

Modes of persuasion toward unanimous consent*

Arjada Bardhi[†] Yingni Guo[‡]

Abstract

A fully committed sender seeks to sway a collective adoption decision through designing experiments. Voters have correlated payoff states and heterogeneous thresholds of doubt. We characterize the sender-optimal policy under unanimity rule for two persuasion modes. Under general persuasion, evidence presented to each voter depends on all voters' states. The sender makes the most demanding voters indifferent between decisions, while the more lenient voters strictly benefit from persuasion. Under individual persuasion, evidence presented to each voter depends only on her state. The sender designates a subgroup of rubber-stampers, another of fully informed voters, and a third of partially informed voters. The most demanding voters are strategically accorded high-quality information.

JEL: D71, D82, D83, G28, K20, O32

Keywords: information design, collective decision-making, unanimity rule, information guard

A tremendous share of decision-making in economic and political realms is made within collective schemes. We explore a setting in which a sender seeks to get the unanimous approval of a group for a project he promotes. Group members care about different aspects of the project and might disagree on whether the project should be implemented. They might also vary in the loss they incur if the project is of low quality in their respective aspects. The sender designs experiments to persuade the members to approve. When deciding as part of a group, individuals understand the informational and payoff interdependencies among their decisions.

*For helpful comments and suggestions, we thank Thomas Mariotti, two anonymous referees, Eddie Dekel, Harry Di Pei, Georgy Egorov, Jeffrey Ely, Emir Kamenica, Navin Kartik, Alessandro Pavan, Nicola Persico, Bruno Strulovici, Jean Tirole, Asher Wolinsky, and seminar participants and conference audiences at University of Chicago, Columbia University, University of Miami, 2017 UK PolEcon Conference (LSE), 2017 NASM (St Louis), 2017 SAET (Faro), 2017 Social Networks and Information Conference (Johns Hopkins University), Summer School of Econometric Society 2017 (Seoul), UBC Summer Theory Conference, Transatlantic Theory Workshop 2016 (Oxford), 27th International Conference in Game Theory (Stony Brook), and EconCon 2016 (Princeton).

[†]Department of Economics, Northwestern University, abardhi@u.northwestern.edu.

[‡]Department of Economics, Northwestern University, yingni.guo@northwestern.edu.

Previous literature has focused mostly on the aggregation and acquisition of (costly) information from exogenous sources in collective decision-making. In contrast, our focus is on optimal persuasion of a heterogeneous group by a biased sender who is able to design the information presented to each group member. We aim to understand the optimal information design, the extent to which group decision-making might be susceptible to information manipulation, and the welfare implications of persuasion for each voter under unanimity rule. Moreover, we contrast optimal persuasion under unanimity with that under nonunanimous rules.

Let us briefly discuss two examples captured by the model. Consider first an industry representative that aims to persuade multiple regulators to approve a project. This representative could be a trade association or an industry-wide self-regulatory authority that interacts directly with regulators.¹ Each regulator is concerned about different but correlated aspects of the project. Typically a successful approval entails the endorsement of all regulators. The representative provides evidence to each regulator by designing informative experiments about the project.

A second example concerns the flow of innovative ideas within organizations. Such ideas are typically born in the R&D department, but they are required to find broad support from other departments, with potentially varied interests, before implementation. The R&D department provides tests to persuade them, and may vary these tests to fit the particular concerns of the department being addressed.

The sender seeks to maximize the ex-ante probability that the project is approved by the entire group.² He establishes and commits to an institutionalized standard for the amount of information to be provided to each voter. Modifying institutional standards on an ad hoc basis is costly and difficult due to legal constraints. At the time of the design the sender is uncertain about the quality of projects to be evaluated using this standard.³

In both examples, the sender is a key source of information for the receivers. In

¹In the context of the U.S. financial industry, such an industry representative is the Financial Industry Regulatory Authority, the private self-regulatory authority created by the industry and serving its internal needs. FINRA provides information to different regulatory agencies such as the Securities and Exchange Commission (SEC), the Consumer Financial Protection Bureau (CFPB), the Federal Deposit Insurance Corporation (FDIC) etc. It arguably has wide authority in determining the precision of the evidence presented to the regulators (McCarty (2015)).

²An example in which majority (rather than unanimous) approval suffices is that of a lobbyist who persuades multiple legislators but only needs a share of them to support the cause.

³There are two sources of commitment in the examples we present. R&D units are naturally uninformed about the quality of their innovation at the time of test design. In contrast, in regulation, the industry representative commits on behalf of the entire industry to certain guidelines of information disclosure for any projects to be presented to regulators in the future. In particular, FINRA has a consolidated rulebook on disclosure rules and standards; for an example, see FINRA's Regulatory Notice 10-41 (<http://www.finra.org/industry/notices/10-41>).

complex policy environments, the regulators are highly dependent on the industry for expertise and knowledge on how to evaluate the project under investigation.⁴ Within innovative organizations, often R&D units exclusively have the required expertise and background information to test the quality of their innovations. They occupy the superior ground of designing tests to persuade other organizational units, which lack the expertise to find independent informational sources.⁵

We take an information design approach in analyzing how the sender optimally persuades a unanimity-seeking group. We consider two scenarios: general persuasion, in which evidence presented to each voter depends on the aspects of all voters; and individual persuasion, in which evidence presented to each voter depends only on her own aspect. We show that the form of the optimal policy as well as the welfare distribution among voters differ drastically under these two persuasion modes.

The two modes describe natural forms of evidence presentation. Within the context of regulation, there is an ongoing debate about the relative effectiveness of comprehensive and targeted evidence to different regulators (Harris and Firestone 2014). General (individual) persuasion is akin to comprehensive (targeted) evidence. Individual persuasion is more natural in regulatory contexts with greater independence and more clearly defined areas of authority across regulators, while general persuasion naturally describes contexts in which each regulator obtains varied evidence on many aspects of a project. Our analysis of the two modes sheds light on the implications of different forms of evidence presentation for optimal information design by the sender and the welfare of different regulators.

Our model features a sender and n voters. A voter's preference is characterized by her binary payoff state. Her payoff from the project is positive if her state is high and negative if the state is low. We assume that the distribution of the voters' payoff states is affiliated and symmetric. With perfectly correlated states, voters agree about the right decision if their states are commonly known. Away from this extreme case, they might disagree about the right decision even if the realized states are commonly known. The magnitude of the loss suffered from approval by a low-state voter differs across voters. We interpret these varying magnitudes as heterogeneous thresholds of doubt; the higher the threshold, the more demanding the voter is. The thresholds and the distribution over state profiles are commonly

⁴See McCarty (2015), Omarova (2011), Woodage (2012). Also, McCarty (2015) and McCarty, Poole, and Rosenthal (2013) have argued that within the financial sector in the United States, federal agencies such as SEC and FDIC are demonstrably reliant on the expertise and information provided by FINRA.

⁵Evaluative authorities might attempt to enforce some minimal standards of disclosure on the sender (e.g. Sarbanes-Oxley Act in the finance sector). Our analysis can be interpreted as the design problem after the sender fulfills the minimal requirement imposed by the authorities.

known by all. Ex ante, neither the voters nor the sender know the realized state profile.

Under general persuasion, the sender designs a mapping from the set of state profiles to the set of distributions over signal profiles. Two features are worth emphasizing: the distribution over the signal profiles is conditioned on the states of all voters; and the signals across voters can be correlated. Under individual persuasion, each voter's signal distribution depends only on her own state. For each voter the sender designs a mapping from the two possible payoff states to the set of distributions over this voter's signal space. We interpret general persuasion as one grand experiment and individual persuasion as voter-specific targeted experiments. Under each mode, we assume that the designed policies are public information, based on the observation that in regulatory settings standards of information are public knowledge either by legal or other institutional requirements. Furthermore, we assume that each voter observes her signal privately. This assumption is based on the observation that confidential information communicated to a specific regulator and pertaining to a specific project is not observable by all parties.⁶ Yet, the results remain valid even if all signals were publicly observed.

Under general persuasion, all players are concerned only with the set of state profiles in which all voters receive a recommendation to approve. In the optimal policy the sender chooses a group of the most demanding voters such that they are made indifferent between approval and rejection whenever the sender recommends approval. The more lenient voters obtain a positive payoff merely due to the presence of the more demanding voters. In the extreme case of perfectly correlated states, this group consists of only the most demanding voter. The sender achieves the same payoff as if he faced this voter alone.

A short detour considers the case in which the sender is required to draw the signals independently across voters, conditional on the entire state profile, which we refer to as *independent general* policies. We show that any general policy can be replicated with an independent general policy for unanimity rule.

Under individual persuasion, each voter learns about her state directly from her policy and indirectly from the conjectured decisions of others. With independent states, no voter learns payoff-relevant information from the approval of others. Hence the optimal policy consists of the single-voter policy for each voter. In contrast, for perfectly correlated states, the sender persuades the most demanding voter

⁶The assumption of private observability of signals is motivated by the fact that regulatory agencies often face legal and bureaucratic obstacles in sharing information smoothly with each other. Moreover, the principle of strict independence of different regulators is often used as justification for the lack of information sharing.

as if he needed only her approval, while all other voters always approve.

When the states are imperfectly correlated, each voter receives the sender's recommendation to approve with certainty if her payoff state is high. Intuitively, a higher probability of approval by a high-state voter benefits the sender while also boosting the beliefs of all other voters about their own states. We show that, when a voter's state is low, the probability that she is asked to approve decreases in her threshold of doubt. Thus, the optimal individual policy provides more precise information to more demanding voters.

The optimal individual policy divides the voters into at most three subgroups: the most lenient voters who rubber-stamp the project, the most demanding voters who learn their states fully, and an intermediate subgroup who are partially-informed. Interestingly, the sender does not persuade all voters to approve as frequently as possible, contrary to the case with only one voter. For moderate correlation of states, the most demanding voter(s) learn their states fully and reject for sure when their state is low. Full revelation of the individual state is more informative than what is necessary to persuade the strictest voter(s), but it allows the sender to persuade other voters more effectively. The most demanding and the least demanding voters might obtain a strictly positive expected payoff. For the former, the payoff is due to the information externality they generate for others by acting as *information guards* for the collective decision. For the latter, the positive expected payoff is due to their willingness to rubber-stamp other voters' informed decisions. The intermediate voters obtain a zero expected payoff.

Under either persuasion mode, the sender prefers smaller to larger groups. Also, he weakly prefers general to individual persuasion: any approval probability attained by an individual policy is also achieved through a general policy. The sender's payoffs across the two modes coincide when the states are sufficiently correlated. The most demanding voter weakly prefers individual persuasion, while the rest of the voters might disagree on the preferred persuasion mode.

When moving away from unanimity rule, the results change drastically. For nonunanimous voting rules, the sender achieves a payoff of one under general and independent general persuasion. The project is approved with certainty, so there is no meaningful check on the adoption decision by the voters.⁷ In contrast, individual persuasion cannot achieve a certain approval. The voters unambiguously prefer individual persuasion to general persuasion because the former allows them to partially discriminate between favorable and unfavorable projects.

⁷For any certain-approval policy, there is a nearby policy such that each voter is pivotal with positive probability and strictly prefers to follow an approval recommendation.

The rest of this section discusses the related literature. Section 1 presents the formal model. Sections 2 through 4 analyze the two persuasion modes. Section 5 compares and contrasts these modes. Section 6 briefly discusses two extensions. In particular, subsection 6.2 characterizes sender-optimal persuasion for nonunanimous rules. We conclude and discuss directions for future research in section 7.

Related literature. This paper is immediately related to the literature on persuasion. Rayo and Segal (2010) and Kamenica and Gentzkow (2011) study optimal persuasion between a sender and a single receiver.⁸ We study the information design problem of a sender who persuades a group of receivers. Bergemann and Morris (2016a, 2016b), Taneva (2014), Mathevet, Peregó, and Taneva (2016), and Bergemann, Heumann, and Morris (2015) also focus on information design with multiple receivers in non-voting contexts. In our setting, voters interact with the sender without any prior private information. The incentive-compatibility constraints for general persuasion characterize the entire set of Bayes correlated equilibria (BCE), while those for individual persuasion characterize a subset of the set of BCE. We identify the sender-optimal BCE within these two sets. In contrast to general persuasion, once attention is restricted to individual persuasion the sender’s problem is no longer a linear program.

