Journal of Economic Behavior and Organization 236 (2025) 107081

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Research paper

Selective exposure reduces voluntary contributions: Experimental evidence from the German Internet Panel

Federico Innocenti *, Linnéa Marie Rohde

University of Mannheim, Department of Economics, L7, 3-5, Mannheim, 68131, Germany

A R T I C L E I N F O

Dataset link: https://doi.org/10.4232/1.13391, https://doi.org/10.4232/1.13464, https://doi.org/10.4232/1.13614, https://doi.org/10.4232/1.13681, https://doi.org/10.4232/1.13682, https://doi.org/10.4232/1.13794

JEL classification:
C99
D12
D61
D83
H41

Keywords:
Experiment
Limited attention
Media bias
Public good
Selective exposure
German Internet Panel

A B S T R A C T

Can strategic information acquisition harm the provision of a public good? We investigate this question in an incentivized online experiment with a large sample of the German population. The marginal returns of the public good are uncertain: it is either socially efficient to contribute or not. In the information treatment, participants can choose between two information sources with opposite biases: one source is more likely to report low marginal returns, whereas the other is more likely to report high marginal returns. We find that information avoidance is a minor phenomenon. Most participants select the source biased towards reporting low marginal returns, independent of their prior beliefs. As a result, the information treatment reduces contributions and increases free-riding. We find that social preferences guide information acquisition: selfish participants are less likely to acquire information, and if they acquire information, they are more likely to select the source biased towards reporting high marginal returns.

1. Introduction

With abundant information available, individuals have to select which sources of information are worthy of their attention. But the media do not necessarily provide accurate information. Hence, individuals might strategically acquire information that confirms their beliefs or aligns with their preferences. Although there is clear evidence of selective exposure in the empirical literature (Del Vicario et al., 2016), its consequences depend on how the acquired information affects an individual's actions. When the information affects an individual's private actions and outcomes, her selective exposure can only affect her own well-being. Instead, when an individual engages in *collective* action, the information she acquires can affect the outcome of all individuals involved and, thus, social welfare. An important area of collective action where information is critical is the provision of public goods. The exact returns of investing in a public good are often uncertain, which can lead to the under-provision of the public good (Levati et al., 2009). At first sight, providing more information about the returns of investing in a public good could reduce uncertainty and mitigate the problem of under-provision. However, when different information sources have opposite claims about the returns of the public good, individuals can strategically select the source that supports their selfish interests and use the information to justify

* Correspondence to: University of Verona, Department of Economics, Via Cantarane 24, Verona, 37129, Italy.
E-mail addresses: f.innocenti93@gmail.com (F. Innocenti), linnea.marie.rohde@gmail.com (L.M. Rohde).

https://doi.org/10.1016/j.jebo.2025.107081
Received 19 June 2024; Received in revised form 20 April 2025; Accepted 19 May 2025

Available online 6 June 2025

lower contributions. In this case, information provision can backfire and, contrary to expectations, further reduce investments in the public good.

In this paper, we study experimentally how strategic information acquisition affects voluntary contributions to public goods. We implement a one-shot Voluntary Contribution Mechanism (VCM) where the public good's marginal per-capita return (MPCR) is uncertain. If the MPCR is high, providing the public good is socially efficient. In contrast, if the MPCR is low, it is socially inefficient. We employ two main treatments, and within each treatment, we vary the assigned prior beliefs about the MPCR. In the *no info* treatment, the participants make their contribution decisions based on their assigned prior beliefs. Instead, in the *info* treatment, participants face two information sources with opposite biases: one supporting the provision of a public good, the other contrasting it. In particular, the high-biased source is more likely to report high MPCR but sometimes reveals that the MPCR is low with certainty. Instead, the low-biased source is more likely to report low MPCR but sometimes reveals that the MPCR is high with certainty. Participants decide whether to receive one unit of costless information about the MPCR and from which source.[1] Then, the participants make their contribution decisions based on their updated beliefs. We provide participants with the correct Bayesian posterior beliefs depending on their information choices.

If participants behave rationally and do not exhibit any social preferences, the equilibrium contribution to the public good in this game is zero, independent of beliefs. In this case, a participant should be indifferent towards information as long as it is costless. Instead, if social preferences play a role, information matters. On the one hand, the direction of optimal information acquisition may depend on prior beliefs (Che and Mierendorff, 2019).[2] On the other hand, participants might strategically avoid information that compels them to be more generous (Golman et al., 2017) and instead seek information that justifies more selfish behavior (Spiekermann and Weiss, 2016).

We conduct our experiment on the German Internet Panel (GIP). The GIP is a long-term online study based on a randomly drawn sample of the general population in Germany. The GIP reaches more than 4000 participants in each wave and covers a multitude of political and economic attitudes. Embedding our experiment in the GIP allows us to relate the participants' behavior in our incentivized experiment to their self-reported field behavior in a public good context, i.e., their willingness to contribute to environmental protection.

Most participants in the *info* treatment (87%) choose to acquire information. Most (65%) of those acquiring information select the low-biased source. In contrast to the theoretical prediction (Che and Mierendorff, 2019), we find no statistically significant differences in the information acquisition decisions between prior beliefs about the MPCR. Selective exposure by the majority to the low-biased source causes beliefs to decline on average. Consequently, the *info* treatment significantly reduces average contributions, while the share of free-riders increases, compared to the *no info* treatment. The treatment effect on efficiency depends on the parameters and the measure of efficiency: under certain conditions, the *info* treatment significantly reduces efficiency compared to the *no info* treatment. Hence, our results show that the provision of information is not always beneficial in the public good context.

We collect self-reported motives for the contribution decision.[3] Based on the elicited motives, we classify participants into different contribution types to investigate how social preferences – as measured by these contribution types – affect information acquisition.[4] We focus on the comparison between participants whom we classify as being "selfish" based on their contribution motives and their contribution behavior compared to those who are more "socially-oriented". We find that, first, selfish participants are less likely to acquire information than more socially-oriented participants. Second, among participants who acquire information, selfish participants are more likely to acquire information from the high-biased source than those who are more socially-oriented. A potential explanation for the observed behavior is that participants seek information confirming with certainty that their individually preferred contribution level is socially desirable.[5] A selfish participant prefers a low contribution and therefore seeks information revealing that the MPCR is low with certainty, which can be provided only by the high-biased source. In contrast, a more socially-oriented participant prefers a high contribution and therefore seeks information revealing that the MPCR is high with certainty, which can be provided only by the low-biased source.

As a robustness check, we replace the endogenous contribution types from our experiment with indexes of actual contribution behavior in the context of environmental protection, which is a public good with uncertain marginal returns.[6] We show that those

---

1 By design, each participant can acquire at most one piece of information; the aim is to mimic limited attention.

2 A participant's prior might determine her perception of the quality of the two information sources (Gentzkow and Shapiro, 2006). Thus, priors might guide her information acquisition choice (Che and Mierendorff, 2019). However, our results contrast with this prediction.

3 Supplementary Material S2 offers an overview of the experimental procedure, and contains the translated instructions and questions as well as screenshots of the original experiment in German.

4 In our experiment, the elicitation of contribution motives occurs only after information acquisition and the contribution decision. To address this endogeneity issue, we use additional data related to real-world public good contributions to environmental protection as a robustness check (see S.1.2.2 in Supplementary Material S1).

5 In this sense, acquiring information from the high-biased source can be interpreted as "looking for low-evidence" and acquiring information from the low-biased source can be interpreted as "looking for high-evidence" (cf. Che and Mierendorff, 2019). An alternative interpretation could be that, by selecting the low-biased source, participants aim to exploit moral wiggle room: If they receive the signal low, they can convince themselves that the true MPCR is low and thus justify not contributing. Some selfish participants may indeed choose the low-biased source for this purpose. However, our analysis of the information acquisition behavior of different contribution types suggests that this mechanism is less common.

6 See S.1.2.2. Table S1.3 shows that environmental contributions are correlated with the contributions to the public good in the experiment in the expected direction: Higher willingness to contribute to environmental protection is associated with a lower probability of contributing zero in the experiment, a higher probability of contributing the entire endowment, and higher average contributions for those who contribute $0 < g_i < 10$. In a previous version of the

participants who are more willing to contribute to environmental protection are also less likely to acquire information from the high-biased source. This result corroborates our experimental finding that social preferences guide strategic information acquisition.

Our findings suggest that strategic information acquisition reduces contributions to the public good through two channels. First, because most participants select the low-biased source, which provides noisy information, posterior beliefs decline on average, leading to lower contributions even from socially-oriented participants. Second, selfish participants are more likely to choose the high-biased source, which allows them to justify their lower contributions.

