"On the veil-of-ignorance principle: welfare-optimal information disclosure in Voting"

Karine Van Der Straeten and Takuro Yamashita

Toulouse
School of
Economics

On the veil-of-ignorance principle: welfare-optimal information disclosure in voting.[*]

Karine Van der Straeten[†]   Takuro Yamashita[‡]

August 2, 2023.

Abstract

Voters' voting decisions crucially depend on their information. Thus, it is an important question how much / what kind of information they should know, as a normative guidance of the optimal extent of transparency. We consider a simple two-alternative majority voting environment, and study the optimal information disclosure policy by a utilitarian social planner. Although full transparency is sometimes (informally) argued as ideal, we show that full transparency is often

strictly suboptimal. This is related to the well-known potential mismatch between what a majority wants and what is socially optimal. Under certain conditions, in order to alleviate this mismatch, the optimal policy discloses just the "anonymized" information about the value of the alternatives to the voters, placing them effectively behind a partial "veil of ignorance".

1   Introduction

Voters' information is crucially important in determining voting outcomes. In some cases, the relevant information is easily accessible to voters. In other cases, acquiring all the relevant information would be prohibitively costly for individual voters, who instead have to rely on other sources – government, media – to inform their choices. These two scenarios represent polar cases, as in practice, for many decisions, both coexist. Consider the example of a structural reform of the labor market proposing changes in employment laws that lower the costs of hiring and firing workers. Each voter, depending on his productivity and switching costs, has some private information about how much he will be directly and immediately affected by the reform. But it may be much more difficult for him to figure out the long-term impact of the reform on his well-being. What will be the consequences of the reform on job creation in the different sectors, growth, the bargaining power of trade unions, tax revenues, the incentives for firms to invest in human capital? The magnitude of these potential changes and how they will specifically affect each voter is extremely difficult to figure out for an individual voter. The government is expected to have better information about this long-term effects.

In order to provide a normative guidance of the optimal extent of transparency in voting, we consider in this paper a simple two-alternative environment ("reform" vs. "status quo") with a simple majority voting rule. We study optimal information disclosure by a (utilitarian) benevolent social

planner, when the planner controls part of the relevant information about the consequences of the alternatives, while each voter possesses their own private information too. Given that the planner is fully benevolent, a natural candidate for an optimal disclosure policy may be full disclosure. This is also often (somewhat informally) argued as ideal as a democratic principle. However, we show that it is in general suboptimal. This is related to the well-known fact that the majority choice may not coincide with what is socially optimal: majority voting only elicits ordinal preferences, while cardinal preferences are relevant for welfare. To provide some idea, imagine a situation where full disclosure of the information makes a minority of voters favor the reform substantially, while it makes the rest slightly worse off. Then, the voting outcome under a simple majority rule is likely the status quo, even if the reform may achieve a higher total welfare. By disclosing information more "partially" (specifically, by hiding "who" would exactly be better off), the voting outcome might be improved. Under some conditions, we are able to fully characterize the optimal disclosure policy.

More specifically, we consider the following problem. A group of voters choose by referendum whether to implement a reform or keep the status quo. The value of the reform to each voter is assumed to be the sum of two orthogonal real-valued components. The first component is each voter's private information: prior to the election, each voter receives a private signal about this first component. These signals are assumed to be i.i.d. across voters; we will refer to this component as the "idiosyncratic" value of the reform to each voter. This information is private in that neither the other voters nor the government have access to it. Going back to the labor market reform example, this first component could capture short term adjustment costs, which are individual-specific, and hard to evaluate for the government. By contrast, the second component is not under the voter's control in that he has no way to learn directly about it. Instead, he has to rely on the inquiries of the government. We assume, as in Kamenica and Gentzkow [2011] or Alonso and Camara [2016], that the government can design a policy experiment that

3

publicly reveals information about a payoff relevant state pertaining to the second component of all voters. For example, this second component could represent the long-term economic benefits and costs of the reform, which might be difficult for the voters to assess. In particular, it might be difficult to assess not only the total average effect of the reform for the economy (is it good or bad on average) but also its distributional consequences (who are the (relative) winners and losers). We assume that the objective of the government when designing the policy experiment is to maximize the expected utilitarian welfare. In that sense, ex ante, the preferences of the government are perfectly aligned with those of the electorate as a whole. After observing the realization of their own private signal and the result of the policy experiment, voters apply Bayes' rule. They then vote and the reform is implemented if and only if it receives the approval of a majority of voters.

As mentioned above, full disclosure will be shown to be strictly suboptimal in general. A non-trivial aspect of the problem is that there are many ways of making the information "partially" disclosed. Indeed, our problem is a complex information design problem in that the planner's information is represented by a multi-dimensional, continuous random vector, and that each voter has his own orthogonal private taste information. Not much is known about the general structures of the solution to this class of problems.[1] We show that, under certain conditions, the optimal disclosure policy only provides *anonymized* information commonly to the voters. That is, if a reform will make some voters better off and others worse off (or some voters much better off and others only slightly better off), then the planner should only disclose the *average* gain of the reform across the electorate. In a sense, voters should be kept under a partial "veil of ignorance". We will show how hiding information that makes voters "asymmetric" proves useful in mitigating the ordinality-cardinality mismatch mentioned above. Although partially dis-

---

[1]See, for example, Malamud and Schrimpf [2021] and Dworczak and Kolotilin [2023], for Bayesian persuasion problems with multi-dimensional, continuous payoff states (but without private taste information); and Kolotilin et al. [2017] for persuasion mechanisms of a privately informed receiver in one-dimensional linear environments.

closing information does not completely resolve this mismatch between the majoritarian choice and the socially optimal alternative, we provide some conditions under which it (optimally) attenuates it.

Among the conditions that imply the optimality of the disclosure of anonymized information, two are key. The first set of conditions is about the distribution of the idiosyncratic preferences. The detail is explained later, but for the moment, let us just note that the conditions, called the *increasing hazard rate* and *decreasing mean residual life* conditions, are satisfied by many popular distributions. The other key condition imposes some constrains on how big the uncertainty regarding the information controlled by the social planner can be relative to the size of the idiosyncratic preferences. This condition is related to the adequacy of the majority voting rule in this environment *given the planner's information at the time of the vote.* Specifically, if the uncertainty regarding the information controlled by the social planner is "large" (in a sense we will make precise), then there will be events where, given her information at the time of the vote and the profiles of individual votes, the social planner will be willing to overrule the majoritarian decision and not to follow the will of the majority (Remember the social planner cannot observe the idiosyncratic preferences, but can learn about them through the individual votes). If this is the case, the disclosure of the anonymized information can be suboptimal. Intuitively, the optimal information disclosure may "try to fix" the inefficiency caused by the inadequate voting rule by creating some asymmetry of the expected gains of the reform across voters (the logic is similar to that in Alonso and Camara [2016]).

Related literature

Our work is related to a series of recent papers applying Bayesian persuasion to voting environments. Extending the framework with a single receiver of Kamenica and Gentzkow [2011], Alonso and Camara [2016] study a model where a group of uninformed voters must vote to choose whether to keep the status quo or implement a proposed reform. Prior to the vote a politician

– who is assumed to be biased in favor of the reform – can influence voters'
choices by strategically designing an experiment that reveals some informa-
tion about voters' payoffs. Bardhi and Guo [2018] study an environment
where a sender can design an experiment to inform voters about their (pos-
sibly) correlated payoff states. Assuming that the sender prefers the reform
to be implemented whatever the state of the world, they characterize the
sender-optimal policy when the voting rule is the unanimity rule. Compared
to these two papers, a key difference of our model is that we assume that
this politician/social planner is fully benevolent and maximizes the voters'
aggregate welfare. Another key difference is that voters have some private
information about how the reform is going to impact them individually. In
our model, information disclosure therefore affects not only the probability
that the reform is implemented, as in the previous papers, but also what is
learned about the voters' idiosyncratic preferences (which is valuable to the
benevolent planner).

With multiple receivers, Bayesian persuasion is an appropriate approach
to study optimal *public* disclosure policy, as assumed in this paper. Given any
public disclosure policy, each voter plays his dominant action, and then our
problem essentially becomes a Bayesian persuasion problem where a sender
controls the informativeness of multi-dimensional continuous random vari-
ables. Rayo and Segal [2010], Malamud and Schrimpf [2021], and Dworczak
and Kolotilin [2023] consider such problems with two or higher-dimensional
random variables, and obtain general properties of the optimal disclosure
policy. The specific structure of our problem enables us to fully characterize
the optimal disclosure policy under certain conditions.[2]

A critical part of our argument rests on the observation that (i) the in-

---

[2]We share with all the aforementioned papers the assumption that information acqui-
sition is costless: our focus is on disclosure only. Another strand of the literature instead
considers settings where agents have to acquire information before making a collective
decision, and that information acquisition is costly. For papers studying the optimal de-
sign of information acquisition, see for example Persico [2004], Caillaud and Tirole [2007],
Gerardi and Yariv [2008], Gershkov and Szentes [2009].

formation disclosure policy affects the voting outcome, and (ii) there can be conflict between cardinal preferences aggregated by simple majority voting and ordinal preferences. We are obviously not the first to make these remarks, and to note that this might create a tension between information revelation and utility maximization. Fernandez and Rodrik [1991] highlight the fact that in the presence of individual-specific uncertainty about the consequences of a reform, whether a majority supports the reform or not can critically depend on what is known about the identity of the winners and losers. In their setting, full information about the consequences of the reform is learned only if (and after) the reform is implemented. If the reform is implemented and a majority learns that it is actually hurt by it, the reform can be repelled. By contrast, if a majority is against the reform ex ante, the reform is not implemented, even if ex post a majority might have turned out to be in favor. This implies a bias towards the status quo. Our information structure is different: we consider situations where the relevant information about the value of the reform can be learned ex ante; our focus is also different in that we look for the optimal information design. Gersbach [1992] studies a setting where society has to decide whether to implement a project or not. Prior to voting on the project, society decides by a vote whether to acquire (full) information about the value of the project. He shows that a majority will always support public disclosure of information, although it might be sub-optimal from a social welfare perspective (See also Gersbach [1995] for similar insights in a setting where the decisions to acquire information are decentralized). Jackson and Tan [2012] study a two-alternative voting environment, where voters have their own, publicly-known biases for each of the two alternatives, but are imperfectly informed about their respective costs. Before the vote takes place, informed (and biased) experts can reveal some information about these costs. They provide an example where welfare would be higher if all information were suppressed (see also Schnakenberg [2015]). Contrary to this setting, we have no biased experts: the social planner is assumed to be fully benevolent and we characterize the

welfare-social information policy. Sun et al. [2021] study the case where the (possibly biased) informed experts can reveal some information about the voters' *common* one-dimensional preference state. It is a common state in the sense that it affects each voter's preference in an homogeneous manner, and in this sense, their information structure is in our *anonymized* class by definition. Our paper considers a more general state space and show that anonymized information is optimal under certain conditions. On the other hand, Sun et al. [2021] allow for more general objectives, and identify an appropriate *single-crossing* condition under which a censorship information policy is optimal.

A more distant literature has studied information aggregation in Condorcet Jury settings. In the jury metaphor, jurors agree that a guilty defendant should be convicted and an innocent acquitted (although they may differ in their thresholds of doubt), and each has some private information about the state of the world.[3] Our informational setting is fundamentally different from that studied in this literature in that we consider a situation where each voter's private information is irrelevant to the other voters' policy-related payoff: his private information does not contain any information about whether the reform is good or bad for other voters. In this sense, our paper also relates to the literature studying optimal voting rules in a private-value setting with cardinal preferences (See Barberà and Jackson [2006] or Azrieli and Kim [2014] in the case of two alternatives, and Gershkov et al. [2017] for any number of alternatives).

---

[3]This simple setting has been extended in many ways to explore a number of questions: information aggregation when decision is reached by strategic voters without (Austen-Smith and Banks [1996], Feddersen and Pesendorfer [1998]) or with a prior stage of deliberation (Austen-Smith and Feddersen [2006], Gerardi and Yariv [2007], when sequential voting is allowed (Dekel and Piccione [2000]), when information acquisition is costly (Persico [2004], Gershkov and Szentes [2009]), when voters can get additional advice from potentially biased experts (Jackson and Tan [2012], Schnakenberg [2015]), or when the decision is delegated to a (better informed) elected politician (Feddersen and Pesendorfer [1996], Feddersen and Pesendorfer [1997], Feddersen and Pesendorfer [1999]).

2 Model

2.1 Environment

A society with N (N odd, $N \geq 3$) voters makes a choice between two alternatives: $k = 1$ ("reform") and $k = 0$ ("status quo"). Each voter i's payoff is normalized to 0 under the status quo, and is $u_i = \theta_i + \varepsilon_i \in \mathbb{R}$ under the reform. Thus, the utilitarian social welfare is given by:

$$\sum_{i=1}^{N} (\theta_i + \varepsilon_i) k$$

if policy k is implemented.

$\theta_i \in [-\bar{\theta}, \bar{\theta}]$ for some $\bar{\theta} > 0$, and the information about θ_i is controlled by the social planner, as we explain below.