More specifically, our paper is closely related to the recent literature on persuading voters. Alonso and Câmara (2016b) explore general persuasion when the sender is restricted to public persuasion. We focus on private persuasion. The differences are threefold. Firstly, we show in subsection 6.1 that whether persuasion is private or public is inconsequential under unanimity rule. Hence, under unanimity our general persuasion setting is a special case of Alonso and Câmara. We strengthen their implications for unanimity by characterizing the general persuasion solution in more detail for a broad class of state distributions. Our main result for general persuasion and unanimity is not implied by their analysis.⁹ Secondly, we also study individual persuasion under unanimity to examine how the optimal policy changes when the evidence presented to each voter depends only on this voter’s state. Lastly, when the voting rule is nonunanimous, we characterize the optimal policy under private persuasion which is drastically different from the optimal policy that Alonso and Câmara identify for public persuasion.

Schnakenberg (2015) shows that the sender can achieve certain approval through

⁸More broadly, the paper is related to the literature on communicating information through cheap talk (Crawford and Sobel (1982)), disclosure of hard evidence (see Milgrom (2008) for a survey), and verification (Glazer and Rubinstein (2004)).

⁹In particular, it is not a special case of their proposition 3 which assumes that all voters rank states in the same order. In our setting, voters might disagree even if the state profile were commonly known.

public persuasion under certain prior distributions if and only if the voting rule is noncollegial. The unanimity rule, which we focus on, is collegial since the approval of all voters is required. Wang (2015) and Chan, Gupta, Li, and Wang (2016) focus on persuading voters who agree under complete information but have different thresholds. In contrast, we examine an environment in which voters might have heterogeneous preferences even under complete information. Moreover, Chan, Gupta, Li, and Wang (2016) characterize the optimal design among information structures that rely on minimal winning coalitions. Arieli and Babichenko (2016) study the optimal group persuasion by a sender who promotes a product. Each receiver makes her own adoption decision so, unlike our setting, there is no payoff externality among receivers.

Another closely related paper is Caillaud and Tirole (2007). In the language of our paper, they also consider individual persuasion. Their setting differs from ours in two aspects: (1) the sender can either reveal or hide a voter’s state perfectly; (2) a voter pays a cost to investigate the evidence if provided by the sender. Due to the cost, only voters with moderate beliefs find investigation worthwhile. They show that the sender optimally provides information to a moderate voter so that a more pessimistic voter, who is not willing to investigate if alone, agrees to rubber-stamp the other’s approval. Both their analysis and ours rely on the observation that each voter learns about her state from the decisions of others. Yet, we are interested in how the sender adjusts the information precision for each individual voter. We show that the sender provides better information to the most demanding voters in order to convince the more lenient ones to follow suit. More surprisingly, he may find it optimal to fully reveal their states to the most demanding voters.¹⁰

Our paper also relates to a large literature on information aggregation in collective decisions with exogenous private information, following Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1996, 1997), as well as on information acquisition in voting games. Li (2001) allows the voters to choose the precision of their signals through costly effort. He argues that groups might choose to commit to more stringent decision-making standards than what is ex-post optimal so as to avoid free riding at the information-acquisition stage. Persico (2004) considers the optimal design of a committee, both in terms of its size and its threshold voting rule, so as to incentivize private acquisition of information. Gerardi and Yariv (2008) and

¹⁰Proposition 6 of Caillaud and Tirole shows that the sender might do better facing two randomly drawn voters than one. The cost of investigation is the key. A single pessimistic voter never investigates or approves. Adding a moderate voter who is willing to investigate might induce the pessimistic voter to rubber-stamp. In our environment, there is no cost and the sender can choose the information precision, so more voters always hurt him.

Gershkov and Szentes (2009) look at a broader class of mechanisms that incentivize costly information acquisition within a committee. Our focus, in contrast, is on the sender's design of the information structure so as to influence the group decision.

1 Model

Players and payoff states. We consider a communication game between a sender (he) and n voters $\{R_i\}_{i=1}^n$ (she). The voters collectively decide whether to adopt a project promoted by the sender. The payoff of each voter from adopting the project depends on her individual payoff state. In particular, R_i 's payoff state, denoted by $\theta_i \in \{H, L\}$, can be either high or low. Let $\theta = (\theta_1, \dots, \theta_n)$ denote the state profile of the group, and let $\Theta := \{H, L\}^n$ denote the set of all such state profiles.

Before the game begins, nature randomly draws the state profile θ according to a distribution f . The realized state profile is initially unobservable to all players and f is common knowledge. Throughout our analysis, we assume that the random variables $(\theta_1, \dots, \theta_n)$ are exchangeable, in the sense that for every θ and for every permutation ρ of the set $\{1, 2, \dots, n\}$, the following holds:

$$f(\theta_1, \dots, \theta_n) = f(\theta_{\rho(1)}, \dots, \theta_{\rho(n)}).$$

The analysis of individual persuasion in section 4 and 6.2 also assumes that the voters' states are affiliated. For any two state profiles $\theta, \theta' \in \Theta$, let $\theta \vee \theta'$ and $\theta \wedge \theta'$ denote the componentwise maximum and minimum state profiles respectively.¹¹ Affiliation of states requires that, for any $\theta, \theta' \in \Theta$,

$$f(\theta \vee \theta')f(\theta \wedge \theta') \geq f(\theta)f(\theta').$$

Let θ^H and θ^L be the state profiles such that $\theta_i^H = H$ for all i and $\theta_i^L = L$ for all i , respectively. The voters' states are perfectly correlated if and only if $f(\theta^H) + f(\theta^L) = 1$. More generally, for $f(\theta^H) + f(\theta^L) < 1$, the states are imperfectly correlated.

Decisions and payoffs. The sender designs an information policy which generates individual signals about the realized state profile. At the time of the design, the sender is uninformed of the realized state profile and fully commits to the chosen policy. The signal intended for R_i is observed only by her.¹²

After observing their signals, voters simultaneously decide whether to approve the project. We let $d_i \in \{0, 1\}$ represent R_i 's approval decision, where $d_i = 1$

¹¹Formally, $\theta \vee \theta' := (\max\{\theta_1, \theta'_1\}, \dots, \max\{\theta_n, \theta'_n\})$ and $\theta \wedge \theta' := (\min\{\theta_1, \theta'_1\}, \dots, \min\{\theta_n, \theta'_n\})$.

¹²Under unanimity, it makes no difference whether the signal profile is public or private, as we show in subsection 6.1. Each R_i behaves as if all other voters had received signals that induce them to approve. The private-observability assumption is crucial when we discuss nonunanimous rules in subsection 6.2.

denotes approval. Analogously, $d \in \{0, 1\}$ denotes the collective adoption decision. Under unanimous consent, the project is adopted ($d = 1$) if and only if $d_i = 1$ for every i .

The sender prefers approval of the project regardless of the realized θ : his payoff is normalized to one if the project is approved and zero otherwise. The payoff of R_i depends only on her own state θ_i . The project yields a payoff of one to R_i if $\theta_i = H$ and $-\ell_i < 0$ if $\theta_i = L$. Here, ℓ_i captures the magnitude of the loss incurred by R_i when the project is adopted and R_i 's state is L . If the project is rejected, the payoffs of all voters are normalized to zero. Without loss, we assume that no two voters are identical in their thresholds: $\ell_i \neq \ell_j$ for $i \neq j$. For notational convenience, the voters are indexed in increasing order of leniency:

$$\ell_i > \ell_{i+1} \text{ for all } i \in \{1, \dots, n-1\}.$$

Each R_i prefers adoption if and only if her state is H . Let $\Theta_i^H := \{\theta \in \Theta : \theta_i = H\}$ and $\Theta_i^L := \{\theta \in \Theta : \theta_i = L\}$ be the set of state profiles with R_i 's state being H and L respectively. Therefore, Θ_i^H and Θ_i^L contain R_i 's favorable and unfavorable state profiles, respectively. For any i , $\Theta_i^H \cup \Theta_i^L = \Theta$, and $\Theta_i^H \cap \Theta_i^L = \emptyset$.

Voter R_i 's prior belief of her state being H is $\sum_{\theta \in \Theta_i^H} f(\theta)$. Due to exchangeability of f , all voters share the same prior belief of their state being H .¹³ We focus on parameter values for which none of the voters prefers approval under the prior belief:

Assumption 1. *For any $i \in \{1, \dots, n\}$, R_i strictly prefers to reject the project under the prior belief, i.e.:*

$$\ell_i > \frac{\sum_{\theta \in \Theta_i^H} f(\theta)}{\sum_{\theta \in \Theta_i^L} f(\theta)}.$$

We make assumption 1 to simplify exposition. It is without loss of generality. We show in online appendix B.2 that, if some voters prefer to approve ex ante, then the sender designs the optimal policy for those who are reluctant to approve. Those who prefer to approve ex ante are always willing to rubber-stamp.

Modes of persuasion. Let S_i denote the signal space for R_i and, correspondingly, $\prod_i S_i$ the space of signal profiles for the entire group. The game has three stages. The sender first commits to an information policy. Subsequently, nature draws θ according to f and then draws the signals $(s_i)_{i=1}^n$ according to the chosen information policy. Finally, each voter receives her signal and chooses d_i .

We consider three major classes of information policies: 1) general policies, 2) independent general policies, and 3) individual policies. Formally, a *general* policy

¹³By assuming exchangeability, we can focus on the impact of different thresholds on the optimal information policy.

is a mapping from the set of state profiles to the set of probability distributions over the signal space:

$$\pi : \Theta \rightarrow \Delta \left(\prod_i S_i \right).$$

For each state profile θ , it specifies a distribution over all possible signal profiles, so the signals sent to the voters could be correlated conditional on the state profile. We let Π^G denote the set of all such policies.

An *independent general* policy specifies, for each voter R_i , a mapping from the set of state profiles to the set of probability distributions over R_i 's signal space:

$$\pi_i : \Theta \rightarrow \Delta(S_i).$$

As in a general policy, each voter's signal distribution depends on the entire state profile. But unlike in a general policy, conditional on θ the signals across voters are independently drawn. Let Π^{IG} denote the set of all independent general policies, with $(\pi_i)_{i=1}^n$ being a typical element of it.

An *individual* policy specifies, for each voter R_i , a mapping from R_i 's state space to the set of probability distributions over S_i :

$$\pi_i : \{H, L\} \rightarrow \Delta(S_i).$$

Unlike in a general or independent general policy, each voter's signal distribution depends only on her own state. Let Π^I denote the set of all individual policies, with $(\pi_i)_{i=1}^n$ being a typical element.

The information policy adopted by the sender determines the information structure of the voting game played among voters. Let Π denote the set of policies available to the sender. If the sender is allowed to use general policies, then $\Pi = \Pi^G$. If the sender is constrained to independent general policies or individual policies, then $\Pi = \Pi^{IG}$ or $\Pi = \Pi^I$ respectively.¹⁴ The strategy σ_i determines the probability R_i approves, given the chosen policy and the realized signal:

$$\sigma_i : \Pi \times S_i \rightarrow [0, 1].$$

Without loss, we focus on *direct obedient policies*, for which (i) S_i coincides with the action space $\{0, 1\}$, and (ii) each voter receives action recommendations, with which she complies.¹⁵ We use $\hat{d}_i \in \{0, 1\}$ to represent the sender's recommendation to R_i , and $\hat{d} = (\hat{d}_i)_i \in \{0, 1\}^n$ to represent the profile of action recommendations.

A direct general policy specifies a distribution over action-recommendation pro-

¹⁴By definition, any individual policy can be replicated by an independent general policy. Any independent general policy can be replicated by a general policy.

¹⁵The first part of online appendix B.3 provides a proof that the restriction to direct obedient policies is without loss of generality.