Our results illustrate how the current media landscape – i.e., biased media and news consumers with limited attention – can impact the provision of public goods with uncertain returns. Environmental protection and COVID-19 containment are two salient examples of such public goods with uncertain returns. First, environmental-friendly behaviors (e.g., switching to electric cars to reduce carbon emissions) have high economic costs and uncertain returns. Citizens can acquire information from heterogeneous sources ranging from environmental activism – such as Fridays for Future – to science denial campaigns, with very different effects on climate change awareness (Björnberg et al., 2017; Baiardi and Morana, 2021). Second, we can interpret social distancing, testing, and vaccinations as contributions to the public good of COVID-19 containment. The returns to these containment measures were initially uncertain since it was not yet clear how the pandemic would evolve. Misleading and false information about the virus and the containment measures spread quickly (Cinelli et al., 2020).

Literature review. Strategic information avoidance and strategic information acquisition have been studied extensively in the dictator game context, providing different explanations for such behavior.[7] Only a few papers study strategic information avoidance and strategic information acquisition in the context of public goods.[8] Aksoy and Krasteva (2020) study a public good game in which participants facing uncertain returns are *exogenously* informed about the true MPCR. They find that participants' contributions react differently to the information depending on participants' generosity and whether they receive "good news" or "bad news" — i.e., whether the true MPCR is above or below the expected MPCR. By contrast, our focus is on the endogenous acquisition of information. We complement their results, showing that social preferences also impact the selection of information sources.[9] Momsen and Ohndorf (2022) study endogenous information acquisition in a framed experiment with repeated carbon-offset purchasing decisions, where the externalities are uncertain. They consider stochastic, potentially unreliable information revelation. They also introduce two information sources, allowing for selective exposure, where participants can acquire one signal from each source.[10] They find evidence for information avoidance but not for selective exposure. Our experiment resembles the "2BR costless" treatment from Momsen and Ohndorf (2022) but differs in several dimensions. First, we study an unframed setting. Second, we allow participants to acquire only one signal to mimic information acquisition under limited attention. Third, we supply participants with the correct updated probability that the MPCR is high after information acquisition. Fourth, we test whether selective exposure depends on the prior probability that the MPCR is high — as suggested by Che and Mierendorff (2019).

2. Experimental design and implementation

2.1. Experimental design

We study a VCM in which the MPCR is stochastic. Participants interact in groups of four. They receive an endowment of 10. Each participant i invests some amount $0 \leq g_i \leq 10$ in a group project, which is the public account. The remaining amount of $10 - g_i$ stays in her private account. The VCM is played only for one round, i.e., participants make exactly one contribution decision. Let ω denote the public good's MPCR, which is the same for all group members. Then, the payoff of participant i is given by

$$\pi_i = 10 - g_i + \omega \sum_{j=1}^{4} g_j \tag{1}$$

such that, if $\omega \in (\frac{1}{4}, 1)$, it is socially efficient to contribute the entire endowment to the public good but individually rational to contribute nothing. With a prior probability of μ, the MPCR is high, ω_h, and with a prior probability of $1 - \mu$, the MPCR is low,

paper (Innocenti and Rohde, 2024), we also showed that the level of contributions to the public good in the experiment is positively correlated with the willingness to voluntarily contribute to (i) environmental protection and (ii) COVID-19 containment. The effect is robust to including controls for socio-demographic variables and comprehension of the experiment. Thus, the contribution behavior observed in the experiment appears to be indicative of actual contributions to a public good, thereby addressing the question of the external validity of our experimental results (Kerschbamer and Müller, 2020).

[7] See Golman et al. (2017) for a comprehensive review of this literature. Dana et al. (2007) introduced the idea that participants exploit a "moral wiggle room" by remaining ignorant about the consequences of their actions to justify selfish behavior. In Spiekermann and Weiss (2016), participants who feel compelled to perform an action implied by a norm can strategically acquire information to manipulate their subjective beliefs, which they use to interpret these normative obligations, thus reducing their subjective normative pressure. The literature has also explored strategic information avoidance in trust games (Van der Weele et al., 2014; Regner, 2018), charitable donations (Exley, 2016), and principal–agent settings (Friedrichsen et al., 2022).

[8] The literature studies the effects of uncertainty about marginal returns on the provision of public goods. See, for instance, Levati et al. (2009) and Fischbacher et al. (2014). We contribute by adding the possibility of receiving information about the MPCR.

[9] Another difference concerns the information revelation process. In Aksoy and Krasteva (2020), informed participants know the true MPCR. Instead, we offer subjects the opportunity to acquire incomplete information, which may reveal the true MPCR.

[10] Thunström et al. (2014) and Momsen and Ohndorf (2020) study information avoidance in similarly framed experiments. In these experiments, there is only one information source that provides a perfectly informative signal that participants can strategically ignore. Thunström et al. (2014) show that participants are less prone to avoid information about social norms than about actual environmental harm. Momsen and Ohndorf (2020) show that participants use information costs as a situational excuse to avoid information that prohibits them from behaving selfishly.

ω_l. We use a value of $\omega_h = 0.5$ for the high MPCR and a value of $\omega_l = 0.1$ for the low MPCR. Thus, the high MPCR ω_h creates a social dilemma: contributing is socially efficient but not individually rational. Instead, contributing is not socially efficient for the low MPCR ω_l; thus, there is no social dilemma. Therefore, selfish and social interests are aligned if the MPCR is low but diverge if the MPCR is high. To study the effect of prior probabilities, we vary $\mu \in \{0.25, 0.5, 0.75\}$ across treatments. For instance, for a risk-neutral participant who makes her contribution decision according to the expected MPCR, full contributions are socially efficient when $\mu = 0.5$ or $\mu = 0.75$, but not when $\mu = 0.25$.[11]

We have two main treatments: *no info* and *info*. In the *no info* treatment, which is our control group, participants do not have the opportunity to acquire further information about the MPCR. They learn about the prior probability (thus, we refer to μ as prior belief later) of the high MPCR and then make their contribution decisions. In the *info* treatment, participants learn about the prior probability of the high MPCR, and they can access one piece of information about the MPCR before making their contribution decisions. They face two information sources with opposite biases, which send one of the two possible signals *high* or *low*.[12] The high-biased source is biased towards sending the signal that the MPCR is high: If the true MCPR is ω_h, the high-biased source always sends signal $S_H = high$. Instead, if the true MPCR is ω_l, the high-biased source sends either signal ($S_H = low$ or $S_H = high$) with equal probability. Analogously, the low-biased source is biased towards sending the signal that the MPCR is low: If the true MCPR is ω_l, the low-biased source always sends signal $S_L = low$. Instead, if the true MPCR is ω_h, the low-biased source sends either signal ($S_L = low$ or $S_L = high$) with equal probability. Participants can acquire one unit of information from only one of the two sources or decide not to acquire any further information about the MPCR. Our information revelation process is an application of Che and Mierendorff (2019) to the special case with one period, no cost of acquiring information, and a probability $\lambda = \frac{1}{2}$ of receiving breakthrough news. In that case, "own-biased learning" is theoretically optimal, i.e., given her prior belief, a participant should acquire information from the source that is biased toward the more likely state of the world.[13]

If a Bayesian participant in the *info* treatment acquires information from the high-biased source and receives signal $S_H = low$ (i.e., breakthrough news), she updates her belief to $\mu'_H = Pr(\omega = \omega_h | S_H = low) = 0$. If she receives signal $S_H = high$, she updates her belief to $\mu'_H = Pr(\omega = \omega_h | S_H = high) = \frac{2\mu}{1+\mu}$. Analogously, when she acquires information from the low-biased source and receives signal $S_L = high$ (i.e., breakthrough news), she updates her belief to $\mu'_L = Pr(\omega = \omega_h | S_L = high) = 1$. If she receives signal $S_L = low$, she updates her belief to $\mu'_L = Pr(\omega = \omega_h | S_L = low) = \frac{\mu}{2-\mu}$.[14] We provide participants with the correct Bayesian posterior beliefs. After acquiring information, the participants in the *info* treatment make their contribution decisions based on their posterior beliefs about the MPCR.

2.2. The German internet panel

The German Internet Panel (GIP) is a long-term online study based on a random probability sample of the general population in Germany aged 16 to 75. The GIP was the central infrastructure project of the Collaborative Research Center (SFB) 884 "Political Economy of Reforms" at the University of Mannheim, which was funded by the German Research Foundation (DFG) from 2010 to 2022. Since 2023, the GIP has been continued as an infrastructure project of the Mannheim University Library. Questionnaire modules are developed by researchers in collaboration with the GIP team. The GIP surveys started in 2012. Panel refreshment samples were recruited in 2014 and 2018, resulting in a current pool of over 6000 potential participants. The online questionnaires are sent to participants on the first day of every other month and remain open for the whole month. The questionnaires take 20-25 min and cover socio-demographic information and various topics, including political attitudes. Until May 2024 (wave 71), the participants received a fixed amount of 4 euros for each completed questionnaire. From then on, the incentive is staggered according to the length of the questionnaire, up to 5 euros for each completed questionnaire. Participants also receive a yearly bonus of up to 10 euros based on completion rates. The GIP data is publicly available in the GIP data archive at the GESIS-Leibniz Institute for the Social Sciences. The original questionnaire documentation for each wave is also available via the GIP data archive.

