

Government Endorsement and Voter Skepticism: An Experiment of Selective Disclosure in Political Communication

Qinyang Yu, Ling Lei, Chenzhe Zhang

ECON 690 Group 3

June 1, 2026

“Persuasion is achieved by the speaker’s personal character when the speech is so spoken as to make us think him credible.”

— Aristotle

Abstract

This paper investigates how explicit knowledge of government bias influences voter skepticism when truthful but selectively disclosed information is presented. We propose a theoretical framework and characterize skeptical and credible Perfect Bayesian Equilibria. We design an individual experiment by randomly assigning participants to interact with either a neutral or biased government with different disclosure strategies. In each round, subjects receive a truthful but potentially ambiguous message, update beliefs, and decide whether to vote for a policy alternative. We test skeptical equilibrium selection under strategic uncertainty and skeptical equilibrium deviation under strategic certainty. Our study contributes to political-communication literature by linking theories of information disclosure and sender credibility to experimental evidence on how bias and transparency jointly shape collective decision-making.

Keywords: Government bias; Voter skepticism; Selective information disclosure; Experimental politics; political communication

1 Introduction

In modern political systems, particularly those dominated by a single ruling party, governments often engage in selective yet truthful information disclosure. The information may or may not be accompanied by explicit policy endorsement, which is defined as public, clear declarations of support for a policy option. Even without lies or fabrications, the decision of

what to disclose—and what to omit—can shape public opinion and ultimately sway voters’ policy choices.

In such contexts, the perceived motive of the ruling party behind the information becomes just as important as its content. Understanding how voters react when they explicitly observe government bias and speculate its strategic motives is essential for evaluating public trust and designing effective state communication strategies.

While prior research has extensively explored how information or sources shape voter preferences and how voter behavior influences government preferences through electoral mechanisms in a multi-party system, less attention has been paid to how known government motives interact with truthful but strategically disclosed information as a common practice in dominant-party systems and how voters learn from such dynamics play out over repeated communications.

To fill this research gap, we build a game-theoretical sender–receiver framework and characterize two types of Perfect Bayesian Equilibria—skeptical and credible—as well as two corresponding forms of voter skepticism: skeptical equilibrium selection (under strategic uncertainty) and skeptical equilibrium deviation (under strategic certainty).

Our central research question asks whether voters are more skeptical of selective information disclosed by a biased government than by a neutral one, and which forms of skepticism they exhibit.

To test our prediction, we conduct a multi-round within-subject experiment implemented in oTree. Each participant is randomly assigned to one of two government types—neutral (main control) or biased (main treatment)—with ambiguous or clear disclosure strategies (finer control and treatment). In each round, the government privately observes the true state of the world and then publicly discloses a truthful but possibly ambiguous subset of states. Subjects update their beliefs via Bayes’ rule and decide whether to vote for a new policy or stick with the status quo policy. Rounds 1 and 20 isolate strategic uncertainty and strategic certainty. We record both voting choices which form the data for our empirical analysis and main results.

Building on our theoretical framework of sender–receiver games, we test our hypotheses in two regression specifications for round 1 and 20. The hypotheses are related to skeptical equilibrium selection, skeptical equilibrium deviation, and selection sensitivity, respectively.

This experiment contributes to the political communication and experimental politics literature. By making government bias explicit and holding message content constant, we isolate the pure effect of perceived bias on belief updating and action. Moreover, our multi-round within-subject design bridges static laboratory games and real-world dynamics, capturing both equilibrium selection under strategic uncertainty and equilibrium deviation under strategic certainty. Beyond academic insight, our findings have policy relevance for any institution—state or private—interested in the ethics and efficacy of strategic disclosure.

2 Literature Review

We relate our study to three main streams of literature: the first examines how source identity shapes public perception; the second investigates how information disclosure influences voter preferences and behavior; and the third explores how voters influence government actions

through electoral channels.

In the realm of political communication, the identity of the information source—such as whether it is a government body or another institution—plays a pivotal role in shaping public responses. For instance, in an authoritarian context like China, [Yu et al. \(2021\)](#) conduct an online survey experiment and find that government sources are perceived as more credible than other media outlets when correcting misinformation. Comparable insights regarding the influence of source identity can be found in other domains, such as food safety ([Rembischevski and Caldas, 2023](#)), and public health advice ([Zarzeczna et al., 2023](#)). These findings suggest that in settings dominated by a central authority, the source identity can exert a strong and direct effect on public perception.

Yet, prior work largely treats source identity as static, comparing different institutions or media outlets; they do not examine situations where the same source may embody multiple potential identities, such as a government acting either as biased or neutral. We instead focus on the same institutional sender whose bias is made explicit to participants, isolating how perceived bias alters responses to exactly the same truthful signals.

The second stream of literature focuses on the impact of the information on voter turnout and voting behavior. [Sun et al. \(2021\)](#) show that biased information can significantly influence short-term voting behavior—especially among young voters ([Carvalho et al., 2023](#))—but has little effect on long-term average vote share, which is determined more by voters’ common prior beliefs. [Arceneaux and Kolodny \(2009\)](#) demonstrate that information from ideologically opposed or untrusted sources increases voter skepticism, particularly among opposition voters. In line with this, [Jackson and Tan \(2013\)](#) develop a model in which strategic information withholding by biased agents affects how others update their beliefs.

