing lists (convenience sample). The benchmark is Loosen et al. (2023), a stratified random sample of active German journalists ($N = 1{,}221$, September 2022–February 2023, HBI Hamburg; Loosen et al.). All figures are percentages. For medium type, Experiment 3 uses a single online category for all digital-native outlets. The Loosen et al. benchmark reports native-online journalists (3.1%) as a distinct population segment; print aggregates newspapers and magazines (57.0%). Broadcaster type is conditional on TV or radio respondents only. Role allows multiple selection; Loosen et al. (2023) report no breakdown by editorial position, so this row appears without a benchmark. Editorial autonomy has no common benchmark: Experiment 3 asks how often the respondent decides which pieces of information appear in a piece (1 = never or almost never to 5 = very often), while Loosen et al. (2023) report freedom in topic selection, so this row appears without a benchmark. For age, Loosen et al. (2023) report a population mean of 45.3 years, Experiment 3 elicited age as a categorical variable. Political affiliation and left-right self-placement were not elicited in Loosen et al. (2023) and appear without a benchmark.

D.2 Targeting

Let $\text{dirR}_{ij} = \text{Dir}_j \times (\text{Stance}_{ij}^{R} - 3)$ be the persuasive room the random assignment of facts to receivers creates, positive when the fact cuts against the receiver's position. Per unit, it raises the expected persuasion return by 0.100 among journalists and 0.015 in the general population's ACTION treatment. With sender and fact fixed effects and a control for political distance, journalists' provision rises by 2.0 pp per unit in immi-

Moderator	Interaction of moderator with:				Obs.
	Align_{ij}	α_{ij}	Δb_{ijk}	$\Delta \pi_{ijk}$	
Distance $\lvert\text{Stance}_i - \text{Stance}_k\rvert$	0.027^{***}	0.043	0.018	0.004	1545
	(0.009)	(0.103)	(0.098)	(0.039)	
Polarization	0.017^{*}	0.086	-0.105	0.025	1540
	(0.010)	(0.086)	(0.096)	(0.037)	
Issue extremity	-0.039	0.085	0.110	0.051	1545
	(0.029)	(0.251)	(0.205)	(0.085)	
Online medium	-0.041	-0.100	-0.016	-0.014	1540
	(0.029)	(0.247)	(0.283)	(0.096)	
Public broadcaster	-0.140^{*}	-0.984^{*}	0.612	-0.147	275
	(0.074)	(0.591)	(0.776)	(0.266)	
National/international reach	-0.054^{*}	-0.248	0.183	0.120	1545
	(0.030)	(0.248)	(0.272)	(0.100)	
Editor role	-0.019	0.159	-0.205	0.063	1545
	(0.030)	(0.253)	(0.272)	(0.096)	
Editorial autonomy	-0.017	-0.150	0.333^{**}	0.062^{*}	1540
	(0.017)	(0.147)	(0.141)	(0.036)	

Notes: The sample is the Experiment 3 journalists. Column (1) is a linear probability model of positional alignment and its interactions. The remaining columns report probit index coefficients of the three-return specification, which are not probability-scale effects (Corollary 1). The dependent variable is the binary provision decision. Each row reports interaction coefficients between one moderator and the selection measures; columns (2)–(4) are from a single regression interacting the moderator with the stance–fact alignment index (α_{ij}), the expected accuracy gain (Δb_{ijk}), and Persuasion ($\Delta \pi_{ijk}$) jointly. All specifications include fact fixed effects and journalist-clustered standard errors. Moderators are centered at the estimation-sample mean. The moderators are sender–receiver distance, polarization ($\lvert$left-right $- 5\rvert$), issue extremity ($\lvert$stance $- 3\rvert$), and the pre-registered journalist background variables, coarsened as online medium (vs. all other medium types), public broadcaster (defined only for television and radio senders), national or international reach (vs. local or regional), any editor role in the multi-select role item, and editorial autonomy (Likert 1–5). Benjamini–Hochberg correction is applied within the pre-registered positional-alignment estimates (column 1; 8 tests) and the channel estimates (columns 2–4; 24 tests). Stars mark unadjusted p-values: $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

Table D.2: Heterogeneity of Selection on Positional Alignment and on the Motive Channels Among Journalists

gration and 2.4 pp in climate (one-sided $p = 0.030$ and 0.020). The general population's provision shows no detectable response ($p = 0.942$). Pooled with the general population's ACTION treatment, journalists' provision rises by 2.2 pp per unit more ($p = 0.025$).

D.3 Heterogeneity and Coefficient Ratios

Table D.2 reports the pre-registered exploratory heterogeneity among journalists. Table D.3 reports, within each sample, the ratio of each determinant's index coefficient to the stance–fact alignment index coefficient, and its difference between journalists and the general population's ACTION treatment.

Table D.3: Ratios to the Stance–Fact Alignment Index Coefficient: Journalists and the General Population

	Ratio to the stance–fact alignment index coefficient		
	Journalists	General population (ACTION treatment)	Difference (1) − (2)
	(1)	(2)	(3)
Persuasion	0.295	0.171	0.124
	[0.068, 0.823]	[0.069, 0.291]	[-0.132, 0.629]
Expected accuracy gain	0.641	0.094	0.547
	[0.073, 1.851]	[-0.150, 0.366]	[-0.074, 1.766]
Political distance	−0.469	−0.441	−0.028
	[-1.126, -0.250]	[-0.614, -0.322]	[-0.685, 0.251]

Notes: Each entry is the ratio of a determinant's probit index coefficient to the stance–fact alignment index coefficient within a sample, from the models of Table 8, Panel A; each sample's latent scale cancels within the ratio. Column (3) is the difference between samples. Brackets are 95% percentile intervals from 1,000 sender-cluster bootstrap replications, senders resampled within each sample. Comparing ratios across samples maintains equal shock shapes, that the imported index measures journalists' alignment, and equal arrival probabilities. Because journalists are matched with certainty, equal weights would make their correction and persuasion ratios larger.

E Experiment 4: Supplementary Results

E.1 Design and Sample Details

Figure E.1 reports receivers' stated policy views on their matched topic, elicited before matching.

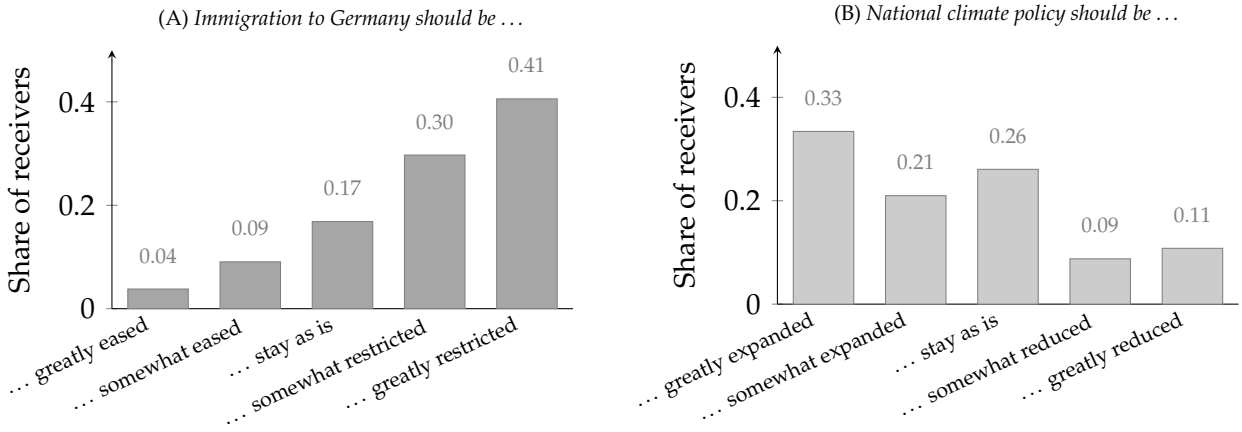


Figure E.1: Stated Policy Views Among Receivers

Notes: The sample is the Experiment 4 receivers (N = 1,043). Each receiver reports a stance on the matched topic on the same five-point item as the senders in Figure C.1, elicited before matching; tick labels complete the stance statement shown in the panel title. The left panel shows receivers matched on immigration (N = 552) and the right panel receivers matched on climate (N = 491).

Belief elicitation layout. The two answer words appear on the same side of the slider throughout for a given participant, and which side is randomized across participants.

Table E.1: Sample Characteristics

	Population (1)	Sample (2)
Gender		
Female	0.499	0.492
Male	0.501	0.508
Diverse	—	0.000
Age band		
18–29	0.179	0.174
30–39	0.178	0.178
40–49	0.170	0.167
50–59	0.193	0.219
60–75	0.279	0.262
Bundesland		
Baden-Württemberg	0.134	0.136
Bayern	0.159	0.158
Berlin	0.045	0.041
Brandenburg	0.030	0.031
Bremen	0.008	0.010
Hamburg	0.023	0.021
Hessen	0.075	0.073
Mecklenburg-Vorpommern	0.019	0.022
Niedersachsen	0.095	0.081
Nordrhein-Westfalen	0.216	0.233
Rheinland-Pfalz	0.050	0.052
Saarland	0.012	0.011
Sachsen	0.047	0.053
Sachsen-Anhalt	0.025	0.019
Schleswig-Holstein	0.035	0.031
Thüringen	0.025	0.030
School-leaving qualification		
Hauptschulabschluss or none	0.297	0.308
Mittlerer Abschluss (Realschule)	0.314	0.269
Fachhochschul-/Hochschulreife or higher	0.390	0.423

Notes: The sample is the $N = 1{,}043$ Experiment 4 receivers, German adults recruited on Bilendi (quota-sampled). Column (1) reports German population proportions, as in Table C.11. For gender, age band, and Bundesland, the benchmark is the population aged 18–75 as of 31 December 2024 (Statistisches Bundesamt, Statistischer Bericht 12411). For school-leaving qualification, the benchmark is the Mikrozensus 2023 (Statistisches Bundesamt, Bildungsstand der Bevölkerung), persons aged 15 and over excluding those still in school. Column (2) reports sample proportions.