Our experiment was part of wave 52 of the GIP (March 2021).[15] We use a question from the same wave to assess how difficult the entire questionnaire, including our experiment, was for the participants. We use data on attitudes towards environmental protection from several past waves of the GIP.[16]

---

[11] Table A.2 in Appendix A shows whether it is socially efficient to contribute to the public good or not for each posterior belief.

[12] The information-generating process mimics the optimal information structure by a biased sender in standard information design problems (Kamenica and Gentzkow, 2011). In other words, it is common for a persuader to send a biased recommendation, which is misleading with a positive but sufficiently small probability.

[13] Thus, for a prior belief $\mu = 0.25$, she should acquire information from the low-biased source, and for $\mu = 0.75$ she should acquire information from the high-biased source. For $\mu = 0.5$, she should be indifferent between the two information sources.

[14] Table A.2 in Appendix A summarizes posterior beliefs by treatment, prior, and revealed signal.

[15] Blom et al. (2021c).

[16] For details on the GIP methodology, see Blom et al. (2015, 2016, 2017), Herzing and Blom (2019) and Cornesse et al. (2020). S1.9 includes a detailed overview of this additional data.

2.3. Implementation of the experiment

We implemented the experiment using five survey questions. In the GIP, participants are not used to incentivized economic experiments like ours. Therefore, we deliberately refrained from using standard elements of public good experiments, such as elicitation of conditional contributions or repetition of the VCM over several rounds. Instead, we simplified the game to a one-shot decision and included it in a single survey question.[17] Moreover, we adapted the instructions to be short and understandable for the general population, as participants may not be familiar with probabilities. Therefore, we presented probabilities in terms of frequencies.[18] We provided the correct Bayesian posterior beliefs to those participants who acquired information to reduce cognitive costs and avoid non-Bayesian updating — which could be a confounding factor for our results.

For the random allocation into treatments, we proceeded as follows: we allocated 75% of the participants to the *info* treatment and the remaining to the *no info* treatment.[19] Within each treatment, we allocated one-third of the participants to each prior probability $\mu \in \{0.25, 0.5, 0.75\}$. For each prior, we randomly allocated a share μ of participants to the high MPCR and a share $1 - \mu$ to the low MPCR. For the information revelation, we proceeded as follows: we randomly allocated 50% of the participants to signal *high* and 50% to signal *low*. This variable then decided which signal the chosen source would reveal in those cases where a piece of breakthrough news was possible — e.g., if the MPCR allocated to the participant was high, and she acquired signal S_L.

To incentivize the experiment, we paid out the payoffs from the game to 50 randomly selected groups, i.e., 200 participants. With an endowment of 10 euros, it was possible to earn up to 25 euros depending on the MPCR and the other group members' decisions. The experiment's endowment and potential payoff were quite sizable compared to the payment of 4 euros for a completed questionnaire or the German minimum hourly wage of 9.50 euros in 2021. On average, those participants randomly selected for payment earned 12.62 euros. The lowest payment was 1.70 euros, while the highest was 24.50 euros.

Our questionnaire contained the following parts[20]: First, we informed the participants about the payment procedure. Second, we explained the VCM. We told the participants that they would receive 10 euros on a virtual account and that they could decide how much of this amount to invest in a group project and how much to keep on their virtual account. To reduce the level of abstraction, we called the group project a "gold" project if the MPCR was $\omega_h = 0.5$ and a "silver" project if the MPCR was $\omega_l = 0.1$. Using an example, we also showed how to calculate the return from the group project. We informed those in the *info* treatment that they would later have the opportunity to reveal – potentially – the type of the group project.

Then, those in the *no info* treatment directly proceeded to the contribution stage. Instead, we informed those in the *info* treatment about the information revelation process. Again, to reduce the level of abstraction and increase plausibility, we presented them with four envelopes, as inspired by Spiekermann and Weiss (2016), whose experiment exploits the same information revelation process as ours in a dictator game. Two envelopes were gold, corresponding to the high-biased source, and two were silver, corresponding to the low-biased source.[21] We carefully explained the interpretation of the envelopes by describing what information can be obtained by opening each envelope. In particular, we told participants that it was possible to infer the type of the group project only from one of the four envelopes. Then, the participants answered a comprehension question about the interpretation of the content of the envelopes and then made their information acquisition decisions.[22] They could choose between opening one of the four envelopes or indicating that they do not want to open any envelope.[23] At the contribution stage, those in the *info* treatment received the information about the content of the envelope and the correct Bayesian posterior.[24] All participants were then asked to decide which amount between 0 and 10 euros they wanted to invest in the group project.

After the contribution decision, we elicited potential contribution types in a multiple-choice question by asking about the motives for the contribution decision.[25] For the answer options, we follow the literature that finds that most participants in public good games are either free-riders, unconditional cooperators, or conditional cooperators (Fischbacher et al., 2001; Fischbacher and Gächter, 2010): Participants could indicate that they wanted to maximize their own payoff, maximize the payoff of the entire group, or that they wanted to contribute neither more nor less than other group members. We also included the option to indicate that they had other reasons.

[17] We also used abstract framing and neutral language, avoiding possibly loaded words such as "public good" or "bias" to be able to study the participants' underlying preferences without an associated context.

[18] Note that since the participants are randomly split into groups of pre-determined size to allocate them into the treatments, using frequencies is mathematically correct and does not constitute deception.

[19] This allocation ensured a sufficiently large number of observations for each posterior belief.

[20] See S2 for the instructions and questions. The codebook to GIP wave 52 is available via the GESIS Data Archive and contains the original programming instructions (see Blom et al., 2021c).

[21] We randomized the order in which the four envelopes were shown to the participants. The order of the envelopes in the information acquisition decision remained identical to the order shown in the instructions.

[22] We randomized the order of the answer options to the comprehension question.

[23] Depending on what they chose, we asked their minimum willingness to pay for the chosen envelope or their minimum willingness to accept to open an envelope if they chose not to. We did not incentivize these questions.

[24] Once a participant reached the contribution stage, she could not go back to the information stage, making it impossible to open more than one envelope. Table A.2 in Appendix A includes a summary of posterior beliefs by prior and revealed signal and their implication for whether contributing is socially efficient.

[25] We randomized the order of answer options to the question about contribution motives.

2.4. Discussion of the framing of information sources

Our results suggest that participants aim to achieve certainty through their information acquisition choice, i.e., they acquire information from the source which might reveal with certainty that their individually preferred contribution level is socially efficient. Therefore, a legitimate question is whether potential framing effects might drive this result. Our instructions for the information stage (see Supplementary Material S2) are inspired by Spiekermann and Weiss (2016), whose experiment exploits the same information revelation process as ours in a dictator game. In their experiment, "the dictators have an unannounced opportunity to acquire information that may inform them about the type of their receiver. [...] If the receiver is a high performer, the information is in one of two envelopes called 'gold envelopes'. If the receiver is a low performer, the information is in one of two envelopes called 'silver envelopes' " (p. 178). In their instructions, Spiekermann and Weiss (2016) focus exclusively on which information reveals the type of the receiver with certainty (see their Web Appendix I, p. 2). In our instructions for the information stage, and especially in the comprehension question, we emphasize how to achieve certainty as well. However, while ensuring that the explanation remains sufficiently simple, we still provide all relevant information to avoid leaving out any part of the story. In particular, we describe in detail what information can be obtained by opening each envelope. First, we explain what card an envelope can contain in each state of the world.[26] Then, we explain how each card-envelope combination can be interpreted.[27] After the information stage, we provide the correct interpretation of the revealed card-envelope combination and the exact Bayesian posterior belief.

We did not run a treatment for different ways of presenting the information, and hence cannot test for framing effects directly. However, if our emphasis on achieving certainty drove participants' behavior, we would expect their contribution behavior to be affected *only* by the revelation of the true state of the world (i.e., to their posterior belief becoming 0 or 1). In contrast, their contribution behavior would not be affected if the revealed information does not lead to certainty about the true state of the world. However, the experimental results suggest this is not the case.

3. Results

A total of 4370 participants took part in GIP wave 52. Of those participants, 100 broke off the survey, and others decided not to participate in our experiment or completed it partially. We dropped all participants who skipped information acquisition or the contribution to the public good, resulting in an overall sample size of 4187 participants. In this sample, the average age is around 52 years, 48% of the participants are female, and 34% have an academic education, i.e., a Bachelor's degree or higher.[28] We now present the results of our experiment in terms of descriptive statistics. Then, we perform a regression analysis that shows how the contribution types elicited in our questionnaire affect information acquisition decisions and how strategic information acquisition, in turn, affects voluntary contributions. We corroborate our findings by investigating whether the participant's information acquisition decisions correlate with the willingness to voluntarily contribute to environmental protection. These additional results are in Supplementary Material S1.2.