However, few studies explore the information channel within the context of government–voter interaction, particularly with respect to how government communication directly shapes voter behavior. [Alonso and Càmarà \(2016\)](#) examine a strategic information design problem where a government seeks to persuade voters through information disclosure. However, in their model, government policy preference is fixed and known, and persuasion is achieved through a full commitment to a communication strategy rather than a dynamic revelation.

The last one is on the relationship between government preferences and voters’ choices, primarily focusing on how voter preferences and behaviors influence government policy through electoral mechanisms. [Guntermann and Persson \(2021\)](#) use Swedish Election Study Panels to show that issue voting enables voters to affect government responsiveness to public policy concerns. Similarly, [Gilens \(2005\)](#) emphasizes that elections act as a key channel linking public preferences to government action and policy implementation, although this responsiveness tends to favor the preferences of the most affluent. Government responsiveness is not limited to the national level. At the local level, city fiscal policies also adapt to constituent preferences ([Einstein and Kogan, 2015](#)).

Nevertheless, the question of how explicitly revealed government preference influences voter behavior is merely addressed. [Blais et al. \(2006\)](#) examine how voters’ impression of potential government coalitions will impact their predictions about policy directions, which in turn influence their voting choices. [Rivard and Lockhart \(2022\)](#) explore how government types, minority or majority, affect voters’ beliefs and electoral decisions. These existing studies operate under the assumption that voters are inferring government preferences based on

structural cues, as the government has not yet been formed at the time of voting. Therefore, no government preference is directly observable to voters, and the voters must predict them before voting. In contrast, little attention has been paid to contexts where the government is already in place, and its policy preferences are made explicitly visible to the public. [Colombo et al. \(2016\)](#) finds empirical evidence that government endorsement of a policy may actually reduce voter support based on survey data from Italy’s Referendum Revolution. However, the study does not provide a theoretical model or explore the underlying mechanisms driving this effect.

Our study contributes to all three streams of literature. First, we extend research on source identity by examining not just differences across types of sources (e.g., government vs. non-government), but how the same government source is perceived differently depending on whether it is seen as neutral or biased. This allows us to investigate how perceptions of bias in a government communicator influence public responses, even when the institutional source remains constant. Second, we contribute to the literature on information disclosure and voter behavior by analyzing how explicit policy endorsements by an already-formed government can provoke skeptical equilibrium selection—particularly when voters suspect selective disclosure of information. Third, we add to the literature on electoral accountability by moving beyond studies that infer government preferences from structural or partisan cues, instead focusing on cases where preferences are clearly stated but still potentially distrusted. The research gap between our question and prior studies is illustrated in Figure 1.




Figure 1: Our Research Contribution to the Existing Literature

3 Theoretical Framework

3.1 Setup

Consider a communication game with two agents: the government and homogeneous voters (a “sender-receiver” setting). The government proposes a binary policy set $X = \{x_0, x_1\}$ to the society, where x_0 represents the status quo or default policy, and x_1 represents the proposed new policy alternative x_1 .

The game follows a sequential structure and involves incomplete information. Before both players move, nature randomly determines the true state of the world and the type of the government, denoted as θ and τ . The true state θ is fully observable to the government but not to the voters, while the type of the government τ is common knowledge. Upon observing θ , the government strategically discloses truthful information to voters. After the message is delivered, voters process information according to Bayes’ rule and decide whether to vote for x_1 .

Despite the similarity in structure, this game differs from Bayesian persuasion (Kamenica and Gentzkow, 2011) because the government lacks full commitment to an information policy. It also differs from the cheap talk game (Farrell and Rabin, 1996), as the government is restricted to disclosing only truthful and verifiable information.

3.2 Belief Space

The complete belief space consists of two components: the true state θ and the government type τ . Let $\Theta = \{a, b, c\}$ denote the finite state space which contains three possible states. Let $\mathcal{T} = \{N, B\}$ denote the type space, where N represents a neutral government and B represents a biased government. The belief space is then defined as $\Omega = \Delta(\Theta \times \mathcal{T})$, where each belief state $\omega(\theta, \tau) \in \Omega$ indicates the probability that the realized state of nature is θ and the government type is τ .

For the prior beliefs about the true state, nature’s random draw assigns a probability of $\omega(\theta, \cdot) = \frac{1}{3}$ for all $\theta \in \Theta$. Since the realization of τ is common knowledge, we can simplify the belief state to $\omega(\theta) = \frac{1}{3}$ for all $\theta \in \Theta$, indicating that voters cannot distinguish between different realizations of θ but can distinguish between N and B .

3.3 Strategies

In the first stage, the government fully observes the realization of θ and selects an information disclosure strategy. The disclosed information will be observed by voters as a signal. We assume that the government must disclose truthful and verifiable information (i.e., the disclosed information set must contain θ). However, it has discretion over how much information to reveal.