Exclusions. 1,081 participants completed the survey. 38 are excluded as speeders under the filed rule (completion time below a quarter of the median). The filed bot exclusion removed no one.

Attrition. Completion differs across reception states: 86.1 percent after receipt, 76.0 percent after failed transmission and 80.5 percent after withholding. Among completers, receipt and failed transmission are balanced on prior loss, initial policy view, alignment with the sender and demographics (joint $p = 0.778$). In the withholding state, completion falls with alignment with the sender ($p = 0.017$) and rises with prior loss ($p = 0.032$). Table E.2 reports trimming bounds for the primary specification.

Table E.2: Trimming Bounds for Differential Completion

Contrast	State trimmed	Trimmed (%)	Estimate	Bounds, trimming ranked on posterior loss	Bounds, trimming ranked on residual loss
Failed transmission vs. receipt (β_F)	R	11.7	0.126	[0.108, 0.186]	[0.063, 0.187]
Withholding vs. failed transmission ($\beta_W - \beta_F$)	W	5.6	0.001	[−0.028, 0.015]	[−0.035, 0.030]
Withholding vs. receipt (β_W)	R	6.5	0.128	[0.119, 0.164]	[0.089, 0.166]

Notes: The analysis sample is Experiment 4's receivers, $N = 1{,}043$. Each row refits equation (7) after trimming the reception state with the higher completion rate, from the top and from the bottom of the ranked outcome, by the completion gap divided by that state's completion rate. Completion is measured among receivers who reached the transmission draw: receipt 373/433, failed transmission 381/501, withholding 327/406. Residual loss is posterior loss net of prior loss and fact fixed effects. Receipt and failed transmission are randomly assigned among receivers whose sender disclosed, so the first row follows Lee (2009) under monotonicity: no receiver who completes after failed transmission would have left after receipt. Because each trimmed sample refits the covariate-adjusted model, the intervals are a trimming sensitivity rather than Lee's sharp bounds. Withholding is not randomly assigned, so the second and third rows are sensitivity ranges. Trimming shares use the completer counts before the filed speeder exclusion.

Robustness. Dropping each of the twelve facts in turn, β_W ranges from 0.110 to 0.135 and β_F from 0.116 to 0.134 (all $p < 0.001$).

Perceived selectivity. The actual disclosure rates are sender-fact-receiver pairing shares and carry binomial sampling error, which attenuates the slope of perceived on actual selectivity. Sampling variance accounts for 22.5 percent of the variance of actual selectivity, so the reliability-corrected slope is $0.066/0.775 = 0.085$. I impute σ_0 as the mean disclosure rate of the opposite-direction facts of the same topic at the same sender and receiver stances, which assumes that disclosure depends on a fact only through its direction relative to the sender's stance. The rule was fixed before the final data were seen, and using any one of these facts instead gives similar slopes of 0.050 to 0.069 and wrong-direction shares of 35.1 to 36.0 percent. Table E.3 reports the pre-registered secondary analyses.

E.2 Sender Forecasts and Receiver Outcomes

I compare senders' forecasts with matched receivers' outcomes. Matched to the receiver they decided about, general-population senders forecast that receipt raises the share answering correctly by 0.158, against a realized 0.217 for receipt against non-receipt, and journalists 0.294, against 0.289. Forecasts of opinion movement are too large in the general population: a shift toward the fact's side of 0.128 scale points against a realized 0.001. Among journalists it is 0.232 against 0.235. Forecasts concern

Table E.3: Updating after Non-Receipt and Moderation by the Fact's Support for the Receiver's View

	Coefficient	SE	p (= 0)	p (= 1)	N
Panel A. Posterior log-odds after non-receipt (Bayes: both coefficients equal one)					
Prior log-odds	0.585	(0.046)	< 0.001	< 0.001	685
Log-likelihood ratio at stated rates	-1.529	(0.387)	< 0.001	< 0.001	685
Panel B. Interaction with the fact's support for the receiver's view					
Belief change after non-receipt	-1.459	(5.825)	0.802		1,043
Perceived selectivity	0.060	(0.037)	0.100		1,043
Policy view after withholding	-0.249	(0.182)	0.170		1,043
Policy view after failed transmission	-0.014	(0.147)	0.924		1,043

Notes: The sample is Experiment 4's receivers. Panel A covers non-receipt states only; it regresses the logit of the posterior on the correct answer on the logit of the prior and the log-likelihood ratio of non-receipt implied by the receiver's own stated disclosure rates and a transmission probability of one half. Beliefs are clamped to $[0.005, 0.995]$ (165 observations). Clamping at 0.01 or 0.001, or dropping the 165 observations, leaves the coefficient on the log-likelihood ratio between -1.81 and -1.01 (all $p < 0.001$ against both zero and one). Bayesian updating at the stated rates implies both coefficients equal one. In Panel B, support is the cross-domain rating of how strongly the fact supports views at the receiver's stance. Belief change and policy view include fact fixed effects; the policy-view row controls for the initial view and demographics. HC3 standard errors.

stance groups, a third or more have no benchmark sender-fact-receiver pairing, and receivers not shown a fact include those whose sender withheld it, so the comparison is descriptive.

F LLM Classification of Stated Motives

F.1 Classification Procedure

I classify the open-ended decision-reason responses of Experiment 1 (1,200 senders, in English). A first-step gate asks whether the response refers to the sharing decision, the receiver, or the information provision, and responses that do not refer to any of these receive no category coding. Categories are not mutually exclusive, and a label is assigned only where the text gives distinct support for it.

Experiment 1 is classified by Claude Sonnet 5, the primary classification, and by GPT-5.6 (`gpt-5.6-terra`) as robustness, under one fixed prompt.

F.2 Validation

For Experiment 1, the two models agree on 91 percent of gate decisions ($\kappa = 0.81$), 94 percent of belief-correction codings ($\kappa = 0.88$) and 98 percent of persuasion codings ($\kappa = 0.73$). Expression is affirmed for twelve and ten responses, and the models agree on six. Under the second classifier every motive coefficient in that table keeps its sign and the persuasion interaction with ACTION remains significant at the five percent level (Appendix Table B.17). Agreement between two models applying one prompt cannot show that either coding is correct.

F.3 Coding Examples

Table F.1 presents representative responses in each category.

Table F.1: Coding Examples: U.S. Senders

Category	Status	Example
Belief correction	Affirmed	*"People should know the truth."* *"I wanted the people who got the questions wrong to know the correct answer. People should be more knowledgeable about things."*
Expression	Affirmed	*"When I thought that the fact was unfavorable towards immigration, I chose to not share the fact."* *"I wanted to show a more positive view of immigrants."*
	Mentioned but rejected	*"Everyone should be informed about facts even if it goes against my beliefs or hiding facts to sway decisions. Nothing but the facts."* *"I think everyone deserves to know the truth and be informed, even if the facts dont align with my personal preferences."*
Persuasion	Affirmed	*"If the question had a direct impact on the way people view immigrants."* *"People who are ill informed, when they are enlightened, may choose to reallocate more funds towards helping immigrants if they know the truth."*
	Mentioned but rejected	*"Everyone should know the truth regardless of how it influences their decisions on immigration."* *"Wanting to be honest instead of manipulative."*
Other stated reason	Affirmed	*"Basically only on whether or not I personally would want/ care to know. For a couple of them, I WOULD want to know. And, for the other two, not as much."*
Referred, no stated reason	Gate outcome	*"Basic knowledge and provided information."* *"The information that was provided at the end."*
Different task only	Gate outcome	*"I made decisions based on what I truly thought was the correct answer and was surprised to find I was mistaken on a few."*
Did not refer to disclosure	Gate outcome	*"My own beliefs and also taking into consideration the information from research."* *"I think that immigration is a valuable thing for the US. I think immigrants help the economy and do important jobs US citizens do not want to do."*

Open-ended responses from Experiment 1 senders, illustrating the boundaries of the coding scheme. Every example shown is coded identically by both classifiers. *Affirmed* means the response gives distinct support for the category; *mentioned but rejected* means the response raises the consideration and states that it did not govern the decision—the two expression examples are senders who name their own position and say they disclosed regardless of it. The last three rows are gate outcomes and receive no category coding: a response may refer to the disclosure decision without giving a reason for it, may describe only a different task in the experiment, or may state a political opinion with no link to the decision.

G Stated Norms and Perceived Design Balance

Most participants perceived no political imbalance in the study: 85.5 percent of the general population and 82.2 percent of journalists (Table G.1, Panel A). These shares

Table G.1: Perceived Design Balance and Normative Items

	General population (Exp. 2)			Journalists (Exp. 3)		
	Mean	SD	% Agree	Mean	SD	% Agree
Panel A: Perceived political imbalance (% of valid answers)						
No imbalance perceived		85.5			82.2	
Imbalance favoring the left		9.5			12.0	
Imbalance favoring the right		5.0			5.8	
Panel B: Normative and policy items (1–5 agreement)						
Share opposing information	4.13	0.90	80.0	4.61	0.67	95.5
Social-media transparency	4.05	0.96	75.7	4.61	0.65	95.1
Viewpoint-diversity obligation	4.17	0.90	80.4	4.21	0.91	82.5

Notes: The data are post-questionnaire self-reports, one per sender, from Experiments 2 and 3. Panel A reports answers to "Did you perceive a political imbalance in the study?", in percent of valid answers; Experiment 2 $N = 1{,}548$, Experiment 3 $N = 309$. Panel B reports agreement on a 1–5 scale (1 = fully disagree, 5 = fully agree) with three items: whoever passes on political information should also pass on factually correct information that contradicts their own position (worded for journalists in Experiment 3); social media should be more transparent about why users see some political information and not others; platforms and media should be obligated to surface relevant information from different perspectives. % Agree is the share answering 4 or 5. Experiment 2 $N = 1{,}548$; Experiment 3 $N = 308$.

barely differ between the treatments. In Experiment 2, 85.7 percent of CONTROL and 85.2 percent of ACTION senders perceived no imbalance ($p = 0.29$), and 1.1 and 0.8 percent named the research question ($p = 0.60$). In Experiment 1, with INFO held off, 89.3 and 85.3 percent perceived no imbalance ($p = 0.16$), and 1.7 and 4.7 percent named it ($p = 0.059$). Without these senders, the ACTION increase in the alignment gap is 5.6 pp (one-sided $p = 0.027$). Panel B reports the three normative items, and Table G.2 their interactions with the selection measures. Among journalists, the positional alignment slope falls from 0.180 among those who agree that one should pass on true information contradicting one's own position to 0.090 among those who strongly agree.