$\varepsilon_i \in [-\bar{\varepsilon}, \bar{\varepsilon}]$ for some $\bar{\varepsilon} > \bar{\theta}$ is i's idiosyncratic payoff term, whose realization is i's private information (that is, i knows its realization at the time of voting, while no one else does). Each ε_i is distributed independently from ε_{-i} and θ. Let G denote its cdf with density g. We assume that g has full-support, is continuously differentiable and symmetric around 0.[4]

2.2 Social Planner and Information Disclosure

The social planner is a benevolent (in the sense of maximizing utilitarian welfare) entity who designs the voters' information about $\theta = (\theta_i)_{i=1}^{N}$. Through this planner's solution, we can provide a normative benchmark regarding the socially desirable information structure.

Assume that θ follows a joint distribution F, which is invariant with respect to permutations ("(N-)exchangeable"). Note that this allows for a variety of (especially positive) correlation. A popular case in the literature is with perfect correlation: $\mathbb{P}(\theta_i = \theta_j, \forall i, j) = 1$, which is our special case. However, we allow for more general (imperfect) correlation structures.

---

[4]The symmetry assumption is just for simplicity. Also, g does not necessarily have a bounded support, we could have g with full-support on $\mathbb{R}$.

The planner's information disclosure policy is denoted as follows:

$$\phi : [-\overline{\theta}, \overline{\theta}]^N \ (\equiv \Theta) \ \to \Delta(X),$$

for some (rich enough) space X, with the interpretation that, once $\theta \in [-\overline{\theta}, \overline{\theta}]^N$ is realized, then the planner publicly discloses a "signal" $x \in X$ to the voters, where x is drawn from distribution $\phi(\theta) \in \Delta(X)$. To the extent that the distribution $\phi(\theta)$ varies with θ, the signal x provides some (possibly noisy / imperfect) information about the realized θ.

By the standard argument in Bayesian persuasion, it is without loss of generality to assume $X = \Theta$, and for each $x = (x_i)_{i=1}^N$, $E[\theta_i|x] = x_i$. That is, the planner's signal x directly tells each i what the expected value of θ_i should be, and each i finds (applying Bayesian updating) that, given signal x, the conditional expected value of θ_i is indeed x_i.

We provide two instances of disclosure policies that play a key role in this paper.

Definition 1. (i) ϕ is a *full disclosure* policy if $\mathbb{P}[x_i = \theta_i, \ \forall i] = 1$.

(ii) ϕ is an *anonymous disclosure* policy if $\mathbb{P}[x_i = x_j, \ \forall i,j] = 1$. In particular, ϕ *fully discloses anonymized* θ if $\mathbb{P}[x_i = \frac{\sum_{j=1}^N \theta_j}{N}, \ \forall i] = 1$.

2.3   Voting

The collective decision is made using a simple majority rule, so that the reform happens if and only if at least $\frac{N+1}{2}$ voters vote for it (recall N is odd). Given a signal x from the planner, voter i votes for the reform if and only if $x_i + \varepsilon_i \geq 0$.

For $z \in [-\overline{\varepsilon}, \overline{\varepsilon}]$, let:

$$
\begin{aligned}
u_+(z) &= \mathbb{E}[z + \varepsilon_i | z + \varepsilon_i > 0], \\
u_-(z) &= \mathbb{E}[z + \varepsilon_i | z + \varepsilon_i < 0]
\end{aligned}
$$

denote a voter's expected payoffs conditional on receiving signal z and his voting for / against the reform, respectively. It is also useful to introduce

10

the probability of voting for / against the reform if the voter has beliefs z:

$$
\begin{aligned}
p_+(z) &= \mathbb{P}[z + \varepsilon_i > 0] = 1 - G(-z) = G(z), \\
p_-(z) &= \mathbb{P}[z + \varepsilon_i < 0] = G(-z).
\end{aligned}
$$

We impose a condition on G that implies useful properties for u_+, u_-, p_+ and p_-.

Assumption 1. G is log-concave.

Assumption 1 together with the symmetry of g guarantee that

(i) p_+ and p_- are log-concave (the "monotone increasing hazard rate" property for G)

(ii) u_+ and u_- are strictly increasing (the "monotone decreasing mean-residual life" property for G)

These last two properties will be extensively used in the sequel. Most standard distributions satisfy this properties (See Bagnoli and Bergstrom [2005]).

3 A simple example: Suboptimality of a full disclosure policy

We first observe that a full disclosure policy can be strictly suboptimal, even in a very simple example.

Example 1. Assume $N = 3$ and $\varepsilon_i \sim U\left(-\bar{\varepsilon}, \bar{\varepsilon}\right)$. Regarding θ, assume that $\theta = (-\bar{\theta}, 0, \bar{\theta})$ or one of its permutations, all with equal probability $\frac{1}{6}$, with $\bar{\theta} \leq \bar{\varepsilon}$. Note that $\sum_j \theta_j = 0$: the aggregate non-idiosyncratic value of the reform is zero.

In the case of a uniform distribution for ε_i, Assumption 1 is satisfied as G is linear. Furthermore:

$$u_+(z) = \frac{z + \overline{\varepsilon}}{2} \;;\; p_+(z) = \frac{z + \overline{\varepsilon}}{2\overline{\varepsilon}}$$
$$u_-(z) = \frac{z - \overline{\varepsilon}}{2} \;;\; p_-(z) = \frac{-z + \overline{\varepsilon}}{2\overline{\varepsilon}}$$

- Consider first a full disclosure policy. The reform is implemented if and only if at least two voters vote in favor of the reform. Therefore expected welfare is:

$$\left[p_+(-\overline{\theta}) * p_+(0) * p_+(\overline{\theta}) \right] * \left[u_+(-\overline{\theta}) + u_+(0) + u_+(\overline{\theta}) \right]$$
$$+ \left[p_+(-\overline{\theta}) * p_+(0) * p_-(\overline{\theta}) \right] * \left[u_+(-\overline{\theta}) + u_+(0) + u_-(\overline{\theta}) \right]$$
$$+ \left[p_+(-\overline{\theta}) * p_-(0) * p_+(\overline{\theta}) \right] * \left[u_+(-\overline{\theta}) + u_-(0) + u_+(\overline{\theta}) \right]$$
$$+ \left[p_-(-\overline{\theta}) * p_+(0) * p_+(\overline{\theta}) \right] * \left[u_-(-\overline{\theta}) + u_+(0) + u_+(\overline{\theta}) \right]$$
$$= \left[\frac{1}{8} - \frac{1}{8}\left(\frac{\overline{\theta}}{\overline{\varepsilon}}\right)^2 \right] * \frac{3}{2}\overline{\varepsilon} + \left[\frac{3}{8} + \frac{1}{8}\left(\frac{\overline{\theta}}{\overline{\varepsilon}}\right)^2 \right] * \frac{1}{2}\overline{\varepsilon}$$

If all three voters vote in favor of the reform, which happens with probability $\frac{1}{8} - \frac{1}{8}\left(\frac{\overline{\theta}}{\overline{\varepsilon}}\right)^2$, one may check that the expected welfare unconditional on the ε_i is $\frac{3}{2}\overline{\varepsilon}$. Thus the first term in the sum in the last line. If two voters exactly vote in favor of the reform, which happens with probability $\frac{3}{8} + \frac{1}{8}\left(\frac{\overline{\theta}}{\overline{\varepsilon}}\right)^2$, expected welfare unconditional of the on ε_i is $\frac{1}{2}\overline{\varepsilon}$ (independently of the identity of these two voters). Thus the second term in the sum in the last line.

- There exists a better disclosure policy. Consider an anonymous disclosure policy with $x_i = 0$ for all i. If all three voters are in favor of the reform, which happens with probability $\frac{1}{8}$, welfare unconditional of the on ε_i is $3u_+(0) = \frac{3}{2}\overline{\varepsilon}$. If exactly two voters out of three are in favor of the reform, which happens with probability $\frac{3}{8}$, welfare unconditional of the on ε_i is $2u_+(0) + u_-(0) = \frac{1}{2}\overline{\varepsilon}$. Therefore, under anonymous

12

disclosure, expected welfare is

$$\frac{1}{8} * \frac{3}{2}\overline{\varepsilon} + \frac{3}{8} * \frac{1}{2}\overline{\varepsilon},$$

which is higher than the expected welfare under full disclosure as soon as $\overline{\theta} > 0$.

Remark 1. The probability of implementing the reform is the same under both information disclosure policies (1/2). The key difference is that the anonymous disclosure policy allows to better detect situations where the idiosyncratic components of the utilities are high (when all three voters approve the reform). Intuitively, the voters' preferences under anonymous disclosure are more aligned than under full disclosure, which turns out to be useful in order to mitigate the mismatch between their ordinal and cardinal preferences. In the next section, our main result will show that if $\overline{\theta} < \frac{1}{3}\overline{\varepsilon}$, anonymous disclosure is actually the optimal information disclosure policy. In this sense, our approach of anonymizing the information may be interpreted as a formalization of the classical "veil-of-ignorance" idea.

Remark 2. Note nevertheless that the anonymous disclosure policy cannot achieve the first best outcome. Indeed, when the aggregate non-idiosyncratic value of the reform is zero, the first best would require to implement the reform if $\sum_i \varepsilon_i > 0$ and not to implement it if $\sum_i \varepsilon_i < 0$. One may check that:

$$\mathbb{P}\left[\sum_i \varepsilon_i \geq 0\right] * \mathbb{E}\left[\sum_i \varepsilon_i \mid \sum_i \varepsilon_i \geq 0\right] = \frac{13}{32}\overline{\varepsilon}$$

while the expected welfare under anonymized disclosure policy is $\frac{3}{8}\overline{\varepsilon}$.[5]

---

[5]The reason why the first best cannot be achieved in this simple example is very general: whatever the information disclosure policy, in some situations (realization of the idiosyncratic preferences) there will be a mismatch between what a majority of voters want, and what is socially optimal. This mismatch, also known as the "intense minority" problem, is well-known: A (one-shot) voting mechanism such as the simple majority rule elicits only ordinal information of the voters' preferences but not cardinal information.

□

In the next section, we provide sufficient conditions under which fully disclosing anonymized θ is the optimal information disclosure policy.

4 Welfare-Optimal Information Disclosure

Our main result is shown in two steps. First, under certain conditions, for any non-anonymous disclosure policy, there is an anonymous policy that is welfare-improving. Second, under additional conditions, fully disclosing anonymized θ is welfare-optimal.

4.1 Expected Welfare

The following expression for the expected welfare given any signal realization x is useful in the proof of our result.

Lemma 1. Given any disclosure policy and any realized x, the expected welfare conditional on x is, denoting $Q = \frac{N+1}{2}$ the threshold for approval under the simple majority rule:

$$
W(x) \equiv \mathbb{P}[k = 1|x] * \mathbb{E}\Big[\sum_{i=1}^{N} u_i|x, k = 1\Big]
$$

$$
= \mathbb{P}(M_{-1,2} = Q - 2)p_+(x_1)p_+(x_2) * \Big[u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\Big]
$$

$$
- \mathbb{P}(M_{-1,2} = Q - 1)p_-(x_1)p_-(x_2) * \Big[u_-(x_1) + u_-(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 1]\Big]
$$

$$
+ \mathbb{P}(M_{-1,2} \geq Q - 1) * \Big[x_1 + x_2 + \mathbb{E}[U_{-1,2}|M_{-1,2} \geq Q - 1]\Big] \tag{1}
$$

where $M_{-i,j}$ is the total number of votes for the reform among the voters except i and j and $U_{-i,j}$ is the total surplus excluding i and j's payoffs. Both depend on $x_{-i,j}$ (which we will omit in the equations whenever there is no risk of confusion).

14

The expression may be interpreted as follows. Note first that the reform can be implemented only if, among the $N-2$ voters other than 1 and 2, at least $Q-2$ vote in favor of the reform (otherwise, the threshold of $Q = \frac{N+1}{2}$ approvals can never be reached).

The first term on the right-hand side of (1) deals with events such that among the $N-2$ voters other than 1 and 2, *exactly* $Q-2$ vote in favor of the reform ($M_{-1,2} = Q-2$). In that case, the reform is implemented if and only if both voters 1 and 2 vote in favor of the reform, in which case the total conditional utility in society is $u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q-2]$.

Consider now events such that among the $N-2$ voters other than 1 and 2, at least $Q-1$ voters vote in favor of the reform. In that case, the reform is always implemented, except if the following two conditions simultaneously hold: (i) among the $N-2$ voters other than 1 and 2, exactly $Q-1$ vote in favor of the reform; (ii) both voters 1 and 2 vote against. The last two terms on the right-hand side of (1) follow this decomposition.

The last term on the right-hand side of (1) can be interpreted as the (hypothetical) welfare that would be generated by events such that among the $N-2$ voters other than 1 and 2, at least $Q-1$ voters vote in favor of the reform *if the reform were implemented whatever the votes of voters 1 and 2.* Under these hypothetical conditions, no information is learnt about ε_1 and ε_2 (as the votes of voters 1 and 2 have no impact on the collective decision), and the total conditional utility in society is $x_1 + x_2 + \mathbb{E}[U_{-1,2}|M_{-1,2} \geq Q-1]$.