Formally, define the government’s disclosure strategy as a mapping $s : \Theta \rightarrow \Delta(2^\Theta)$, where $s(\theta)$ represents the set of possible states that the government may disclose given the true state θ . The power set 2^Θ contains all subsets of Θ , capturing different levels of disclosure. To summarize, the government’s pure disclosure strategies can be classified into

three types:

$$s(\theta) = \begin{cases} \{\theta\} & \text{(full disclosure)} \\ S \subsetneq \Theta, \text{ where } \theta \in S, |S| > 1 & \text{(partial disclosure)} \\ \Theta & \text{(no disclosure)} \end{cases}$$

Note that $s(\theta) = \Theta$ conveys no new information, so it is equivalent to $s(\theta) = \emptyset$ and the latter is omitted. As an illustrating example, if the true state is b , the government's pure information disclosure follows the mapping:

$$s(b) = \begin{cases} \{b\} & \text{(full disclosure)} \\ \{a, b\} \text{ or } \{b, c\} & \text{(partial disclosure)} \\ \{a, b, c\} & \text{(no disclosure)} \end{cases}$$

In the second stage, based on $s(\theta)$ and $\omega(\theta)$, the voters update their beliefs according to Bayes' rule and form a posterior belief about the true state θ :

$$\mu(\theta|s(\theta)) = \frac{\Pr(s(\theta)|\theta)\omega(\theta)}{\sum_{\theta' \in \Theta} \Pr(s(\theta)|\theta')\omega(\theta')}$$

Based on the posterior belief, the voters then calculate their expected payoff and decide whether to vote for x_0 or x_1 . Their strategy is defined as a mapping from the belief space Ω to the action set $A = \{\text{vote}, \text{not vote}\}$ in favor of x_1 . Since we assume all voters are homogeneous, their decisions are the same and directly determine which policy will be implemented.

3.4 Utility

Assume both the government and voters are risk-neutral and their preference can be characterized by a von Neumann-Morgenstern utility function.

The government's preferences over the two policies are state-independent but vary by types. A neutral government N is indifferent between the two policies, while a biased government B strictly prefers the new policy over the default policy. Formally, the government's utility function $U : X \rightarrow \mathbb{R}$ is given by:

$$U_N(x_0) = U_N(x_1), \quad U_B(x_0) < U_B(x_1).$$

The voters' preferences over policies are represented by the utility function $V : X \times \Theta \rightarrow \mathbb{R}$. The status quo policy, x_0 , always yields a zero payoff to all voters, while the new policy, x_1 , generates a state-dependent payoff. Specifically, the new policy results in a positive return v_a in state a , a small negative return $-v_b$ in state b , and a large negative return $-v_c$ in state c . The payoff values satisfy the assumption $v_c > v_a > v_b$ for $v_a, v_b, v_c \in \mathbb{R}^+$. The values of $V(x, \theta)$ are summarized in Table 1. An intuition behind this assumption is that it creates opportunities for a biased government B to strategically manipulate information disclosure by disclosing $\{a, b\}$ in state b .

Nature	State _a	State _b	State _c
x_0 (Status Quo)	0	0	0
x_1 (New Policy)	v_a	$-v_b$	$-v_c$

Table 1: Voters’ Payoff Structure

3.5 Equilibrium

For simplicity, we focus on pure-strategy Perfect Bayesian Equilibria (PBE) in this study. The game admits multiple pure-strategy PBEs, including some that are trivial or exhibit anomalous strategic behavior. Among these, we specifically focus on equilibria in which the disclosure $\{a, b\}$ occurs, as our primary interest lies in voter behavior conditioned on this ambiguous information.

The rationale for this equilibrium refinement is that the disclosure $\{a, b\}$ introduces strategic uncertainty for voters. While the informational content of this disclosure guarantees a positive expected payoff, the government’s strategic response may vary—potentially resulting in either a favorable or adversarial outcome for voters.

To illustrate this, we first analyze the case where the government is biased. In any PBE, the government must disclose $\{a, b\}$ in state b ; otherwise, $\{a, b\}$ could only be disclosed in state a and voters would thus vote upon receiving it. Anticipating this, the government would have an incentive to deviate in state b . Additionally, the government avoids disclosing $\{a, c\}$ or $\{a, b, c\}$ in state a , since doing so either mimics disclosures from state c , leading to profitable deviation in state a , or it differs from state c and becomes distinguishable, leading to profitable deviation in state c . Given these constraints, we end up with two PBEs, as summarized in the final two columns of Table 2. We characterize these two PBEs as the **skeptical PBE** and the **credible PBE** based on voters’ responses upon receiving $\{a, b\}$. In the skeptical PBE, the government discloses $\{a\}$ in state a , so voters interpret $\{a, b\}$ cautiously as a manipulative intent and refrain from voting upon receiving $\{a, b\}$. In the credible PBE, the government discloses $\{a, b\}$ in state a , so voters interpret $\{a, b\}$ as a trustworthy signal and vote upon receiving $\{a, b\}$.

When the government is neutral, one can readily verify the existence of both the skeptical and credible PBEs. In addition to these, there also exists a trivial PBE in which the government discloses $\{a, b\}$ in state a but not in state b , and voters nonetheless choose to vote upon receiving the message $\{a, b\}$. These PBEs are summarized in the first three columns of Table 2.

In both the neutral and biased government scenarios, multiple PBEs can arise. However, the skeptical PBE is a unique channel in that it leads to voter inaction—that is, a reluctance to vote following the ambiguous disclosure $\{a, b\}$. We refer to this phenomenon as **skeptical equilibrium selection**—the idea that when voters face an uncertain disclosure, they tend to believe they are more likely in the worse equilibrium, as they would have received a more precise disclosure if the true state were more favorable, making the skeptical PBE more likely to emerge.