Two research assistants independently coded each participant's open-ended answer on what the study was about as a correct guess if it named the selective passing on of information depending on the sender's own position. Table G.3 reports the shares. The assistants agree on 96 to 99 percent of answers in Experiments 1, 2 and 4 and on 78 percent among journalists. In Experiment 4, 34 receivers' answers were exported before coding; counting all of them as correct guesses would raise the share to at most 4 percent.

Experimenter demand and social desirability most plausibly work against these findings. Participants could expect a study of information sharing to reward impartiality, so both pressures push toward passing on facts regardless of one's own position (de Quidt et al. 2018, Bursztyn et al. 2025). Between 82.2 and 86.6 percent of senders perceived no political imbalance in the study, and few identified its research question: 0.9 to 3.2 percent of general-population senders and 18.1 percent of journalists when

Table G.2: Normative Items and Selection: Channel × Item Interactions, General Population and Journalists

	Interaction of item with:			
Item	Align_{ij}	α_{ij}	$\Delta\pi_{ijk}$	Obs.
Panel A: General population (Experiment 2)				
Share opposing information	0.006	0.027	0.006	7740
	(0.010)	(0.021)	(0.007)	
Social-media transparency	0.008	0.003	0.005	7740
	(0.009)	(0.019)	(0.006)	
Viewpoint-diversity obligation	0.017*	0.054***	0.010	7740
	(0.009)	(0.020)	(0.007)	
Panel B: Journalists (Experiment 3)				
Share opposing information	−0.090***	−0.110*	−0.040	1540
	(0.021)	(0.058)	(0.026)	
Social-media transparency	−0.054**	−0.009	−0.050**	1540
	(0.024)	(0.047)	(0.024)	
Viewpoint-diversity obligation	0.002	0.026	−0.011	1540
	(0.022)	(0.042)	(0.016)	

Notes: The samples are the Experiment 2 general population and the Experiment 3 journalists. The dependent variable is the binary provision decision. Each cell reports the interaction coefficient between the row item (centered at 3, the scale midpoint) and one selection measure, from a separate regression of provision on the measure, the item, and their interaction, with fact fixed effects and sender-clustered standard errors (in parentheses). The selection measures are positional alignment (Align_{ij}), the stance–fact alignment index (α_{ij}), and Persuasion ($\Delta\pi_{ijk}$). Items as in Table G.1, Panel B. Stars mark unadjusted p-values: $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.1$.

both research assistants agree (Table G.3).

Table G.3: Perceived Bias and Guessed Purpose of the Study

	Exp. 1 (U.S. senders)	Exp. 2 (German senders)	Exp. 3 (journalists)	Exp. 4 (receivers)
Perceived no political bias in the study	86.6	85.5	82.2	87.5
Correct guess, both research assistants	3.2	0.9	18.1	0.2
Correct guess, at least one research assistant	7.1	1.9	39.8	0.8
Agreement between research assistants	96.2	99.0	78.2	99.4
Participants	1,200	1,548	309	1,043

Notes: Shares are in percent. "Perceived no political bias" is the answer "no imbalance" to the closing question whether the study was politically imbalanced. The guessed purpose is coded from the open-ended closing question on what the study was about by two research assistants independently, from the text alone. An answer is a correct guess if it names the selective passing on of information depending on the sender's own position; in Experiment 1, the category "selective information provision" of a six-category scheme. In Experiment 4, the guessed purpose is coded for 1,009 of the 1,043 receivers, whose answers were exported during fielding.

H Experimental Instructions

Instructions of each experiment are shown in abbreviated form. Repeated screens are shown once, and Experiments 2–4 are translated from German.

H.1 Experiment 1

Participants saw the pages below in this order. The four treatments cross CONTROL/ACTION with INFO off/on; the treatment changes only Page 17 (shown with INFO only) and the wording on Pages 23–27. The order of the four statements, of the two organizations and of the four persons is randomized. Participants who refuse consent, fail the captcha, exceed a time limit or switch browser tabs more than twice are routed to an exit page.

Page 1: Informed consent. Standard informed-consent text [...]. Substantive sentences: the study is run by researchers at the Max Planck Institute for Research on Collective Goods; participants "read instructions, answer questions, and make decisions that affect your own and other participants' payment"; they receive $2.50 for completing the study, "any part may be relevant for determining your final payment" and they "should make decisions according to your genuine views and preferences"; participants are "expected to dedicate your full attention to the study" and "must complete comprehension checks, attention checks, and respond within the time limits shown at the top of certain pages"; failing these checks or exceeding a time limit terminates the study without payment. Response: "I consent" / "I do not consent".

Page 2: Captcha. "A quick check to confirm you are human." Participants type the letters shown in an image and tick: "I hereby confirm that I am human. I am aware that using a bot can lead to exclusion from this and any further studies without payment."

Page 3: Schedule for today's study. Welcome! This study consists of three parts and will take about 15 minutes. You will receive $2.50 for participating in the study. Additionally, you can earn a bonus of $1.50 in this study. You will receive the bonus within seven days of completing the study. Your bonus will be based on one of the three parts of the study. The following table shows the probabilities: Part 1, 45%; Part 2, 45%; Part 3, 10%.

Since each part may determine your bonus, it is in your best interest to always act as if that part will determine your bonus.

We expect you to dedicate your full attention to the study.

All information presented in this study is truthful.

Page 4: Instructions (Part 1). You are now starting Part 1 of the study!

Your Task. You'll see four sentences about immigration. For each one, do one of the following: select the word that best completes the sentence; or choose "I don't know" if you're unsure.

You have 30 seconds to answer each question. Please respond within the time limit to avoid being excluded from the survey.

Important: Please do not look up answers online or switch browser tabs. Doing so will result in exclusion from the survey.

Your Bonus Payment. In Part 1, each question is equally likely to be selected to determine your bonus payment. Therefore, you should treat each question as if it's the one that will determine your bonus payment. How your chance of winning the bonus is calculated: your probability of receiving the bonus payment of $1.50 is: if you are right, 100% chance; if you are wrong, 0% chance; if you say you don't know, 25% chance.

Pages 5–8: Prior-knowledge elicitation (four pages, one per statement, in randomized order). The elicitation and its payment rule are summarized in Appendix A.2; the page wording follows. Example (crime statement); the other three are indicated below. Page title: "Your assessment of the crime rate". Some people believe immigrants are more likely to commit violent crimes than U.S.-born Americans, while others believe they are less likely to do so. Please complete the sentence below so that it is true.

The crime rate of immigrants is *[higher / lower / I don't know]* than that of U.S.-born Americans.

Labor: "Some people believe immigrants are less likely to work than U.S.-born Americans, while others believe they are more likely to work." / "The labor participation rate of immigrants is *[higher / lower / I don't know]* than that of U.S.-born Americans." Border: "Some people believe the number of illegal border crossings is historically high, while others believe it is not." / "The number of unauthorized border crossings has been *[higher / lower / I don't know]* in the past five years compared to historical averages." Fiscal: "Some people believe immigrants contribute more in taxes than they use in government services, while others believe they use more in government services than they contribute in taxes." / "When comparing what first-generation immigrants use in government services versus what they pay in taxes, the net cost is *[higher / lower / I don't know]* than the same calculation for U.S.-born Americans at the state and local levels." [30-second limit per page; a selection is required.]

Page 9: Your Decision (donation). [Mobile users: "Please rotate your device to landscape mode to improve readability of this page."] The description text and both organizations' mission statements are reproduced in Appendix A.2. Further instruction text

on the page:

You will decide how to split $500 between two organizations with different positions on immigration.

All participants' decisions from the current series of surveys will be collected. One decision will be randomly selected and implemented with real money. At the end of the study, you will receive a link to a page with the donation receipt to verify the donation was placed. The page will be available for 6 months.

Since your decision might be the one selected, you should decide as if it will definitely determine how the $500 is allocated.

Please indicate below how you want the $500 to be split between these two organizations:

[Slider from "AIC" to "NumbersUSA", $0–$500, initial position $250, with the left–right order of the endpoints following the randomized display order. Confirmation screen: "Do you want to confirm your decison [sic]? If you continue, you cannot change it later." 6-minute limit.]

Page 10: The research findings. The next section presents current research findings on the issues you just considered. Please read the information carefully.

Pages 11–14 (four pages): Research findings (four pages, one per statement, in randomized order). The four statements, values and sources are reproduced in Appendix A.1 (paragraph "Sources"). Each page repeats the participant's own answer (e.g. "You indicated that the crime rate of immigrants is **higher** than that of U.S.-born Americans." / "You indicated that you were **unsure** about [...]") and then shows the headline sentence of the finding: *crime*: "Recent comprehensive research shows: Immigrants have consistently lower crime rates than U.S.-born Americans across multiple studies and time periods." *Labor*: "Federal labor statistics consistently show: Immigrants have higher labor force participation rates than U.S.-born Americans." *Border*: "Border monitoring data indicates: Unauthorized border crossings have reached historically high levels in recent years." *Fiscal*: "Comprehensive fiscal analysis shows: First-generation immigrants have higher net costs (they are more expensive to the taxpayer) at state and local levels." [The next button appears after 10 seconds.]