But as noted above, when $M_{-1,2} \geq Q-1$, the reform is in fact not always implemented. To get the actual welfare, one has to substract events where $M_{-1,2} = Q-1$ and both voters 1 and 2 vote against the reform. This corresponds to the second term on the right-hand side of (1). In that case, the total conditional utility in society is $u_-(x_1) + u_-(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q-1]$.

In this decomposition, the first term pertains to events such that voters 1 and 2 are jointly pivotal in getting the reform implemented, while the second term pertains to events such that voters 1 and 2 are jointly pivotal in blocking the reform.

4.2 Criterion for Optimality of Anonymous Policy

We first provide sufficient conditions under which, for any non-anonymous disclosure policy, there is an anonymous policy that is welfare-improving.

Roughly, anonymizing information means that we send a less dispersed signal. For example, instead of sending a signal $x = (x_1, x_2, x_3, \ldots, x_N)$ with $x_1 < x_i < x_2$ for all $i \neq 1, 2$, imagine that we send $(x_1 + \delta, x_2 - \delta, x_3, \ldots, x_N)$, for a small positive δ.[6] The expected welfare changes by:

$$\delta\Big(W_1(x) - W_2(x)\Big),$$

where $W_i(x) = \frac{\partial W(x)}{\partial x_i}$. Thus, if this is positive, anonymizing information improves the expected welfare.

Based on this simple criterion, we provide sufficient conditions where an anonymous policy is optimal. Recall that from Lemma 1, total welfare $W(x)$ can be decomposed according to (1). Regarding the derivatives of W with respect to x_1 and x_2, two remarks directly follow from this expression. First, the third term being linear in $x_1 + x_2$, its derivatives with respect to x_1 and x_2 are equal, and thus they cancel out in the difference $W_1 - W_2$. Second, given the symmetry of g the first and second lines should deliver qualitatively similar effects (as will be checked in the proof). Thus, in order to provide the intuition, we focus here on the first line. Denoting:

$$H(x) := p_+(x_1)p_+(x_2) * \Big[u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\Big], \quad (2)$$

we have:

$$
\begin{aligned}
H_1(x) &= p'_+(x_1)p_+(x_2) * \Big[u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\Big] \\
&\quad + p_+(x_1)p_+(x_2) * u'_+(x_1)
\end{aligned}
$$

---

[6]Of course, we must do it simultaneously appreciating the Bayes update constraint: $\mathbb{E}[\theta_i|x] = x_i$ for all i. This can be done by changing the policy for some other x' at the same time, and is taken care of in the formal proof of Theorem 1. Here, we omit this complication. Notice that, without the Bayes update constraint, the problem is analogous to the one where the planner promises some "transfer" from / to each agent, in a balanced-budget manner.

and hence,

$$H_1(x) - H_2(x)$$

$$= p_+(x_1)p_+(x_2)\left(\frac{p'_+(x_1)}{p_+(x_1)} - \frac{p'_+(x_2)}{p_+(x_2)}\right) * \left[u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\right]$$

$$+ p_+(x_1)p_+(x_2)) * \left[u'_+(x_1) - u'_+(x_2))\right]. \tag{3}$$

The first line on the right-hand-side can be interpreted as the *probability effect*. By increasing x_1 and decreasing x_2 by the same amount, the probability that voters 1 and 2 both vote for the reform increases. Indeed, by the log-concavity of p_+ (increasing hazard rate of G), $p'_+(x_1)p_+(x_2) - p_+(x_1)p'_+(x_2) > 0$. The probability effect is positive if the utility in society conditional on voters 1 and 2 being pivotal is positive, that is, $u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2] > 0$. We now discuss a condition that guarantees that this is the case.

Define

$$m^* := \max\left\{m \in [0, \bar{\varepsilon}] : \text{For all } (z_1, \ldots, z_N) \in [-m, m]^N, \right.$$

$$\left. \sum_{i=1}^{\frac{N+1}{2}} u_+(z_i) + \sum_{i=\frac{N+1}{2}+1}^{N} u_-(z_i) \geq 0\right\}. \tag{4}$$

Notice $m^* > 0$, as $\frac{N+1}{2}u_+(0) + \frac{N-1}{2}u_-(0) = u_+(0) > 0$.

Condition 1. $\bar{\theta} \leq m^*$.

Condition 1 stipulates that whatever the voters' signal, the conditional expected welfare when there is a one-vote margin in favor of the reform is positive. When it is satisfied, $u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2] > 0$ for all x, and the probability effect (first line on the right-hand side of (3)) is positive.

The second line of (3) can be interpreted as the *information effect*. By changing the disclosure policy, one changes what is learnt about the idiosyncratic component of the preferences of the voters when they both vote in favor

17

of the reform. In case of a uniform distribution, u_+ is linear and hence this effect does not exist. In the general case, although u'_+ itself is always positive (Assumption 1), by the potential non-linearity of u_+, this information effect may be positive or negative.

The discussion suggests that anonymity is more likely be optimal with stronger monotonicity in hazard rate and larger conditional postive gain from the reform (and in this sense, the probability effect is positive and more significant), and with more regulated non-linearity (so that the potentially negative information effect is less significant).

Proposition 1 provides a formalization of this intuition. For $z \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, let

$$\alpha(z) := \frac{u''_+(z)}{\left(\frac{g(z)}{G(z)}\right)'} + 2u_+(z) \tag{5}$$

The function $\alpha(.)$ captures information about the curvature of u_+ and the monotonicity of hazard rate. Define

$$\hat{m} := \max \left\{ m \in [0, \bar{\varepsilon}] : \text{ For all } z \in [-m, m], \right.$$
$$\left. \alpha(z) + \frac{N-3}{2} * u_+(-m) + \frac{N-1}{2} * u_-(-m) \geq 0 \right\} \tag{6}$$

Notice that when $(x_3, \ldots, x_N) \in [-m, m]^{N-2}$, $(Q-2)u_+(-m)+(N-Q)u_-(-m)$ can be interpreted as the largest externality that voters 1 and 2 can impose on the rest of the electorate when they are jointly pivotal in getting the reform implemented.

We show in Appendix A that Assumption 1 guarantees that $\hat{m} > 0$.

Proposition 1. Assume Assumption 1. If $\bar{\theta} < \hat{m}$, then for all $x \in [-\bar{\theta}, \bar{\theta}]^N$:

$$x_1 < x_2 \Rightarrow W_1(x) > W_2(x).$$

We prove Proposition 1 in Appendix A.

4.3 Full disclosure of anonymized information

With additional conditions, full disclosure of anonymized θ is optimal. The additional conditions, though having different expressions, are qualitatively similar to the previous conditions for Proposition 1.

Let $W^*(z) = W(z, \ldots, z)$ for $z \in [-\bar{\theta}, \bar{\theta}]$ denote the expected welfare under any anonymous disclosure policy, when z is announced as the average value of θ. We identify sufficient conditions guaranteeing that W^* is convex, and hence full disclosure of anonymized θ is optimal.

Observe that:

$$\frac{\partial^2 W^*(z)}{\partial z^2} = \sum_{i,j} W_{i,j}(z, \ldots, z)$$

$$= W_{11}(z, \ldots, z) + (N-1)W_{12}(z, \ldots, z).$$

Recall the three terms in the decomposition of $W(x)$ (See (1)). The third line is linear in $x_1 + x_2$, and hence its second (own or cross) derivative is zero. The first and second lines are again qualitatively similar to each other, so we focus on $H(x)$ (see (2)). That is, defining

$$H^*(z) := H(Z, \ldots, z),$$

we seek some conditions under which

$$\frac{\partial^2 H^*(z)}{\partial z^2} = H_{11}(z, \ldots, z) + (N-1)H_{12}(z, \ldots, z) > 0,$$

where H_{11}, H_{12} are the own and cross second derivatives of H.

As seen above, the first derivative of H with respect to x_1 is:

$$H_1(x) = p'_+(x_1)p_+(x_2)\Big(u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q-2]\Big)$$
$$+ p_+(x_1)p_+(x_2)u'_+(x_1).$$

Therefore:

$$H_{12}(x) = p'_+(x_1)p'_+(x_2)\Big(u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q-2]\Big)$$
$$+ p_+(x_1)p'_+(x_2)u'_+(x_1) + p'_+(x_1)p_+(x_2)u'_+(x_2) > 0.$$

19

and

$$H_{11}(x) = p''_+(x_1)p_+(x_2)\Big(u_+(x_1) + u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\Big)$$
$$+ 2p'_+(x_1)p_+(x_2)u'_+(x_1) + p_+(x_1)p_+(x_2)u''_+(x_1).$$

It follows that

$$\frac{\partial^2 H^*(z)}{\partial z^2} = N(p'_+(z))^2 w(z) + \Big(p_+(z)p''_+(z) - (p'_+(z))^2\Big)w(z)$$
$$+ 2Np_+(z)p'_+(z)u'_+(z) + p_+^2(z)u''_+(z)$$

where

$$w(z) := \frac{N+1}{2} * u_+(z) + \frac{N-1}{2} * u_-(z) \tag{7}$$

is the conditional expected welfare when all voters have z and there is a one vote margin in favor of the reform. Noticing that $u''_+ = -\left(\frac{g}{G}\right)' u_+ - \frac{p'_+}{p_+} u'_+$ and rearranging terms, we have:

$$\frac{\partial^2 H^*(z)}{\partial z^2} = p'_+(z))^2 * \beta(z)$$

where

$$\beta(z) := Nw(z) + (2N - 1)\left(\frac{G(z)}{g(z)}\right)u'_+(z) - \left(\frac{G(z)}{g(z)}\right)'\Big(w(z) - u_+(z)\Big)$$

Define

$$\tilde{m} := \max\Big\{m \in [0, \bar{\varepsilon}] : \text{ for any } z \in [-m, m], \beta(z) \geq 0\Big\}. \tag{8}$$

As $w(0) = u_+(0) > 0$, Assumption 1 implies that when z is close enough to 0, $\beta(z)$ is strictly positive, and therefore $\tilde{m} > 0$.

Proposition 2. Assume Assumption 1. If $\bar{\theta} < \tilde{m}$, then W^* is convex.

Proof. The proof directly follows from the discussion above and from the definition of $\tilde{m}$. A similar logic applies to the second line of $W^*(x)$. □

Propositions 1 and 2 provide the key steps towards our main result, stated in Theorem 1:

Theorem 1. Assume Assumption 1. If $\bar{\theta} < \min\{\hat{m}, \tilde{m}\}$, then it is an optimal disclosure policy to fully disclose the anonymized information.

The proof of Theorem 1 is provided in Appendix B. To provide an illustration of our main result in the case of some familiar distributions, consider the case where the idiosyncratic component of the utility is distributed according to a uniform or a normal distribution.

Claim 1. (i) Assume $\varepsilon_i \sim U(-\bar{\varepsilon}, \bar{\varepsilon})$. Assumption 1 is satisfied and:

$$\hat{m} = \frac{1}{N}\bar{\varepsilon} = m^*$$
$$\tilde{m} = \frac{3N-1}{N(N+1)}\bar{\varepsilon} \in \left(\hat{m}, \bar{\varepsilon}\right)$$

As discussed in Section 4.2, in the case of a uniform distribution, there is no information effect, and the probability effect is positive as soon as Condition 1 is satisfied. This explains why $m^* = \hat{m}$ in that case.

(ii) Assume $\varepsilon_i \sim \mathcal{N}(0, \sigma)$. Assumption 1 is satisfied, and when $N = 3$:

$$\hat{m} = \sigma * 0.5606 < m^* = \sigma * 0.7678$$
$$\tilde{m} = \sigma * 1.05427$$

In the case of a normal distribution, the information effect is negative as u_+ is convexe. This explains why $\hat{m} < m^*$ in that case.

The proof of Claim 1 is provided in Appendix C.

Remark 3. Theorem 1 provides an upper bound for $\bar{\theta}$ given the distribution of each ε_i fixed. An alternative interpretation of this result is to take $\bar{\theta}$ as given, and imagine that g is in a scale-family of another "baseline" density γ:

$$g(\varepsilon_i) = \frac{1}{\sigma}\gamma\left(\frac{\varepsilon_i}{\sigma}\right)$$

for a scale parameter $\sigma > 0$. In particular, if this baseline distribution is chosen to be standardized, σ is the standard deviation of g. In Appendix D, we provide an equivalent statement of Theorem 1 that concerns a lower bound of σ, the scale parameter for the distribution of each ε_i, given $\bar\theta$ fixed (See Corollary 1 in Appendix D).

5   Discussion

In this section, we first illustrate the critical role of Assumption 1, and of having θ small enough. We then briefly discuss the cases of a priori asymmetric voters and of non-separable idiosyncratic noises. Last, we touch upon the question of whether the social planner could improve welfare by using more complex information disclosure mechanisms.