In addition, even along the equilibrium paths of both the skeptical and credible PBEs, voters may still deviate from equilibrium behavior due to perceived government bias. We refer to this phenomenon as **skeptical equilibrium deviation**—the idea that, even when the

disclosure strategy is fully known and consistent with equilibrium play, voters may choose not to follow the equilibrium response because they distrust the government’s motive. In this case, the deviation is not due to ambiguity about which equilibrium is being played, but rather to a lack of trust in the credibility or intent behind the disclosed information itself. Skeptical equilibrium deviation thus reflects a behavioral failure of equilibrium implementation, driven by perceived manipulation rather than informational ambiguity. However, we are interested the sensitivity of this deviation to informational ambiguity.

Our experiment is designed to examine whether government bias induces voter skepticism—either as skeptical equilibrium selection under strategic uncertainty or as skeptical equilibrium deviation under strategic certainty—and whether this skepticism ultimately leads to a reduction in overall voting rates.

Table 2: Equilibrium Classification Based on Information Disclosure $\{a, b\}$

Strategy	Neutral Government			Biased Government	
	Skeptical PBE	Credible PBE	Trivial PBE	Skeptical PBE	Credible PBE
Government ($\theta = a$)	$\{a\}$	$\{a, b\}$	$\{a, b\}$	$\{a\}$	$\{a, b\}$
Government ($\theta = b$)	$\{a, b\}$	$\{a, b\}$	$s(b) \setminus \{a, b\}$	$\{a, b\}$	$\{a, b\}$
Government ($\theta = c$)	$s(c)$	$s(c)$	$s(c)$	$s(c)$	$s(c)$
Voters (receive $\{a, b\}$)	not vote	vote	vote	not vote	vote

4 Experiment Design

The experiment is conducted individually without requiring coordination among participants. Each participant plays 20 rounds of a voting game, making a binary choice in each round: to vote or not to vote. Prior to beginning the experiment, participants must pass a comprehension test to ensure they understand the instructions and structure of the game. Upon completion of the 20 rounds, participants are asked to complete a follow-up survey. Figure 2 presents the flow chart summarizing the full sequence of the experiment. The entire session is expected to take approximately 20 minutes per participant.

4.1 Preparation

The experiment begins with a preparation phase consisting of three sequential components. First, participants are presented with an informed consent form outlining the purpose of the study and their rights as participants (Figure 4).

Second, participants receive detailed instructions that explain the key components of the experiment, including the payoff structure (Figure 5), the government’s information disclosure mechanisms (Figure 6), and the type of government they will interact with—either neutral for the control group (Figure 7) or biased for the treatment group (Figure 8).

Third, participants must complete a comprehension test consisting of five questions designed to assess their understanding of: (1) the probability distributions over policy states




Figure 2: Experiment Flow Chart

(Figure 9); (2) the government’s disclosure policies across different states (see example in Figure 10); and (3) the underlying payoff structure (Figure 11). Participants are allowed up to three attempts per question. Failure to answer any question correctly within the allowed attempts results in exclusion from the experiment. This ensures that only participants with a clear grasp of the experimental framework proceed to the main task.

4.2 Random Assignment

Upon successfully completing the comprehension test, participants are first randomly assigned to one of two groups: the control group (C) or the treatment group (T). Participants in the control group are informed that they are interacting with a neutral government (N), while those in the treatment group are told they are interacting with a biased government (B).

Following the first round of the game, participants in both the control and treatment groups are further randomized into subgroups: C_1 or C_2 for the control group, and T_1 or T_2 for the treatment group. These subgroups correspond to different information disclosure strategies—either high clarity or low clarity. Within each pair (C_1/T_1 and C_2/T_2), the control and treatment groups are exposed to the same disclosure strategy, allowing us to isolate the effect of government type from the effect of information clarity.

Table 4 provides a detailed overview of the disclosure strategies assigned to each subgroup.

Table 3: Experiment Design

Round	Control Group C		Treatment Group T	
1	$\{a, b\}$		$\{a, b\}$	
	Control C_1	Control C_2	Treatment T_1	Treatment T_2
2	Information generated by disclosure strategies in Table 4			
$\vdots$				
$n - 1$				
n	$\{a, b\}$	$\{a, b\}$	$\{a, b\}$	$\{a, b\}$

Table 4: Classification of Disclosure Strategies

True State	C_1/T_1	C_2/T_2
a	$\{a\}$	$\{a, b\}$
b	$\{a, b\}$	$\{a, b\}$
c	$s(c)$	$s(c)$

4.3 20-Round Voting Game

The third section of the experiment consists of a 20-round voting game. Each participant completes all 20 rounds independently. In each round, participants must decide whether to vote or not (Figure 12). After submitting their decision, they receive immediate feedback displaying both their payoff outcome and the true policy state (Figure 13). This process repeats across all rounds until completion.

Throughout the game, participants have continuous access to: (1) the assigned government type¹, (2) the current round’s disclosed information set, and (3) the payoff structure. Starting from round 2, participants can also review their own voting history and the true policy state history from previous rounds (Figure 14).

Notably, the final round (round 20) applies a standardized disclosure strategy that deviates from the group’s assigned treatment in earlier rounds. Specifically, all participants receive the same ambiguous information set, $\{a, b\}$ ², regardless of group assignment. This design allows us to elicit posterior beliefs. After round 20, participants are asked: “After completing 20 rounds of voting and observing that the disclosed information in the final round was ‘the true state is either good or bad,’ what do you think is the probability that the true state is good?”

4.4 Follow Up Questions

The final component of the experiment consists of a series of follow-up questions aimed at collecting demographic and background information from participants. These questions cover variables such as gender, education level, political attitudes, and other relevant characteristics. Figure 15 presents the interface used for the demographic questionnaire.