Page 15: Attention check. Title: "Your assessment of the most popular sport". The next question is about the following problem. In questionnaires like ours, sometimes there are subjects who do not carefully read the questions and just quickly click through the survey. This means that there are a lot of random answers which compromise the results of research studies. To show that you read our questions carefully, please choose

"Badminton" as your answer to the next question, even if you believe another answer is more accurate.

What do you think is the most popular sport among immigrants in the United States? (a) Soccer (b) Basketball (c) Baseball (d) American football (e) Tennis (f) Ice hockey (g) Cricket (h) Ultimate frisbee (i) Badminton (j) Lacrosse (k) Pickleball. [correct: Badminton]

Page 16: Instructions (Part 2) and the two predictions. The Haaland–Roth study description, the two incentivized predictions and their scoring are reproduced in Appendix A.3 (Stage 1). Page title "Instructions". *Your Task.* You will now be asked to estimate how the research information affected the attitudes and behavior of the participants in the study. Specifically, you will estimate:

1. Compared to participants who received no information, by what percentage did learning about the research findings increase or decrease signing rates for a White House petition on increased immigration?

2. Did people who initially believed (falsely) that immigration harmed U.S.-born workers change their attitudes towards immigrants more, less, or the same amount as people who already held the correct belief initially?

Important: Before answering, please do not look up answers online or switch browser tabs. Doing so will result in exclusion from the survey. A link to the study will be provided at the end.

Your Bonus Payment. In Part 2, both estimates are equally likely to be selected to determine your bonus payment. Therefore, you should try to be as accurate as possible on every single estimate. How your chance of winning the bonus is calculated: probability of receiving the bonus payment of $1.50: (1) If the first question is selected, you have a 100% chance of receiving the bonus if you guess the right percentage. Your chances decrease the more your answer deviates from the true percentage. Should you be off by more than 20 percentage points, you have a 0% chance of receiving the bonus. (2) If the second question is selected, you have a 100% chance of receiving the bonus if you are right and a 0% chance of receiving the bonus if you are wrong.

Question 1 label (verbatim): "By what percentage did learning about the research findings change people's petition signing rates?" with the example "If you think the information doubled the amount of people who signed the petition, i.e., caused a 100% increase in signing rates, enter 100%. If you think the information meant nobody signed the petition anymore, i.e., caused a 100% decrease enter -100%." (integer, -100 to 200). Question 2: "After learning about the research findings, who do you think changed their attitudes toward immigration more?" (a) People who initially believed immigra-

tion harmed U.S.-born workers (b) People who initially believed immigration didn't harm U.S.-born workers (c) Both groups changed about the same amount. [correct: a]

Page 17 (INFO only, treatments T3 and T4): The actual results from the study. The feedback screen is reproduced in Appendix A.3 (Stage 2). Not shown in INFO-off treatments (T1, T2). Additional page text: "Study Summary: Receiving information strongly changed behavior: Signature rates changed dramatically. Receiving information had the strongest effect on people who initially held false beliefs."

Page 18: Instructions (expected attitude change). Your Task. You will estimate how ten participants who hold false beliefs or do not know the answer would change their attitudes about immigration if they are provided with research findings. In particular, you will estimate how many of them ... become less supportive of immigration; don't change their opinion; become more supportive of immigration. These ten participants are illustrated below [ten person icons]. You will assign all 10 participants to one of the three categories. Please give careful consideration to each of your estimates. You should try to be as accurate as possible on every single estimate. (The task is also summarized in Appendix A.3, Stage 3.)

Pages 19–22: Expected attitude change (four pages, one per statement, in randomized order). Page titles: "Your first estimate", "Your second estimate", "Your third estimate", "Your fourth estimate". Example (crime): Consider ten participants who *[falsely believed that immigrants were more likely to commit crimes than U.S.-born Americans / did not know if immigrants were more or less likely to commit crimes than U.S.-born Americans]*. These participants then read these research findings: *[the finding box of Pages 11–14]*. How will these findings affect these ten participants attitudes on **immigration**? Assigned: *[n]* of 10 participants assigned. Responses (number of participants, each between 0 and 10, set with plus/minus buttons): Become less supportive of immigration; Don't change their opinion; Become more supportive of immigration. If the total is not 10, "You have to assign exactly 10 participants!". Confirmation screen as on Page 9 ("decison [sic]"). The false-belief/don't-know description is randomized per statement (the other statements use the belief descriptions of Pages 24–27 below).

Page 23: Instructions (Part 3). You are now starting Part 3 of the study!

Your Decisions. You will now make decisions about how much four other participants in this study will know. Here are illustrations of these four people: [four colored icons labeled Person A, B, C, D]. Each of these individuals is a real, distinct person!

The People for Whom You Decide. Persons A through D each completed one of the four immigration-related questions that you were asked earlier. They answered false

or said they did not know the answer. However, they do not know that their answer was not correct. You must decide for each of them whether they should be informed about the research findings. Persons A through D have no knowledge that you exist or that someone else is making decisions about them. From their perspective, when they receive the research findings, it will seem like a standard part of the study. If they do not receive the research findings, they have no idea that they were even available.

What Happens Next?

CONTROL: Each person will either receive or not receive information about the research findings. After this, they will not take any further action, and the study ends for them.

ACTION: Each person will either receive or not receive information about the research findings. After that, they will also allocate $500 between two immigration-related organizations, as you did earlier.

ACTION: Reminder: *[one-sentence description of the two organizations, in the participant's randomized order, as in the donation-page description text of Appendix A.2]*. One donation decision from this study will be selected and implemented with real money. At the end of the study, you will receive a link to a page with the donation receipt to verify the donation was placed. The page will be available for 6 months. Important: This means the person's decision will directly affect how likely it is that each organization receives a donation.

Your Bonus Payment. In Part 3, you will always receive a bonus payment of $1.50. Therefore, your decisions here do not affect your bonus payment in any way. Please make your decisions based on your genuine preferences.

For each of these four people, one participant from the current series of surveys will be selected to determine how much they know. Therefore, your decisions could be implemented. ACTION: Your decisions today may have real effects—on both real people and real-world advocacy efforts. CONTROL: Your decisions today may have real effects on real people. Please take your decisions seriously.

Pages 24–27: Information decision (shown 4 times, once per person A–D; treatment-specific wording in one sentence). Title: "Your decision regarding Person A" (B, C, D). Example (crime statement; the other statements differ only in the belief description): Person A *[falsely believed that immigrants were more likely to commit crimes than U.S.-born Americans / did not know if immigrants were more or less likely to commit crimes than U.S.-born Americans]*.

CONTROL: You can either provide Person A either only their own answer (Option A) or their own answer and the *research findings* (Option B).

ACTION: You can either provide Person A either only their own answer (Option A) or their own answer and the *research findings* (Option B) before their donation decision.

Click on the buttons for a preview! [Buttons "Option A" and "Option B" reveal a preview: Option A shows "You said that [...]" (the person's own answer, e.g. "You said that you did not know whether the crime rate of immigrants is higher or lower than that of U.S.-born Americans."); Option B shows the same plus the finding box of Pages 11–14.] Which feedback screen should Person A see? (a) Option A (their own answer) (b) Option B (their own answer + research findings). [Selection required. After the fourth decision the bonus is drawn.]

Page 28: Post-questionnaire (study). Items, in order: (1) What factors or considerations influenced your decisions during the study? [open text]. (2) Should people be corrected when they are wrong about facts? [slider 1–5, step 0.5, "Never corrected"–"Always corrected"]. (3) How important is it to know accurate facts when forming political opinions? [slider 1–5, "Not important"–"Extremely important"]. (4) Should social media companies fact-check information posted on their platforms? [slider 1–5, "Should not fact-check"–"Should fact-check"]. (5) When people disagree with me on important issues, it's usually because we have different: [slider 1–5, "Facts"–"Views"]. (6) Did you find the provided research findings believable? [Yes/No]. If the answer indicates no: "Please elaborate which research findings you did not find believable." Did you find the provided research findings on (a) the effect of information on White House petition signatures; (b) the crime rates of immigrants; (c) the number of unauthorized border crossings; (d) the labor participation rates of immigrants; (e) the net fiscal costs of new immigrants believable? [Yes/No each]. (7) Do you feel this survey was biased in any way? (a) Yes, left-wing biased (b) Yes, right-wing biased (c) No, unbiased. (8) Do you have any additional comments or feedback about this study? [open text, optional]. (9) What do you think the researchers were investigating in this study? [open text]. [The four sliders must be touched before the page can be submitted.]

Page 29: Post-questionnaire (demographics). "Please provide the following demographic information." Age [18–85]; Gender (Male / Female / Non-binary/ Third gender / Prefer not to say); Ethnicity (White / Black or African American / Hispanic or Latino / Asian or Asian American / American Indian or Alaska Native / Native Hawaiian or Other Pacific Islander / Other); Annual household income before taxes (Less than $25,000 / $25,000 to $49,999 / $50,000 to $74,999 / $75,000 to $99,999 / $100,000 to $149,999 / $150,000 or more / Prefer not to say); Highest level of education completed (Less than high school / High school graduate / Some college / Associate's degree / Bachelor's degree / Master's degree / Professional degree (JD, MD, etc.) / Doctoral degree); Current employment status (Employed full-time / Employed part-time / Self-employed / Unemployed / Student / Retired / Other); Marital status (Single / Married / Divorced / Other); ZIP code of your current residence [optional];

Generally speaking, do you usually think of yourself as: (Republican / Democrat / Independent / Other); On a scale from 1 (very liberal) to 5 (very conservative), where would you place yourself politically? (Very liberal / Somewhat liberal / Moderate / Somewhat conservative / Very conservative); Do you think the number of immigrants coming to the United States each year should be increased, reduced, or remain the same? (It should be increased a lot / increased a little / remain as it is / decreased a little / decreased a lot).