Modes	Definition	Notation
General	$\pi : \Theta \rightarrow \Delta(\{0, 1\}^n)$	$(\pi(\cdot \theta))_{\theta \in \Theta}$
Independent general	$\pi_i : \Theta \rightarrow \Delta(\{0, 1\}), \forall i$	$(\pi_i(\theta))_{\theta \in \Theta}, \forall i$
Individual	$\pi_i : \{H, L\} \rightarrow \Delta(\{0, 1\}), \forall i$	$(\pi_i(H), \pi_i(L)), \forall i$

Table 1: Modes of persuasion

files as a function of the realized state profile, i.e., it specifies $\pi(\cdot|\theta) \in \Delta(\{0, 1\}^n)$ for each θ . A direct independent general policy specifies for each R_i and for each θ the probability $\pi_i(\theta)$ with which R_i is recommended to approve, since the action space is binary. For individual persuasion, any direct policy specifies $(\pi_i(H), \pi_i(L)) \in [0, 1]^2$ for each R_i , where $\pi_i(\theta_i)$ is the probability that the sender recommends approval to R_i when her state is θ_i . Table 1 summarizes the definition and notation for these three different modes of persuasion.

Refinement. We allow for any policy that is the limit of a sequence of direct obedient policies with full support. The full-support requirement demands that for any state profile, all possible recommendations are sent with positive probabilities.¹⁶ We impose this refinement so that along the sequence (i) no recommendation is off the equilibrium path; and (ii) a voter is always pivotal with positive probability.¹⁷ For each persuasion mode, we solve for the sender-optimal policy.

2 General persuasion

2.1 General formulation

In the regulatory process for complex industries, the most general form that the evidence presented to the regulators can take is as an experiment that generates correlated action recommendations conditional on the realized states of all regulators. The structure of such an experiment is formally captured by a *general policy*.

Recall that $\pi(\hat{d}|\theta)$ is the probability that the recommendation profile \hat{d} is sent, given the state profile θ . If R_i rejects, the project is definitively rejected. If R_i approves, the project is collectively approved if and only if all other voters approve as well; this is the only event in which R_i 's decision matters. Let \hat{d}^a be the recommendation under which all voters receive a recommendation to approve. We refer to \hat{d}^a as the unanimous (approval) recommendation. Voter R_i obeys a recommendation

¹⁶The second part of online appendix B.3 provides a formal definition of full-support policies.

¹⁷For unanimity rule, we can easily construct a sequence of full-support direct obedient policies that approaches the optimal policy. Therefore, we do not invoke this refinement explicitly in the main discussion. However, when we discuss nonunanimous rules in subsection 6.2, we construct explicitly the sequence of full-support policies that approaches the optimal policy.

to approve if

$$\sum_{\theta \in \Theta_i^H} f(\theta) \pi(\hat{d}^a | \theta) - \ell_i \sum_{\theta \in \Theta_i^L} f(\theta) \pi(\hat{d}^a | \theta) \geq 0. \quad (\text{IC}^a\text{-}i)$$

Let $\hat{d}^{r,i}$ be the recommendation under which all voters except R_i receive a recommendation to approve, i.e., $\hat{d}_i^{r,i} = 0$ and $\hat{d}_j^{r,i} = 1$ for all $j \neq i$. Voter R_i obeys a recommendation to reject if

$$\sum_{\theta \in \Theta_i^H} f(\theta) \pi(\hat{d}^{r,i} | \theta) - \ell_i \sum_{\theta \in \Theta_i^L} f(\theta) \pi(\hat{d}^{r,i} | \theta) \leq 0. \quad (\text{IC}^r\text{-}i)$$

The sender chooses $(\pi(\cdot | \theta))_{\theta \in \Theta}$ so as to maximize the probability of the project being approved, $\sum_{\theta \in \Theta} f(\theta) \pi(\hat{d}^a | \theta)$, subject to (i) approval and rejection IC constraints, and (ii) the following feasibility constraints: $\pi(\hat{d} | \theta) \geq 0$, for all \hat{d}, θ , and $\sum_{\hat{d} \in \{0,1\}^n} \pi(\hat{d} | \theta) = 1$ for all θ .

We first analyze the relaxed problem in which IC^r constraints are ignored:

$$\begin{aligned} \max_{(\pi(\hat{d}^a | \theta))_{\theta \in \Theta}} \quad & \sum_{\theta \in \Theta} f(\theta) \pi(\hat{d}^a | \theta), \\ \text{subject to} \quad & (\text{IC}^a\text{-}i)_i \quad \text{and} \quad \pi(\hat{d}^a | \theta) \in [0, 1], \quad \forall \theta \in \Theta. \end{aligned} \quad (1)$$

A solution to this relaxed problem specifies, for each θ , only the probability that all voters are recommended to approve, i.e., $(\pi(\hat{d}^a | \theta))_{\theta \in \Theta}$. For any such solution, we can easily construct the probabilities $(\pi(\hat{d}^{r,i} | \theta))_{i,\theta}$ so that the IC^r constraints are satisfied as well.¹⁸ Therefore, focusing on the relaxed problem is without loss.

2.2 Characterization of the optimal policy

Each R_i learns about her state from the relative frequency with which the unanimous recommendation \hat{d}^a is generated in Θ_i^H rather than in Θ_i^L . The posterior belief that R_i holds about her state being H conditional on \hat{d}^a having been drawn is:

$$\Pr(\theta_i = H | \hat{d}^a) = \frac{\sum_{\theta \in \Theta_i^H} f(\theta) \pi(\hat{d}^a | \theta)}{\sum_{\theta \in \Theta} f(\theta) \pi(\hat{d}^a | \theta)}.$$

Each $\text{IC}^a\text{-}i$ can be rewritten in terms of this posterior belief as:

$$\Pr(\theta_i = H | \hat{d}^a) \geq \frac{\ell_i}{1 + \ell_i}.$$

This posterior belief has to be sufficiently high for R_i to obey an approval recommendation. The cutoff value for this posterior belief increases in the threshold of

¹⁸Lemma 2.2 establishes that for any policy that satisfies the approval IC constraints and for any R_i there exists $\theta' \in \Theta_i^L$ such that $\pi(\hat{d}^a | \theta') < 1$. The sender can specify $\pi(\hat{d}^{r,i} | \theta') = \varepsilon$ and $\pi(\hat{d}^{r,i} | \theta) = 0$ for all $\theta \neq \theta'$. Such a specification guarantees that $\text{IC}^r\text{-}i$ holds for each R_i .

doubt: naturally, the larger the loss a voter experiences if her state is L , the higher the posterior belief about $\theta_i = H$ needed in order for this voter to prefer approval.

We first examine the optimal policy for perfectly correlated states. Only two state profiles, θ^H and θ^L , are possible to realize.

Proposition 2.1 (Perfect correlation). *Suppose the voters' states are perfectly correlated. The unique optimal policy, for which only IC^a-1 binds, is given by:*

$$\pi(\hat{d}^a|\theta^H) = 1, \quad \pi(\hat{d}^a|\theta^L) = \frac{f(\theta^H)}{f(\theta^L)} \frac{1}{\ell_1}.$$

The unanimous recommendation is sent with certainty given θ^H . The probability of unanimous approval in θ^L is determined only by the threshold of the most demanding voter R_1 . Due to perfect correlation, all voters share the same posterior belief about their respective states being H . The highest cutoff on this posterior belief is imposed by R_1 . Thus, the sender provides sufficiently accurate recommendations so as to leave R_1 indifferent between approval and rejection. Being more lenient than R_1 , all other voters receive a strictly positive expected payoff.

We now generalize our discussion to imperfectly correlated states. Our first observation is that the sender recommends approval with certainty to all voters when all their respective states are high. The intuition is straightforward: increasing $\pi(\hat{d}^a|\theta^H)$ strengthens the posterior belief $\Pr(\theta_i = H|\hat{d}^a)$ of each voter R_i , while also strictly improving the probability of a collective approval.

Lemma 2.1 (Certain approval for θ^H). *The sender recommends with certainty that all voters approve when every voter's state is high, i.e., $\pi(\hat{d}^a|\theta^H) = 1$.*

We next show that, for any R_i , an optimal policy does not set $\pi(\hat{d}^a|\theta) = 1$ for all $\theta \in \Theta_i^L$. If there existed such a voter who knows that the group receives \hat{d}^a for sure whenever her state is L , her posterior belief conditional on \hat{d}^a would be lower than the prior belief. She would not be willing to approve given assumption 1.

On the other hand, there does not exist a voter who, given that \hat{d}^a is sent, is fully confident that her state is H . Put differently, every voter mistakenly obeys a unanimous recommendation for a project for which her state is low with positive probability. If indeed some voter R_i learned her state fully, her IC^a- i would be slack. Moreover, full revelation would require that $\pi(\hat{d}^a|\theta) = 0$ for $\theta \in \Theta_i^L$ such that $\theta_j = H$ for all $j \neq i$. By increasing the frequency with which \hat{d}^a is generated in this state profile, the sender improves his payoff while still satisfying the previously slack IC^a- i and strictly relaxing all other IC^a constraints.

Lemma 2.2 (No certain approval or rejection for Θ_i^L). *For each i , there exist $\theta, \theta' \in \Theta_i^L$ such that $\pi(\hat{d}^a|\theta) < 1$ and $\pi(\hat{d}^a|\theta') > 0$.*

A natural next question concerns the pattern of binding and slack IC^a constraints across voters. Under the unanimous rule with an outside option normalized at zero, a binding IC^a constraint implies that the corresponding voter's expected payoff is exactly zero. Hence, an analysis of the subset of binding IC^a constraints has immediate implications for the welfare of the voters. The following proposition establishes that there exists an index $i' \geq 1$ such that all voters who are more demanding than $R_{i'}$ have binding IC^a constraints and a zero expected payoff in any optimal policy. The only voters who might obtain a positive expected payoff are the most lenient voters in the group.¹⁹

Proposition 2.2 (The strictest voters' IC^a constraints bind). *Suppose f is exchangeable. In any optimal policy, a subgroup of the strictest voters' IC^a constraints bind, i.e., IC^a - i binds iff $i \in \{1, \dots, i'\}$ for some $i' \geq 1$.*

This proposition holds as long as f is exchangeable. Ex ante, the voters differ only in their thresholds. The proof makes use of the dual problem corresponding to (1). Thinking of each voter's IC^a as a resource constraint, we show that granting positive surplus to a tough voter is more expensive than to a lenient one. Intuitively, the voters with the highest thresholds are the hardest to persuade; hence, the sender provides sufficiently precise information about the strictest voters' states in order to leave them indifferent between approval and rejection.

Let us now briefly touch upon the multiplicity of optimal policies that arises under general persuasion. The dual problem corresponding to (1) identifies the set of binding IC^a constraints. It also identifies the state profiles for which the project is approved or rejected for sure. These conditions pin down the sender's payoff. As a result, he has flexibility in designing how frequently approval is recommended to all voters in other state profiles. For voters with slack IC^a constraints, their expected payoffs vary across different optimal policies. The following example with $n = 3$ and independent payoff states illustrates this multiplicity.

Example 1 (Multiplicity of optimal policies). *We suppose that (ℓ_1, ℓ_2, ℓ_3) equals $(20, 15, \frac{2291}{229})$. Voters' states are independent. The state of each voter is H with probability $9/10$. By examining the dual problem (which is a linear program), we determine that in any optimal policy:*

$$\pi(\hat{d}^a|HHH) = \pi(\hat{d}^a|HHL) = 1, \quad \pi(\hat{d}^a|LLH) = \pi(\hat{d}^a|LLL) = 0.$$

Moreover, both IC^a -1 and IC^a -2 bind. These pin down the sender's payoff. Subject to the binding IC^a for R_1 and R_2 , the feasibility constraints, and IC^a -3, the sender

¹⁹We say that R_i 's IC^a constraint binds if the dual variable associated with this constraint is strictly positive.

has flexibility in specifying the rest of the unanimous recommendation probabilities. Among the optimal policies, R_3 's payoff is the highest if

$$\pi(\hat{d}^a|HLH) = \frac{210}{299}, \pi(\hat{d}^a|LHH) = \frac{160}{299}, \pi(\hat{d}^a|HLL) = \pi(\hat{d}^a|LHL) = 0.$$

R_3 receives a strictly positive payoff from this policy. R_3 's payoff is the lowest if

$$\pi(\hat{d}^a|HLH) = \frac{47}{69}, \pi(\hat{d}^a|LHH) = \frac{160}{299}, \pi(\hat{d}^a|HLL) = \frac{57}{299}, \pi(\hat{d}^a|LHL) = 0,$$

which grants R_3 a zero expected payoff. \square

In this example, the optimal policy in which R_3 receives a zero payoff is Pareto dominated by the optimal policy in which R_3 receives a positive payoff. The sender may very well choose an optimal policy that is also Pareto efficient. The set of optimal and Pareto efficient policies is given by maximizing the weighted sum of voters' payoffs subject to their IC^a constraints and the constraint that the sender obtains her optimal payoff.