At the conclusion of the experiment, each participant’s final compensation is determined by randomly selecting one of the 20 rounds and adding that round’s payoff to a fixed participation fee.

4.5 Hypothesis and Predictions

Strategic communication by governments often leaves room for voters to question the credibility and intent behind disclosed information. In our experiment, we explicitly vary both the type of government (neutral vs. biased) and the clarity of the information provided, allowing us to observe how these features influence voter behavior. We expect a set of behavioral hypotheses that reflect different equilibrium responses:

- H_1 : Under strategic uncertainty, there exists skeptical equilibrium selection, such that government bias increases the likelihood of selecting the skeptical PBE. Specifically:

$$\text{Voting}_{\{a,b\}}^T < \text{Voting}_{\{a,b\}}^C$$

¹This assignment remains fixed for each participant throughout all 20 rounds.

²In the actual experiment, “good” corresponds to state a , “bad” to b , and “horrible” to c .

- H_2 : Under strategic certainty, there exists skeptical equilibrium deviation in both skeptical and credible PBEs. Specifically:

$$\text{Voting}_{\{a,b\}}^{T_1} - \text{Voting}_{\{a,b\}}^{C_1} < 0, \quad \text{Voting}_{\{a,b\}}^{T_2} - \text{Voting}_{\{a,b\}}^{C_2} < 0$$

- H_3 : Under strategic certainty, skeptical equilibrium deviation is more sensitive to government bias in the credible PBE. Specifically:

$$|\text{Voting}_{\{a,b\}}^{T_2} - \text{Voting}_{\{a,b\}}^{C_2}| > |\text{Voting}_{\{a,b\}}^{T_1} - \text{Voting}_{\{a,b\}}^{C_1}|$$

Specifically, H_1 tests skeptical equilibrium selection under strategic uncertainty, H_2 tests skeptical equilibrium deviation under strategic certainty, and H_3 tests the sensitivity of such deviation. Figure 3 plots the hypothetical conditions in these three hypotheses for illustration.




Figure 3: An Illustration for Hypothesis Testing

5 Implementation

5.1 Software and Platform

We employ oTree as our primary experimental platform and plan to conduct our study through Prolific. oTree offers several advantages as a Python-based software solution—it features an intuitive interface for experiment design and also provides customization options. Prolific serves as an ideal recruitment platform for our experiment, offering precise demographic targeting capabilities that align with our research requirements.

5.2 Sample Collection

We aim to recruiting 1200 participants. We will first divide them into two big control and treatment groups, and then divide each big group into 2 subgroups (C_1/C_2 vs. T_1/T_2). Since this is an individual experiment, we do not impose any minimum group size requirement to initiate a session. Participants can be randomly assigned and proceed individually, allowing for flexible scheduling and efficient data collection without the need for coordinated group starts.

5.3 Technique Details

Our experiment employed a two-stage randomization process. In the first stage, participants were randomly assigned to either the control group C or the treatment group T , which determined the government type they received either being “Neutral” or “Biased”. In the second stage, within each of these groups, participants were further randomized into subgroups: C_1 or C_2 within the control condition, and T_1 or T_2 within the treatment condition. These subgroups were exposed to different disclosure strategies across subsequent rounds, as detailed in Table 3.

Subgroups C_1 and T_1 were assigned to a disclosure strategy that selectively revealed clearer information: when the true state was a , only $\{a\}$ was disclosed; for states b or c , the disclosure was either $\{a, b\}$ or a stochastic signal $s(c)$, respectively. This design reflects a setting of high disclosure clarity, where the government does reveal clearer information. In contrast, subgroups C_2 and T_2 were exposed to a disclosure strategy with lower transparency: both states a and b resulted in the same ambiguous disclosure $\{a, b\}$, while state c still triggered $s(c)$. This structure allows us to isolate how voters respond to differences in information clarity (C_1/T_1 vs. C_2/T_2), conditional on whether the government is perceived as neutral or biased.

In each round, the true policy state is independently drawn by the oTree software. From round 2 to round 19, the state was selected with equal probability from three possible outcomes: “ a ”, “ b ”, or “ c ”. In contrast, round 1 and 20 were restricted to only two possible states: “good” and “bad”, each with 50% probability. This restriction in the final round was intentionally implemented to facilitate the elicitation of participants’ posterior beliefs—particularly how they update their expectations about the likelihood of the “ a ” state when presented with disclosed information $\{a, b\}$. By holding the underlying distribution of states constant across participants and rounds, this design ensures that observed differences in voting behavior can be attributed to the experimental treatments rather than to systematic variation in the state-generating process.

6 Data Analysis Plan

The resulting dataset includes three primary categories of variables. First, we collected comprehensive demographic information including participant ID, age, gender, nationality, education level, employment status, political affiliation, and government trust level. Second, we recorded key experimental outcomes throughout the 20 rounds, tracking each participant’s voting decisions, round-specific payoffs, the information disclosed by the government, and

the true state revealed after each voting decision. Third, our primary outcome variable is the voting result in round 1 and 20. Based on the dataset and variables, we outline our main empirical analysis models in the remainder of this section.