Page 30: End of the survey. Thank you for participating in our study! The experiment is now complete, and you have successfully completed all required tasks. You will receive $4.00 (if the bonus was received) or $2.50 in total for this study. A table lists the payout for participation ($2.50) and for Parts 1–3 ("Not selected for bonus" for the unselected parts), followed by a description of the selected part: in Part 1 the participant's answer on the drawn statement and whether it was correct (100% chance), "I don't know" (25%) or incorrect (0%); in Part 2 the drawn estimate and the resulting chance; in Part 3 "you received the bonus regardless of your choices"; and whether the $1.50 bonus was received. Participants are given a Prolific completion link, a link to the Haaland–Roth study, and a link to verify that a donation was placed. Exit pages for time-outs, tab switching, ineligibility and refused consent exist and pay nothing.

H.2 Experiments 2 and 3

The instrument was administered in German and the text below is translated. Experiment 3 differs only as listed in the final block. The worked example is the immigration domain; the climate domain is analogous, with the climate facts (Table 4) and the climate stance scale. Where text differs by treatment it is labelled CONTROL or ACTION. Text that varies with the fact, the receiver or the sender's earlier answers is given in square brackets. The twelve facts and their sources are reproduced in Tables 4 and A.4; the two organizations' descriptions and positions are reproduced in Appendix A.5. Participants who refuse consent, live outside Germany, or hit a full quota are routed out of the study by Bilendi screen-out pages.

Experiment 2

Page 1: Consent. Title: "Study information and consent". Standard informed-consent text [...]. Substantive statements: the study has four parts and a final questionnaire and takes about 20 minutes; participants have a chance of a bonus payment of €10; there is no deception: "All statements about other participants and about the procedures in this study are true." Response: (a) "Yes, I agree and wish to participate." (b) "No, I do not wish to participate."

Page 2: Personal information. Title: "Information about you". "Please first answer a few questions about yourself." Age (18–29, 30–39, 40–49, 50–59, 60–75 years); gender (male, female, diverse, no answer); highest educational attainment (no school-leaving qualification; Hauptschulabschluss; Mittlere Reife/Realschulabschluss; Fachhochschulreife; Abitur; Bachelor or equivalent; Master, Diplom, Magister or equivalent; doctorate or higher); federal state (the 16 Bundesländer, or "I do not live in Germany").

"Please state your personal view on the following topics."

- "Migration and asylum policy: Which of the following positions comes closest to your view?" (1) Immigration to Germany should be made much easier. (2) Immigration to Germany should be made somewhat easier. (3) Immigration to Germany should stay about as it is. (4) Immigration to Germany should be made somewhat harder. (5) Immigration to Germany should be made much harder.
- "Climate policy: Which of the following positions comes closest to your view?" Analogous five options (climate protection measures in Germany should be expanded considerably / somewhat / stay about as they are / be reduced somewhat / considerably).

The page contains no other items.

Conditional page: quota check. Shown after Page 2 only as an exit screen (duplicate participant, screened out, or quota full); otherwise participants pass through without seeing it.

Page 3: Study schedule. "This study consists of **three parts** and takes **about 20 minutes**." Title: "Study schedule". "Among all participants, **25 persons are randomly selected** who can receive a **bonus payment of €10**." "The bonus is paid **in addition** at the end of the whole series of studies." "If you are selected, it is determined at random which part of the study counts for your bonus."

Part	Probability
Part 1 (donation decision)	0 %
Part 2 (decisions)	50 %
Part 3 (estimates)	50 %

"All information in this study is **true and is shown with its source**. Both you and **all other participants in our studies know** that the **information shown is always true**." "We ask you to give this study your **full attention**."

Page 4: Part 1, donation decision. Title: "Part 1: Donation decision". The description sentence, the sentence introducing the statements and the organizations' statements

are as reproduced in Appendix A.5. Then: "You now **divide €250 between these two organizations**." Box "Important!": "One decision from this study is selected and **actually implemented with real money**. Your decision today can therefore have **real financial consequences** for these organizations." Below, a slider has the two organizations' names at its left and right ends; underneath, a readout lists "[organization]: [amount] €" for each. The slider starts at the midpoint (€125 each); moving it towards one organization raises that organization's amount and lowers the other's, in steps of €1. The left–right placement of the organizations is randomized.

Page 5: Part 2, decisions (instructions). "You now make **decisions that concern persons from another study**." "Here are depictions of these persons:" [Persons A to E, each shown as a coloured person icon.]

Your task.

- You decide for each person A to E whether this person **receives a piece of information on [migration and asylum policy]**—or **not**.
- **Exactly one of these persons** from the other study **can be actually assigned** to you.[1]

"Some people send information, others do not. Both decisions can have good reasons. There is no right or wrong decision."

Box: "Your decision for the **person assigned to you** is **implemented with 50 % probability**. Otherwise the person assigned to you receives **no information**—regardless of how you decided."

What the person assigned to you learns.

1. "She always learns **your opinion on [migration and asylum policy]**—but **nothing about your identity**."
2. "She knows that **chance co-determines** whether your decision is implemented."

What does the person assigned to you do afterwards?

- CONTROL: "She will **state her opinion on [migration and asylum policy] again**." "She will make **no donation decision**."
- ACTION: "She will **state her opinion on [migration and asylum policy] again**." "She will **make a donation decision**. She will divide 250 euros between **[pro organization]** and **[anti organization]**." Box: "One decision from this study is selected and **actually implemented with real money**. The decision of the person assigned to you can therefore have **real financial consequences** for these organizations."

Your bonus. "On the next page you answer short comprehension questions on the rules of Part 2. If this part is selected for your bonus payment, you **receive the bonus of €10** if you answer all questions **correctly on the first attempt**. Otherwise you receive no bonus."

[1] "This happens in at least 4 out of 10 cases." (Footnote shown in Experiment 2 only.)

Page 6: Part 2, comprehension questions. Each question allows several answers ("Please select all correct answers"); a button "Help" opens a reminder box that restates the Page 5 rules without adding anything new. Participants cannot proceed until every question is answered correctly; a submission with an unanswered question is rejected without counting as an attempt. After one wrong attempt the correct answers are highlighted. The bonus of Page 5 requires all correct on the first attempt.

Question 1. What does the person assigned to you learn about you? (Please select all applicable answers.) (a) She always learns your political opinion, whatever you decide. (b) She does not learn who you are. (c) She learns your political opinion only if you send a piece of information. [correct: a, b]

Question 2. How is your decision for the person assigned to you implemented? (Please select all applicable answers.) (a) If you **send** the information, it is **always displayed**. (b) If you **send** the information, it is **sometimes displayed**. (c) If you **do not send** the information, it is **never displayed**. [correct: b, c]

Question 3. Which decisions does the person assigned to you make during the rest of the study? (Please select all applicable answers.) (a) She states her political opinion on [migration and asylum policy] again. (b) She makes a decision on how 250 euros are divided between [pro organization] and [anti organization]. (c) She in any case receives another piece of information. [correct: ACTION: a, b; CONTROL: a]

Page 7: Part 2, sending decision (shown 5 times, once per fact and person A–E). Title: "Your decision: Person [A]". [Person icon.] "Person [A] is of the opinion: **'[stance sentence, one of the five scale sentences of the domain, e.g. Immigration to Germany should be made somewhat easier.]**'" Then a fact card, in the immigration example:

> *Skilled-labour shortage means that in certain occupations, such as food production, finishing and drywall work, or bus and tram transport, there are too few suitably trained workers.* [context line, shown for some facts only] The share of people with an immigration history in occupations with a skilled-labour shortage is **higher** than their share of the population. [fact statement, with the correct answer in bold]

Statistisches Bundesamt (2025): "Fachkräftemangel: Anteil ausländischer Beschäftigter in Engpassberufen gestiegen"

Buttons: "Not send" / "Send". [Reading gate: on loading, everything except the fact card and the button "I have read the information" is blurred and the continue button is disabled; after the click, everything except the receiver's icon and opinion is blurred for 3 seconds. The participant can then choose and continue.] The five facts and five stances are assigned as described in Section 4.1; the page text is identical in the two treatments.

Page 8: Attention check (Experiment 2 only). Title: "Interim question". "In studies like this, participants are sometimes inattentive and answer questions without reading them completely. To show that you read this question carefully, please select 'No, I skipped them' as your answer below, even if that is not true." **Have you read the instructions for the last part of the study carefully?** (1) No, I skipped them. [correct] (2) I read them partly. (3) Yes, I read the instructions. (4) I am not sure. [Not shown to participants: any answer other than (1) excludes the participant from the bonus.]

Page 9: Part 3, estimates (instructions). "In the other study there are **several persons with the same opinion** on [migration and asylum policy]. **Persons from your study** with the opinion '[stance sentence]' **decided** on sending the information."

Your task. "You now give two estimates for each of these persons."

1. What share of these persons knows that the **information is true**, if … (i) the information is **not displayed**? (ii) the information is **displayed**?
2. What is the **average opinion** of these persons, if … (i) the information is **not displayed**? (ii) the information is **displayed**?