5.1   The critical role of Assumption 1

Theorem 1 states that if Assumption 1 is satisfied, then it is an optimal disclosure policy to fully disclose the anonymized information if $\bar\theta$ is smaller than a threshold (which depends on g and on the number of voters). In this section, we provide an example showing that if u_+ is strictly decreasing on some interval $[-\kappa, \kappa]$ in the support of g (and thus violates Assumption 1), then we can exhibit a family of distributions for θ, $\{F_{\bar\theta}\}_{\bar\theta \in (0,\kappa)}$, such that the full disclosure information policy strictly dominates the anonymyzed disclosure policy for all $\bar\theta$.

Example 2. Assume $N = 3$ and, as in Example 1, $\theta = (-\bar\theta, 0, \bar\theta)$ or one of its permutations, all with equal probability $\frac{1}{6}$. For the time being, we make no specific assumptions on g (other than assuming that g is symmetric around 0 and that its support includes $[-\bar\theta, \bar\theta]$).

Using the decomposition of expected welfare provided in Lemma 1, the

expected welfare under full disclosure can be written as:

$$W(-\overline{\theta}, 0, \overline{\theta}) = p_-(0) * p_+(-\overline{\theta})p_+(\overline{\theta}) * \left(u_+(-\overline{\theta}) + u_+(\overline{\theta}) + u_-(0) \right)$$
$$-p_+(0) * p_-(-\overline{\theta})p_-(\overline{\theta}) * \left(u_- + (-\overline{\theta}) + u_-(\overline{\theta}) + u_+(0) \right)$$
$$+p_+(0) * \left((-\overline{\theta}) + (\overline{\theta}) + u_+(0) \right)$$

while the expected welfare under anonymous disclosure is:

$$W(0, 0, 0) = p_-(0) * p_+^2(0) * \left(2u_+0) + u_-(0) \right)$$
$$-p_+(0) * p_-^2(0) * \left(2u_-(0) + u_+(0) \right)$$
$$+p_+(0) * \left((0) + (0) + u_+(0) \right)$$

Therefore full disclosure is better than anonymous disclosure iff

$$p_+(-\overline{\theta})p_+(\overline{\theta}) * \left(u_+(-\overline{\theta}) + u_+(\overline{\theta}) + u_-(0) \right) > p_+^2(0) * \left(2u_+(0) + u_-(0) \right) \quad (9)$$

We recognize on the left-hand side $H(-\overline{\theta}, \overline{\theta}, 0)$, and on the right-hand side $H(0, 0, 0)$. Using the property that $\frac{\partial}{\partial z}\left(p_+(z)u_+(z) \right) = p_+(z)$, one may check that

$$\frac{\partial}{\partial z}\left[H(-z, z, 0) \right] = p_+(z)p_+'(-z) * \left(u_+(0) - u_+(z) \right)$$
$$+ p_+'(z)p_+(-z) * \left(u_+(-z) - u_+(0) \right)$$

At this point, we can explain the importance of Assumption 1. Suppose that Assumption 1 is violated so that we have a reverse monotonicity around zero. Specifically, assume that there exists $\kappa > 0$ such that u_+ is strictly decreasing on $[-\kappa, \kappa]$ in the support of g. Then if $\overline{\theta} < \kappa$, for $0 < z \leq \overline{\theta}$, $u_+(z) < u_+(0) < u_+(-z)$. In that case, the expression above is strictly positive, and inequality (9) holds. That is, the anonymous disclosure policy is strictly worse than full disclosure.

As mentioned when introducing Assumption 1 earlier in the text, most standard distributions do satisfy Assumption 1. It is neverthemess possible

23

to exhibit some "pathological" distributions such that u_+ is locally decreasing around 0. For example, consider the following truncated Cauchy distribution: For $z \in [100, 100]$,

$$g(z) = \frac{1}{\int_{-100}^{+100} \frac{1+z^2}{1+t^2} dt}$$

One may check that u_+ is strictly decreasing on $[-\frac{1}{2}, \frac{1}{2}]$.

$\square$

5.2 The importance of an upper bound on $\overline{\theta}$

Theorem 1 states that under Assumption 1, full disclosure of the anonymized information is optimal provided that $\overline{\theta}$ is not too large. Below, we provide an example illustrating the importance of this upper bound on $\overline{\theta}$ for the result.

Example 3. Assume $N = 3$; $\varepsilon_i \sim U(-\overline{\varepsilon}, \overline{\varepsilon})$ for some $\overline{\varepsilon} > 0$; $\theta = (\overline{\theta}, \overline{\theta}, -\frac{1}{3}\overline{\theta})$ or any of its permutations with equal probability $\frac{1}{6}$ ('good' states of the world), and $\theta = (-\overline{\theta}, -\overline{\theta}, \frac{1}{3}\overline{\theta})$ or any of its permutations with equal probability $\frac{1}{6}$ ('bad' states of the world), where $\overline{\theta} = \frac{9}{10}\overline{\varepsilon}$.

From Claim 1, in the uniform case, $\min\{\hat{m}, \tilde{m}\} = \hat{m} = m^* = \frac{1}{N}\overline{\varepsilon}$. Therefore in this example, $\overline{\theta} = \frac{9}{10}\overline{\varepsilon} > \min\{\hat{m}, \tilde{m}\} = m^* = \frac{1}{3}\overline{\varepsilon}$. In particular, Condition 1 is not satisfied. The expected welfare when voters gets (x_1, x_2, x_3) and k voters exactly vote in favor of the reform is $\frac{1}{2}\sum_i x_i + \frac{1}{2}(2k - N)\overline{\varepsilon}$. In the good states of the world $(\sum_i x_i = 3\overline{\varepsilon}/2)$, the majority rule is 'too restrictive', as the expected welfare when one voter exactly approves the reform is positive $(3\overline{\varepsilon}/4 - \overline{\varepsilon}/2 = \overline{\varepsilon}/4)$ and still the reform is rejected. Symmetrically, in the bad states of the world, the majority rule is 'too lenient', as the expected welfare when two voters exactly approve the reform is negative.

Under the anonymous disclosure policy, $x_i = \frac{1}{3}\sum_j \theta_j$ for all i, that is $x_i = \overline{\varepsilon}/2$ in good states and $-\overline{\varepsilon}/2$ in bad states. Therefore expected welfare is $\frac{1}{2}W(\overline{\varepsilon}/2, \overline{\varepsilon}/2, \overline{\varepsilon}/2) + \frac{1}{2}W(-\overline{\varepsilon}/2, -\overline{\varepsilon}/2, -\overline{\varepsilon}/2)$, while under full disclosure,

expected welfare is: $\frac{1}{2}W(\bar{\theta}, \bar{\theta}, -\frac{1}{3}\bar{\theta}) + \frac{1}{2}W(-\bar{\theta}, -\bar{\theta}, \frac{1}{3}\bar{\theta})$, where $\bar{\theta} = \frac{9}{10}\bar{\varepsilon}$. One may check that the difference between the two is

$$\frac{1}{2}\left[0.421875 - 0.315875\right]*(9\bar{\varepsilon}/4)$$

$$+\frac{1}{2}\left[0.421875 - 0.619875\right]*(5\bar{\varepsilon}/4))$$

$$+\frac{1}{2}\left[0.015625 - 0.001625\right]*(3\bar{\varepsilon}/4)$$

$$+\frac{1}{2}\left[0.140625 - 0.062625\right]*(-\bar{\varepsilon}/4)$$

$$=(0.726563 - 0.735562)\bar{\varepsilon} < 0$$

The first line pertains to situations where the state of world is good and the reform is unanimously approved. In that case, the conditional expected welfare is $9\bar{\varepsilon}/4$, the probability of approval is 0.421875 under anonymous disclosure and 0.315875 under full disclosure. Similarly, the second line pertains to situations where the state of world is good and there is a one-vote margin in favor of the reform. The third and fourth lines pertain to situations where the state of the world is bad, and there is respectively unanimity and a one-vote margin in favor of the reform.

In this example, welfare is strictly higher with full disclosure. The intuition is the following. Condition 1 is violated and the simple majority voting does not select the welfare-maximizing policy in any occasion. In particular, it is "too lenient" in the bad states, as when exactly two voters are in favor of the reform, welfare (conditional on the profile of votes) is negative $(-\bar{\varepsilon}/4)$, and yet the reform is approved. Note that under both disclosure policies, the probability of approving the reform is the same (1/2). Compared to full disclosure, the anomymized policy increases the probability that voters unanimously approve the reform (which is good), but also increases the probability of a one-vote victory when the state of the world is bad (which is bad). Under the assumptions made on the parameters, the latter effect dominates, and full disclosure outperforms anonymous disclosure. $\square$

5.3 A priori heterogeneous voters

So far, we have considered an environment in which voters are a priori symmetric: F is N-exchangeable and all voters have the same g. Studying in full generality the case of heterogeneous voters is beyond the scope of this paper. Still, we provide in this section an example showing that some of the results might remain valid, although for different reasons.

With symmetric voters, as discussed in section 4.2, a key observation is that, whenever $\frac{p'(x)}{p(x)}$ is decreasing in x, moving two voters closer one to the other in terms of beliefs increases the probability that they jointly approve the reform, which, together with Condition 1 ($\bar{\theta} < m^*$), guarantees that the probability effect is positive.

This is no longer the case if voters have different distributions of idiosyncratic shocks. Intuition suggests that to increase the probability of joint approval the planner would like to "compensate" the voter with a priori bad idiosyncratic shocks by giving him better news about θ. If Bayesian plausibility prevents this, no disclosure might still be optimal.

We provide a simple example illustrating this point.

Example 4. Assume two groups of voter. One group ($t = p$) is a priori in favor of the reform: the idiosyncratic shock ϵ_i is uniformly distributed on $[a_p - \bar{\epsilon}, a_p + \bar{\epsilon}]$, with $a_p \geq 0$. The other group ($t = n$) is a priori against the reform: the idiosyncratic shock ϵ_i is uniformly distributed on $[a_n - \bar{\epsilon}, a_n + \bar{\epsilon}]$, with $a_n \leq 0$. The identity of voters within each group are known to the Planner, with N_t voters in group t ($N_p + N_n = N$).

We keep the assumption that in expectation, the sum of the idiosyncratic shocks in the electorate is zero, that is, $\sum_t N_t a_t = 0$.

As for θ, we assume that there are two equiprobable states of the world: one state of the world is good news for group-p voters, while the other is good news for group-n voters. Specifically, in the former state, voters in group p get $\theta_i = b_p > 0$ while voters in group n get $\theta_i = -b_n < 0$. In the latter state, voters in group p get $\theta_i = -b_p < 0$ while voters in group n get

26

$\theta_i = b_n > 0$. We assume that $N_p b_p = N_n b_n$. These assumptions imply that for each voter, absent any additional information, $\mathbb{E}[\theta_i] = 0$, and that in both states of the world, $\sum_i \theta_i = 0$.

Last, we assume that for both groups: $|a_t| + b_t < \bar{\epsilon}$. This assumption implies that whatever his beliefs about θ_i, voter i might still vote for or against the reform depending on the realization of his idiosyncratic shock.[7]

If beliefs are x such that $\sum_i x_i = 0$, one may check that the expected welfare if exactly k voters vote in favor of the reform is $\left(k - \frac{N}{2}\right)\bar{\epsilon}$, independently of the identity of these k voters. Therefore, the simple majority rule always selects the alternative (reform or status quo) with the highest conditional welfare.

As an example, consider the case $N = 3; N_p = 1; N_n = 2$. We show that in that case, no disclosure is the optimal disclosure policy.

We follow the proof in Kamenica and Gentzkow (2011). If the belief that the state is $(b_p, -b_n, -b_n)$ is μ (good news for group p), one may check that the expected welfare is:

$$V(\mu) = \left(\frac{1}{2} + y_p\right)\left(\frac{1}{2} + y_n\right)^2 * \left(\frac{3}{2}\bar{\epsilon}\right)$$
$$+ \left[\left(\frac{1}{2} - y_p\right)\left(\frac{1}{2} + y_n\right)^2 + 2\left(\frac{1}{2} + y_p\right)\left(\frac{1}{2} + y_n\right)\left(\frac{1}{2} - y_n\right)\right] * \left(\frac{1}{2}\bar{\epsilon}\right)$$

where $y_p = \frac{a_p + (2\mu - 1)b_p}{2\bar{\epsilon}}$ and $y_n = \frac{a_n + (1 - 2\mu)b_n}{2\bar{\epsilon}}$. The first line in the expression of $V(\mu)$ corresponds to events where all three voters unanimously vote in favor of the reform, the second line to events where two voters exactly are in favor.

Rearranging terms and using the property that $y_p + 2y_n = 0$, one gets:

$$V(\mu) = \left(\frac{1}{4} - y_n^2\right) * \left(\frac{3}{2}\bar{\epsilon}\right),$$

Therefore $\frac{dV}{d\mu}(\mu) = 3b_n y_n$; $\frac{d^2V}{d\mu^2}(\mu) = -\frac{3b_n^2}{\bar{\epsilon}} < 0$, which shows that V is strictly concave in μ: no disclosure is uniquely optimal.