3.1. Descriptive results

Information acquisition. Most participants in the *info* treatment (87%) choose to acquire a signal from either of the two sources; only a small share (13%) chooses not to acquire any information. Among those participants who do acquire information, the majority (65%) choose the low-biased source.[29] The finding that signal S_L is the most frequent information acquisition choice aligns with the results of Spiekermann and Weiss (2016). Fig. 1 shows that signal choices do not differ in a statistically significant manner between the three prior probabilities. How information acquisition choices affect posterior beliefs helps to understand how information acquisition affects voluntary contributions compared to the *no info* treatment. The beliefs of 41% of the participants in the *info* treatment decline compared to their prior. Only 8% of the participants found out that the true MPCR is low with certainty, while 15% revealed that the true MPCR is high with certainty.[30]

Voluntary contributions. We study how the information treatment affects three main features of the distribution of voluntary contributions to the public good: average contributions, the share of free riders, and the share of participants who contribute their entire endowment.

In the *no info* treatment, participants contribute on average 6.94 euros to the public good. The *info* treatment significantly reduces the average contributions to 6.13 euros (Wilcoxon rank sum test, $p < 0.0001$), which corresponds to a reduction by 8.1% of the endowment. The treatment effect is statistically significant for all prior beliefs, and average contributions do not differ significantly between prior beliefs in either treatment.[31]

[26] For instance, we write: "Only if the group project is GOLD, exactly one of the two silver envelopes contains a gold card and hence reveals the type of the group project. Otherwise, the silver envelopes always contain a silver card."

[27] For instance, we write: "Only if you find a gold card in a silver envelope, you can be completely certain that the group project is a GOLD project. If you find a gold card in a gold envelope, you can be more certain that it is a GOLD project than without this information, but you cannot be completely certain."

[28] The demographic characteristics are balanced across treatments. See Table A.1 in Appendix A.

[29] We report the results concerning the participants' willingness to pay/accept in S1.2.1.

[30] Figure S1.1 in S1.1 displays the changes in the posterior beliefs across treatments. Table A.2 in Appendix A summarizes the posterior beliefs by treatment, prior, and revealed signal, and how many participants ended up with a given posterior belief.

[31] Figure S1.2 in S1.1 displays average contributions for the three different prior beliefs.

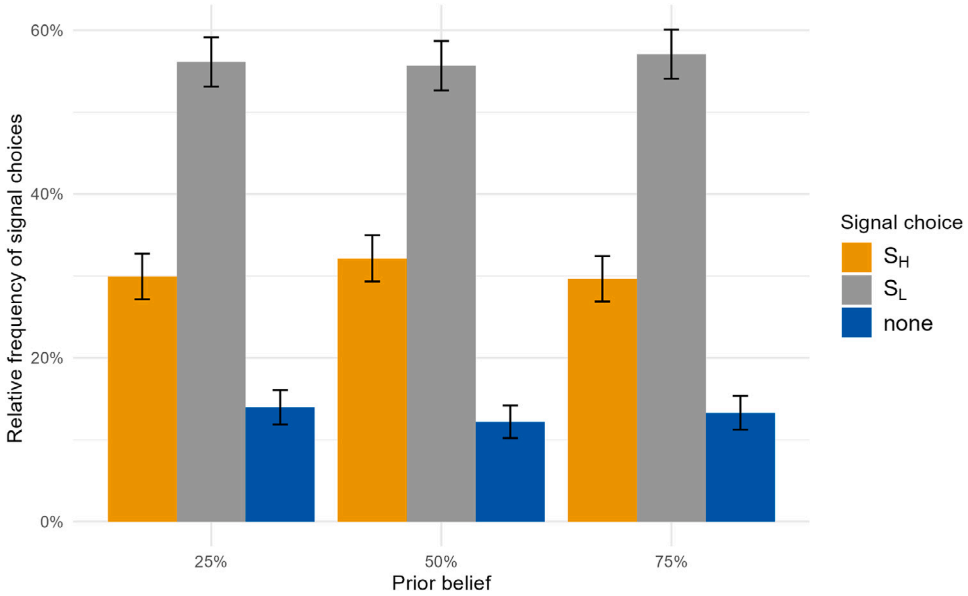

Fig. 1. Information acquisition choices for the different prior beliefs μ. Error bars represent 95% confidence intervals.

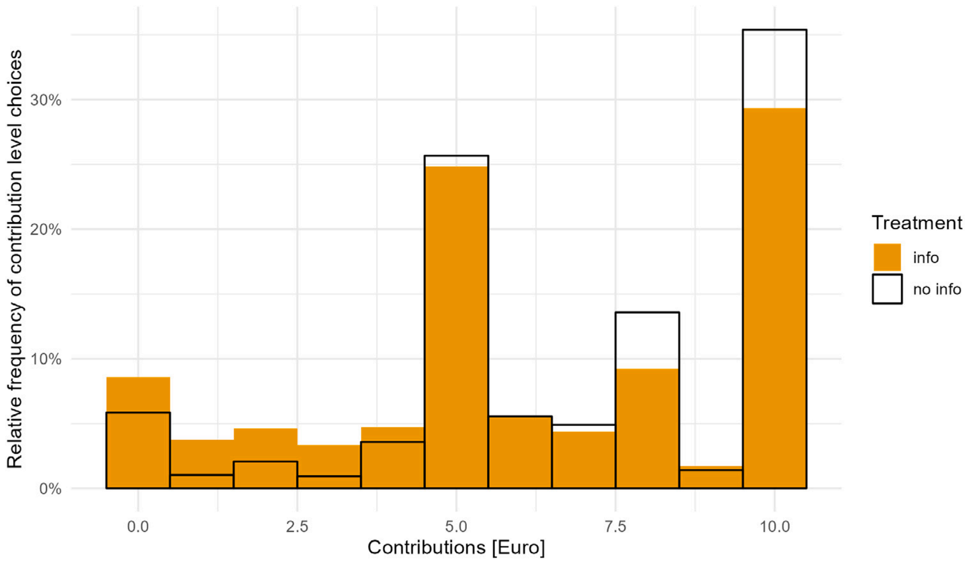

Fig. 2. The distribution of contributions to the public good in the two treatments.

Fig. 2 displays the distribution of voluntary contributions to the public good in the two treatments. In both treatments, the most frequently chosen contribution levels are 10 euros – the entire endowment – and 5 euros — half of the endowment. We compare the distribution of contributions in the two treatments and observe a shift to the left following information acquisition, resulting in lower contribution levels having higher frequencies.[32]

In particular, only 6% of the participants contribute zero in the *no info* treatment. This share increases to 9% in the *info* treatment, which is a significant difference (two-proportions z-test, $p = 0.0066$). Similarly, the share of participants who contribute their entire endowment of 10 euros significantly decreases from 35% in the *no info* treatment to 29% in the *info* treatment (two-proportions z-test, $p = 0.0003$).[33]

[32] Figure S1.6 in S1.1 shows that the same shift to the left of the distribution of contributions in the *info* treatment compared to the *no info* treatment occurs independently of whether the true MPCR is high or low.

The reduction in contributions in the *info* treatment, compared to the *no info* treatment, is driven by the signal choices as well as revealed information. Those participants who do not acquire information contribute only 4.45 euros on average, compared to an average of 6.94 euros in the *no info* treatment, which is a significant difference (Wilcoxon rank sum test, $p < 0.0001$). Those who select the high-biased source contribute 6.53 euros on average, which is driven by those who reveal that the true MPCR is low: they significantly reduce their contribution to 5.79 euros on average (Wilcoxon rank sum test, $p < 0.0001$). Instead, the average contribution of those who receive the signal $S_H = high$ amounts to 6.79 euros, which is not a significant difference to the *no info* treatment (Wilcoxon rank sum test, $p = 0.2413$). Those who select the low-biased source contribute 6.30 euros on average, which is driven by those who receive the signal $S_L = low$: they significantly reduce their contribution to 5.82 euros on average (Wilcoxon rank sum test, $p < 0.0001$). In contrast, the average contribution of those who reveal that the true MPCR is high significantly increases to 7.62 euros (Wilcoxon rank sum test, $p < 0.0001$). The former effect dominates because the signal $S_L = low$ is obtained more frequently. Average contributions of participants who do not reveal the true MPCR do not differ significantly between posterior beliefs, neither among those who receive the signal $S_H = high$, nor among those who receive the signal $S_L = low$.[34]

Concerning the motives behind their contribution decisions, most participants indicated exactly one motive[35]: 12% want to maximize their own payoff, 45% want to maximize the payoff of the entire group, 21% want to contribute neither more nor less than other group members, and 13% had "other reasons".[36] Among those who indicated exactly one of the three main motives, the participants interested in maximizing the group's payoff contribute the largest amount on average, are least likely to contribute zero, and are most likely to contribute the entire endowment.[37] Interestingly, those who claim to be interested in contributing neither more nor less than other group members appear to be slightly more selfish than those who indicate that they want to maximize their own payoff: They contribute less on average and are less likely to contribute the entire endowment.[38] Therefore, we merge these two types of participants into a single "selfish" category.[39] In the following analysis, we will thus focus on these "selfish" types, compared to those who indicated that they want to maximize the group's payoff as their unique answer, whom we interpret as being the more "socially-oriented" types.