Your bonus. "If Part 3 is selected for you: each question can determine your bonus. The closer your estimate is to the data in the other study, the higher your chance of the **bonus of €10**. **You always achieve the best result by stating your actual assessment.**" Expandable "Additional explanation: how your chance of the bonus is computed in detail":

- Question 1: $B = \max(0,\, 1 - (v - w)^2/400)$, where v is your prediction in percentage points and w the value in the other study. Deviation of 5 points → bonus chance 94 %; 10 points → 75 %; 20 points → 0 %.
- Question 2: $B = \max(0,\, 1 - (v - w)^2/4)$, where v is your prediction and w the average value in the other study. Deviation of 0.5 scale points → 94 %; 1.0 → 75 %; 2.0 → 0 %.

- "You always achieve the best result if you state your honest assessment." "One of the questions for which, in the other study, at least one person received the information and at least one did not is selected with equal probability. Chance also determines which of the estimates counts."

Page 10: Part 3, estimates (shown 5 times, once per receiver position and fact). [Row of icons for persons A to E's group.] "Persons [A] were initially of the opinion: **'[stance sentence]'**"

1. "What share of these persons knows that: [fact statement, with the correct answer in bold]" … "if the information is **not displayed** to them?" … "if the information is **displayed** to them?" (two sliders from 0 % to 100 %; each starts at the midpoint, shows "– %" until moved, and must be moved before the participant can continue).

2. "What is the **average opinion** of these persons, …" Instruction (red): "If you do not move the slider, you state that the persons do not change their opinion. If you move it to the right, you expect a change towards '[much harder]'; to the left, towards '[much easier]'. The text below the slider shows which statement corresponds to the slider position." … "if the information is **not displayed** to them?" … "if the information is **displayed** to them?" (two continuous sliders from 1 to 5, each starting at the position of the persons' initial opinion, which is also marked by a tick; the endpoints read, for immigration, "ease considerably" and "make considerably harder"; below each slider the scale sentence for the current position is displayed). Unmoved percentage sliders trigger "Please move the slider to give an estimate." [For 3 seconds after loading, all parts of the page except the icons, the initial opinion and the fact card are blurred.]

Pages 11–16: Final questionnaire. Title on every page: "Final questionnaire". A fourth part and an additional perception task were removed from the sequence and were not shown.

Page 11. (1) "How did you proceed in your sending decisions in Part 2?" (open text). (2) "How did you proceed in your estimates of the persons in Part 3 and 4?" (open text). (3) "What do you think this study was about?" (open text). (4) "Did you perceive a political imbalance in the study?" (a) No, the study was balanced. (b) Yes, rather in favour of left-wing positions. (c) Yes, rather in favour of right-wing positions. If (b) or (c): "What makes you think so?" (open text).

Page 12. "Please briefly assess the following pieces of information on [climate policy]. There are no right or wrong answers." (Title "Final questionnaire".) (The two items are one pro and one anti fact of the topic not assigned to the sender; the example below is therefore the climate domain.) For each: [fact statement with source]. (i) "Does this information show that your opinion ('[the participant's own stated position on that other topic]') is correct? We are **not asking whether your opinion is correct in**

itself or whether this information **persuades others**." No / Yes. (ii) "Do you think this information is true?" Not true / True.

Page 13. "State your agreement with these statements." Each on a five-point scale (1 strongly disagree, 2 disagree, 3 neither, 4 agree, 5 strongly agree): (i) "Whoever passes on political information should also pass on factually correct information that contradicts their own political position." (ii) "Social media should make more transparent why users see certain political information and not other." (iii) "Platforms or media should be obliged to make relevant information from different perspectives visible on political topics." Then: "If there were a federal election next Sunday, which party would you vote for?" (CDU/CSU; AfD; SPD; Bündnis 90/Die Grünen; Die Linke; BSW; FDP; other; I would not vote/no answer). Then: "Many people use the terms 'left' and 'right' to describe political attitudes. When you think of your own political views: where would you place yourself on a scale from 0 to 10, where 0 means 'left' and 10 means 'right'?" (slider, 0–10, starting at 5).

Page 14. "Do you have any additional comments or feedback on this study?" (open text, optional).

Page 15. "The study is complete. The following information is **voluntary** and has **no effect on your payment** or the evaluation of your answers." "Note on the supplementary analysis: As part of this study we optionally collect your social-media profiles in order to analyse publicly accessible posts in addition. Since your profile names could be identifying, we store them separately from your survey answers and analyse the posts anonymously. We do not pass them on to third parties and delete them after the analysis is complete. In publications we report only aggregated results—individual profiles or user names never appear there." "Do you agree that publicly accessible posts on the profiles given below are analysed anonymously for a supplementary scientific analysis?" Yes/No; if yes, optional fields for an X/Twitter, Bluesky and Mastodon profile.

Page 16: End screen. "Thank you for your participation! You have successfully completed the survey." "Please click the following button to return to Bilendi and complete your participation." Button: "Back to Bilendi".

Experiment 3

Experiment 3 uses the same templates and the same sequence as Experiment 2 with the ACTION text throughout (every journalist is assigned to ACTION). Receivers are called "readers" (*Leser:innen*) instead of "persons" in all texts ("Reader A", "Readers A"), and the group in Page 9 is "journalists" instead of "persons from your study". Differences by page:

- **Screener and exit.** A screening page precedes consent: "In this study we are interested in the decisions and assessments of journalists. We would therefore first like to know whether you work as a journalist. On the next page we then explain the details of the study." "Are you currently working, or have you worked in the last 12 months, as a journalist as your main occupation?" (Yes / No). A conditional exit page appears only for a participant who answered No on the screener ("Unfortunately you cannot take part in this study, since it is designed exclusively for persons who currently work, or in the last 12 months worked, as journalists as their main occupation").
- **Page 1 (Consent).** The summary adds that the study is about the decisions and assessments of journalists, and states: "You receive **€15 for your participation** and have a chance of a **bonus payment of €30**. Payment is made, as you wish, by PayPal or bank transfer." It says the study has *three* parts and a final questionnaire; the optional data block on the final page concerns name and outlet rather than social-media profiles; the Bilendi withdrawal sentence is omitted.
- **Page 2 (Personal information).** Identical.
- **Page 3 (Schedule).** "10 persons are randomly selected" who can receive a **bonus of €30** (instead of 25 persons and €10). Added: "You receive the base payment for participation **within seven working days** after you complete the study. The bonus is paid **in addition** at the end of the whole series of studies." Same table and truthfulness paragraph. The same €30 replaces €10 in the bonus sentences of Pages 5 and 9.
- **Page 4 (Donation decision).** Identical.
- **Page 5 (Part 2 instructions).** "readers" replaces "persons"; the matching is certain, not probabilistic: "You are **actually matched** with **exactly one of these readers** from the other study." The footnote "This happens in at least 4 out of 10 cases" is absent. The 50 % implementation rule is unchanged.
- **Page 6 (Comprehension questions).** Same questions with "reader" wording and the ACTION answer key (Question 3: a, b). The Help box has the same matching sentence as Page 5.
- **Pages 7 and 10 (decisions, estimates).** Identical except "reader" for "person" (title "Your decision: Reader [A]").
- **Page 8 (Attention check).** Not shown to journalists.
- **Page 9 (Part 3 instructions).** "Journalists with the opinion '[stance sentence]' **decided** on sending the information" in place of "persons from your study".
- **Final questionnaire.** Page 12 (ratings of other-topic facts) is not shown. Page 11, item 2 reads "… in your estimates of the readers in Part 3" (without "and 4").

Page 13: the first agreement item reads "Journalists should also pass on factually correct information that contradicts their own political position."; the vote question has a separate response list (CDU/CSU; AfD; SPD; Bündnis 90/Die Grünen; Die Linke; BSW; FDP; other; I would not vote; no answer). Page 14 first asks the journalist background items, then the open comment: (a) "In which type of medium do you mainly work?" (print; online; television; radio/audio/podcast; news agency; other; no answer); (b) only if television or radio/audio/podcast: "Do you mainly work for public-service or private broadcasting?" (public-service; private); (c) "What reach does your main medium mainly have?" (local; regional; national; international; no answer); (d) "Which role do you take in your journalistic work?" (select all: reporter; editor (Redakteur/in); desk editor/chief of desk/copy editor; correspondent; columnist/commentator; fact-checking/research; investigative journalist; other; at least one required); (e) "How often do you decide yourself which pieces of information appear in a piece?" (never or almost never; rarely; sometimes; often; very often). Page 15 has the same opening sentence, then the notice: "Note on the supplementary analysis: As part of this study we optionally collect your name and your medium. These data serve exclusively scientific purposes: they allow us to analyse publicly accessible articles in addition and to place the survey results in a broader context. All data are treated strictly confidentially, not passed on to third parties and irrevocably deleted after the analysis is complete. Publications report only anonymized findings—your name and your medium never appear in publications." "Do you agree that publicly accessible articles that appeared under your name are analysed anonymously for a supplementary scientific analysis?" (yes/no); if yes: "Under which name do you publish journalistic work?" and "For which main medium or editorial office do you mainly publish?" (free text). No favourite-colour item or other extra item is displayed in Experiment 3.

- **End screen.** Instead of the Bilendi redirect: "Thank you for your participation! … Your personal payout code is: [8-character code]. Please be sure to write down this code! You need it to receive your payment." Button "To the payout form" linking to the external payout form. The bonus draw is between the sending decisions ("provision") and the estimates ("delta").

H.3 Experiment 4

The task was administered in German and the text below is translated. Participants are German residents recruited through Bilendi. Text in square brackets varies by receiver (topic, fact, sender's stance, matched treatment).

Page 1: Consent. Standard informed-consent text [...]. Substantive sentences: "Duration: about 5 minutes. Payment: In addition to the base payment you have a chance of a bonus payment of €10." "Deception: You will not be deceived in this study. All statements about other participants and about the procedure of the study are true." Options: (a) "Yes, I agree and wish to participate." (b) "No, I do not wish to participate."

Page 1b: Consent refused. Participants who decline see an exit page and are returned to Bilendi (conditional).