To test H_1 , we focus on the voting result in round 1. The regression model is as follows:

$$Y_i = \beta_0 + \beta_1 T_i + \mu_i, \quad (1)$$

where

$$T_i = \begin{cases} 0 & \text{if } i \text{ is in the control group,} \\ 1 & \text{if } i \text{ is in the treatment group,} \end{cases}$$

and

$$Y_i = \begin{cases} 0 & \text{if } i \text{ does not vote,} \\ 1 & \text{if } i \text{ votes.} \end{cases}$$

The conditional expectations under each treatment condition are:

$$\text{Voting}_{\{a,b\}}^T = \mathbb{E}[Y_i \mid T_i = 1] = \beta_0 + \beta_1,$$

$$\text{Voting}_{\{a,b\}}^C = \mathbb{E}[Y_i \mid T_i = 0] = \beta_0.$$

Based on these expectations, to test whether $\text{Voting}_{\{a,b\}}^T < \text{Voting}_{\{a,b\}}^C$ for H_1 , it is equivalent to test whether $\beta_1 < 0$ is statistically significant, suggesting skeptical equilibrium selection under strategic uncertainty. Specifically,

$$H1 : \text{Voting}_{\{a,b\}}^T < \text{Voting}_{\{a,b\}}^C \Leftrightarrow \beta_1 < 0.$$

To test H_2 and H_3 , we focus on the voting result in round 20. The regression model is as follows:

$$Y_i = \beta_0 + \beta_1 T_i + \beta_2 D_i + \beta_3 T_i D_i + \mu_i, \quad (2)$$

where

$$T_i = \begin{cases} 0 & \text{if } i \text{ is in the control group } (C_1 \text{ or } C_2), \\ 1 & \text{if } i \text{ is in the treatment group } (T_1 \text{ or } T_2), \end{cases}$$

$$D_i = \begin{cases} 0 & \text{if } i \text{ is in pair } (C_1, T_1), \\ 1 & \text{if } i \text{ is in pair } (C_2, T_2), \end{cases}$$

and

$$Y_i = \begin{cases} 0 & \text{if } i \text{ does not vote,} \\ 1 & \text{if } i \text{ votes.} \end{cases}$$

Similarly, we list all the conditional expectations under each treatment condition, which are summarized in Table 5. Based on these expectations, the conditions in H_2 can be tested

through regression coefficients based on the following relationships:

$$\begin{aligned}\text{Voting}_{\{a,b\}}^{T_1} - \text{Voting}_{\{a,b\}}^{C_1} &= \mathbb{E}[Y_i \mid T = 1, D = 0] - \mathbb{E}[Y_i \mid T = 0, D = 0] \\ &= (\beta_0 + \beta_1) - \beta_0 \\ &= \beta_1,\end{aligned}$$

$$\begin{aligned}\text{Voting}_{\{a,b\}}^{T_2} - \text{Voting}_{\{a,b\}}^{C_2} &= \mathbb{E}[Y_i \mid T = 1, D = 1] - \mathbb{E}[Y_i \mid T = 0, D = 1] \\ &= (\beta_0 + \beta_1 + \beta_2 + \beta_3) - (\beta_0 + \beta_2) \\ &= \beta_1 + \beta_3.\end{aligned}$$

In other words, we can simplify H_2 as:

$$H2(1) : \text{Voting}_{\{a,b\}}^{T_1} - \text{Voting}_{\{a,b\}}^{C_1} < 0 \Leftrightarrow \beta_1 < 0,$$

$$H2(2) : \text{Voting}_{\{a,b\}}^{T_2} - \text{Voting}_{\{a,b\}}^{C_2} < 0 \Leftrightarrow \beta_1 + \beta_3 < 0.$$

For the first part of H_2 , we cannot reject the hypothesis if $\beta_1 < 0$. For the second part of H_2 , we cannot reject the hypothesis if $\beta_1 + \beta_3 < 0$. Both indicate a negative impact of government bias on voting rates along the two equilibrium paths, reflecting skeptical equilibrium deviation.

The last hypothesis H_3 can be tested through regression coefficients based on the following relationship:

$$H3 : |\text{Voting}_{\{a,b\}}^{T_2} - \text{Voting}_{\{a,b\}}^{C_2}| - |\text{Voting}_{\{a,b\}}^{T_1} - \text{Voting}_{\{a,b\}}^{C_1}| < 0 \Leftrightarrow (-\beta_1 - \beta_3) - (-\beta_1) = \beta_3 < 0.$$

It simplifies to testing $\beta_3 < 0$. Therefore, we cannot reject H_3 if β_3 is significantly negative, suggesting that the equilibrium deviation is more sensitive to government bias in the credible PBE compared to the skeptical one.

For all regressions, we also estimate using a logit model, include control variables, add individual fixed effects, and cluster standard errors at the subject level for robustness check.

Group	T_i	D_i	$T_i D_i$	Expected Value of Y_i
C_1	0	0	0	$\mathbb{E}[Y_i \mid T = 0, D = 0] = \beta_0$
C_2	0	1	0	$\mathbb{E}[Y_i \mid T = 0, D = 1] = \beta_0 + \beta_2$
T_1	1	0	0	$\mathbb{E}[Y_i \mid T = 1, D = 0] = \beta_0 + \beta_1$
T_2	1	1	1	$\mathbb{E}[Y_i \mid T = 1, D = 1] = \beta_0 + \beta_1 + \beta_2 + \beta_3$

Table 5: Expected values of Y_i under treatment and disclosure conditions

7 Limitations

While our design effectively isolates the impact of perceived government bias on voter behavior under conditions of strategic uncertainty and certainty, a key limitation lies in the

exogeneity of perception formation. Specifically, we assign the perception of government bias through experimental treatments rather than allowing it to evolve endogenously through participants’ interpretation of disclosure patterns. In real-world political environments, voter beliefs about whether a government is biased often emerge through indirect cues such as media framing, partisan alignment, prior trust, or iterative experiences—not simply by labeling the government as biased or neutral. Although our multi-round setup incorporates feedback and allows subjects to update their beliefs about disclosure strategies, the experiment does not model how those beliefs about the government’s intent are formed in the first place, nor how they may vary across individuals.