Social efficiency. We investigate how the information treatment affects the social efficiency of contributions. Recall that if the true MPCR is high (ω_h), it is socially efficient to contribute the entire endowment to the public good. Instead, if the true MPCR is low (ω_l), it is socially efficient to contribute nothing. Therefore, we define the ex-post efficiency E_p of a contribution $g \in [0, 10]$ as

$$E_p(g, \omega) = \begin{cases} 1 - \frac{g}{10} \text{ if } \omega = \omega_l \\ \frac{g}{10} \text{ if } \omega = \omega_h \end{cases}$$

where $E_p(\cdot) \in [0, 1]$. We find that the average ex-post efficiency level is 0.51 in the *no info* treatment and 0.54 in the *info* treatment, where the difference is significantly different from zero (Wilcoxon rank sum test, $p = 0.0148$). Hence, overall, the information treatment increases ex-post efficiency.[40] However, the effect of the information treatment on the ex-post efficiency of contributions is heterogeneous and depends on the nature of the public good. Crucially, the information treatment leads to lower contributions and, thus, can increase ex-post efficiency only if the true MPCR is low. Indeed, we find that the treatment effect on ex-post efficiency is only significantly positive when the true MPCR is low.[41]

Because ex-post efficiency is affected by noise due to the realization of the state of the world, we also consider ex-ante expected efficiency. Given posterior belief μ', the expected MPCR of the public good is $\mathbb{E}(\omega) = \mu'\omega_h + (1 - \mu')\omega_l$. Then, the expected payoff

33 Comparing our results to the literature, we find that our sample from the general population seems more altruistic than the typical sample of students in the laboratory. See, for instance, Fischbacher et al. (2001).

34 Since average contributions do not differ significantly between prior beliefs in the *no info* treatment either, this result could indicate that participants only care about being uncertain, but not about the exact probability of the true MPCR being high.

35 Among those who indicated more than one of the three main motives, maximizing their own and the group's payoff is the most frequent combination. However, the contribution behavior of those who indicated more than one motive is erratic and cannot be rationalized. For example, we observe that those answering that they want to maximize their own and the group's payoff contribute more on average than those who only want to maximize the group's payoff. The most likely explanation is confusion and lack of understanding. We keep these contribution types in our sample for completeness but refrain from interpreting their effect on information acquisition in the following regression analysis.

36 We included an open-answer field for those having "other reasons" behind their contribution decisions. Some participants indicate either risk-averse behavior (not investing because of the uncertainty about the returns) or risk-seeking behavior (investing the entire endowment in gambling). There is also a tendency to evenly split the money between the private and public accounts, which might explain the high share of investments of 5 euros. However, most open answers show confusion and a lack of comprehension. Therefore, we will not focus on the "other reasons" category in the following analysis.

37 Figure S1.3 in S1.1 shows how the contribution decisions differ by contribution motive.

38 When designing the question to elicit contribution types, we thought of "contributing neither more nor less than other group members" as reciprocity concerns, in the sense of conditional cooperation. However, we cannot identify reciprocating behavior with our data because we neither elicit participants' beliefs about others' contributions nor their conditional contributions in a contribution table (Fischbacher et al., 2001). Potentially, the reciprocity motive might have been chosen as an excuse by selfish participants who do not want to appear so.

39 We classify all participants who answered that they are interested in maximizing their own payoff, or in contributing neither more nor less than other group members, or in both as "selfish".

40 In a previous version of the paper, we also studied ex-post social welfare (Innocenti and Rohde, 2024). We found that the treatment effect on social welfare is overall positive but heterogeneous: when the true MPCR is low, social welfare increases by up to 12.4%, whereas when the true MPCR is high, social welfare decreases by up to 5.3%.

41 Figure S1.4 in S1.1 shows the treatment effect on ex-post efficiency by MPCR for each prior belief.

of participant i is given by

$$\mathbb{E}(\pi_i) = 10 - g_i + \mathbb{E}(\omega) \sum_{j=1}^{4} g_j \tag{2}$$

such that, if $\mathbb{E}(\omega) \in (\frac{1}{4}, 1]$, it is socially efficient to contribute the entire endowment, whereas if $\mathbb{E}(\omega) \in [0, \frac{1}{4}]$ it is socially efficient to contribute nothing. Thus, the ex-ante expected efficiency E_a of a contribution $g \in [0, 10]$, given posterior belief μ', is

$$E_a(g, \mu') = \begin{cases} 1 - \frac{g}{10} \text{ if } \mu' \in [0, \frac{3}{8}] \\ \frac{g}{10} \text{ if } \mu' \in (\frac{3}{8}, 1] \end{cases}$$

where $E_a(\cdot) \in [0, 1]$.[42] We find that the average ex-ante efficiency level is 0.57 in the *no info* treatment and 0.56 in the *info* treatment, where the difference is not significantly different from zero (Wilcoxon rank sum test, $p = 0.2147$). Hence, overall, the information treatment does not significantly affect ex-ante efficiency. However, the effect of the information treatment on the ex-ante efficiency of contributions is heterogeneous and depends on the prior belief μ. The reduction of contributions in the *info* treatment increases ex-ante efficiency only for the prior belief $\mu = 0.25$, whereas it significantly reduces ex-ante efficiency for $\mu = 0.5$ and $\mu = 0.75$.[43]

3.2. Regression analysis

We now use regression analysis to answer two questions about the interplay of selective exposure and voluntary contributions. How do contribution types affect information acquisition decisions? How does strategic information acquisition affect voluntary contributions in the *info* treatment compared to the *no info* treatment?

Information acquisition. From an econometric perspective, we study information acquisition as if it consists of two separate decisions. First, each participant decides whether to acquire a signal or not. Second, if she acquires a signal, she chooses between S_H and S_L. Thus, we estimate two probit regressions that model these two decisions separately.[44] Table 1 presents the probit estimates of the marginal effects of priors and contribution types on the decision about whether to acquire information or not. Table 2 presents the same effects on the decision about which signal (S_H or S_L) to acquire among those who acquired information.[45] For the contribution motives, the reference category is those who want to maximize the group's payoff.

Table 1 shows that those classified as being "selfish" are less likely to acquire information than those who want to maximize the group's payoff. Table 2 shows that, if they acquire information, they are more likely to acquire signal S_H. Both effects remain significant, at least at the 5% level, when controlling for the comprehension and perceived difficulty of the experiment. Priors affect neither information acquisition decision in a statistically significant manner.[46]

So far, we studied the effect of social preferences – measured by the contribution types – on information acquisition based on the participants' *stated* preferences elicited through our questionnaire. Because the classification of participants into different types occurs only after information acquisition, this measure suffers from endogeneity problems. To address this issue, we study the same questions based on the participants' *revealed* preferences demonstrated by their behavior in real-world public good contexts. To do so, we return to one of the two salient examples of public goods with uncertain marginal returns introduced at the beginning: environmental protection.[47] We exploit several questions from previous GIP waves to measure participants' willingness to contribute to environmental protection (see Supplementary Material S.1.2.2 for the analysis and Supplementary Material S1.9 for a description of the variables and variable selection). We find that, although the willingness to contribute to environmental protection does not affect the decision to acquire information, among those who acquire information, a higher willingness to contribute to environmental protection correlates with the decision to acquire information from the low-biased source. This finding corroborates our result that underlying social preferences guide strategic information acquisition. Overall, we show that more socially-oriented (selfish) participants – either identified by their stated preferences or their higher (lower) willingness to contribute to environmental protection – are more likely to acquire information from the low-biased (high-biased) source.

---

42 Table A.2 in Appendix A shows whether contributing to the public good is ex-ante efficient for each posterior belief obtainable in the different treatments.

43 Figure S1.5 in S1.1 shows the average ex-ante efficiency level in the two treatments by prior belief.

44 Alternatively, we can study the overall decision problem between the three options of acquiring no signal, acquiring S_H, or acquiring S_L using multinomial logit regression. This alternative approach produces results that are consistent with those of our main specification (see Table S1.13 in S1.3).

45 The full regression tables, including the coefficients for the other contribution motives, comprehension, and difficulty, are in S1.4.

46 We conduct several robustness checks to test whether potential comprehension problems could drive our results. First, we re-run the regressions on the subsample of those participants who indicated they did not find the questionnaire difficult. Second, we consider the response times — i.e., the time a participant spent on each question page, including the instructions. We drop the top 10% and the bottom 10% of the distribution concerning the time spent on the instructions for the public good game. Third, we use the subsample of those who correctly answered the comprehension question about the information revelation process. S1.8 includes all tables for these robustness checks. The two main findings are robust to these modifications.

47 In a previous version of the paper, we also considered the example of COVID-19 containment (Innocenti and Rohde, 2024). However, the correlation between information acquisition and voluntary contributions to COVID-19 containment is less significant than for environmental protection. The reason could be that these two public goods are very different: In particular, the willingness and ability to contribute to COVID-19 containment are affected by more external factors and are tangled together with health concerns and personal liberty issues, which makes it a more imperfect measure of one's tendency to contribute to public goods.