---

[7]Indeed, it implies that $[-b_t, b_t] \subset [a_t - \bar{\epsilon}, a_t + \bar{\epsilon}]$ for $t = p, n$.

Notice that $\frac{dV}{d\mu}(\mu) = 0$ iff $y_n = y_p = 0$, that is, $\mu = \frac{1}{2} - \frac{1}{2}\frac{a_p}{b_p}$. When $a_p = a_n = 0$ (all voters are identical regarding the distribution of their idiosyncratic shock), $V(\mu)$ is maximized when $\mu = 1/2$ (no disclosure). If $a_p > 0$, maximizing $V(\mu)$ would require to give voter with group n better news (than $\mu = 1/2$). Intuitively, the planner would like to "compensate" a voter with bad a priori idiosyncratic shocks by giving him better news about his θ_i. Bayesian plausibility prevents this. In this example, the payoff function is concave, and the planner cannot improve on no disclosure. $\qquad\square$

In Example 4, $\sum_i \theta_i = 0$ in both states of the world. Therefore, in that case, anomymized disclosure simply coincides with no disclosure. In Appendix E, we propose a generalization of Example 4, where besides uncertainty about the distributional consequences of the reform, there is also some uncertainty about a common shock affecting all voters in the same way: the total value of the reform (that is, the part controlled by the social planner $\sum_i \theta_i$) can be positive or negative. We show that in that case, disclosure of the anonymized information is still the optimal policy. These examples show that anomymous disclosure might still be optimal with a priori heterogeneous voters. The logic is nevertheless somewhat different. With asymmetric voters, the planner would like to "compensate" the voters with a priori bad idiosyncratic shocks by giving them better news about θ. If Bayesian plausibility prevents this, anonymous disclosure might still be optimal.

5.4 Non-separable idiosyncratic noise

In the main model, each voter i's idiosyncratic type enters additively separably to his payoff function. This implies that the marginal increase in θ_i and that in θ_j are treated equally in the welfare. Thus, disclosing only the (non-weighted) average is a natural way of anonymizing the information.

This assumption would be satisfied, for example, if θ_i refers to i's gain in the monetary term and all the voters have quasilinear payoffs in the monetary term. Even if they do not have quasilinear payoffs, our result holds as long

as θ_i and ε_i are separable: for example, let $v(\theta_i) + \varepsilon_i$ be i's payoff under the reform, where v is a concave function (risk aversion / non-quasilinearity). Then, we can redefine $\hat{\theta}_i = v(\theta_i)$ as the state, and apply the previous analysis. In this case, each i would be informed of the average of $v(\theta_i)$.

Without additive separability, the analysis would be more complicated. Let $u(\theta_i, \varepsilon_i)$ be i's payoff under the reform (and 0 be the status-quo payoff). In this case, what would be the analogue of the anonymized information? For example, imagine a two-voter situation ($i = 1, 2$) with $\theta \in \{(0,0), (1,-1), (-1,1)\}$. If u is sufficiently concave in θ_i ("risk aversion" for the reform), then the social planner would be more willing to implement the reform with $\theta = (0,0)$, while much less so otherwise. In that case, revealing (only) $x = \frac{\sum_i \theta_i}{n}$, which is equivalent to no disclosure, may not be the right idea, and it can be better to reveal "whether $\theta = (0,0)$ or $\theta \in \{(1,-1), (-1,1)\}$". On the other hand, it may still be a good idea not to separate $(1,-1)$ and $(-1,1)$, in order to mitigate the ordinal-cardinal mismatch. As suggested in this example, the right analogue may be to reveal the *histogram* of θ (i.e., "how many would enjoy $\theta_i = x$, y, etc.", without revealing "who would enjoy those x, y, etc.).

The precise conditions under which such a histogram-revealing policy is optimal would look different from those in the previous sections, but they would be based on the similar intuition about the probability effect and information effect. For example, imagine two voters again, and consider unanimity as a voting mechanism. The expected social welfare is given by:

$$p_+^1(q)p_+^2(q)(u_+^1(q) + u_+^2(q)),$$

where $q \in \Delta(\Theta)$ is the posterior for θ given the disclosed information by the planner; $p_+^i(q)$ is the probability that i votes for the reform:

$$p_+^i(q) = 1 - G(\varepsilon_i(q)),$$

where $\varepsilon_i(q)$ is the threshold type given q:

$$\mathbb{E}_{\theta_i \sim q}[u(\theta_i, \varepsilon_i(q))] = 0;$$

29

and $u_+^i(q)$ is i's expected payoff conditional on his voting for the reform:

$$\frac{1}{p_+^i(q)} \int_{\varepsilon_i \geq \varepsilon_i(q)} \mathbb{E}_{\theta_i \sim q}[u(\theta_i, \varepsilon_i)]dG.$$

Note that, if the planner only discloses the histogram of θ, then we have $\varepsilon_1(q) = \varepsilon_2(q)$, and hence $p_+^1(q) = p_+^2(q)$ and $u_+^1(q) = u_+^2(q)$. As in the previous sections, equalizing the voters' threshold values of ε is the key effect of anonymization. We thus conjecture that, under appropriate conditions, a similar logic based on probability effects and information effects would imply the conditions under which disclosing the histogram is optimal.

5.5 More complex mechanisms

So far, we have considered simple mechanisms ("experiments") which publicly disclose information independent of the voters' idiosyncratic preferences. Could the social planner do better by using more complex mechanisms? For example, can the social planner do better by using a mechanism that first elicits some information about the voters' idiosyncratic preferences, and then tailors information disclosure to voters' report? We provide in Appendix F a simple example showing that this can indeed be the case. Finding the optimal information design in this broader class of mechanisms for general distribution of θ and ϵ_i is beyond the scope of this paper. Similarly, we only consider here public disclosure of information ; it remains an open question to know if the planner could improve welfare by *privately* disclosing information to voters.

Although the class of mechanism we consider in this paper is admittedly limited, we believe it is an important benchmark case when looking for some normative guidance about the optimal extent of transparency in a democracy. Besides, this class of mechanism has some desirable properties in terms of implementation. In Yamashita and Van der Straeten [2022], we consider a general information design problem with multiple agents, where each agent has initial private information, orthogonal to the principal's information.

We ask which distribution jointly over the state and action profile is implementable *in an ex post dominant-strategy equilibrium*, i.e., an equilibrium such that each agent, given the signal from the principal and his own private information, finds an optimal action regardless of the other agents' information and play. We show that any ex post dominant-strategy implementable outcome is implementable by a mechanism where the principal only provides a public signal about her information.

6 Conclusion

In this paper, we show that, under certain conditions, fully disclosing the anonymized information is optimal in terms of utilitarian expected welfare. Two key conditions are identified: the first imposes the monotonicity of each voter's expected payoff conditional on voting for or against the reform; the second imposes some constrains on how big the uncertainty regarding the information controlled by the social planner can be compared to the size of the idiosyncratic preferences.

References

Ricardo Alonso and Odilon Camara. Persuading voters. *American Economic Review*, 106(11):3590–3605, November 2016.

David Austen-Smith and Jeffrey S. Banks. Information aggregation, rationality, and the condorcet jury theorem. *The American Political Science Review*, 90(1):34–45, 1996.

David Austen-Smith and Timothy J. Feddersen. Deliberation, preference uncertainty, and voting rules. *The American Political Science Review*, 100 (2):209–217, 2006.

Yaron Azrieli and Semin Kim. Pareto efficiency and weighted majority rules. *International Economic Review*, 55(4):1067–1088, 2014.

Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. *Economic Theory*, 26(2):445–469, 2005.

Salvador Barberà and Matthew O. Jackson. On the weights of nations: Assigning voting weights in a heterogeneous union. *Journal of Political Economy*, 114(2):317–339, 2006.

Arjada Bardhi and Yingni Guo. Modes of persuasion toward unanimous consent. *Theoretical Economics*, 13(3):1111–1149, 2018.

Bernard Caillaud and Jean Tirole. Consensus building: How to persuade a group. *American Economic Review*, 97(5):1877–1900, December 2007.

Eddie Dekel and Michele Piccione. Sequential voting procedures in symmetric binary elections. *Journal of Political Economy*, 108(1):34–55, 2000.

Piotr Dworczak and Anton Kolotilin. The persuasion duality, 2023. mimeo.

Timothy Feddersen and Wolfgang Pesendorfer. Voting behavior and information aggregation in elections with private information. *Econometrica*, 65(5):1029–1058, 1997.

Timothy Feddersen and Wolfgang Pesendorfer. Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. *The American Political Science Review*, 92(1):23–35, 1998.

Timothy J. Feddersen and Wolfgang Pesendorfer. The swing voter's curse. *The American Economic Review*, 86(3):408–424, 1996.

Timothy J. Feddersen and Wolfgang Pesendorfer. Abstention in elections with asymmetric information and diverse preferences. *The American Political Science Review*, 93(2):381–398, 1999.

Raquel Fernandez and Dani Rodrik. Resistance against reform: Status quo bias in the presence of individual speci ? c uncertainty. *American Economic Review*, 81:1146–55, 02 1991.

Dino Gerardi and Leeat Yariv. Deliberative voting. *Journal of Economic Theory*, 134(1):317–338, 2007.

Dino Gerardi and Leeat Yariv. Information acquisition in committees. *Games and Economic Behavior*, 62(2):436–459, 2008.

Hans Gersbach. Allocation of information by majority decisions. *Journal of Public Economics*, 48(2):259–268, 1992. ISSN 0047-2727.

Hans Gersbach. Information efficiency and majority decisions. *Social Choice and Welfare*, 12(4):363–370, 1995.

Alex Gershkov and Balázs Szentes. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1):36–68, 2009.

Alex Gershkov, Benny Moldovanu, and XianWen Shi. Optimal voting rules. *The Review of Economic Studies*, 84(2 (299)):688–717, 2017.

Matthew Jackson and Xu Tan. Deliberation, disclosure of information, and voting. *Journal of Economic Theory*, 148, 06 2012.

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

Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964, 2017.

Semyon Malamud and Andreas Schrimpf. Persuasion by dimension reduction, 2021.

Nicola Persico. Committee design with endogenous information. *The Review of Economic Studies*, 71(1):165–191, 2004.

Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, 2010.

Keith E. Schnakenberg. Expert advice to a voting body. *Journal of Economic Theory*, 160:102–113, 2015.

Junze Sun, Arthur Schram, and Randolph Sloof. Public persuasion in elections: Single-crossing property and the optimality of censorship, 2021. mimeo.

Takuro Yamashita and Karine Van der Straeten. Dominant-strategy information design, 2022. mimeo.

APPENDIX

A   Proof of Proposition 1

The proof proceeds in three steps. First, we show that if Assumption 1 holds, then $\hat{m}$ defined in (6) is positive. Second, we show that if Assumption 1 holds and $\bar{\theta} < \hat{m}$, then for all $x \in [-\bar{\theta}, \bar{\theta}]^N$: $x_1 < x_2 \Rightarrow H_1(x) > H_2(x)$. Third, we show that if Assumption 1 holds and $\bar{\theta} < \hat{m}$, then for all $x \in [-\bar{\theta}, \bar{\theta}]^N$: $x_1 < x_2 \Rightarrow W_1(x) > W_2(x)$.

Step 1. Assume Assumption 1 holds. Let us show that $\hat{m} > 0$.
By definition (see (6)):

$$\hat{m} := \max \left\{ m \in [0, \bar{\varepsilon}] : \text{ For all } z \in [-m, m], \right.$$

$$\left. \alpha(z) + \frac{N-3}{2} u_+(-m) + \frac{N-1}{2} u_-(-m) \geq 0 \right\}$$

To show that $\hat{m} > 0$, it is sufficient to show that $\alpha(0) + \frac{N-3}{2} u_+(0) + \frac{N-1}{2} u_-(0) > 0$, that is, $\alpha(0) > u_+(0)$.

We now show that for all $z \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, $\alpha(z) > u_+(z)$. Notice that

$$(p_+(z) u_+(z))' = p_+(z).$$

Therfeore, we have

$$\frac{g(z)}{G(z)} = \frac{p'_+(z)}{p_+(z)} = \frac{1 - u'_+(z)}{u_+(z)}$$

and

$$\left(\frac{g(z)}{G(z)} \right)' = -\frac{u''_+(z)}{u_+(z)} - \frac{u'_+(z)(1 - u'_+(z))}{u_+^2(z)} < -\frac{u''_+(z)}{u_+(z)}$$

where the last inequality follows from the observation that, under Assumption 1, $u'_+(z)(1 - u'_+(z)) > 0$. Therefore,

$$\frac{u''_+(z)}{\left(\frac{g(z)}{G(z)} \right)'} > -u_+(z)$$

35

and for all $z \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, $\alpha(z) > u_+(z)$.

Step 2. Assume that Assumption 1 holds. We show that if $\bar{\theta} < \hat{m}$, then for all $x \in [-\bar{\theta}, \bar{\theta}]^N$: $x_1 < x_2 \Rightarrow H_1(x) > H_2(x)$.