Example 1 also illustrates that even when the voters' payoff states are entirely independent, the general persuasion problem faced by the sender is not separable across voters. For example, the probability with which R_3 is recommended to approve the project when her state is L depends on the realized states of the other voters: in any optimal policy in example 1, $\pi(HHL) = 1$ but $\pi(LLL) = 0$.²⁰

3 Equivalence of general and independent general policies

In the context of our motivating example, the regulatory process might require the sender to conduct an independent experiment for each regulator. If the sender is allowed to condition the recommendations to each regulator on the entire state profile, the sender designs for each regulator a mapping from the set of state profiles to the set of distributions over this regulator's signal space. Formally, such a profile of independent experiments is an *independent general policy*. We show in this section that under unanimity, for any general policy, there is an independent general policy that achieves the same payoffs for all players.

Recall that $\pi_i(\theta)$ denotes the probability of an approval recommendation made to R_i when the realized state profile is θ . The approval and rejection incentive-compatibility constraints have to be slightly modified for an independent general

²⁰This observation stands in sharp contrast to the case of individual persuasion, under which no voter can possibly receive a strictly positive payoff when the states are independent. In section 4, we show that the optimal individual policy exploits the lack of information externalities between voters to set all IC^{a-i} constraints binding.

policy $(\pi_i)_{i=1}^n$:

$$\sum_{\theta \in \Theta_i^H} f(\theta) \prod_{j=1}^n \pi_j(\theta) - \ell_i \sum_{\theta \in \Theta_i^L} f(\theta) \prod_{j=1}^n \pi_j(\theta) \geq 0, \quad (2)$$

$$\sum_{\theta \in \Theta_i^H} f(\theta) (1 - \pi_i(\theta)) \prod_{j \neq i} \pi_j(\theta) - \ell_i \sum_{\theta \in \Theta_i^L} f(\theta) (1 - \pi_i(\theta)) \prod_{j \neq i} \pi_j(\theta) \leq 0. \quad (3)$$

Similarly, the objective of the sender changes to

$$\sum_{\theta \in \Theta} f(\theta) \prod_{i=1}^n \pi_i(\theta).$$

We argue that a payoff is attainable under independent general persuasion if and only if it is attainable under general persuasion. One direction of this statement is trivial: any independent general policy can be formulated as a general policy. The next propositions shows that the other direction holds as well.

Proposition 3.1 (Equivalence of general/independent general policies). *Under unanimity, the set of attainable payoffs for the sender and the voters is the same under general policies and independent general policies.*

The key assumption for this result is the unanimity rule. The distribution f need not be exchangeable or affiliated. Moreover, the result holds even if each voter's state space is not binary. Under unanimity, the sender cares only about the event in which all voters receive an approval recommendation; so does each voter when she contemplates obeying an approval recommendation. Hence, for each θ , the sender can choose the approval probability for each voter so that the product of all these probabilities equals the unanimous approval probability in a general policy. When we consider nonunanimous rules, the equivalence does not hold any more: In subsection 6.2, we show that the sender is strictly worse off under independent general persuasion when voters' states are perfectly correlated. Moreover, this equivalence result is not expected to hold in other group persuasion games beyond voting games.²¹

4 Individual persuasion

4.1 General formulation

Restrictive regulatory processes might require that a regulator be provided only evidence directly pertaining to her area of interest. Such a requirement is often justified on grounds of protecting the independence of different regulatory agencies

²¹See Arieli and Babichenko (2016).

in their evaluations. Other times law assigns separate and disjoint areas of authority to different regulators: they decide based on evidence pertaining to their area of authority. Targeted experiments allow each regulator to focus her limited resources and have full authority over the evaluation of one aspect in parallel regulatory processes (i.e. processes that involve more than one regulator). Such experiments are formally captured by an *individual policy*.

This section characterizes optimal individual persuasion. We show that a more demanding voter enjoys a more informative policy. The sender essentially divides the group into (at most) three subgroups: (i) the most demanding voters fully learn their states; (ii) the intermediate voters are partially informed; and (iii) the most lenient voters rubber-stamp. We further show that only the extreme voters might obtain a positive payoff: the most demanding voters due to their role as informational guards, and the least demanding voters due to their willingness to rubber-stamp.

Recall that $\pi_i(\theta_i)$ denotes the probability that R_i receives an approval recommendation when her state is θ_i . Let $\Pr(\theta_i = H | R_{-i} \text{ approve})$ denote the probability that $\theta_i = H$, conditional on all voters other than R_i approving:

$$\Pr(\theta_i = H | R_{-i} \text{ approve}) = \frac{\sum_{\theta \in \Theta_i^H} f(\theta) \prod_{j \neq i} \pi_j(\theta_j)}{\sum_{\theta \in \Theta} f(\theta) \prod_{j \neq i} \pi_j(\theta_j)}.$$

R_i 's incentive-compatibility constraint when she receives an approval recommendation is:

$$\Pr(\theta_i = H | R_{-i} \text{ approve})\pi_i(H) - \ell_i(1 - \Pr(\theta_i = H | R_{-i} \text{ approve}))\pi_i(L) \geq 0. \quad (\text{IC}^a\text{-}i)$$

The sender maximizes the probability of a collective approval:

$$\sum_{\theta \in \Theta} f(\theta) \prod_{i=1}^n \pi_i(\theta_i). \quad (\text{OBJ})$$

We focus on the relaxed problem of maximizing (OBJ), subject to the set of $\text{IC}^a\text{-}i$ constraints.²² Focusing on such a relaxed problem is without loss: we show later in lemma 4.1 that recommendation to reject in any optimal policy for the relaxed problem is conclusive news that the voter's own state is low; hence she always obeys a rejection recommendation.

²²Under optimal individual persuasion, every voter approves with positive probability, i.e., $\pi_i(H) + \pi_i(L) > 0$, $\forall i$. Otherwise, the sender's payoff is zero. The sender could do better by fully revealing her payoff state to each voter. Hence, $\Pr(\theta_i = H | R_{-i} \text{ approve})$ is well defined.

4.2 Characterization of the optimal policy

The IC^a-*i* constraint, rewritten in the following form, emphasizes the informational externalities among the voters' decisions:

$$\Pr(\theta_i = H | R_{-i} \text{ approve}) \geq \frac{\ell_i \pi_i(L)}{\ell_i \pi_i(L) + \pi_i(H)}.$$

The left-hand side is R_i 's belief that her state is high when she conditions on the others' approvals. This belief depends on the policies of all voters other than R_i .²³ Due to affiliation, an increase in $\pi_j(H)$ of another voter R_j boosts the posterior belief of R_i , while an increase in $\pi_j(L)$ makes R_i more pessimistic about her state. The right-hand side depends only on R_i 's own policy and her threshold of doubt. It decreases in $\pi_i(H)$ and increases in $\pi_i(L)$ and ℓ_i . The more likely that R_i receives an approval recommendation when her state is high, the easier it is to induce compliance. The more frequently R_i receives an approval recommendation when her state is low or the more demanding R_i is, the more difficult it is to induce compliance. So, an increase in $\pi_i(H)$ relaxes not only IC^a-*i* but also IC^a-*j* for all other R_j , while an increase in $\pi_i(L)$ increases the cutoff for R_i and lowers the posterior belief of all other voters, thus tightening all IC^a constraints.

Before approaching the more general problem of imperfectly correlated states, we first solve the polar cases of perfectly correlated and independent states.

Proposition 4.1 (Perfectly correlated or independent states). *Suppose the voters' states are perfectly correlated. Any optimal policy is of the form:*

$$\pi_i(H) = 1, \quad \forall i, \quad \text{and} \quad \prod_{i=1}^n \pi_i(L) = \frac{f(\theta^H)}{f(\theta^L)} \frac{1}{\ell_1}.$$

Suppose the states are independent. The policy for each voter is the same as if the sender were facing only this voter:

$$(\pi_i(H), \pi_i(L)) = \left(1, \frac{\sum_{\theta \in \Theta_i^H} f(\theta)}{\sum_{\theta \in \Theta_i^L} f(\theta)} \frac{1}{\ell_i} \right) \text{ for all } i.$$

For perfectly correlated states, the project is approved for sure when every voter's state is H , while the probability of approval in state L is chosen so that only IC^a-1 binds. One such optimal policy is the one in which the sender persuades the most demanding R_1 and recommends that all other voters rubber-stamp the decision of R_1 .²⁴ For this policy, $\pi_1(L) = f(\theta^H)/(f(\theta^L)\ell_1)$ and $\pi_i(L) = 1$ for all $i \geq 2$.

For independent states, each voter R_i 's posterior belief $\Pr(\theta_i = H | R_{-i} \text{ approve})$ equals her prior belief of her state being H regardless of the other voters' policies.

²³With slight abuse of the term "policy," we call $(\pi_i(H), \pi_i(L))$ the policy of R_i under individual persuasion.

²⁴A voter rubber-stamps if she approves with probability one in both states.

In the absence of information externalities across voters, the sender sets $\pi_i(L)$ as high as IC^a - i allows for each voter R_i .

We next show that for imperfectly correlated states it continues to be the case that a high-state voter obtains an approval recommendation with probability one.

Lemma 4.1 (High states approve for sure). *Suppose that f is exchangeable and affiliated. In any optimal policy, the sender recommends approval to each R_i with probability one when $\theta_i = H$, i.e., $\pi_i(H) = 1$ for every i .*

The main step of the proof relies on the Ahlswede-Daykin inequality, through which we show that increasing $\pi_i(H)$ for R_i relaxes the IC^a constraints for voters other than R_i as well. The interests of the sender and all voters are aligned when it comes to an increase of $\pi_i(H)$ for any R_i : such an increase improves the sender's payoff, makes R_i more compliant with an approval recommendation, and makes all other voters more optimistic given that R_i approves. Thus, lemma 4.1 reduces the problem to simply choosing $(\pi_i(L))_{i=1}^n$.

Lemma 4.2 (At least one IC^a binds). *Suppose that f is exchangeable and affiliated. In any optimal policy, $\pi_i(L) < 1$ for some i . Moreover, at least one voter's IC^a constraint binds, so $\pi_j(L) > 0$ for some j .*

Lemma 4.2 rules out the possibility that all voters approve without any additional information from the policy, as a direct consequence of assumption 1. Moreover, any optimal policy accords a zero expected payoff to at least one voter. If that were not the case, the sender could strictly improve the approval probability by slightly increasing $\pi_i(L)$ for some R_i . Hence, lemma 4.2 rules out the possibility of the optimal policy being fully revealing about all states. Yet, it may well be the case that the optimal policy is fully revealing about the states of some voters. Indeed, we present in example 2 an optimal policy for which $\pi_i(L) = 0$ for some R_i .

Example 2 (Full revelation to some voter). *Let f be: $f(HHH) = 6/25$, $f(HHL) = 1/250$, $f(HLL) = 7/750$, and $f(LLL) = 18/25$. The thresholds are $(\ell_1, \ell_2, \ell_3) = (41, 40, 39)$. Based on lemma 4.1, the sender chooses $(\pi_i(L))_{i=1}^3$ to maximize the probability of unanimous approval subject to IC^a constraints. The solution is:*

$$(\pi_1(L), \pi_2(L), \pi_3(L)) = (0, 0.606, 0.644).$$

The most demanding voter R_1 learns her state fully. The project is never approved when $\theta_1 = L$, so R_1 's IC^a is slack. R_1 takes the role of a very accurate veto player: the more lenient voters depend on R_1 to veto a bad project. The sender could recommend the low-state R_1 to approve more frequently, but he optimally chooses to fully reveal θ_1 so that he can persuade R_2 and R_3 more effectively.