Second, while we document behavioral changes in response to perceived bias, we do not disentangle the underlying psychological or informational channels through which this bias operates. Specifically, it remains unclear whether the observed voter skepticism is driven by a first-impression effect (where early exposure to government types) or by a credibility filter (where voters actively discount the information due to strategic distrust in the source). Our current design identifies the presence of skepticism but does not distinguish which pathway is responsible for it.

8 Conclusion

This paper investigates how voters respond to policy endorsements in political communication, with a focus on the role of perceived bias and strategic information disclosure. Through a multi-round experimental design, we independently vary the government’s perceived type—neutral or biased—and the clarity of the information it provides. This structure allows us to isolate the behavioral impact of government source identity and information disclosure strategy on voting behavior under conditions of strategic uncertainty and certainty.

Our findings aim to contribute to three key strands of literature. First, we extend the literature on source identity by examining not just who communicates, but how perceived bias within a single source influences credibility. Second, we contribute to work on strategic information disclosure by studying how voters respond to ambiguous versus precise signals. Third, we provide new insights into voter decision-making in the presence of explicit government preferences, enriching our understanding of accountability mechanisms in political communication.

Future extensions could endogenize perceptions of bias—allowing participants to infer government type over time—incorporate richer communication channels (e.g., media intermediaries or group deliberation), or explore cross-cultural variation in baseline trust. By elucidating the mechanisms through which bias and transparency interact, this study provides both theoretical insight and practical guidance for policymakers and communicators seeking to maintain credibility in selective disclosure environments.

References

- Alonso, R. and Câmara, O. (2016). Persuading voters. *American Economic Review*, 106(11):3590–3605.
- Arceneaux, K. and Kolodny, R. (2009). Educating the least informed: Group endorsements in a grassroots campaign. *American Journal of Political Science*, 53(4):755–770.
- Blais, A., Aldrich, J. H., Indridason, I., and Levine, R. (2006). Do voters vote for government coalitions? *Party Politics*, 12:691 – 705.
- Carvalho, B. P., Custódio, C., Geys, B., Mendes, D., and Peralta, S. (2023). Information, perceptions, and electoral behaviour of young voters: A randomised controlled experiment. *Electoral Studies*, 84:102625.
- Colombo, C., De Angelis, A., and Morisi, D. (2016). New survey evidence: Renzi’s support is damaging the chances of a yes vote in Italy’s referendum. *LSE European Politics and Policy (EUROPP) Blog*.
- Einstein, K. and Kogan, V. (2015). Pushing the city limits. *Urban Affairs Review*, 52:3 – 32.
- Farrell, J. and Rabin, M. (1996). Cheap talk. *Journal of Economic Perspectives*, 10(3):103–118.
- Gilens, M. (2005). Inequality and democratic responsiveness. *Public Opinion Quarterly*, 69:778–796.
- Guntermann, E. and Persson, M. (2021). Issue voting and government responsiveness to policy preferences. *Political Behavior*, 45:561–584.
- Jackson, M. O. and Tan, X. (2013). Deliberation, disclosure of information, and voting. *Journal of Economic Theory*, 148(1):2–30.
- Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

- Rembischevski, P. and Caldas, E. (2023). Consumers’ trust in different sources of information related to food hazards and their judgment of government performance—a cross-sectional study in brazil. *Foods*, 12.
- Rivard, A. B. and Lockhart, M. (2022). Government preferences, vote choice and strategic voting in canada. *Canadian Journal of Political Science*, 55(3):721–733.
- Sun, J., Schram, A., and Sloof, R. (2021). “why biased endorsements can manipulate elections. *Available at SSRN 3420703*.
- Yu, W., Shen, F., and Min, C. (2021). Correcting science misinformation in an authoritarian country: An experiment from china. *Telematics Informatics*, 66:101749.
- Zarzeczna, N., Hanel, P. H. P., Rutjens, B., Bono, S. A., Chen, Y.-H., and Haddock, G. (2023). Scientists, speak up! source impacts trust in health advice across five countries. *Journal of experimental psychology. Applied*.

A Experiment Interface

Consent Form

Principal Investigator:

Supervisor: Prof. Pellumb Reshidi - Department of Economics, Duke University;

Organizer: Qinyang Yu, Ling Lei, Chenzhe Zhang - Department of Economics, Duke University

Brief Introduction:

This research aims at helping economists and political scientists understand how voters process information from the government. Participation is voluntary, and you can withdraw at any time without penalty. The session will last approximately 20 minutes, during which you will make decisions as a voter that affect your final payment. You will receive detailed instructions at the beginning of the session on how payments are calculated based on your decisions. There are no known risks beyond those of everyday life associated with this study. While there are no direct benefits to you beyond the financial compensation, your participation may help in the development of social science policies.

Confidentiality:

Your personal data will remain confidential. We do not have access to identifiable information. All records will be securely stored on password-protected computers, accessible only to the research team.