The difference $H_1(x) - H_2(x)$ is given by Equation (3). Noticing that $p'_+ u_+ + p_+ u'_+ = p_+$, we have

$$H_1(x) - H_2(x) = p'_+(x_1)p_+(x_2)\Big(u_+(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\Big)$$
$$- p_+(x_1)p'_+(x_2))\Big(u_+(x_1) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\Big)$$

Therefore, $H_1(x) - H_2(x) > 0$ whenever $x_1 < x_2$ if

$$z \to \frac{G(z)}{g(z)}\Big(u_+(z) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]\Big)$$

is increasing in z whatever $\mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]$, that is (as Assumption 1 holds), for all $z \in [-\bar{\theta}, \bar{\theta}]$,

$$\left(\frac{G(z)u_+(z)}{g(z)}\right)' + \left(\frac{G(z)}{g(z)}\right)'\left(\frac{N-3}{2}u_+(-\bar{\theta}) + \frac{N-1}{2}u_-(-\bar{\theta})\right) > 0$$

or equivalently,

$$u''_+(z) + \left(\frac{g(z)}{G(z)}\right)'\left(2u_+(z) + \frac{N-3}{2}u_+(-\bar{\theta}) + \frac{N-1}{2}u_-(-\bar{\theta})\right) < 0 \quad (10)$$

Indeed,

$$\left(\frac{G(z)u_+(z)}{g(z)}\right)' = \left(\left(\frac{G(z)}{g(z)}\right)^2\left(\frac{g(z)u_+(z)}{G(z)}\right)\right)'$$
$$= 2u_+(z)\left(\frac{G(z)}{g(z)}\right)' + \left(\frac{G(z)}{g(z)}\right)^2\left(\frac{g(z)u_+(z)}{G(z)}\right)'$$
$$= 2u_+(z)\left(\frac{G(z)}{g(z)}\right)' - \left(\frac{G(z)}{g(z)}\right)^2 u''_+(z)$$

where the last line follows from the observation that $\frac{p'_+}{p_+}u_+ + u'_+ = 1$.

36

It follows from the definition of $\hat{m}$ that, if $\bar{\theta} < \hat{m}$ then for all $z \in [-\bar{\theta}, \bar{\theta}]$ and all $\mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2]$, Condition (10) is satisfied. Therefore, if $x_1 < x_2$, $H_1 > H_2$.

Step 3. To conclude the proof of Proposition 1, denote

$$K(x) = -p_-(x_1)p_-(x_2)\Big(u_-(x_1) + u_-(x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 1, (x_3, ..., x_N)]\Big)$$

Because of the symmetry of g, $p_-(z) = p_+(-z)$ and $u_-(z) = -u_+(-z)$. Therefore for all x,

$$K(x)$$
$$= p_+(x_1)p_+(x_2)\Big(u_+(-x_1) + u_+(-x_2) + \mathbb{E}[U_{-1,2}|M_{-1,2} = Q - 2, (-x_3, ..., -x_N)]\Big)$$
$$= H(-x)$$

and $K_1(x) - K_2(x) = -H_1(-x) + H_2(-x)$.

Take some x such that $x_1 < x_2$. Then $-x_1 > -x_2$ and the property of H that was established in Step 2 of the proof implies that $H_1(-x) < H_2(-x)$. Which completes the proof of the proposition.

B   Proof of Theorem 1

Assume Assumption 1, and $\bar{\theta} < \min\{\hat{m}, \tilde{m}\}$.

We prove Theorem 1 in a number of steps.

First, we consider an auxiliary environment where all the voters *know* $\frac{\sum_{i=1}^{N} \theta_i}{N}$ (but not individual θ_i). Applying Proposition 1, we show that disclosing no additional information is optimal in this hypothetical environment (Lemma 2, see below). Hence, for any x with $\frac{\sum_{i=1}^{N} x_i}{N} = \frac{\sum_{i=1}^{N} \theta_i}{N}$, we have:

$$W^* \left(\frac{\sum_{i=1}^{N} \theta_i}{N} \right) \geq W(x)$$

where $W(x)$ denotes the expected total welfare if the voters are informed of x and $W^*(z) = W(z, ..., z)$.

Applying Proposition 2, $W^*(m)$ is convex for $m \in [-\bar{\theta}, \bar{\theta}]$.

For $x \in [-\bar{\theta}, \bar{\theta}]$, let

$$\hat{W}(x) := W^* \left(\frac{\sum_{i=1}^N x_i}{N} \right).$$

Then, $\hat{W}(x) \geq W(x)$ for all $x \in [-\bar{\theta}, \bar{\theta}]^N$, and moreover, it is convex. Indeed, for all $x, x' \in [-\bar{\theta}, \bar{\theta}]^N$, $\lambda \in [0, 1]$:

$$\begin{aligned}
\hat{W}(\lambda x + (1 - \lambda x')) &= W^* \left(\frac{\sum_{i=1}^N (\lambda x_i + (1 - \lambda) x_i'))}{N} \right) \\
&= W^* \left(\lambda \frac{\sum_{i=1}^N x_i}{N} + (1 - \lambda) \frac{\sum_{i=1}^N x_i'}{N} \right) \\
&\leq \lambda W^* \left(\frac{\sum_{i=1}^N x_i}{N} \right) + (1 - \lambda) W^* \left(\frac{\sum_{i=1}^N x_i'}{N} \right) \\
&= \lambda \hat{W}(x) + (1 - \lambda x') \hat{W}(x')
\end{aligned}$$

where the inequality on the third line follows from the convexity of W^*.

The planner's problem is

$$\max_{\xi} \int_x W(x) d\xi,$$

which is not larger than

$$\max_{\xi} \int_x \hat{W}(x) d\xi.$$

Note that:

$$\max_{\xi} \int_x \hat{W}(x) d\xi = \int_\theta \hat{W}(\theta) dF,$$

because $\hat{W}$ is convex. However, $\hat{W}(\theta)$ is precisely the total welfare for each θ under the policy of fully disclosing the anonymized information. Therefore, it must be an optimal disclosure policy.

Statement and proof of Lemma 2

Consider an auxiliary environment where all the voters know $m = \frac{\sum_{i=1}^{N} \theta_i}{N}$ (but not individual θ_i). Applying Proposition 1, we show that disclosing no additional information is optimal.

Lemma 2. Assume Assumption 1. Consider an auxiliary environment where all the voters know $m = \frac{\sum_{i=1}^{N} \theta_i}{N}$ (but not individual θ_i). If $\overline{\theta} < \hat{m}$, then it is optimal to disclose no additional information.

Proof. Fix any $m \in [-\overline{\theta}, \overline{\theta}]$. Let $\Theta^m = \{\theta | \frac{\sum_{i=1}^{N} \theta_i}{N} = m\}$. Let $\xi^m \in \Delta(\Theta^m)$ be any disclosure policy (more formally, its equivalent, as discussed in Section 2). Letting $W(x)$ denote the expected total welfare if the voters are informed of $x = (x_i)_{i=1}^{N}$, our problem is:

$$\max_{\xi^m} \int_x W(x) d\xi^m.$$

If no additional information is disclosed, its induced expected total welfare is:

$$W(m, \ldots, m).$$

Thus, it suffices to show:

$$W(m, \ldots, m) \geq \int_x W(x) d\xi^m.$$

Fix any policy ξ^m with $\Pr(x_i \neq x_j, \ \exists i, j) > 0$. Fix $\lambda \in (0, 1)$, and define an alternative disclosure policy, $\tilde{\xi}^m$, by:

$$\tilde{\xi}^m(A) = \xi^m(\{x' | \exists x \in A; x' = (1 - \lambda)x + \lambda(m, \ldots, m)\}),$$

for each (measurable) $A \subseteq [-\overline{\theta}, \overline{\theta}]$. $\tilde{\xi}^m$ is less convex than ξ^m in the sense of the convex stochastic ordering, and hence, is feasible. Equivalently, $\tilde{\xi}^m$ can be described as a probability distribution such that, first, x is drawn according to ξ^m, and then the actual value realized is given by $x' = (1 - \lambda)x + \lambda(m, \ldots, m)$.

This alternative description is useful in that the expected total welfare given $\tilde{\xi}^m$ is given by:

$$\int_x W((1-\lambda)x + \lambda(m,\ldots,m))d\xi^m,$$

and to show that $\tilde{\xi}^m$ is an improvement over ξ^m, it suffices to show that

$$W((1-\lambda)x + \lambda(m,\ldots,m)) > W(x)$$

whenever $x_i \neq x_j$ for some i, j.[8]

Hence, from here on, we fix an arbitrary x such that $x \neq (m,\ldots,m)$. For small λ, we have

$$W((1-\lambda)x + \lambda(m,\ldots,m)) - W(x) \simeq \lambda \sum_{i=1}^N (m - x_i)W_i(x),$$

where $W_i(x) = \frac{\partial W}{\partial x_i}(x)$ for all $i = 1,\ldots,N$. Given that $m = \frac{1}{N}\sum_{j=1}^N x_j$,

$$
\begin{aligned}
\sum_{i=1}^N (m - x_i)W_i(x) &= \frac{1}{N}\sum_{i=1}^N \sum_{j=1}^N (x_j - x_i)W_i(x) \\
&= \frac{1}{N}\sum_{i=1}^N \sum_{j=i}^N (x_j - x_i)(W_i(x) - W_j(x))
\end{aligned}
$$

Thus, it suffices to show that $W_i(x) - W_j(x) > 0$ if $x_j > x_i$, for any given i, j and x. This is precisely the Condition obtained in Proposition 1. $\qquad\square$

C Proof of Claim 1

C.1 Uniform distribution

Assume $\varepsilon_i \sim U\left(-\bar{\varepsilon}, \bar{\varepsilon}\right)$.

---

[8]More precisely, for each $n \in \mathbb{N}$, let

$$X_n = \{x \mid W((1-\lambda)x + \lambda(m,\ldots,m)) - W(x) > \frac{1}{n}\},$$

and let $X_\infty = \bigcup_{n\in\mathbb{N}} X_n$. Then $\xi^m(X_\infty) > 0$ by assumption, which implies $\xi^m(X_n) > 0$ for some n by continuity of a measure with respect to countable set operations.

G is linear, which implies that Assumption 1 is satisfied.

In the uniform case:

$$u_+(z) = \frac{z + \overline{\varepsilon}}{2} \ ; \ u_-(z) = \frac{z - \overline{\varepsilon}}{2}.$$

Determination of m^*. The definition of m^* is given in (4). When Assumption 1 holds, u_+ and u_- are increasing, therefore

$$m^* = \max\left\{m \in [0, \overline{\varepsilon}] : w(-m) \geq 0\right\}.$$

where w is defined in (7). As $w(z) = \frac{Nz + \overline{\varepsilon}}{2}$, we have $m^* = \frac{1}{N}\overline{\varepsilon}$.

Determination of $\hat{m}$. The definition of $\hat{m}$ is given in (6). In the uniform case, $u'' = 0$, therefore $\alpha(z) = 2u_+(z)$, which is increasing in z, therefore

$$\hat{m} = \max\left\{m \in [0, \overline{\varepsilon}] : \ 2u_+(-m) + \frac{N-3}{2} * u_+(-m) + \frac{N-1}{2} * u_-(-m) \geq 0\right\},$$

which is precisely the definition of m^*.

Determination of $\tilde{m}$. The definition of $\tilde{m}$ is given in (8). In the uniform case, $\frac{G(z)}{g(z)} = z + \overline{\varepsilon} = 2u_+(z)$, therefore

$$\beta(z) = (N-1)w(z) + 2Nu_+(z) = \frac{N(N+1)z + (3N-1)\overline{\varepsilon}}{2},$$

which is increasing in z, therefore $\tilde{m} = \frac{3N-1}{N(N+1)}\overline{\varepsilon}$. Notice $\tilde{m} \in \left(\frac{1}{N}\overline{\varepsilon}, \overline{\varepsilon}\right)$.

C.2 Normal distribution

Assume $\varepsilon_i \sim \mathcal{N}(0, \sigma)$.

The density g is logconcave, which implies the log-concavity of G (see Bagnoli and Bergstrom [2005]). Therefore Assumption 1 is satisfied.

We first consider the case $\sigma = 1$.

Determination of m^*. When Assumption 1 holds, u_+ and u_- are increasing, therefore m^* is the unique solution of the equation $w(-m) = 0$. One may check numerically that, for $N = 3$, $m^* \approx 0.767847$.

Determination of $\hat{m}$. When $\varepsilon_i \sim \mathcal{N}(0,1)$:

$$
\begin{aligned}
u_+(z) &= z + \frac{g(z)}{G(z)} \\
u'_+(z) &= 1 + \Big(\frac{g(z)}{G(z)}\Big)' = 1 - \frac{g(z)}{G(z)}u_+(z) \\
u''_+(z) &= -\Big(\frac{g(z)}{G(z)}\Big)'u_+(z) - \frac{g(z)}{G(z)}u'_+(z)
\end{aligned}
$$

Therefore

$$
\alpha(z) = z + \frac{1}{u_+(z)}
$$

One may check that $\alpha(z)$ is increasing in z on $\mathbb{R}$, therefore $\hat{m}$ is the unique positive solution (if it exists) of $\alpha(-\hat{m}) + u_-(-\hat{m}) = 0$. One finds numerically that $\hat{m} \approx 0.5606$.