Table 1
Probit model for the decision to acquire information.

	Dependent variable:			
	Acquired information			
	Probit			
	(1)	(2)	(3)	(4)
Prior = 0.25	−0.018	−0.012	−0.011	−0.011
	(0.015)	(0.014)	(0.014)	(0.014)
Prior = 0.75	−0.011	−0.013	−0.008	−0.008
	(0.015)	(0.014)	(0.014)	(0.014)
Selfish		−0.093***	−0.070***	−0.069***
		(0.014)	(0.013)	(0.013)
Constant	–	–	–	–
Other motives	No	Yes	Yes	Yes
Comprehension	No	No	Yes	Yes
Difficulty	No	No	No	Yes
Observations	3127	3111	3111	3100
Log likelihood	−1216.005	−1131.738	−1028.666	−1023.624

Note: *p < 0.1; **p < 0.05; ***p < 0.01.
All columns report marginal effects, with robust standard errors in parentheses. The sample is the subsample of those in the *info* treatment. The dependent variable is a binary indicator variable which takes the value 1 if the participant chose to acquire either of the two signals, and the value 0 if the participant did not acquire any signal. *Prior* is a categorical variable with 0.5 as the omitted reference category. *Selfish* and *other motives* belong to the same categorical variable which captures the motives behind the contribution decision, with *group payoff* as the omitted reference category. The control variable *comprehension* captures whether the participant answered the comprehension question correctly, and *difficulty* captures the perceived difficulty of the entire questionnaire. The number of observations in columns 2–4 is reduced because some participants did not answer the question about the contribution motives or the question about the difficulty of the questionnaire.

Table 2
Probit model for the decision to acquire signal S_H among those who acquire information.

	Dependent variable:			
	Acquired signal S_H			
	Probit			
	(1)	(2)	(3)	(4)
Prior = 0.25	−0.018	−0.017	−0.018	−0.020
	(0.023)	(0.023)	(0.022)	(0.022)
Prior = 0.75	−0.024	−0.023	−0.028	−0.030
	(0.022)	(0.022)	(0.022)	(0.022)
Selfish		0.059***	0.049**	0.054***
		(0.021)	(0.021)	(0.021)
Constant	–	–	–	–
Other motives	No	Yes	Yes	Yes
Comprehension	No	No	Yes	Yes
Difficulty	No	No	No	Yes
Observations	2716	2707	2707	2697
Log likelihood	−1761.147	−1748.796	−1701.577	−1688.004

Note: *p < 0.1; **p < 0.05; ***p < 0.01.
All columns report marginal effects, with robust standard errors in parentheses. The sample is the subsample of those who acquired information. The dependent variable is a binary indicator variable which takes the value 1 if the participant acquired signal S_H, and the value 0 if the participant acquired signal S_L. *Prior* is a categorical variable with 0.5 as the omitted reference category. *Selfish* and *other motives* belong to the same categorical variable which captures the motives behind the contribution decision, with *group payoff* as the omitted reference category. The control variable *comprehension* captures whether the participant answered the comprehension question correctly, and *difficulty* captures the perceived difficulty of the entire questionnaire. The number of observations in columns 2–4 is reduced because some participants did not answer the question about the contribution motives or the question about the difficulty of the questionnaire.

Voluntary contributions. To analyze how strategic information acquisition affects voluntary contributions in the *info* treatment compared to the *no info* treatment, we performed several regressions with the signal choices and the revealed information as explanatory variables. Fig. 2 shows that the distribution of contributions is roughly continuous, with significant mass at the endpoints. Therefore, we study three main features of the distribution of contributions: the probability of contributing zero, the

Table 3
Three-part model for contributions.

	Dependent variable:								
	Zero contribution			Contributions			Full contribution		
	Probit			*Tobit*			*Probit*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Info	0.026***			−0.648***			−0.061***		
	(0.009)			(0.083)			(0.017)		
Prior = 0.25	0.029***	0.019**	0.019**	0.030	0.100	0.153*	−0.012	−0.0001	0.010
	(0.010)	(0.009)	(0.009)	(0.094)	(0.090)	(0.089)	(0.017)	(0.017)	(0.017)
Prior = 0.75	0.018*	0.013	0.016*	0.145	0.168*	0.119	0.031*	0.034**	0.021
	(0.010)	(0.009)	(0.009)	(0.094)	(0.088)	(0.088)	(0.018)	(0.017)	(0.017)
Acquired signal S_H		0.002			−0.463***			−0.003	
		(0.010)			(0.103)			(0.020)	
Acquired signal S_L		−0.001			−0.605***			−0.045***	
		(0.009)			(0.088)			(0.017)	
No signal acquired		0.159***	0.159***		−0.986***	−0.992***		−0.030	−0.033
		(0.019)	(0.019)		(0.161)	(0.160)		(0.027)	(0.027)
Posterior = 1			−0.006			0.005			0.083***
			(0.014)			(0.143)			(0.026)
Posterior = 0			0.047**			−0.812***			−0.029
			(0.019)			(0.183)			(0.032)
Posterior increased			−0.017			−0.344***			0.005
			(0.010)			(0.110)			(0.022)
Posterior reduced			0.0002			−0.759***			−0.095***
			(0.010)			(0.092)			(0.018)
Constant				5.729***	6.236***	6.232***			
				(0.087)	(0.121)	(0.120)			
Motives	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Difficulty	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Observations	4187	4153	4153	2567	2544	2544	4187	4153	4153
Log likelihood	−1141.922	−874.169	−867.264	−5364.466	−5164.037	−5145.304	−2577.495	−2345.903	−2318.268

Note: *p < 0.1; **p < 0.05; ***p < 0.01.
Robust standard errors in parentheses. Columns 1–3 and 7–9 report marginal effects. *Zero contribution* is a binary indicator variable. *Contributions* is the level of contributions (in euros) for the subset of participants who contributed $0 < g_i < 10$. *Full contribution* is a binary indicator variable. The truncated normal model in columns 4–6 is estimated on the subsample of those who contributed $0 < g_i < 10$. *Prior* is a categorical variable with 0.5 as the omitted reference category. *Signal choice* and *posterior* are categorical variables with "no info treatment" as the omitted reference category. *Motives* captures the different contribution motives, and *difficulty* captures the perceived difficulty of the entire questionnaire.

probability of contributing the entire endowment, and the average contribution for those who contribute $0 < g_i < 10$. We use a three-part model to model these three distribution features separately. This model provides the highest possible flexibility by allowing separate mechanisms to determine the three decisions of interest.[48] Table 3 summarizes the three-part model.[49] We estimate a probit regression to model the decision to contribute zero, a truncated normal model for the contribution level on the subsample of participants who contribute $0 < g_i < 10$, and another probit regression to model the decision to contribute the entire endowment.

For each part, we report three different specifications of the explanatory variables. First, we consider the overall effect of the *info* treatment on the three decisions, compared to the *no info* treatment (columns 1, 4, and 7). Second, to gain insight into the mechanisms behind this treatment effect, we include the signal choices (columns 2, 5, and 8) and the changes in the posterior beliefs (columns 3, 6, and 9).[50] Because the contribution types affect both the signal choice and the contribution decisions, we include them as control variables. We additionally control for the perceived difficulty of the questionnaire.

The three-part model highlights several results. The probability of contributing zero is higher in the *info* treatment than in the *no info* treatment, while both the amount contributed among those with $0 < g_i < 10$ and the probability of contributing the entire endowment are smaller in the *info* treatment than in the *no info* treatment. Those who did not acquire information drive the increase in zero contributions in the *info* treatment, whereas those who acquire signal S_L drive the decrease in full contributions. Among

[48] There exist alternative models potentially suitable for our data. The two-limit Tobit model (Table S1.16 in S1.3) takes into account the pileups at the endpoints. However, it does not allow separate mechanisms to determine the different decisions. The two-part hurdle model (Tables S1.14 and S1.15 in S1.3) models only the participation decision separately from the amount decision, but it does not consider the decision to contribute the entire endowment. Our main results are robust when using these alternative models. Comparing the values of the log-likelihood function reveals that the three-part model reported in this section provides the best model fit. S1.6 includes details about the model selection process.

[49] The full regression tables for all three parts, including the coefficients for the contribution motives and difficulty, are in S1.4.