Page 2: About you. "Please first answer a few questions about yourself." Age (18–29, 30–39, 40–49, 50–59, 60–75), gender (male, female, diverse, no answer), highest educational degree (eight categories from no degree to doctorate), federal state (16 states, "I do not live in Germany"). "Please give your personal assessment on the following topics." *Migration and asylum policy:* "Which of the following positions comes closest to your view?"

1. Immigration to Germany should be made much easier.
2. Immigration to Germany should be made somewhat easier.
3. Immigration to Germany should stay about as it has been.
4. Immigration to Germany should be made somewhat more difficult.
5. Immigration to Germany should be made much more difficult.

Climate policy: same question, with the options: (1) Climate protection measures in Germany should be expanded considerably. (2) ... should be expanded somewhat. (3) ... should stay about as they have been. (4) ... should be reduced somewhat. (5) ... should be reduced considerably.

Page 3: Quota check. Participants whose quota cell is full, duplicates and screened-out participants are redirected to Bilendi (conditional).

Page 4: Study schedule. "This study consists of **four parts** and takes **about 5 minutes**."

"Among all participants, **10 persons are randomly selected for a possible bonus payment of €10**. For each of these persons one of the four parts is drawn. Each part is selected with a probability of 25%. Whether a drawn person receives the bonus depends on their answers in the selected part."

Part	Probability
Part 1	25%
Part 2	25%
Part 3	25%
Part 4	25%

"A possible bonus is paid **in addition to the base payment** after the entire study series has been completed."

"All information in this study is **true and is shown with a source**. **All other participants of our studies** know this as well."

"Please work through the study **attentively**."

Page 5: Part 1, Your assessment. "**Your task:** In Part 1 you indicate which of the two words **correctly completes** the sentence."

"**Your bonus:** If Part 1 is selected for the bonus calculation, your chance of the bonus depends on your assessment. **You achieve the best result if you give your honest assessment.**" Expandable details: "Let p be your entered assessment for '[word at right end]' divided by 100 (e.g. 75% $\rightarrow$ $p = 0.75$), and let $x = 1$ if '[word at right end]' is correct and $x = 0$ if '[word at left end]' is correct. Your bonus chance is calculated as: $B = 1 - (p - x)^2$. 75% for '[right]' corresponds to 25% for '[left]'; 25% for '[right]' corresponds to 75% for '[left]'. If '[right]' is correct: assessment 50% for '[right]' $\rightarrow$ bonus chance 75%; 75% $\rightarrow$ 94%; 100% $\rightarrow$ 100%. If '[left]' is correct: assessment 50% for '[right]' $\rightarrow$ 75%; 25% for '[right]' (= 75% for '[left]') $\rightarrow$ 94%; 0% for '[right]' (= 100% for '[left]') $\rightarrow$ 100%. You always achieve the best result if you give your honest assessment."

Warning box: "Please answer with your **current knowledge** and do not switch the browser tab."

Fact screen (the receiver's matched fact; example, immigration, fact `imm_pro_1`; the facts are listed in Table 4): *[background sentence, if any] [introductory sentence: "Some people think that … Others think that …"]* "The share of people with an immigration background in occupations with a skilled-labor shortage is … than their share of the population." (the blank shows the word currently favored by the slider, "…" before any answer, and "—" at the midpoint together with the note "Both answers are equally likely.") "**Which word correctly completes the sentence?**" Before an answer: "Click or tap on the scale to give your assessment." The slider has no starting position: no handle is shown until the participant clicks or taps the scale, and a submission without an answer is rejected ("Please click or tap on the scale to give an assessment."). Scale ends: "higher" and "lower" (which word is on the left is randomized once per participant and kept on all later screens showing the two answer words), center label "Both equally likely". Live feedback sentence: "You believe that the correct word is [X]% '[left word]'

and the correct word is [Y]% '[right word]'." (literally, "Sie glauben, dass das richtige Wort zu [X]% '[left]' und das richtige Wort zu [Y]% '[right]' ist"); at 50: "You think that both words are equally likely to be correct."; at an end: "You are sure that the correct word is '[word]'."

Page 6: Part 2, Information. "In the second part you can find out which word **correctly completes** the sentence from the previous page. Whether you receive the information depends on the **decision of a person from another study** and on **chance**."

"How it works." (i) "The person decides whether to **send you the information**." (ii) "If the person does **not** send the information, you receive **no information**." (iii) "If the person sends the information, a **coin toss** additionally decides: **Heads**: you do not receive the information. **Tails**: you receive the information." (iv) "The person knew your opinion on [Migration and asylum policy / Climate policy]."

"If you receive no information, there can be two reasons: (1) The person **decided against sending**. (2) The **coin toss** showed **heads**."

"**Your task:** On the next page you answer two **comprehension questions** on this part." "**Your bonus:** If Part 2 is selected for the bonus calculation, you receive the bonus if you answer both questions correctly **at the latest on the second attempt**."

When the participant submits this page, the transmission coin is drawn (50% Tails): the receiver *receives* the fact only if the sender's decision for the receiver's stance was "send" and the coin shows Tails.

Page 7: Part 2, Comprehension questions. **Question 1:** "If you **do not receive** the information: What reasons can this have?" (select all that apply) (a) The person **did not send** the information. (b) There was no information on the question from Part 1. (c) The person **sent** the information, but the coin toss showed **heads**. (d) I answered the question from Part 1 incorrectly. [correct: a, c]

Question 2: "If you **receive** the information: What must have happened?" (select all that apply) (a) All participants receive the information. (b) The person **sent** the information and the coin toss showed **tails**. (c) The person **did not send** the information, but the coin toss showed **tails**. (d) The coin toss alone decides – the person's decision does not matter. [correct: b]

Empty answers are rejected ("Please answer question 1." / "Please answer both questions.") and do not count as an attempt. After a wrong submission: "Please check your answers: Question 1 is wrong. Question 2 is wrong." The participant repeats until both are correct; after two failed attempts the correct options are marked "Correct. Please select." A "Hilfe" button opens a reminder repeating the box of Page 6 (how it works, two reasons).

Page 8: Part 2, Information (reception). Header: Part 2: Information. In both states the page first shows: "The person assigned to you is of the opinion: *'[sender's stance sentence]'*." (one of the five scale sentences of Page 2 for the matched topic). The rest of the page depends on the state.

Receipt (sender sent and coin showed Tails): "**You receive the following information:**" *[background sentence, if any]* "The share of people with an immigration background in occupations with a skilled-labor shortage is **higher** than their share of the population." (the correct word in bold; example fact `imm_pro_1`), followed by the source, e.g. "**Statistisches Bundesamt (2025)** *Fachkräftemangel: Anteil ausländischer Beschäftigter in Engpassberufen gestiegen*".

Non-receipt (sender did not send, or sent and coin showed Heads): the stance sentence is followed only by "You receive **no information**." The two causes are not distinguished on the screen.

Page 9: Your opinion. Header: Your opinion. "Now we are interested in **only your opinion**." *Immigration:* "Some people think that immigration to Germany should be made easier. Others think that it should be made more difficult." *Climate:* "Some people think that climate protection measures in Germany should be expanded. Others think that they should be reduced." (matched topic) "**Which of the following positions comes closest to your view?**" The five options are identical to those of Page 2 for that topic.

Page 10: Your donation decision. Shown only to receivers matched to a sender in the ACTION treatment; all other receivers skip this page.

"To conclude this part we introduce **two organizations** that deal with the topic **[Migration and asylum policy / Climate policy]**." "The two organizations hold different positions. On one side is *[organization]*, which advocates *[position]*. On the other side is *[organization]*, which advocates *[position]*." (order randomized per participant) "Here you find some statements of the two organizations in their own words." The statements are those reproduced in Appendix A.5 (the quotes are identical; the introductory sentences above differ from the ones printed there). "You now divide €250 **between these two organizations**." Box: "**Important!** A donation decision from this study is randomly selected and **implemented with real money**. Your decision today can therefore have **real financial consequences** for these organizations." Decision: a slider from 0 to 250 whose two ends are labeled with the two organizations (left–right order randomized); below it the euro amount for each organization is displayed. The slider starts at €125 each.

Page 11: Part 3, Your assessment. Identical to Page 5 for the same matched fact, with "Part 3" in place of "Part 1" and the Part 3 bonus box, except that the background sentence is not repeated and the warning box reads "Please do not switch the browser tab." The answer words keep their left–right positions from Page 5.

Page 12: Part 4, Further assessments. "People from our earlier studies knew the correct answer to the question. In the following we call these people **senders**. Each sender decided whether to send the correct answer to another person. We call this other person the **recipient**."

"**Your task:** The question has two possible answers. For each answer you estimate separately how many of the senders sent the information if this answer is correct."

"Both estimates refer to **the same** N **[randomly selected] senders**." ($N = 10$ in most cells; in a few cells N is 8 or 9 and all those senders are used, and "randomly selected" is omitted.)

"**Your bonus:** If Part 4 is selected for the bonus calculation, only your estimate for the answer that is actually correct counts. **You achieve the best result if you give your honest assessment.**" Expandable details: "For the bonus calculation **one of the** N **persons is selected at random**. Your statement 'v of N' is evaluated as the probability $p = v/N$. If the selected person sent the information, $x = 1$. If the person did not send the information, $x = 0$. Your bonus chance is then calculated as: $B = 1 - (p - x)^2$. Example: '6 of 10' corresponds to a probability of 60% that a randomly selected person sent the information. You always achieve the best result if you give your honest assessment."

Box "For both estimates:" *Opinion of the senders: "[sender's stance sentence]"*. *Opinion of the recipients: "[receiver's own stance sentence]"*.