Determination of $\tilde{m}$. Recall

$$
\beta(z) = Nw(z) + (2N-1)\Big(\frac{G(z)}{g(z)}\Big)u'_+(z) - \Big(\frac{G(z)}{g(z)}\Big)'\Big(w(z) - u_+(z)\Big)
$$

Using the fact that $\Big(\frac{G}{g}\Big)' = \frac{G}{g}u_+$, we have:

$$
\beta(z) = Nw(z) + (2N-1)\Big(\frac{1}{u_+(z) - z} - u_+(z)\Big) - \frac{u_+(z)(w(z) - u_+(z))}{u_+(z) - z}
$$

One may check numerically that for $N = 3$, $\tilde{m} = 1.05427$.

Consider now the general case $\sigma > 0$. We show in the Appendix D that the critical thresholds m^*, $\hat{m}$ and $\tilde{m}$ are proportional to the standard deviation of the density, which concludes the proof of Claim 1.   □

D An alternative interpretation of Theorem 1

Theorem 1 provides an upper bound for $\overline{\theta}$ given the distribution of each ε_i fixed. An alternative interpretation of this result is to take $\overline{\theta}$ as given, and imagine that g is in a scale-family of another "baseline" density γ:

$$g(\varepsilon_i, \sigma) = \frac{1}{\sigma} \gamma \left(\frac{\varepsilon_i}{\sigma} \right)$$

for a scale parameter $\sigma > 0$. In particular, if this baseline distribution is chosen to be standardized, σ is the standard deviation of g. Corollary 1 states an equivalent statement of Theorem 1, which concerns a lower bound of σ, the scale parameter for the distribution of each ε_i (given $\overline{\theta}$ fixed). To make this statement precise, assume that γ has full-support on $[-\overline{e}, \overline{e}]$, for some $\overline{e} \geq \overline{\theta}$. In that case, g has full-support on $[-\sigma\overline{e}, \sigma\overline{e}]$, and we assume that $\sigma > \frac{\overline{\theta}}{\overline{e}}$.

Denote μ^* , $\hat{\mu}$ and $\tilde{\mu}$ the solution of equations (4), (6) and (8), respectively, when the density is γ, and let

$$\hat{\sigma} = \frac{\overline{\theta}}{\hat{\mu}} \; ; \; \tilde{\sigma} = \frac{\overline{\theta}}{\tilde{\mu}}.$$

Then the following holds:

Corollary 1. Assume γ satisfies Assumption 1. If $\sigma > \max\{\hat{\sigma}, \tilde{\sigma}\}$, then it is an optimal disclosure policy to fully disclose the anonymized information.

Proof. When m^* , $\hat{m}$ and $\tilde{m}$ are the solution of equations (4), (6) and (8), respectively, when the density is $g(., \sigma)$, we prove that:

$$\frac{m^*}{\mu^*} = \frac{\hat{m}}{\hat{\mu}} = \frac{\tilde{m}}{\tilde{\mu}} = \sigma,$$

which will show the result as a direct application of Theorem 1. Denote by Γ the cumulative of the baseline density γ, and let

$$\nu_+(z) := z + \frac{\int_{-z}^{\infty} t\gamma(t)dt}{\Gamma(z)}.$$

Notice:

$$u_+(z) = \sigma * \nu_+\left(\frac{z}{\sigma}\right)$$

$$u'_+(z) = \nu'_+\left(\frac{z}{\sigma}\right)$$

$$u''_+(z) = \frac{1}{\sigma} * \nu''_+\left(\frac{z}{\sigma}\right)$$

$$G(z) = \Gamma\left(\frac{z}{\sigma}\right)$$

which implies

$$\left(\frac{G(z)}{g(z)}\right)' = \left(\frac{\Gamma(z/\sigma)}{\gamma(z/\sigma)}\right)'$$

$$\left(\frac{g(z)}{G(z)}\right)' = \frac{1}{\sigma^2} * \left(\frac{\gamma(z/\sigma)}{\Gamma(z/\sigma)}\right)'$$

The result follows from the definition of m^*, $\hat{m}$ and $\tilde{m}$ in (4), (6) and (8).

$\square$

E   An example with heterogeneous voters

In Example 4, $\sum_i \theta_i = 0$ in both states of the world. Therefore, in that case, anomymized disclosure coincides with no disclosure. We now propose an extension of Example 4, where, besides some uncertainty about the distributional consequences of the reform, there is some uncertainty about a common shock affecting all voters in the same way: the total value of the reform (that is, the part controlled by the social planner, $\sum_i \theta_i$) can be positive or negative. We show that in that case, disclosure of the anonymized information is optimal.

Specifically, we make the following assumptions:

Example 5. As for the ϵ_i: we keep the same assumptions as in Example 4. We still assume that $N = 3; N_p = 1; N_n = 2$. As for θ, we assume that there are four possible states of the world:

$$\theta = \tilde{C}(c, c, c) + \tilde{D}(b_p, -b_n, -b_n)$$

where b_p, b_n have been defined in Example 4, c is a non-negative real number, and $\tilde{C}$ and $\tilde{D}$ are two independent binary variables. Variable $\tilde{C}$, which refers to the common shock, takes the value 1 with probability $\lambda_0 = \frac{1}{2}$ (good news for the whole economy) and value -1 with the remaining probability, while $\tilde{D}$, which refers to the distributional shock, takes the value 1 with probability $\mu_0 = \frac{1}{2}$ (good news for group-p voters) and value -1 with the remaining probability.

Last, we assume that for $t = p, n$:

$$|a_t| + b_t + c < \bar{\epsilon}$$

Notice that if $c = 0$ (no aggregate shock), the setting is exactly that of Example 4.

If the belief that $\tilde{C} = 1$ (good news about the common shock) is λ, and the belief that $\tilde{D} = 1$ (good news for group-p voters) is μ, one may check that the expected welfare if exactly k voters vote in favor of the reform is

$$\frac{N}{2}(2\lambda - 1)c + \left(k - \frac{N}{2}\right)\bar{\epsilon}$$

independently of the identity of these k voters. Note that these quantities are independent of μ (as in Example 4), but depend on λ. The simple majority rule selects the alternative (reform or status quo) with the highest conditional expected welfare on any occasion iff for all $\lambda \in [0, 1]$, $\frac{3}{2}(2\lambda - 1)c + \frac{1}{2}\bar{\epsilon} > 0$ and $\frac{3}{2}(2\lambda - 1)c - \frac{1}{2}\bar{\epsilon} < 0$, that is, $3c < \bar{\epsilon}$. We assume in the sequel that this condition holds.

Expected welfare when beliefs are (λ, μ) is:

$V(\lambda, \mu)$
$$= \left(\frac{1}{2} + z_p\right)\left(\frac{1}{2} + z_n\right)^2 * \left(\frac{3}{2}(2\lambda - 1)c + \frac{3}{2}\bar{\epsilon}\right)$$
$$+ \left[\left(\frac{1}{2} - z_p\right)\left(\frac{1}{2} + z_n\right)^2 + 2\left(\frac{1}{2} + z_p\right)\left(\frac{1}{2} + z_n\right)\left(\frac{1}{2} - z_n\right)\right] * \left(\frac{3}{2}(2\lambda - 1)c + \frac{1}{2}\bar{\epsilon}\right)$$

45

where $z_p = \frac{1}{2\bar{\epsilon}}\left(a_p + (2\mu - 1)b_p + (2\lambda - 1)c\right)$ and $z_n = \frac{1}{2\bar{\epsilon}}\left(a_n + (1 - 2\mu)b_n + (2\lambda - 1)c\right)$.
The first line in the expression of $V(\lambda, \mu)$ corresponds to events where all three voters unanimously vote in favor of the reform, and the second line to events where two voters exactly are in favor.

We show that:

1. For any λ, V is strictly concave in μ ;

2. For any μ, V is strictly convex in λ ;

which implies that full disclosure on $\tilde{C}$ and no disclosure on $\tilde{D}$ is uniquely optimal.

Proof: Rearranging terms in V, one may check that

$$V(\lambda, \mu) = \left(\frac{1}{2} + \frac{1}{2}(z_p + 2z_n) - 2z_p z_n^2\right) * \frac{3}{2}(2\lambda - 1)c$$
$$+ \left(\frac{3}{4} + (z_p + 2z_n) + 2z_p z_n + z_n^2\right) * \frac{\bar{\epsilon}}{2}$$

- We first show that V is strictly concave in μ. Using the property that $z_p + 2z_n = \frac{3(2\lambda - 1)c}{2\bar{\epsilon}}$ and $\frac{\partial z_p}{\partial \mu} + 2\frac{\partial z_n}{\partial \mu} = 0$, we have:

$$\frac{\partial}{\partial \mu}V(\lambda, \mu) = \left[3(2\lambda - 1)c\left(z_p z_n - z_n^2\right) + \frac{\bar{\epsilon}}{2}(z_n - z_p)\right]\frac{\partial z_p}{\partial \mu}$$
$$\frac{\partial^2}{\partial \mu^2}V(\lambda, \mu) = 3\left[(2\lambda - 1)c\left(2z_n - \frac{1}{2}z_p\right) - \frac{1}{4}\bar{\epsilon}\right]\left(\frac{\partial z_p}{\partial \mu}\right)^2$$

It follows that V is strictly concave in μ iff

$$2\left(4z_n - z_p\right)(2\lambda - 1)c - \bar{\epsilon} < 0$$

which is equivallent to

$$\left(\frac{(2\lambda - 1)c}{\bar{\epsilon}}\right)^2 - \frac{a_p + (2\mu - 1)b_p}{\bar{\epsilon}} * \frac{(2\lambda - 1)c}{\bar{\epsilon}} < \frac{1}{3}$$

One may check that this inequality is satisfied for all $(\lambda, \mu) \in [0, 1]^2$ iff $a_p + b_p + c < \frac{\bar{\epsilon}^2}{3c}$. Conditions $a_p + b_p + c < \bar{\epsilon}$ and $3c < \bar{\epsilon}$ together imply that this inequality holds.

- We now show that V is strictly convex in λ. Straightforward computation shows that:

$$\frac{\partial}{\partial \lambda} V(\lambda, \mu) = \left(\frac{3}{2} - 2z_n^2 - 4z_p z_n \right) * \frac{3}{2} (2\lambda - 1) c * \frac{\partial z_n}{\partial \lambda}$$
$$+ \left(3 + \frac{5}{2} (z_p + 2z_n) - 6z_p z_n^2 \right) * \overline{\epsilon} * \frac{\partial z_n}{\partial \lambda}$$

and

$$\frac{\partial^2}{\partial \lambda^2} V(\lambda, \mu)$$
$$= \left(-(z_p + 2z_n)^2 + 3(1 - z_n^2 - 2z_p z_n) * 4\overline{\epsilon} * \left(\frac{\partial z_n}{\partial \lambda} \right)^2 \right.$$
$$= \left(\left(3z_n - \frac{3(2\lambda - 1)c}{2\overline{\epsilon}} \right)^2 + 3 - 2 \left(\frac{3(2\lambda - 1)c}{2\overline{\epsilon}} \right)^2 \right) * 4\overline{\epsilon} * \left(\frac{\partial z_n}{\partial \lambda} \right)^2$$

The condition $c < 3\overline{\epsilon}$ implies that $3 - 2 \left(\frac{3(2\lambda-1)c}{2\overline{\epsilon}} \right)^2 > 0$, which shows that V is strictly convex in λ.

F An example with more complex mechanisms

Our focus in this paper is on simple information disclosure policies, with no prior communication between the planner and the voters (that is, before the planner discloses information). In this section, we exhibit an example showing that the social planner could do better by using a mechanism that first elicits some information about the voters' idiosyncratic preferences, and then tailors information disclosure to voters' reports.

Example 6. Let's consider the setting of Example 1: Assume $N = 3$ voters; $\varepsilon_i \sim U(-\overline{\varepsilon}, \overline{\varepsilon})$; $\theta = (-\overline{\theta}, 0, \overline{\theta})$ or one of its permutation with probability $\frac{1}{6}$. Additionally, assume that $\overline{\varepsilon} > N\overline{\theta}$. Theorem 1 states that within the

47

class of mechanism studied in the paper, the optimal disclosure policy is no disclosure at all.

We show in this section that the social planner can stricty improve welfare by using the following mechanism:

- Stage 1: The planner asks voters to report the sign of their ε_i. Denote $\hat{s}_i = 1$ if voter i reports a positive sign, $\hat{s}_i = 0$ otherwise.

- Stage 2: The planner publicly discloses the number of voters who reported a positive sign and the number of voters who reported a negative sign, as well as $\sum_{i:\hat{s}_i=0} \theta_i$ (when $\{i : \hat{s}_i = 0\}$ is not empty), and $\sum_{i:\hat{s}_i=1} \theta_i$ (when $\{i : \hat{s}_i = 1\}$ is not empty) .