Example 2 highlights the crucial role of the most demanding voter(s): they serve as information guards vis-a-vis the more lenient voters who either partially or fully rubber-stamp the project. This example also suggests a monotonicity feature of the optimal policy $(\pi_i(L))_{i=1}^n$ with respect to the threshold: for $\ell_i > \ell_j$, $\pi_i(L) \leq \pi_j(L)$. More demanding voters are less likely to receive an approval recommendation when their states are low. Hence, the approval recommendations made to more demanding voters are more informative of their respective states being H . Any voter is more optimistic about her state being high when taking into account the approval of another more demanding voter than the approval of a more lenient voter. The following proposition establishes this monotonicity property.

Proposition 4.2 (Monotonicity of persuasion). *Suppose that f is exchangeable and affiliated. There exists an optimal policy in which more demanding voters' policies are more informative:*

$$\pi_i(L) \leq \pi_{i+1}(L) \text{ for all } i \in \{1, \dots, n-1\}.$$

Moreover, in any optimal policy in which R_i 's IC^a constraint binds, those who are more demanding than R_i must have strictly more informative policies, i.e., $\pi_j(L) < \pi_i(L)$ for all $j < i$.

Proposition 4.2 states that the sender essentially divides the group into (at most) three subgroups. The most demanding voters learn their states fully. The intermediate voters are partially manipulated. The most lenient voters rubber-stamp the collective decision.

To prove the first part of proposition 4.2, we show that if there is a pair of voters for which the more lenient voter enjoys a more informative policy, we can swap the individual policies between the two. The new policy remains incentive-compatible. Intuitively, when the stricter voter is compliant with a less informative policy, she continues to comply when assigned the more informative policy. After the swap, the stricter voter's belief that her state is H is weakened when conditioning only on the approval of all other voters, but the more accurate information acquired through her own policy offsets this increased pessimism.²⁵ After the swap, the more lenient voter is assigned the less informative policy accorded previously to the stricter voter. Since the stricter voter was willing to comply with an approval recommendation from this policy, the more lenient voter is willing to do so as well. Therefore, there always exists an optimal policy such that stricter voters have more informative policies. The second part of proposition 4.2 follows naturally: if IC^{a-i}

²⁵After all, each voter learns about her own state indirectly from the others' approval decisions but she learns about her own state directly from her own policy.

binds for some R_i , any voter who is more demanding than R_i must have a strictly more informative policy. Otherwise, this stricter voter's IC^a is not satisfied.

Based on the monotonicity property and the previous examples, a natural conjecture is that any voter with a slack IC^a constraint either fully learns her state or rubber-stamps. This conjecture is not true, as the following example 3 demonstrates. Nonetheless, among those voters who are partially informed (i.e., $\pi_i(L) \in (0, 1)$), at most one of them has a slack IC^a . If such a voter exists, she must be the strictest among those voters with partially informative policies. We state this result formally in proposition 4.3 below.

Example 3 (Slack IC^a with interior $\pi_i(L)$). *Let f be the same as in example 2 and the threshold profile be $(31, 30, 24)$. The optimal policy is given by:*

$$(\pi_1(L), \pi_2(L), \pi_3(L)) = (0.003, 0.516, 1).$$

Only IC^a -1 is slack. The sender provides very precise information to R_1 so as to be able to recommend approval more frequently to R_2 and R_3 . Decreasing $\pi_1(L)$ further does not benefit the sender once the recommendation to R_3 becomes fully uninformative, i.e., $\pi_3(L) = 1$. \square

Proposition 4.3. *Suppose distribution f is exchangeable and satisfies strict affiliation for any 3-voter subgroup. Among those voters who are partially informed, at most one has a slack IC^a constraint. If such a voter exists, she is the strictest voter among those who are partially informed.*

Taken together, propositions 4.2 and 4.3 establish that only the extreme voters might obtain a positive payoff from persuasion: the most demanding voters due to their role as informational guards, and the least demanding voters due to their willingness to rubber-stamp.

We conclude this section by discussing the proof of proposition 4.3. Suppose, to the contrary, that we can find two partially informed voters with thresholds $\ell_i < \ell_j$ whose IC^a constraints are slack. From proposition 4.2, it is without loss that $\pi_i(L) \geq \pi_j(L)$. Due to the slack IC^a - i and IC^a - j , the sender can slightly increase $\pi_i(L)$ and decrease $\pi_j(L)$ so that, given R_i and R_j 's approvals, every other voter's posterior belief of her respective state being H is at least as high as prior to the change. Because $\pi_i(L)$ is greater than $\pi_j(L)$ to start with, the boost in another voter's belief from a lower $\pi_j(L)$ has a stronger effect than the drop in the belief from a higher $\pi_i(L)$. Therefore, only a small decrease in $\pi_j(L)$ is required to offset the change in $\pi_i(L)$. Importantly, this necessary decrease is sufficiently small so that the sender's payoff strictly improves from this policy perturbation.

Hence, at most one partially informed voter has a slack IC^a . The second part of proposition 4.3 strengthens this observation: if some partially informed voter receives a positive payoff, she must be the strictest voter among all those who are partially informed. If the sender ever provides to some voter more precise information than what is required by her IC^a constraint, he prefers to do so with a voter who will necessarily be assigned a low $\pi_i(L)$. Reducing an already low $\pi_i(L)$ generates a stronger optimism boost among other voters.

4.3 When do some voters fully learn their states?

Individual persuasion introduces the possibility that the strictest voter(s) learn their own states fully. This stands in contrast to general persuasion. Lemma 2.2 established that no voter fully learns her state under general persuasion. The following discussion identifies necessary conditions for full revelation to the strictest voters to be optimal under individual persuasion. More importantly, we characterize the parameter region under which full revelation to some voters arise for a subclass of state distributions.

First, the group size must be at least three. We show in online appendix B.4 that when facing two voters, the sender sets $\pi_1(L)$ and $\pi_2(L)$ as high as their IC^a constraints permit; he never fully reveals the state to any voter. When the sender provides more precise information to some voter than what this voter's IC^a constraint requires, the probability that this voter approves the project is reduced. On the other hand, the sender can persuade the other voters to approve in their low state more frequently. This information externality is a public good so the benefit is larger for a larger group. The group size has to be at least three for the sender to find it worthwhile to provide more precise information than necessary.

Secondly, when the states are sufficiently independent or correlated, the sender does not fully reveal to any voter her state. Propositions 4.4 and 4.5 establish this fact. If the states are sufficiently independent, the information externality is not sufficiently strong for the benefits of full revelation to offset its cost in terms of forgone approval probability. On the other hand, if the states are sufficiently correlated, full revelation is not necessary: due to the strong correlation, the sender does not need to reduce by much the probability with which the strictest voter approves in her low state before the other voters are willing to rubber-stamp her decision.²⁶

²⁶In a sense, proposition 4.4 establishes a conceptual continuity between (i) full revelation to some voter for imperfectly correlated states, and (ii) lack of full revelation to any voter for perfectly correlated states.

Proposition 4.4 (No full revelation if states are sufficiently correlated). *For a fixed threshold profile $(\ell_i)_{i=1}^n$, there exists a critical degree of correlation above which full revelation to any voter is not optimal:*

$$(\pi_j(H), \pi_j(L)) \neq (1, 0) \text{ for any } j.$$

Proposition 4.5 (No full revelation if states are sufficiently independent). *For sufficiently independent states, all IC^a constraints bind in the optimal individual policy.*

Lastly, we fully characterize the parameter region under which some voters learn their own states fully for a relatively broad class of state distributions and a group in which all voters have the same threshold $\ell > 1$.²⁷

Each distribution in this class is parametrized by $\lambda_1 \in [1/2, 1]$, which measures the degree of correlation among the voters' states. More specifically, let there be a grand state $\omega \in \{G, B\}$, for which $\Pr(\omega = G) = p_0$. The state of each voter is drawn conditionally independently according to the following probabilities: $\Pr(H|G) = \Pr(L|B) = \lambda_1$. In particular, $\lambda_1 = 1/2$ corresponds to independent payoff states, and $\lambda_1 = 1$ to perfectly correlated payoff states. Without loss, we assume $\pi_i(L) \leq \pi_j(L), \forall i < j$.

Figure 1 summarizes the optimal policy when the group size n is sufficiently large. When the states are sufficiently independent (i.e. $\lambda_1 \in (1/2, \ell/\ell + 1]$), the optimal policy is symmetric across all voters (Proposition B.4). No voter learns her state fully for such low λ_1 . Moreover, every voter's IC^a constraint binds. If the states are sufficiently correlated (i.e., $\lambda_1 \in (\lambda_1^*, 1)$), the sender provides one voter with more precise information than her IC^a constraint requires so that every other voter is willing to rubber-stamp (Proposition B.2). Due to sufficiently high correlation, the better informed voter does not have to learn her state fully before the others are willing to rubber-stamp. When the state correlation is intermediate (i.e. $\lambda_1 \in (\ell/\ell + 1, \lambda_1^*]$), the optimal policy is to have some voters who are fully revealed their own states, one partially informed voter, and all other voters as rubber-stampers (Proposition B.3). This analysis shows that full revelation to some voters is optimal for a broad range of parameter values.

5 Comparison of persuasion modes

In this subsection we compare the different persuasion modes analyzed in sections 2 through 4. Proposition 3.1 has shown that independent general persuasion and

²⁷The formal analysis and results for this case can be found in the online appendix B.5.




Figure 1: Optimal individual persuasion with homogeneous thresholds

general persuasion are equivalent under unanimity. Therefore, we focus on the comparison of general and individual persuasion.

Under both general and individual persuasion, more voters hurt the sender. We show that, under both modes, any approval probability that can be achieved by n voters can also be achieved after a voter is removed from the group. Hence, the sender is weakly better off after the removal of a voter. Intuitively, due to the required unanimity for a collective approval, a greater number of voters can only hurt the probability of a unanimous approval.

Lemma 5.1 (More voters hurt the sender). *For any fixed threshold profile, the sender attains a weakly higher payoff with $(n - 1)$ voters than with n voters under both general and individual persuasion.*

When the states are sufficiently correlated, the sender attains the same payoff under general and individual persuasion. Under the latter, the sender persuades only the most demanding voter R_1 , whereas all other more lenient voters are willing to rubber-stamp R_1 's decision. This obviously induces the highest attainable payoff for the sender since persuading more voters than just R_1 can only leave him worse off. Therefore, general persuasion cannot improve upon the payoff from individual persuasion. There exists an optimal policy under general persuasion such that the approval probability in each state profile is the same across the two modes.

Lemma 5.2 (Equivalence with sufficiently high correlation). *Given a perfect correlation distribution f' , there exists $\varepsilon > 0$ such that for any f with $\|f - f'\| < \varepsilon$, the approval probability for any $\theta \in \Theta$ is the same across the two modes, i.e.,*

$$\pi(\theta) = \prod_{i=1}^n \pi_i(\theta_i),$$

where π and $(\pi_i)_{i=1}^n$ denote, respectively, an optimal general policy and an optimal individual policy associated with f .

As the correlation among the states weakens, the sender does strictly better under general persuasion than under individual persuasion. The sender is able to pool different state profiles more freely under general persuasion in order to obtain a higher payoff. In contrast, the strictest voter always weakly prefers individual persuasion. This is because any optimal general policy accords the strictest voter a zero payoff, while the optimal individual policy might accord her a strictly positive payoff. Moreover, if IC^a - i binds for $i \in \{1, \dots, i'\}$ in any optimal general policy, the strictest voters $\{R_1, \dots, R_{i'}\}$ are better off with individual persuasion.

While the strictest voter(s) unequivocally prefer individual persuasion, the ranking of the two persuasion modes by more lenient voters can go either way. For instance, in example 1, R_3 obtains a positive payoff under general persuasion when the sender chooses an optimal and Pareto efficient policy. Due to independent states, R_3 's payoff must be zero under individual persuasion. In this example, the more lenient voter prefers general persuasion. We can find another example in which the most lenient voter obtains a strictly higher payoff under individual persuasion than under general persuasion, even if we restrict attention to optimal general policies that are Pareto efficient. As a result, the more lenient voters might disagree on the preferred persuasion mode.