[50] To test whether the effects of information on the contribution decisions differs by prior belief, we also estimated models for all three parts in which we included interactions between prior beliefs and signal choices, or prior beliefs and posterior beliefs (Tables S1.6 - S1.11 in S1.3). Our main results are robust to including these interaction effects. In each case, a likelihood ratio test fails to reject the null hypothesis that the more complex model, including the interaction effects, fits the data as well as the nested model without the interactions. Therefore, we conclude that adding the interaction terms does not improve the model, so we focus on the simpler model here.

those who contribute $0 < g_i < 10$, both those who acquire any signal and those who do not reduce their contributions compared to those in the *no info* treatment. The changes in posterior beliefs mainly affect the contribution decisions in the expected direction. Obtaining a posterior belief of $\mu'_L = 1$ (i.e., revealing that the true MPCR is high) increases the probability of contributing the entire endowment compared to the *no info* treatment. Obtaining a posterior belief of $\mu'_H = 0$ (i.e., revealing that the true MPCR is low) increases the probability of contributing zero and reduces contributions among those with $0 < g_i < 10$, compared to the *no info* treatment. A reduced posterior $0 < \mu < \mu'_L$ reduces the probability of contributing the entire endowment and reduces contributions among those with $0 < g_i < 10$, compared to the *no info* treatment. Only the negative effect of an increased posterior $\mu < \mu'_H < 1$ on the level of contributions among those with $0 < g_i < 10$ is unexpected.[51] This effect is most likely caused by the selection at the information stage — because those who acquire signal S_H are generally less willing to contribute than those in the *no info* treatment, independent of their beliefs.[52]

We again repeat the analysis using the participants' *revealed* preferences demonstrated by their behavior in the context of voluntary contributions to environmental protection instead of their *stated* preferences elicited through our questionnaire as a control variable (see Table S1.3 in Supplementary Material S1.2). The results based on environmental contributions align with the results in Table 3. The unexpected negative effect of an increased posterior remains, highlighting that both elicited contribution types as well as environmental contributions are imperfect measures of social preferences in the context of our experiment.

To summarize, the information treatment reduces the level of contributions independently of the acquired signal. Because our measures of social preferences are imperfect, we cannot fully disentangle the mechanisms driving the interplay between strategic information acquisition and contribution behavior. However, our results suggest that the mechanisms driving the reduction in contributions differ depending on the signal choice. Acquiring information from the low-biased source reduces posterior beliefs on average. Participants acquiring information from this source are more likely to be socially-oriented and, hence, less willing to contribute if they are less confident that the MPCR is high. Therefore, they contribute less when they receive signal low. However, they contribute more when they receive the signal high from the low-biased source – and, hence, become sure that the MPCR is high – but this happens less frequently. The sum of these effects explains the reduction in contributions compared to the *no info* treatment. Acquiring information from the high-biased source increases posterior beliefs on average. Participants acquiring information from this source are more likely to be selfish and, hence, less willing to contribute to the public good independent of their beliefs. Indeed, since the ultimate goal of selfish participants is to avoid contributing and they process information strategically to justify this behavior, they do not respond to higher beliefs – following signal high – by increasing their contributions. However, they contribute less when they receive the signal low from the high-biased source — and, hence, become sure that the MPCR is low. Therefore, although posterior beliefs are higher on average, contributions are generally lower than those in the *no info* treatment.

4. Conclusion

In this paper, we investigate whether strategic information acquisition can harm the provision of public goods in a setting where decision-makers with limited attention choose between biased information sources. Most participants acquire information from the low-biased source. In particular, our analysis suggests that participants sort into the audiences of different information sources following their social preferences. Selfish participants are less likely to acquire information. If they acquire information, they are more likely to choose the high-biased source, which might reveal that the true MPCR is low, justifying low contributions. In contrast, socially-oriented participants are more likely to acquire information from the low-biased source – if they acquire information – for the opposite reason. The fact that selfish participants strategically seek information that justifies selfish behavior – or avoid information that compels them to behave more generously – has already been documented in the context of dictator games. Observing the same behavior in a public good game has more far-reaching consequences because it impacts social efficiency.

Our results show that more information is not always better: the possibility of accessing noisy information from biased sources can be harmful in the public good context. Compared to the case where no further information is available, strategic information acquisition leads to lower contributions and more free-riding, which can reduce social efficiency. Therefore, a policymaker concerned with providing a public good that requires citizens' investments, such as improving environmental quality or containing a virus, should consider the information environment. The previous considerations leave several open questions for future research: How can desirable collective outcomes, such as the provision of a public good, be reached despite strategic information acquisition? Moreover, it might be the case that a policymaker is more informed about the actual state of the world than the citizens – e.g., because she is directly in contact with scientists – and that she might want to persuade citizens of her belief. How can she credibly convey her information when other sources might make different, unreliable claims? This question is especially relevant during times of low trust in governments and general skepticism towards science.

[51] We estimate the three-part model on the two subsamples of those who acquired signal S_H and those who acquired signal S_L separately, using priors and changes in posterior beliefs as explanatory variables (Table S1.12 in S1.3). Then, the information revelation is exogenous and random in each subsample by construction. The results show that the participants react in the expected direction when they reveal the true MPCR.

[52] Another potential explanation might be confusion among the participants concerning the information received. Our robustness checks address this potential problem. First, we re-run the regression analysis using the subsample of participants who did not find the questionnaire difficult (Table S1.32 in S1.8). Second, we make use of the response times contained in the dataset: we drop from the sample the bottom 10% and top 10% of the distribution concerning the time spent on the instructions for the public good game (Table S1.35 in S1.8). In both cases, the sign and significance of the coefficients remain the same.

Table A.1
Demographic characteristics across treatments.

Treatment	Prior belief	n	Average age	Female	Academic education
No info	0.25	352	52	47%	34%
No info	0.5	357	53	49%	34%
No info	0.75	351	51	47%	38%
Info	0.25	1046	52	48%	34%
Info	0.5	1042	52	49%	35%
Info	0.75	1039	52	49%	32%

n is the number of participants in each treatment arm. The average age is calculated from an age variable denoting the midpoint (in years) of the six age categories. An academic education is classified as having at least a Bachelor's degree.

Table A.2
Posterior beliefs.

Treatment	Prior belief	Signal	Posterior belief	n	Is contributing ex-ante efficient?
No info	0.25	n/a	0.25	352	No
No info	0.5	n/a	0.5	357	Yes
No info	0.75	n/a	0.75	351	Yes
Info	0.25	None	0.25	145	No
Info	0.25	$S_H = low$	0	112	No
Info	0.25	$S_H = high$	0.4	201	Yes
Info	0.25	$S_L = low$	0.14	508	No
Info	0.25	$S_L = high$	1	76	Yes
Info	0.5	None	0.5	127	Yes
Info	0.5	$S_H = low$	0	94	No
Info	0.5	$S_H = high$	0.66	240	Yes
Info	0.5	$S_L = low$	0.33	431	No
Info	0.5	$S_L = high$	1	149	Yes
Info	0.75	None	0.75	138	Yes
Info	0.75	$S_H = low$	0	41	No
Info	0.75	$S_H = high$	0.86	266	Yes
Info	0.75	$S_L = low$	0.6	349	Yes
Info	0.75	$S_L = high$	1	242	Yes

n is the number of participants that ended up with the respective posterior belief.

Funding sources

Federico Innocenti gratefully acknowledges funding by the German Research Foundation (DFG) through CRC TR 224 (Project B05). Linnéa Marie Rohde gratefully acknowledges funding by the German Research Foundation (DFG) through SFB 884 (Project number 139943784, Project A7).

Declaration of competing interest

The authors declare no conflict of interest.

Acknowledgments

We are grateful to the GIP team for their collaboration in conducting the experiment. We thank Hans Peter Grüner for his encouragement and support. We thank Duk Gyoo Kim and Wladislaw Mill for their helpful comments. We thank seminar participants in Mannheim and Verona as well as at the 15th RGS Doctoral Conference, the 2022 MaCCI Annual Conference, the 9th CRC TR 224 Conference, the 2022 North American Summer Meetings of the Econometric Society and the XI IBEO Workshop.

Appendix A. Description of the sample

Table A.1 shows how many participants were allocated to each treatment arm and their demographic characteristics captured by age, gender, and academic education.

Table A.2 summarizes the posterior beliefs by treatment, prior, and revealed signal. It also shows how many participants ended up with a given posterior belief and for which posterior beliefs it is socially efficient to contribute to the public good.

Appendix B. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jebo.2025.107081.

Data availability

This article uses data from waves 38, 41, 44, 48, 49, and 52 of the German Internet Panel (GIP; DOIs: https://doi.org/10.4232/1.13391, https://doi.org/10.4232/1.13464, https://doi.org/10.4232/1.13614, https://doi.org/10.4232/1.13681, https://doi.org/10.4232/1.13682, https://doi.org/10.4232/1.13794), (Blom et al., 2019, 2020a,b, 2021a,b,c). The data is available for academic research and teaching via the GIP data archive at the GESIS-Leibniz Institute for the Social Sciences, with the data depositors written authorization. The response times used for our robustness checks are only available via On-Site Data Access (ODA). A study description can be found in Blom et al. (2015). The GIP was funded by the German Research Foundation (DFG) as part of the Collaborative Research Center 884 (SFB 884; Project Number 139943784; Project Z1).

References

Aksoy, B., Krasteva, S., 2020. When does less information translate into more giving to public goods? Exp. Econ. http://dx.doi.org/10.1007/s10683-020-09643-1.

Baiardi, D., Morana, C., 2021. Climate change awareness: Empirical evidence for the European Union. Energy Econom. 96, 105163. http://dx.doi.org/10.1016/j.eneco.2021.105163.