Two blocks follow, one per answer word (order randomized, tied to the left–right randomization of Pages 5 and 11). Each reads: "**Suppose the correct answer is:**" *[the fact sentence completed with that word]* (preceded by the background sentence if the fact is not the receiver's matched fact) "How many of the N senders sent this information?" Before an answer the display reads "Click or tap on the scale or drag the slider"; the slider runs from 0 to N and has no starting position. After an answer it reads "v of N", and two boxes of person icons are shown: "Not sent" (left) with the count "$N - v$ of N" and "Sent" (right) with the count "v of N"; the icon numbers follow the slider. A missing answer is rejected ("Please click or tap on the scale or drag the slider to give an estimate."). Receivers who did not receive the information are asked about their own matched fact; receivers who did receive it are asked about a randomly drawn other fact on the same topic.

Page 13: Final questionnaire (1). "1. What do you think this study was about?" (open text). "2. Did you have the impression that the study favored a political direction?" (a) No, the study was balanced. (b) Yes, rather left positions. (c) Yes, rather right positions. If (b) or (c): "What makes you think so?" (open text).

Page 14: Final questionnaire (2). (i) "Which news outlet have you read, watched or heard political or social news from most often in the past seven days?" with the note "Please give the news outlet the content comes from – not the way you access it. Example: if you see *Tagesschau* posts on Instagram, give 'ARD-Tagesschau'." (free text, placeholder "e.g. ARD-Tagesschau"; checkbox: "I have not used any political or social news in the past seven days."). If an outlet is named: "Your assessment of: [outlet]" – "How often are **important pieces of information left out** there that do not fit the political orientation of this news outlet?" (1) Never (2) Rarely (3) Occasionally (4) Often (5) Very often. (ii) "The statistical information on the topic **[topic]** that you received or could have received comes from the following source:" *[source of the matched fact]* "**How credible do you find this source for statistical information on the topic [topic]?**" (slider; 0 – not at all credible, 10 – very credible; checkbox: "I did not know this source before the study."). (iii) "**The information concerned: [short description of the matched fact].**" "**How important is the content of this information for political decisions on [topic]?**" (slider; 0 – not at all important, 10 – very important).

Page 15: Final questionnaire (3). "Where would you place yourself politically?" (slider; endpoints "0 – Left", "5 – Center", "10 – Right"). "If a federal election were held next Sunday, which party would you vote for?" (CDU/CSU; AfD; SPD; Bündnis 90/Die Grünen; Die Linke; BSW; FDP; Other; I would not vote / no answer). "Do you have any comments on the study?" (open text).

Page 16: End. "Thank you for your participation! You have successfully completed the survey. Please click the following button to return to Bilendi and complete your participation."

I Derivations

I.1 Sender Framework

Throughout, $U_A(x) = \tau\,\theta_B(A)\,h_B(\Delta b) + \theta_E(A)\,h_E(\alpha) + \tau\,\theta_P(A)\,w(\Delta\pi) - \theta_R(A)\,D(\rho) + \kappa A$, where h_B, h_E, w and D are differentiable with strictly positive derivatives. Equation (2) is the case in which all four are the identity. The marginal utilities are

$$\begin{aligned}
\frac{\partial U_A}{\partial \Delta b} &= \tau\,\theta_B(A)\,h_B'(\Delta b),\\
\frac{\partial U_A}{\partial \alpha} &= \theta_E(A)\,h_E'(\alpha),\\
\frac{\partial U_A}{\partial \Delta \pi} &= \tau\,\theta_P(A)\,w'(\Delta\pi),\\
\frac{\partial U_A}{\partial \rho} &= -\theta_R(A)\,D'(\rho),
\end{aligned} \tag{8}$$

each a weight times a strictly positive factor, with a minus sign for distance.

Proof of Proposition 1. Under Assumption 1, F_A does not depend on x, so differentiating Equation (3) gives $\partial \Pr(d = 1 \mid x, A)/\partial x_j = f_A(U_A)\,\partial U_A/\partial x_j$. Since $f_A > 0$, this has the sign of $\partial U_A/\partial x_j$, which by Equation (8) is the sign of $\theta_B(A)$, $\theta_E(A)$, $\theta_P(A)$ and $-\theta_R(A)$ for Δb, α, $\Delta\pi$ and ρ. □

Proof of Proposition 2. Differencing Equation (8) across A at the same determinant values gives $\tau[\theta_B(1) - \theta_B(0)]\,h_B'(\Delta b)$ for Δb and $\tau[\theta_P(1) - \theta_P(0)]\,w'(\Delta\pi)$ for $\Delta\pi$. The differences for α and ρ are zero by Assumption 3(ii)–(iii). Under the identity transformations these are the changes stated in the proposition, and otherwise they have the same signs. Parts (i)–(iii) follow. Under Assumption 3(i), these values describe the same fact and receiver. □

Derivation of Equation (4). Differencing the probability response in the proof of Proposition 1 across A at the same x gives Equation (4). For $\Delta\pi$ it equals $\tau\,w'(\Delta\pi)\,[f_1(U_1)\,\theta_P(1) - f_0(U_0)\,\theta_P(0)]$, and analogously for Δb with θ_B and h_B'. If $\theta_P(1) > \theta_P(0) \geq 0$ and $f_1(U_1) \geq f_0(U_0)$ it is positive; if $\theta_P(1) > 0$ it is negative whenever $f_1(U_1)/f_0(U_0) < \theta_P(0)/\theta_P(1)$. The two densities need not coincide and are evaluated at different points, so their ratio is unrestricted and the sign of Equation (4) need not match the change in marginal utility. For α, under Assumption 3(ii), it equals $\theta_E\,h_E'(\alpha)\,[f_1(U_1) - f_0(U_0)]$, which is nonzero whenever $\theta_E \neq 0$ and $f_1(U_1) \neq f_0(U_0)$. □

Proof of Corollary 1. Under Assumption 4(i), $\Pr(d = 1 \mid x, A) = G\big((U_A(x) - \mu_A)/s_A\big)$. Since G is strictly increasing, $G^{-1}\big(\Pr(d = 1 \mid x, A)\big) = (U_A(x) - \mu_A)/s_A$, and since μ_A and s_A do not depend on x (Assumption 1), $\beta_j(A) = s_A^{-1}\,\partial U_A/\partial x_j$. As $s_A > 0$, $\beta_E(A)$ has the sign of $\theta_E(A)$.

(i) With $s_1 = s_0 = s$, $\beta_j(1) - \beta_j(0)$ is s^{-1} times the corresponding difference in the proof of Proposition 2, and $s > 0$ preserves its sign.

(ii) The scale cancels in $\beta_j(A)/\beta_E(A) = (\partial U_A/\partial x_j)/(\partial U_A/\partial \alpha)$, defined when $\theta_E(A) \neq 0$. Under the identity transformations, Equation (8) gives $\beta_m(A)/\beta_E(A) = \tau\,\theta_m(A)/\theta_E(A)$ for $m \in \{B, P\}$ and $\beta_\rho(A)/\beta_E(A) = -\theta_R(A)/\theta_E(A)$. Under Assumption 3, $\theta_E(A) = \theta_E$ and $\theta_R(A) = \theta_R$, so differencing across A gives $\tau[\theta_m(1) - \theta_m(0)]/\theta_E$ for $m \in \{B, P\}$, which has the sign of $\theta_m(1) - \theta_m(0)$ when $\theta_E > 0$, and zero for distance. Under general transformations the ratios carry the factors $h'_B(\Delta b)/h'_E(\alpha)$, $w'(\Delta\pi)/h'_E(\alpha)$ and $D'(\rho)/h'_E(\alpha)$, which are positive and, evaluated at the same determinant values, common to both values of A, so every sign is unchanged. With heterogeneous weights or scales the cancellation holds sender by sender, and a pooled coefficient ratio need not equal the average of individual ratios. □

Remark. By Equation (8), $\partial \Pr(d = 1 \mid x, A)/\partial \Delta\pi = f_A(U_A)\,\tau\,\theta_P(A)\,w'(\Delta\pi)$. If $\theta_P(A) > 0$, disclosure rises with the expected strength of persuasion when $\Delta\pi > 0$ and falls with it when $\Delta\pi < 0$.

I.2 Receiver Benchmark

A sender who observes a binary state $\omega \in \{0, 1\}$ discloses it with state-contingent probabilities $\sigma_1 = \Pr(d = 1 \mid \omega = 1)$ and $\sigma_0 = \Pr(d = 1 \mid \omega = 0)$. Disclosure is truthful, and a disclosed state reaches the receiver with probability τ. A received message reveals the state. For a receiver with prior $p \in (0, 1)$ who knows the disclosure rule, Bayes' rule gives after non-receipt, whenever it has positive probability,

$$b_{\varnothing}^{\text{Bayes}} = \Pr(\omega = 1 \mid m = \varnothing) = \frac{p\,(1 - \tau\sigma_1)}{p\,(1 - \tau\sigma_1) + (1 - p)\,(1 - \tau\sigma_0)}. \tag{9}$$

Non-receipt is informative only if $\sigma_1 \neq \sigma_0$: at $\sigma_1 = \sigma_0$ the posterior equals the prior, and a smaller τ moves it toward p. This fixes a reference and assumes nothing about behavior: a receiver may match it, fall short of it, or move against it. Identifying σ_0 from disclosure observed at the realized state needs a restriction linking the sender's behavior across states.

The benchmark at the actual rates. Evaluating Equation (9) at the senders' disclosure rates in the sender-fact-receiver pairing of each receiver who saw no fact, with σ_0 set to the sender-fact-receiver pairing's rate for facts of the opposite direction, with the receiver's own prior, prescribes a mean absolute belief movement of 1.95 points on the 0–100 scale. Across the 685 receivers who saw nothing, mean squared loss is 0.312 at the prior, 0.301 at the reported posterior and 0.310 at the Bayesian posterior, so Bayes'

rule would lower loss after non-receipt by 0.002. Selectivity measured this way runs in opposite directions across sender stances (mean -0.024).