6 Extensions

6.1 Public/sequential persuasion and observability of policies

This subsection discusses the robustness of the results as we relax two assumptions of the benchmark model: the private observability of recommendations, and the simultaneous structure of voting. We argue that the relaxation of either assumption does not affect the results from sections 2 through 5. Moreover, we briefly discuss the implications of privately observed policies under individual persuasion.

Let us first assume that the recommendations are announced publicly. Given any optimal policy under the previous assumption of private communication, the voting game that follows after voters publicly observe the same recommendations admits an obedient equilibrium in which all voters follow their respective recommendations. This is true for any of the three persuasion modes. If the public recommendation profile is the unanimous recommendation \hat{d}^a , the previous IC^a constraints ensure that all voters comply with the approval recommendation. If $\hat{d}^{r,i}$ is realized instead, the previous IC^r - i ensures that R_i prefers to reject. The other voters comply as well

since they are no longer pivotal. If two or more voters receive a recommendation to reject, no voter can overturn the collective rejection given that all other voters follow their individual recommendations.²⁸ The following proposition summarizes this reasoning.

Proposition 6.1 (Public persuasion). *Fix an optimal general or individual policy. Suppose that the recommendation profile \hat{d} is observable by all voters. There exists an equilibrium in which each voter complies with the recommendation.*

Suppose now that the signals are privately observed but the sender encounters the voters sequentially in a particular order, one at a time. Each voter perfectly observes the decisions of preceding voters. We analyze how well the sender performs, compared to simultaneous voting, and whether the payoff achieved by the sender depends on the particular order in which the voters are approached.

Under unanimity rule, the decision of R_i is significant for the collective decision only if all preceding voters have already approved and all succeeding voters will approve as well. This observation—that in sequential voting too each voter decides as if all other voters have already approved the project—suggests that sequential voting is equivalent to simultaneous voting.²⁹ Indeed, the set of incentive-compatible policies under sequential voting is the same as that under simultaneous voting. Moreover, the order in which the voters are encountered is immaterial for the sender’s payoff.

Proposition 6.2 (Sequential persuasion). *Fix an optimal general or individual policy with simultaneous voting. This policy is also optimal under sequential voting for any order of voters.*

Furthermore, it follows from propositions 6.1 and 6.2 combined that the set of incentive-compatible policies remains the same if we relax the assumptions of private communication and simultaneous voting concurrently.

Finally, we discuss how results change if under individual persuasion, each voter privately observes her own policy. First of all, if we use Nash equilibrium as our solution concept, the optimal policy stays the same. This is because if the sender deviates and alters the equilibrium policy for R_i , then R_i will reject for sure. However, rejecting for sure on R_i ’s side regardless of the deviation policy and the realized signal is not sequentially rational. After the deviation, R_i needs to form a belief

²⁸Moreover, if two or more voters receive a recommendation to reject, it cannot be the case that all voters prefer to approve. If that were so, the sender could recommend that all voters approve under private communication and increase his payoff. This contradicts the presumption that the policy is optimal under private communication.

²⁹The reasoning clearly does not extend to nonunanimous voting rules.

about the sender's policies toward other voters and their behavior, in order to form a belief about their respective states. The most pessimistic belief R_i can hold is that all other voters' states are L upon approval. If the deviation policy toward R_i is incentive compatible under this most pessimistic belief, then R_i is willing to follow the recommendation. This gives an upper bound on how informative R_i 's equilibrium policy can be. If $\pi_i(H) = 1$, this gives a positive lower bound on how low $\pi_i(L)$ can be in order for the policy to be sustained in equilibrium. The sender's problem is essentially the same as before except that each $\pi_i(L)$ now must stay above its lower bound.

By this reasoning, the optimal policy identified in example 2 is not an equilibrium. The sender prefers to deviate from the fully revealing policy π_1 and increase $\pi_1(L)$ by a very small amount: for any belief R_1 might hold, there exists a sufficiently small $\pi_1(L)$ for which R_1 continues to obey the recommendation. Full revelation does not arise with privately observed policies.

However, the optimal policy identified in example 3 is credible under the most pessimistic off-path beliefs described above. R_1 obeys a recommendation of an off-path policy $\tilde{\pi}_1$ with $\tilde{\pi}_1(H) = 1$ if and only if $\tilde{\pi}_1(L) < f(HLL)/(\ell_1 f(LLL)) = 0.00042$. This imposes a lower bound on any credible $\pi_1(L)$. Hence, the sender does not deviate from the optimal policy $\pi_1^*(L) = 0.003$. Similarly, the sender prefers not to deviate from π_2^* and π_3^* as well. This example shows that even with privately observed individual policies, there are environments $(f, (\ell_i)_{i=1}^n)$ for which the sender provides to the strictest voters more precise information than what is needed to persuade them in order to persuade others more often.

6.2 Nonunanimous decision rules

This subsection considers nonunanimous voting rules, i.e., when $k < n$ votes suffice for project adoption. We show that the project is approved with certainty under general persuasion, and with probability strictly less than one under individual persuasion. Voters, on the other hand, strictly prefer individual to general persuasion.

A general policy that trivially achieves a certain approval is one that always recommends all voters to approve for all state profiles. No voter is ever pivotal, so each voter trivially follows the approval recommendation. This construction, however, relies on the failure to be pivotal, so each voter is indifferent when asked to approve. We take a different route.

As noted in section 1, we allow only for any policy that is the limit of a sequence of full-support incentive-compatible policies. For each limiting policy, there exists a nearby full-support policy for which each voter is pivotal with positive probabil-

ity for both approval and rejection recommendations, and each voter’s incentive-compatibility constraints are strictly satisfied.³⁰ We can construct a sequence of such full-support policies that achieves a payoff arbitrarily close to one for the sender. This shows that the certain-approval result does not rely on the voters’ failure to be pivotal.

This certain-approval result stands in contrast to the optimal policy under unanimity where the sender’s payoff is strictly below one. The key observation is that the sender benefits from the event that *at least* k voters approve, whereas each voter cares about the event that *exactly* k voters (including herself) receive an approval recommendation. When $k < n$, the sender can recommend most of the time that at least $k + 1$ voters approve, without jeopardizing the IC constraints of the voters. This observation explains the discontinuity between the unanimous and nonunanimous rules.

Proposition 6.3 (Certain approval under general policies).³¹ *Suppose that the sender needs $k \leq n - 1$ approvals. Under general persuasion, the sender’s payoff is one.*

$\theta \backslash \hat{d}$	(0,0)	(1,0)	(0,1)	(1,1)
<i>HH</i>	ε^2	ε	ε	$1 - 2\varepsilon - \varepsilon^2$
<i>HL</i>	ε^2	ε^2	ε^2	$1 - 3\varepsilon^2$
<i>LH</i>	ε^2	ε^2	ε^2	$1 - 3\varepsilon^2$
<i>LL</i>	ε	ε^2	ε^2	$1 - \varepsilon - 2\varepsilon^2$

Table 2: General policy for $n = 2, k = 1$

To illustrate the proof idea, we construct a full-support policy for $n = 2$ and $k = 1$ in table 2. Each row summarizes the recommendation distribution for a fixed state profile. Upon receiving an approval recommendation, R_1 is pivotal only if recommendation $(1, 0)$ has realized. This recommendation is sent much more frequently in state profile *HH* than in other state profiles; therefore, R_1 is confident that her state is *H* when conditioning on being pivotal. The reasoning for R_2 is similar. Therefore, the constructed policy guarantees compliance with approval recommendations. On the other hand, the only recommendation for which R_i is pivotal upon being recommended to reject is $(0, 0)$. Since this recommendation

³⁰Suppose that the sender does not provide any information and voters do not use weakly dominated strategies. For any voting rule $1 \leq k \leq n$, all voters reject in the unique correlated equilibrium.

³¹The proofs for propositions 6.3 and 6.4 can be found in online appendix B.1. Chan et al. (2016) independently reach a similar result to proposition 6.3 for the case of perfectly correlated states.

profile is suggestive of the unfavorable state profile LL more than of any other state profile, both voters are willing to follow a rejection recommendation. Therefore, this sequence of policies (indexed by ε) are obedient and attain a payoff for the sender that is arbitrarily close to one.

The certain-approval result is further strengthened by the observation that the sender achieves a certain approval even when constrained to independent general persuasion.

Proposition 6.4 (Certain approval under independent general policies). *Suppose that f has full support and the sender needs $k \leq n - 1$ approvals. Under independent general persuasion, the sender's payoff is one.*

To illustrate the proof idea, we revisit the two-voter example with $k = 1$. The left-hand side of table 6.2 demonstrates the independent policy for each voter. The right-hand side shows the corresponding probability of each recommendation profile for each state profile.³²

$\theta \backslash \pi_i$	$\pi_1(\cdot)$	$\pi_2(\cdot)$
HH	$1 - \varepsilon_1$	$1 - \varepsilon_1$
HL	$1 - \varepsilon_2$	ε_3
LH	ε_3	$1 - \varepsilon_2$
LL	$1 - \varepsilon_4$	$1 - \varepsilon_4$

$\theta \backslash \hat{d}$	$(0,0)$	$(1,0)$	$(0,1)$
HH	ε_1^2	$\varepsilon_1(1 - \varepsilon_1)$	$\varepsilon_1(1 - \varepsilon_1)$
HL	$\varepsilon_2(1 - \varepsilon_3)$	$(1 - \varepsilon_2)(1 - \varepsilon_3)$	$\varepsilon_2\varepsilon_3$
LH	$\varepsilon_2(1 - \varepsilon_3)$	$\varepsilon_2\varepsilon_3$	$(1 - \varepsilon_2)(1 - \varepsilon_3)$
LL	ε_4^2	$\varepsilon_4(1 - \varepsilon_4)$	$\varepsilon_4(1 - \varepsilon_4)$

Table 3: Independent general policy for $n = 2, k = 1$

In state profiles with exactly one high-state voter, this voter is recommended to approve with very high probability, while the other low-state voter is recommended to reject with very high probability. Hence, when R_1 is recommended to approve and she conditions on $(1, 0)$ being sent, the state profile is most likely to be HL . Because HL is favorable for R_1 , she is willing to approve. By the same logic, upon being recommended to reject, R_1 conditions on $(0, 0)$ being sent. If ε_4 is chosen so as to shrink to zero at a much slower rate than ε_1 and ε_2 , R_1 is sufficiently assured that the state profile is the unfavorable LL . Hence, she obeys the rejection recommendation. The reasoning for R_2 is similar. Therefore, the policies in the sequence are obedient and attain an arbitrarily high payoff for the sender.

More generally under independent general persuasion, whenever a state profile with exactly k high-state voters realizes, the high-state voters are very likely to be recommended to approve and the low-state voters very likely to be recommended to reject. Therefore, when R_i is recommended to approve, the pivotality of her decision

³²The distribution for recommendation profile $(1, 1)$ has been omitted for ease of exposition.

suggests that the realized state profile very probably admits exactly k high-state voters, with R_i being among them. This construction has the flavor of “targeted persuasion” as in Alonso and Câmara (2016b), since the sender targets the voters who benefit from the project. By making the state profiles with exactly k high-state approvers the most salient event for each voter when her vote is pivotal, the sender is able to achieve a guaranteed approval.

Our construction relies on the existence of state profiles with $k > 0$ high-state voters and $(n - k) > 0$ low-state voters. So such state profiles have to occur with positive probability in order for the argument to hold. Lemma B.1 in online appendix B.1 shows that for any affiliated and exchangeable state distribution f , either (i) f has full support on Θ , or (ii) the states are perfectly correlated. If the latter holds, the problem of the sender reduces to one of individual persuasion, under which a certain approval is never attained, as shown in the next proposition. Hence, a full-support state distribution is both necessary and sufficient for certain approval under independent general persuasion.

Whenever each voter’s recommendation can depend on the entire profile θ , a requirement for nonunanimous consent effectively imposes no check on the adoption of the project, as approval is guaranteed. We next show that when the sender can condition individual recommendations to a voter only on her state, the project is never approved for sure.