Compensation:

Your final payment will consist of a \$5 participation fee for completing the session, plus an additional payoff based on the decisions you make during the experiment. Your decisions could result in either a bonus or a penalty beyond the completion fee, while the final payment is guaranteed to be positive. By paying close attention during the experiment, you increase your chances of earning a bonus and boosting your final payment.

Figure 4: Consent Form

Instructions

Your Decisions & Payoffs:

Beyond a \$5 completion fee, you will earn a (positive or negative) payoff based on your voting decisions during the experiment.

The **old policy** always generates **\$0** payoff. The **new policy**, however, introduces more variability depending on the true state of the world:

- If the true state is **good**, the payoff is **\$2**.
- If the true state is **bad**, the payoff is **-\$1**.
- If the true state is **horrible**, the payoff is **-\$3**.

Each state happens with **equal probability 1/3**.

In other words, while a **good** state significantly benefits you under the new policy, a **bad** and **horrible** state pose moderate and severe financial risks, respectively.

Here is the summary table of your possible payoff according to each true state and your choice:

True State	Not Vote	Vote
Good	0	2
Bad	0	-1
Horrible	0	-3

Next

Figure 5: Payoff Instruction

Instructions

Government Information Disclosure:

The government will directly observe whether the true state is **good**, **bad**, or **horrible** and selectively disclose information to you.

All information must be **truthful and verifiable**, ensuring that the government **cannot** provide false information about the true state.

For example, if the true state is **good**, the government has the following options for disclosure:

- **Full Disclosure:** The true state is **good**.
- **Partial Disclosure:**
 - The true state is either **good** or **bad**.
 - The true state is either **good** or **horrible**.
- **No Disclosure:** The true state is either **good**, **bad**, or **horrible**.

The information may help you better guess which state is the more likely to be the true state. As a result, you can better balance potential risks and benefits to make more informed decisions.

Next

Figure 6: Government Information Disclosure

Instructions

Neutral Government Preference:

The government's information disclosure strategy is closely related to its own preference. Unlike you, the government **does not** have specific preferences towards the two policies.

In other words, the government is completely **neutral**. It is indifferent between the **old policy** and the **new policy**. Your voting decision **will not affect** government welfare.

You should assume this preference throughout the experiment.

Next

Figure 7: Government Type for Control Group

Instructions

Biased Government Preference:

The government's information disclosure strategy is closely related to its own preference. Just like you, the government **also** has specific preferences towards the two policies.

In other words, the government is completely **biased**. Specifically, it strongly prefers the **new policy** to the old policy. Your voting decision **will significantly affect** government welfare.

You should assume this preference throughout the experiment.

Next

Figure 8: Government Type for Treatment Group

Comprehension Test 1

Are the states **good**, **bad**, and **horrible** equally likely to occur?

Yes

No

Check Answer

Figure 9: Comprehension Test 1

Comprehension Test 2

Assume the true state is **good**, then what type of selective information can the government disclose? (please select all the correct answers)

- ☐ The true state is **good**.
- ☐ The true state is **bad**.
- ☐ The true state is **horrible**.
- ☐ The true state is either **good** or **bad**.
- ☐ The true state is either **good** or **horrible**.
- ☐ The true state is either **bad** or **horrible**.
- ☐ The true state is either **good**, **bad**, or **horrible**.

Check Answer

Figure 10: Comprehension Test 2

Comprehension Test 5: Payoff Structure

Please answer the following questions about the payoff structure.

1. What is your additional payoff if you **vote** and the true state is **good**?

- ☐ \$-3
- ☐ \$-1
- ☐ \$0
- ☐ \$2

2. What is your additional payoff if you do **not vote** and the true state is **good**?

- ☐ \$-3
- ☐ \$-1
- ☐ \$0
- ☐ \$2

3. What is your additional payoff if you **vote** and the true state is **bad**?

- ☐ \$-3
- ☐ \$-1
- ☐ \$0
- ☐ \$2

4. What is your additional payoff if you do **not vote** vote and the true state is **bad**?

- ☐ \$-3
- ☐ \$-1
- ☐ \$0
- ☐ \$2

Figure 11: Comprehension Test 5

Round 1 of 20: Voting Decision

Round Information

The government holds a **Biased** preference and discloses the following verifiable information about the true state:

The true state is either **good** or **bad**.

Make Your Decision

Do you choose to vote or not?

☐ Vote

☐ Not Vote

Next

Figure 12: Voting Game Interface

Round 1 Results

Round 1 Outcome

The true state was: **good**

Your choice: **Vote**

Your payoff this round: **\$2 points**

Payoff explanation:

You voted and the state was **good**, so you earned \$2.

Get ready for the next round!

Next

Figure 13: Voting Result Check

Previous Rounds Summary				
Round:	1	2	3	4
Disclosure:	The true state is either good or bad .	The true state is either good or bad .	The true state is good .	The true state is either
True_State:	good	bad	good	horribl

Payoff Structure			
State	Not Vote	Vote	
Good	0	2	
Bad	0	-1	
Horrible	0	-3	

Figure 14: Cumulative Voting Record

Survey Questions

Demographic Information

What is your name?

What is your age?

What is your gender?

☐ Male
☐ Female
☐ Non-binary
☐ Prefer not to say

What is your nationality?

What is your highest level of education?

☐ High school
☐ Bachelor's
☐ Master's
☐ PhD
☐ Other

Figure 15: Follow Up Questions