- Stage 3: Vote under the majority rule. The reform is implemented iff at least two voters are in favor.

With no prior elicitation: Theorem 1 states that the optimal disclosure policy is no disclosure at all. The reform is implemented if and only if at least two voters have a positive ε_i.

With prior elicitation as described in the mechanism above: Assume for now that voters report truthfully at the first stage.

- If $\varepsilon_i > 0$ for all i: no additional information is disclosed, and the reform is implemented.

- If $\varepsilon_1 > 0, \varepsilon_2 > 0, \varepsilon_3 < 0$ (or one of its permutation): If $\theta_3 \leq 0$, the news disclosed by the planner "goes in the same direction" as what is learnt from the idiosyncratic shocks, and the reform is always implemented. If $\theta_3 = \bar{\theta}$ (in which case voter 3 learns that $\theta_3 = \bar{\theta}$ while voters 1 and 2 share the same expectation about their θ_i equal to $-\bar{\theta}/2$), the reform is rejected with a positive probability. Specifically, it is rejected if and only if one of the three following situations arises:

(i) all three voters unanimously vote against the reform, in which case the expected welfare of the reform (conditional on the information available at the time of the vote and on the vote profile) is

$$m\left(0, \frac{\bar{\theta}}{2}\right) + m\left(0, \frac{\bar{\theta}}{2}\right) + m\left(-\bar{\varepsilon}, -\bar{\theta}\right) = -\frac{\bar{\varepsilon}}{2} < 0$$

where for $z_1, z_2 \in [-\bar{\varepsilon}, \bar{\varepsilon}]$, $z_1 < z_2$:

$$m(z_1, z_2) := \mathbb{E}\left[\varepsilon_i | \varepsilon_i \in [z_1, z_2]\right]$$

or (ii) Voters 1 and 2 vote against the reform and voter 3 votes in favor, in which case the expected welfare of the reform is:

$$m\left(0, \frac{\bar{\theta}}{2}\right) + m\left(0, \frac{\bar{\theta}}{2}\right) + m\left(-\bar{\theta}, 0\right) = 0$$

or (iii) Voters 1 and 2 split their votes and voter 3 votes against, in which case the expected welfare of the reform is:

$$m\left(\frac{\bar{\theta}}{2}, \bar{\varepsilon}\right) + m\left(0, \frac{\bar{\theta}}{2}\right) + m\left(-\bar{\varepsilon}, -\bar{\theta}\right) = 0$$

Compared to the scenario with no disclosure, the additional information revealed by the planner is strictly welfare improving. In particular, it allows to reject the reform in situations such as case (i) where a priori positive voters (voters 1 and 2 in the example) have only mildly positive shocks, while voter 3 is badly hurt by the reform (very negative ε_3), so that the total expected value of the reform negative. In that case, voters unanimously reject the reform (which would have been approved should no additional information had been disclosed).

- If $\varepsilon_1 > 0, \varepsilon_2 < 0, \varepsilon_3 < 0$ for all i: the reform is not implemented except if $\theta_1 = -\bar{\theta}$ and a majority vote in favor. By symmetry with the case $\varepsilon_1 > 0, \varepsilon_2 > 0, \varepsilon_3 < 0$, the value of the reform if implemented is non-negative.

- If $\varepsilon_i < 0$ for all i: the reform is never implemented (same as with anomymous disclosure)

So far, we have shown that if all voters truthfully report the sign of their ε_i, the mechanism allows to strictly improve welfare compared to the no disclosure policy. Last, let us show that the situation where all voters report truthfully is indeed an equilibrium. Consider voter 1 when all the other two voters report truthfully. Wlog, assume that $\varepsilon_1 > 0$. One may check that the difference between his expected utility if he reports $m_1 = 1$ (reports truthfully) and his expected utility if he reports $m_1 = 0$ is:

$$
= \mathbb{P}\Big[\theta_1 = 0\Big] * \frac{1}{4}\Big[\frac{\bar{\theta}}{2} + \Big(\frac{3\bar{\theta}}{2} - \bar{\theta}^2\Big)\mathbf{1}_{0<\varepsilon_1<\frac{\bar{\theta}}{2}}\Big] * [0 + \varepsilon_1] \tag{11}
$$

$$
+ \mathbb{P}\Big[\theta_1 = \bar{\theta}\Big] * \frac{1}{4}\Big[\bar{\theta} - \frac{\bar{\theta}^2}{4}\Big] * [\bar{\theta} + \varepsilon_1]
$$

$$
+ \mathbb{P}\Big[\theta_1 = -\bar{\theta}\Big] * \frac{1}{4}\Big[\Big(-\frac{\bar{\theta}}{2} + \frac{\bar{\theta}^2}{4}\Big) + \Big(\frac{3\bar{\theta}}{2} - \bar{\theta}^2\Big)\mathbf{1}_{0<\varepsilon_1<\frac{\bar{\theta}}{2}} + \Big(\bar{\theta} - \frac{\bar{\theta}^2}{2}\Big)\mathbf{1}_{\varepsilon_1>\bar{\theta}}\Big] * [-\bar{\theta} + \varepsilon_1]
$$

To understand the intuition behind the expression above, consider for example the case where $\bar{\theta}_1 = 0$ (which corresponds to the first line in the expression above):

- If $\varepsilon_2 > 0, \varepsilon_3 > 0$: whatever his report, no additional information is disclosed, voters 2 and 3 vote in favor, and the reform is implemented no matter what 1 votes.

- If $\varepsilon_2 < 0, \varepsilon_3 < 0$: whatever his report, no additional information is disclosed, voters 2 and 3 vote against, and the reform is not implemented no matter what 1 votes.

- If ε_2 and ε_3 have different signs: voter 1's report impacts what he learns and what the other two voters learn and votes.

  – Consider first the case where ε_2 and ε_3 have different signs, and for $i = 2, 3$, θ_i and ε_i have the same sign (which happens with probability 1/4). Then no matter what 1 reports, the posterior of

50

voters 2 and 3 will have the same sign as their ε_i, so that they will vote in opposite directions and voter 1 is pivotal. In that case, if he reports truthfully, he will share his beliefs with the voter having the positive ε_i, and therefore he will vote in favor of the reform (as by assumption $\varepsilon_1 > 0$), which will be implemented. If he does not report truthfully, he will share his beliefs with the voter having the negative ε_i, and therefore his posterior will be $-\frac{\bar{\theta}}{2}$ and he will vote in favor of the reform iff $\varepsilon_1 > \frac{\bar{\theta}}{2}$.

– Consider now the case where where ε_2 and ε_3 have different signs, for $i = 2, 3$, θ_i and ε_i have opposite signs (which happens with probability $1/4$). For example, consider the case $\theta_2 < 0 < \varepsilon_2$. If voter 1 report truthfully, he will share his beliefs with voter 2, and the posterior beliefs are $\left(-\frac{\bar{\theta}}{2}, -\frac{\bar{\theta}}{2}, \bar{\theta} \right)$, while if voter 1 lies, he will share his beliefs with voter 3, and the posterior beliefs are $\left(\frac{\bar{\theta}}{2}, -\bar{\theta}, \frac{\bar{\theta}}{2} \right)$. In the former case, the reform is implemented with probability

$$
\mathbb{P}\left[\varepsilon_2 > \frac{\bar{\theta}}{2}, \varepsilon_3 > -\bar{\theta} \middle| \varepsilon_3 < 0 < \varepsilon_2 \right]
$$
$$
+ \left(\mathbb{P}\left[\varepsilon_2 < \frac{\bar{\theta}}{2}, \varepsilon_3 > -\bar{\theta} \middle| \varepsilon_3 < 0 < \varepsilon_2 \right] + \mathbb{P}\left[\varepsilon_2 > \frac{\bar{\theta}}{2}, \varepsilon_3 < -\bar{\theta} \middle| \varepsilon_3 < 0 < \varepsilon_2 \right] \right) * \mathbf{1}_{\varepsilon_1 > \frac{\bar{\theta}}{2}}
$$

while in the latter, it is implemented with probability

$$
\mathbb{P}\left[\varepsilon_2 > \bar{\theta}, \varepsilon_3 - \frac{\bar{\theta}}{2} \middle| \varepsilon_3 < 0 < \varepsilon_2 \right]
$$
$$
+ \mathbb{P}\left[\varepsilon_2 < \bar{\theta}, \varepsilon_3 > -\frac{\bar{\theta}}{2} \middle| \varepsilon_3 < 0 < \varepsilon_2 \right] + \mathbb{P}\left[\varepsilon_2 > \bar{\theta}, \varepsilon_3 < -\frac{\bar{\theta}}{2} \middle| \varepsilon_3 < 0 < \varepsilon_2 \right]
$$

The difference between the former and the latter probability is

$$
\left[\left(1 - \frac{\bar{\theta}}{2} \right)\bar{\theta} - (1 - \bar{\theta})\frac{\bar{\theta}}{2} \right] + \left[\frac{\bar{\theta}^2}{2} + \left(1 - \frac{\bar{\theta}}{2} \right)(1 - \bar{\theta}) \right] * \left[\mathbf{1}_{\varepsilon_1 > \frac{\bar{\theta}}{2}} - 1 \right]
$$
$$
= \frac{\bar{\theta}}{2} + \left[-1 + \frac{3\bar{\theta}}{2} - \bar{\theta}^2 \right] * \mathbf{1}_{\varepsilon_1 < \frac{\bar{\theta}}{2}}
$$

Collecting terms, conditional on $\theta_1 = 0$, the difference in the probability of the reform being implemented if voter 1 reports truthfully and if he lies is

$$\frac{1}{4} * \mathbf{1}_{\varepsilon_1 < \frac{\bar{\theta}}{2}} + \frac{1}{4} * \left[\frac{\bar{\theta}}{2} + \left(-1 + \frac{3\bar{\theta}}{2} - \bar{\theta}^2\right) * \mathbf{1}_{\varepsilon_1 < \frac{\bar{\theta}}{2}}\right] = \frac{1}{4} * \left[\frac{\bar{\theta}}{2} + \left(\frac{3\bar{\theta}}{2} - \bar{\theta}^2\right) * \mathbf{1}_{\varepsilon_1 < \frac{\bar{\theta}}{2}}\right],$$

which one can recognize on the first line of Expression (11). The algebra is similar when $\theta_1 = \bar{\theta}$ and $\theta_1 = -\bar{\theta}$ (leading to the second and third line of (11), respectively).

Using Expression (11), one can now check that truthfully reporting the sign of ε_1 is a best response if other voters are truthful. Indeed:

- If $\varepsilon_1 > \bar{\theta}$, Expression (11) gives

$$\frac{1}{12} * \left[\frac{\bar{\theta}}{2}\right] * [\varepsilon_1] + \frac{1}{12} * \left[\bar{\theta} - \frac{\bar{\theta}^2}{4}\right] * [\bar{\theta} + \varepsilon_1] + \frac{1}{12} * \left[\frac{\bar{\theta}}{2} - \frac{\bar{\theta}^2}{4}\right] * [-\bar{\theta} + e_1]$$

$$= \frac{1}{12} * \left[2\bar{\theta} - \frac{\bar{\theta}^2}{2}\right] * [\varepsilon_1] + \frac{1}{12} * \left[\frac{\bar{\theta}}{2}\right] * [\bar{\theta}]$$

  which is strictly positve

- If $\frac{\bar{\theta}}{2} < \varepsilon_1 < \bar{\theta}$, the difference is equal to

$$\frac{1}{12} * \left[\frac{\bar{\theta}}{2}\right] * [\varepsilon_1] + \frac{1}{12} * \left[\bar{\theta} - \frac{\bar{\theta}^2}{4}\right] * [\bar{\theta} + \varepsilon_1] + \frac{1}{12} * \left[-\frac{\bar{\theta}}{2} + \frac{\bar{\theta}^2}{4}\right] * [-\bar{\theta} + \varepsilon_1]$$

$$= \frac{1}{12} * \left[\bar{\theta}\right] * [\varepsilon_1] + \frac{1}{12} * \left[\frac{3\bar{\theta}}{2} - \frac{\bar{\theta}^2}{2}\right] * [\bar{\theta}]$$

  which is strictly positve

- If $0 < \varepsilon_1 < \frac{\bar{\theta}}{2}$, the difference is equal to

$$\frac{1}{12} * \left[2\bar{\theta} - \bar{\theta}^2\right] * [\varepsilon_1] + \frac{1}{12} * \left[\bar{\theta} - \frac{\bar{\theta}^2}{4}\right] * [\bar{\theta} + \varepsilon_1] + \frac{1}{12} * \left[\bar{\theta} - \frac{3\bar{\theta}^2}{4}\right] * [-\bar{\theta} + \varepsilon_1]$$

$$= \frac{1}{12} * \left[4\bar{\theta} - 2\bar{\theta}^2\right] * [\varepsilon_1] + \frac{1}{12} * \left[\frac{\bar{\theta}^2}{2}\right] * [\bar{\theta}]$$

  which is strictly positve

$\square$

52