Proposition 6.5 (No certain approval under individual policies). *Suppose that the sender needs $k \leq n - 1$ approvals. Under individual persuasion, the sender’s payoff is strictly below one. The payoff of each voter is strictly higher than that under general and independent general persuasion.*

Under individual persuasion, each voter approves more frequently under a high state than a low state. If the project is approved with certainty, it must be approved with certainty under all possible state profiles. There must then exist a coalition of at least k voters who approve the project regardless of their individual states. However, any voter in this coalition does not become more optimistic about her state being high when she receives a recommendation to approve and considers her decision to be pivotal. Due to assumption 1, she is not willing to follow the approval recommendation. Hence, the project cannot be approved with certainty.

We further establish that there exists at least one voter who approves strictly more frequently under a high state than under a low state and she is pivotal with positive probability. Due to affiliation, all other voters benefit from such a voter’s decision. Therefore, for any nonunanimous voting rule, all voters are strictly better

off under individual persuasion than under either general or independent general persuasion.

The role of communication. For a fixed information structure π or $(\pi_i)_{i=1}^n$, a stage of pre-voting communication can be added to this voting game of incomplete information. The set of equilibria depends on the communication protocol. For instance, we can add one round of public cheap-talk communication: all voters first observe their own private signal and then simultaneously announce a public cheap-talk message. Afterwards they vote. For this particular communication protocol, there is always an equilibrium in which each voter votes informatively based on her private signal generated according to π or $(\pi_i)_{i=1}^n$.³³ This is clearly the optimal equilibrium for the sender among all possible communication equilibria.³⁴ Moreover, if all voters *share a consensus* given π or $(\pi_i)_{i=1}^n$, this communication protocol admits also an equilibrium in which all voters announce their signals truthfully and vote according to the aggregated information.³⁵ This might lead to a lower payoff for the sender than the previous equilibrium. In general, when voters do not share such a consensus, there might not exist such a full-revelation equilibrium.

To sum up, for any information structure π or $(\pi_i)_{i=1}^n$, the equilibrium outcome when the voters can communicate depends not only on the communication protocol but also on the equilibrium selection criterion. We view our result as an important and useful benchmark when the sender-optimal equilibrium is selected.

7 Concluding remarks

This paper analyzed the problem faced by an uninformed sender aiming to persuade a unanimity-seeking group of heterogeneous voters through the design of informative experiments. We characterized the optimal policies for two main modes of persuasion, —general and individual persuasion—and explored their implications for the players’ welfare. Returning to our motivating examples, our results clarify who benefits from particular forms of evidence: the most demanding regulators prefer targeted evidence, while comprehensive evidence benefits, besides the sender,

³³There is always a babbling equilibrium in which each voter sends an uninformative message and then votes according to her private signal.

³⁴In fact, for any direct obedient π or $(\pi_i)_{i=1}^n$, there exists a communication equilibrium in which voters follow their private signals. See Gerardi and Yariv (2007) for details.

³⁵As defined in Coughlan (2000) and Austen-Smith and Feddersen (2006), a consensus exists if, given the public revelation of all private information, all voters always agree on whether to approve. If voters do not share such a consensus, a full-revelation equilibrium might not exist. Austen-Smith and Feddersen (2006) present an example in which voters have the same state but different thresholds. Given the information structure in that example, there is no full-revelation equilibrium.

only the least demanding ones. Moreover, providing evidence that is comprehensive yet independent across regulators does not constrain the sender at all in his ability to persuade. When restricted to targeted evidence, —for instance, in highly specialized regulatory evaluations, — the sender might find it optimal to accurately inform the strictest regulators of the quality of their aspect of interest. An institutional arrangement that is based on targeted evidence designates a strong informational role to the most demanding regulators: the sender leverages the fully revealing recommendations provided to them to persuade others more easily. When the voters' states are affiliated and exchangeable, the possibility of full revelation is exclusive to targeted evidence.

Under unanimity rule, the characterized optimal policies remain optimal even when the voters publicly observe the recommendation profile and/or take turns to cast their votes. For nonunanimous rules, the sender obtains a sure approval if allowed to offer a general policy or an independent general policy to the group. In the context of regulation, this is effectively as if there were no regulatory check on the proposals of the industry as long as it can provide comprehensive evidence regarding all payoff-relevant aspects. We interpret this result as advocating for the institutionalization of unanimity rule whenever the sender (i) communicates with voters in private, and (ii) offers comprehensive evidence. Moreover, the fact that all voters prefer individual to general persuasion under any nonunanimous rule supports the institutionalization of targeted evidence in environments governed by such rules.

Our analysis of unanimity rule invites future work on the full characterization of optimal individual persuasion for nonunanimous rules. We have taken a first step by analyzing some features of individual persuasion under such rules. The effect of communication among voters on the persuasion effort of the sender also remains largely unexplored. Another natural question concerns the implications of sequential persuasion in settings with and without payoff externalities among receivers. For sequential collective decision-making under different persuasion modes, a departure from the unanimity rule complicates the analysis substantially. Moreover, we leave for future work a systematic examination of differences among persuasion modes when the receivers vary in their informational importance for the group, that is, when the assumption of exchangeability is dropped. The exploration of these questions promises to shed further light on the dynamics of group persuasion.

A Appendix: Proofs for sections 2 to 4

Proof for proposition 2.1. The policy only specifies $\pi(\hat{d}^a|\theta^H)$ and $\pi(\hat{d}^a|\theta^L)$. First, $\pi(\hat{d}^a|\theta^H) = 1$. Otherwise the objective can be improved without hurting any IC^a constraints. Secondly, in order for $\pi(\hat{d}^a|\theta^L)$ to be as high as possible without hurting any IC^a constraints, the sender optimally sets it so as to make IC^a-1 bind (because $\ell_1 = \max_i \ell_i$). Hence, $f(\theta^H) - f(\theta^L)\pi(\hat{d}^a|\theta^L)\ell_1 = 0$. \square

Proof for lemma 2.1. Suppose that $\pi(\hat{d}^a|\theta^H) < 1$. Increasing $\pi(\hat{d}^a|\theta^H)$ relaxes IC^a- i for all i as $\theta^H \in \Theta_i^H$ for all i . It also strictly improves the payoff of the sender. Hence $\pi(\hat{d}^a|\theta^H) < 1$ cannot be optimal. \square

Proof for lemma 2.2. Suppose to the contrary that there exists i such that $\pi(\hat{d}^a|\theta) = 1$ for all $\theta \in \Theta_i^L$. Then IC^a- i requires that $\left(\sum_{\theta \in \Theta_i^H} f(\theta)\pi(\hat{d}^a|\theta)\right) / \left(\sum_{\theta \in \Theta_i^L} f(\theta)\right) \geq \ell_i$. But assumption 1 tells us that

$$\ell_i > \frac{\sum_{\theta \in \Theta_i^H} f(\theta)}{\sum_{\theta \in \Theta_i^L} f(\theta)} \geq \frac{\sum_{\theta \in \Theta_i^H} f(\theta)\pi(\hat{d}^a|\theta)}{\sum_{\theta \in \Theta_i^L} f(\theta)}$$

for any specification of $\pi(\hat{d}^a|\theta)$ for $\theta \in \Theta_i^H$. We have thus reached a contradiction.

Suppose to the contrary that for some i , $\pi(\hat{d}^a|\theta) = 0$ for all $\theta \in \Theta_i^L$. Then IC^a- i is slack. Consider $\theta' \in \Theta_i^L$ such that $\theta'_j = H$ for all $j \neq i$. Consider increasing $\pi(\hat{d}^a|\theta')$ to some strictly positive value so as to still satisfy IC^a- i . This increase: i) strictly improves the payoff of the sender, ii) relaxes all IC^a- j for $j \neq i$. \square

Proof for proposition 2.2. Slightly reformulated, the sender's problem is

$$\begin{aligned} & \min_{\pi(\hat{d}^a|\theta) \geq 0} - \sum_{\theta \in \Theta} f(\theta)\pi(\hat{d}^a|\theta) \\ \text{s.t. } & \pi(\hat{d}^a|\theta) - 1 \leq 0, \forall \theta, \quad \text{and} \quad \sum_{\theta \in \Theta_i^L} f(\theta)\pi(\hat{d}^a|\theta)\ell_i - \sum_{\theta \in \Theta_i^H} f(\theta)\pi(\hat{d}^a|\theta) \leq 0, \forall i. \end{aligned}$$

Let γ_θ be the dual variable associated with $\pi(\hat{d}^a|\theta) - 1 \leq 0$ and μ_i the dual variable associated with IC^a- i . Since the constraints of the primal are inequalities, for the dual we have the constraints that $\gamma_\theta \geq 0$ for all θ and $\mu_i \geq 0$ for all i . The dual is:

$$\min_{\gamma_\theta \geq 0, \mu_i \geq 0} \sum_{\theta \in \Theta} \gamma_\theta, \quad \text{s.t.} \quad \gamma_\theta \geq f(\theta) \left(1 + \sum_{i:\theta_i=H} \mu_i - \sum_{i:\theta_i=L} \mu_i \ell_i \right), \quad \forall \theta \in \Theta.$$

In the dual problem, there is an inequality constraint for each state profile θ . The associated primal variable for each inequality constraint is $\pi(\hat{d}^a|\theta)$.

We first show that it cannot be that $\mu_i = 0$ for all i . Suppose that was indeed the case. Then $\sum_{\theta \in \Theta} \gamma_\theta = \sum_{\theta \in \Theta} f(\theta) = 1$. Consider increasing μ_1 by a small

amount $\varepsilon > 0$. The dual objective changes by $\varepsilon \sum_{\theta \in \Theta_1^H} f(\theta) - \varepsilon \ell_1 \sum_{\theta \in \Theta_1^L} f(\theta) < 0$ by assumption 1. Hence the dual objective can be improved.

We next show that there exists an optimal solution to the dual such that if $\mu_j = 0$, then $\mu_{j'} = 0$ for any $j' > j$. Suppose that $\mu_j = 0$ and $\mu_{j'} > 0$ for some $j' > j$. We can rewrite the constraint associated with θ as:

$$\gamma_\theta \geq f(\theta) \left(1 + \sum_{\substack{i \neq j, j' \\ \theta_i = H}} \mu_i - \sum_{\substack{i \neq j, j' \\ \theta_i = L}} \mu_i \ell_i + \sum_{i \in \{j, j'\}} \mu_i (\mathbf{1}_{\theta_i = H} - \mathbf{1}_{\theta_i = L} \ell_i) \right).$$

We construct a new set of $\tilde{\mu}_i$ and $\tilde{\gamma}_\theta$ such the the dual objective stays constant and the dual constraints are satisfied. We let $\tilde{\mu}_{j'} = \mu_j = 0$, $\tilde{\mu}_j = \mu_{j'} > 0$ and $\tilde{\mu}_i = \mu_i$ for $i \neq j, j'$. To construct $\tilde{\gamma}_\theta$, we consider two cases:

1. If the state profile θ is such that $\theta_j = \theta_{j'}$, then we let $\tilde{\gamma}_\theta = \gamma_\theta$. Since $\ell_j > \ell_{j'}$ and $\tilde{\mu}_j > \tilde{\mu}_{j'} = 0$, the RHS of the inequality associated with $\tilde{\gamma}_\theta$ is weakly lower than before, so the constraint is satisfied.
2. If the state profile θ is such that $\theta_j \neq \theta_{j'}$, then there exists another state profile θ' such that θ' is the same as θ except for θ'_j and $\theta'_{j'}$. In this case, we let $\tilde{\gamma}_\theta = \gamma_{\theta'}$ and $\tilde{\gamma}_{\theta'} = \gamma_\theta$. The inequalities associated with $\tilde{\gamma}_\theta$ and $\tilde{\gamma}_{\theta'}$ are both satisfied. This is easily verified given that $\ell_j > \ell_{j'}$, $f(\theta) = f(\theta')$, and the presumption that $\mu_{j'} > \mu_j = 0$ (or equivalently, $\tilde{\mu}_j > \tilde{\mu}_{j'} = 0$).

This shows that there exists an optimal solution to the dual such that μ_i 's for the most demanding regulators are positive. Moreover, based on case (ii), if $\mu_{j'} > \mu_j = 0$ and $j' > j$, it must be true that $\gamma_\theta = 0$ for any θ such that $\theta_j = H, \theta_{j'} = L$. Otherwise, after the exchange operation in case (ii), we can lower γ_θ (or equivalently, $\tilde{\gamma}_{\theta'}$) without violating any constraint.