Björnberg, K.E., Karlsson, M., Gilek, M., Hansson, S.O., 2017. Climate and environmental science denial: A review of the scientific literature published in 1990–2015. J. Clean. Prod. 167, 229–241. http://dx.doi.org/10.1016/j.jclepro.2017.08.066.

Blom, A.G., Bosnjak, M., Cornilleau, A., Cousteaux, A.-S., Das, M., Douhou, S., Krieger, U., 2016. A comparison of four probability-based online and mixed-mode panels in europe. Soc. Sci. Comput. Rev. 34 (1), 8–25. http://dx.doi.org/10.1177/0894439315574825.

Blom, A.G., Fikel, M., Friedel, S., Höhne, J.K., Krieger, U., Rettig, T., Wenz, A., SFB 884 'Political Economy of Reforms', Universität Mannheim, 2019. German internet panel, wave 38 (november 2018). GESIS Data Archive, Cologne. ZA6958 Data file Version 1.0.0, https://doi.org/10.4232/1.13391.

Blom, A.G., Fikel, M., Friedel, S., Höhne, J.K., Krieger, U., Rettig, T., Wenz, A., SFB 884 'Political Economy of Reforms', Universität Mannheim, 2020a. German internet panel, wave 41 (may 2019). GESIS Data Archive, Cologne. ZA7590 Data file Version 1.0.0, https://doi.org/10.4232/1.13464.

Blom, A.G., Fikel, M., Friedel, S., Höhne, J.K., Krieger, U., Rettig, T., Wenz, A., SFB 884 'Political Economy of Reforms', Universität Mannheim, 2020b. German internet panel, wave 44 (november 2019). GESIS Data Archive, Cologne. ZA7593 Data file Version 1.1.0, https://doi.org/10.4232/1.13614.

Blom, A.G., Fikel, M., Friedel, S., Krieger, U., Rettig, T., SFB 884 'Political Economy of Reforms', Universität Mannheim, 2021a. German internet panel, wave 48 (july 2020). GESIS Data Archive, Cologne. ZZA7645 Data file Version 1.0.0, https://doi.org/10.4232/1.13681.

Blom, A.G., Fikel, M., Friedel, S., Krieger, U., Rettig, T., SFB 884 'Political Economy of Reforms', Universität Mannheim, 2021b. German internet panel, wave 49 - core study (september 2020). GESIS Data Archive, Cologne. ZA7646 Data file Version 1.0.0, https://doi.org/10.4232/1.13682.

Blom, A.G., Gathmann, C., Krieger, U., 2015. Setting up an online panel representative of the general population: The german internet panel. Field Methods 27 (4), 391–408. http://dx.doi.org/10.1177/1525822X15574494.

Blom, A.G., Gonzalez Ocanto, M., Fikel, M., Friedel, S., Krieger, U., Rettig, T., SFB 884 'Political Economy of Reforms', Universität Mannheim, 2021c. German internet panel, wave 52 (march 2021). GESIS Data Archive, Cologne. ZA7760 Data file Version 1.0.0, https://doi.org/10.4232/1.13794.

Blom, A.G., Herzing, J.M.E., Cornesse, C., Sakshaug, J.W., Krieger, U., Bossert, D., 2017. Does the recruitment of offline households increase the sample representativeness of probability-based online panels? Evidence from the german internet panel. Soc. Sci. Comput. Rev. 35 (4), 498–520. http://dx.doi.org/10.1177/0894439316651584.

Che, Y.-K., Mierendorff, K., 2019. Optimal dynamic allocation of attention. Am. Econ. Rev. 109 (8), 2993–3029. http://dx.doi.org/10.1257/aer.20171000.

Cinelli, M., Quattrociocchi, W., Galeazzi, A., Valensise, C.M., Brugnoli, E., Schmidt, A.L., Zola, P., Zollo, F., Scala, A., 2020. The COVID-19 social media infodemic. Sci. Rep. 10 (1), 1–10. http://dx.doi.org/10.1038/s41598-020-73510-5.

Cornesse, C., Blom, A.G., Dutwin, D., Krosnick, J.A., De Leeuw, E.D., Legleye, S., Pasek, J., Pennay, D., Phillips, B., Sakshaug, J.W., Struminskaya, B., Wenz, A., 2020. A review of conceptual approaches and empirical evidence on probability and nonprobability sample survey research. J. Surv. Stat. Methodol. 8 (1), 4–36. http://dx.doi.org/10.1093/jssam/smz041.

Dana, J., Weber, R.A., Kuang, J.X., 2007. Exploiting moral wiggle room: experiments demonstrating an illusory preference for fairness. Econom. Theory 33 (1), 67–80. http://dx.doi.org/10.1007/s00199-006-0153-z.

Del Vicario, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., Stanley, H.E., Quattrociocchi, W., 2016. The spreading of misinformation online. Proc. Natl. Acad. Sci. 113 (3), 554–559. http://dx.doi.org/10.1073/pnas.1517441113.

Exley, C.L., 2016. Excusing selfishness in charitable giving: The role of risk. Rev. Econ. Stud. 83 (2), 587–628. http://dx.doi.org/10.1093/restud/rdv051.

Fischbacher, U., Gächter, S., 2010. Social preferences, beliefs, and the dynamics of free riding in public goods experiments. Am. Econ. Rev. 100 (1), 541–556. http://dx.doi.org/10.1257/aer.100.1.541.

Fischbacher, U., Gächter, S., Fehr, E., 2001. Are people conditionally cooperative? Evidence from a public goods experiment. Econom. Lett. 71 (3), 397–404. http://dx.doi.org/10.1016/S0165-1765(01)00394-9.

Fischbacher, U., Schudy, S., Teyssier, S., 2014. Heterogeneous reactions to heterogeneity in returns from public goods. Soc. Choice Welf. 43 (1), 195–217. http://dx.doi.org/10.1007/s00355-013-0763-x.

Friedrichsen, J., Momsen, K., Piasenti, S., 2022. Ignorance, intention and stochastic outcomes. J. Behav. Exp. Econ. 100, 101913. http://dx.doi.org/10.1016/j.socec.2022.101913.

Gentzkow, M., Shapiro, J.M., 2006. Media bias and reputation. J. Political Econ. 114 (2), 280–316. http://dx.doi.org/10.1086/499414.

Golman, R., Hagmann, D., Loewenstein, G., 2017. Information avoidance. J. Econ. Lit. 55 (1), 96–135. http://dx.doi.org/10.1257/jel.20151245.

Herzing, J.M.E., Blom, A.G., 2019. The influence of a person's digital affinity on unit nonresponse and attrition in an online panel. Soc. Sci. Comput. Rev. 37 (3), 404–424. http://dx.doi.org/10.1177/0894439318774758.

Innocenti, F., Rohde, L.M., 2024. Selective exposure reduces voluntary contributions: Experimental evidence from the german internet panel. Available at SSRN: http://dx.doi.org/10.2139/ssrn.4892090.

Kamenica, E., Gentzkow, M., 2011. Bayesion persuasion. Am. Econ. Rev. 101, 2590–2615. http://dx.doi.org/10.1257/aer.101.6.2590.

Kerschbamer, R., Müller, D., 2020. Social preferences and political attitudes: An online experiment on a large heterogeneous sample. J. Public Econ. 182, 104076. http://dx.doi.org/10.1016/j.jpubeco.2019.104076.

Levati, M., Morone, A., Fiore, A., 2009. Voluntary contributions with imperfect information: An experimental study. Public Choice 138 (1), 199–216. http://dx.doi.org/10.1007/s11127-008-9346-2.

Momsen, K., Ohndorf, M., 2020. When do people exploit moral wiggle room? An experimental analysis of information avoidance in a market setup. Ecol. Econom. 169, 106479. http://dx.doi.org/10.1016/j.ecolecon.2019.106479.

Momsen, K., Ohndorf, M., 2022. Information avoidance, selective exposure, and fake (?) news: Theory and experimental evidence on green consumption. J. Econ. Psychol. 88, 102457. http://dx.doi.org/10.1016/j.joep.2021.102457.

Regner, T., 2018. Reciprocity under moral wiggle room: Is it a preference or a constraint? Exp. Econ. 21 (4), 779–792. http://dx.doi.org/10.1007/s10683-017-9551-2.

Spiekermann, K., Weiss, A., 2016. Objective and subjective compliance: A norm-based explanation of 'moral wiggle room'. Games Econom. Behav. 96 (C), 170–183. http://dx.doi.org/10.1016/j.geb.2015.11.007.

Thunström, L., van't Veld, K., Shogren, J.F., Nordström, J., 2014. On strategic ignorance of environmental harm and social norms. Rev. Econ. Polit. 124 (2), 195–214. http://dx.doi.org/10.3917/redp.242.0195.

Van der Weele, J.J., Kulisa, J., Kosfeld, M., Friebel, G., 2014. Resisting moral wiggle room: how robust is reciprocal behavior? Am. Econ. J.: Microecon. 6 (3), 256–264. http://dx.doi.org/10.1257/mic.6.3.256.