We next show that in *any* optimal solution of the dual, if $\mu_j = 0$, then $\mu_{j'} = 0$ for any $j' > j$. Suppose not. Suppose that $\mu_j = 0$ and $\mu_{j'} > 0$ for some $j' > j$. This implies that $\gamma_\theta = 0$ for any θ such that $\theta_j = H, \theta_{j'} = L$. In particular, $\gamma_\theta = 0$ when $\theta_j = H, \theta_{j'} = L$ and all other states are H . The constraint with respect to this state profile is $\gamma_\theta = 0 \geq f(\theta) \left(1 + \sum_{i \neq j, j'} \mu_i - \mu_{j'} \ell_{j'} \right)$. This imposes a lower bound on $\mu_{j'}$: $\mu_{j'} \geq \left(1 + \sum_{i \neq j, j'} \mu_i \right) / \ell_{j'}$. We next show that it is uniquely optimal to set $\mu_{j'}$ to be this lower bound. Firstly, this lower bound of $\mu_{j'}$ ensures that for any θ such that $\theta_{j'} = L$, we can set γ_θ to be zero. For any θ such that $\theta_{j'} = H$, lowering $\mu_{j'}$ makes the constraint associated with θ easier to be satisfied. In particular, for θ such that all states are H , lowering $\mu_{j'}$ strictly decreases the lower bound on γ_θ and thus strictly lower the dual objective. This shows that it is uniquely optimal to set $\mu_{j'}$ to be the lower bound $(1 + \sum_{i \neq j, j'} \mu_i) / \ell_{j'}$.

Lastly, we show that we can strictly improve by constructing a new set of $\tilde{\mu}_i$ and $\tilde{\gamma}_\theta$. We let $\tilde{\mu}_{j'} = 0$, $\tilde{\mu}_j = (1 + \sum_{i \neq j, j'} \mu_i) / \ell_j < (1 + \sum_{i \neq j, j'} \mu_i) / \ell_{j'}$ and $\tilde{\mu}_i = \mu_i$ for $i \neq j, j'$. For any θ such that $\theta_j = L$, we can set $\tilde{\gamma}_\theta$ to be zero. For any θ such that $\theta_j = H$, there exists a state profile θ' such that $\theta'_{j'} = \theta_j$, $\theta'_j = \theta_{j'}$, and $\theta'_i = \theta_i$ for $i \neq j, j'$. We let $\tilde{\gamma}_\theta = \gamma_{\theta'}$. Since $\tilde{\mu}_j$ is strictly smaller than $\mu_{j'}$, the constraint associated with $\tilde{\gamma}_\theta$ is satisfied. Moreover, the lower $\tilde{\mu}_j$ ensures that we can further lower $\tilde{\gamma}_{\theta^H}$ (which was equal to γ_{θ^H}). This shows that we can strictly decrease the objective. We have thus reached a contradiction.

Therefore, it cannot be true that $\mu_j = 0$ and $\mu_{j'} > 0$ for some $j' > j$. Moreover, at least one μ_i is strictly positive. This shows that there exists an $i' \geq 1$ such that $\mu_i > 0$ for $i \leq i'$. Therefore, in any optimal policy, the strictest regulators' IC^a constraints bind. \square

Proof for proposition 3.1. Consider a general policy $\pi = \left(\pi(\hat{d}|\theta) \right)_{\hat{d}, \theta}$. We want to construct an independent general policy $(\pi_i(\theta))_{i, \theta}$ that implements the same payoff for the sender. We construct $(\pi_i(\theta))_{i, \theta}$ such that $\prod_{i=1}^n \pi_i(\theta) = \pi(\hat{d}^a|\theta)$ for any $\theta \in \Theta$. IC^a constraints (2) are satisfied automatically since the right-hand side of (2) depends only on the product $\prod_{i=1}^n \pi_i(\theta)$. We next show that we can choose $(\pi_i(\theta))_{i, \theta}$ to satisfy IC^r constraints (3) as well. First, if $\pi(\hat{d}^a|\theta) = 0$, we set $\pi_i(\theta)$ to be zero for all i . Such state profiles contribute zero to the left-hand side of (3). Secondly, if $\pi(\hat{d}^a|\theta) = 1$, we have to set $\pi_i(\theta) = 1$ for all i . Again, such state profiles contribute nothing to the left-hand side of (3). Lastly, if $\pi(\hat{d}^a|\theta) \in (0, 1)$, we set $\pi_i(\theta)$ to be one if $\theta_i = H$ and $\pi_i(\theta)$ to be interior if $\theta_i = L$. Such state profiles contribute a weakly negative term to the left-hand side of (3). Thus, we have shown that (3) is satisfied. \square

Proof for proposition 4.1. Under perfect correlation, IC^a constraint can be rewritten compactly as: $f(\theta^H) \prod_{j=1}^n \pi_j(H) \geq f(\theta^L) \ell_i \prod_{j=1}^n \pi_j(L)$, while the objective is: $f(\theta^H) \prod_{j=1}^n \pi_j(H) + f(\theta^L) \prod_{j=1}^n \pi_j(L)$. Increasing $\prod_j \pi_j(H)$ relaxes all approval IC constraints and also benefits the sender, therefore the sender sets $\pi_j(H) = 1$ for all j . The sender's payoff increases in $\prod_j \pi_j(L)$, so he sets it to be the highest possible value allowed by the IC^a constraints: $\prod_{j=1}^n \pi_j(L) = f(\theta^H) / (f(\theta^L) \ell_1)$. This policy satisfies IC^r constraints as well. If R_i receives a rejection recommendation, she knows for sure that her state is L because $\pi_i(H) = 1$, hence she strictly prefers to reject the project.

If states are independent, R_i receives no additional information from conditioning on the approval of R_{-i} . IC^a- i is simply: $\sum_{\theta \in \Theta_i^H} f(\theta) \pi_i(H) - \sum_{\theta \in \Theta_i^L} f(\theta) \pi_i(L) \ell_i \geq 0$. The policy offered to one voter does not affect the sender's ability to persuade

other voters. So the problem of the sender is separable across voters. Therefore, the following single-voter policy is optimal: $\pi_i(H) = 1$, and

$$\pi_i(L) = \frac{\sum_{\theta \in \Theta_i^H} f(\theta)}{\ell_i \sum_{\theta \in \Theta_i^L} f(\theta)}.$$

□

Proof for lemma 4.1. Suppose $\pi_j(H) < 1$ for some j . It is easily verified that (OBJ) increases in $\pi_j(H)$, and increasing $\pi_j(H)$ weakly relaxes IC^a- j . It remains to be shown that increasing $\pi_j(H)$ weakly relaxes IC^a- i for any $i \neq j$. Consider IC^a- i rewritten in the following form:

$$\frac{\pi_j(H) \sum_{\Theta_i^H \cap \Theta_j^H} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k) + \pi_j(L) \sum_{\Theta_i^H \cap \Theta_j^L} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k)}{\pi_j(H) \sum_{\Theta_i^L \cap \Theta_j^H} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k) + \pi_j(L) \sum_{\Theta_i^L \cap \Theta_j^L} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k)} \geq \ell_i \frac{\pi_i(L)}{\pi_i(H)}.$$

For this step, note that if the denominator of the LHS equals zero, then IC^a- i holds for any $\pi_j(H)$, and if the denominator of the RHS equals zero, then $\pi_i(L)$ has to be zero so IC^a- i holds for any $\pi_j(H)$ as well. Therefore, we focus on the case in which neither denominator equals zero.

The derivative of the LHS with respect to $\pi_j(H)$ is positive if:

$$\left(\sum_{\Theta_i^H \cap \Theta_j^H} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k) \right) \left(\sum_{\Theta_i^L \cap \Theta_j^L} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k) \right) - \left(\sum_{\Theta_i^H \cap \Theta_j^L} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k) \right) \left(\sum_{\Theta_i^L \cap \Theta_j^H} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k) \right) \geq 0. \quad (4)$$

To prove this inequality, we will use the Ahlswede-Daykin inequality (Ahlswede and Daykin, 1978): Suppose (Γ, \succ) is a finite distributive lattice and functions $f_1, f_2, f_3, f_4 : \Gamma \rightarrow \mathbb{R}^+$ satisfy the relation that $f_1(a)f_2(b) \leq f_3(a \wedge b)f_4(a \vee b)$, $\forall a, b \in \Gamma$. Then $f_1(A)f_2(B) \leq f_3(A \wedge B)f_4(A \vee B)$, $\forall A, B \subset \Gamma$, where $f_k(A) = \sum_{a \in A} f_k(a)$ for all $A \subset \Gamma$, $k \in \{1, 2, 3, 4\}$, and $A \vee B = \{a \vee b : a \in A, b \in B\}$, $A \wedge B = \{a \wedge b : a \in A, b \in B\}$.

First, we let Γ be the set of state profiles excluding R_i and R_j . That is, $\Gamma = \Theta_{-ij}$. The lattice Γ is finite. Secondly, the lattice is distributive, i.e. for any $\theta, \theta', \theta'' \in \Theta_{-ij}$ $\theta_{-ij} \vee (\theta'_{-ij} \wedge \theta''_{-ij}) = (\theta_{-ij} \vee \theta'_{-ij}) \wedge (\theta_{-ij} \vee \theta''_{-ij})$. To see this, consider the state of $R_k \neq R_i, R_j$ in both sides. If $\theta_k = H$, then the state of R_k is H in both sides. If $\theta_k = L$ and $\theta'_k = \theta''_k$, then the state of R_k is θ'_k in both sides. If $\theta_k = L$, and $\theta'_k \neq \theta''_k$, the state of R_k is L in both sides. Therefore the lattice is distributive. We write $f(\theta) = f(\theta_i \theta_j, \theta_{-ij})$. Define functions:

$$f_1(\theta_{-ij}) = f(HL, \theta_{-ij})\pi_{-ij}(\theta_{-ij}), \quad f_2(\theta_{-ij}) = f(LH, \theta_{-ij})\pi_{-ij}(\theta_{-ij}),$$

$$f_3(\theta_{-ij}) = f(LL, \theta_{-ij})\pi_{-ij}(\theta_{-ij}), \quad f_4(\theta_{-ij}) = f(HH, \theta_{-ij})\pi_{-ij}(\theta_{-ij}),$$

where $\pi_{-ij}(\theta_{-ij}) := \prod_{k \neq i,j} \pi_k(\theta_k)$. Due to affiliation and the easily verified equality that $\pi_{-ij}(\theta_{-ij})\pi_{-ij}(\theta'_{-ij}) = \pi_{-ij}(\theta_{-ij} \wedge \theta'_{-ij})\pi_{-ij}(\theta_{-ij} \vee \theta'_{-ij})$, the premise of the theorem holds:

$$f(HL, \theta_{-ij})\pi_{-ij}(\theta_{-ij})f(LH, \theta'_{-ij})\pi_{-ij}(\theta'_{-ij}) \leq \\ f(LL, \theta_{-ij} \wedge \theta'_{-ij})\pi_{-ij}(\theta_{-ij} \wedge \theta'_{-ij})f(HH, \theta_{-ij} \vee \theta'_{-ij})\pi_{-ij}(\theta_{-ij} \vee \theta'_{-ij}).$$

Take subsets of the lattice $A = B = \Gamma$, so $A \wedge B = A \vee B = \Gamma$. It follows from Ahlswede-Daykin inequality that

$$\left(\sum_{\theta_{-ij} \in A} f(HL, \theta_{-ij})\pi_{-ij}(\theta_{-ij}) \right) \left(\sum_{\theta_{-ij} \in B} f(LH, \theta_{-ij})\pi_{-ij}(\theta_{-ij}) \right) \leq \\ \left(\sum_{\theta_{-ij} \in A \wedge B} f(LL, \theta_{-ij})\pi_{-ij}(\theta_{-ij}) \right) \left(\sum_{\theta_{-ij} \in A \vee B} f(HH, \theta_{-ij})\pi_{-ij}(\theta_{-ij}) \right).$$

This is precisely the inequality we wanted to show. This concludes the proof. \square

Proof for lemma 4.2. We first argue that there exists a voter R_i such that $\pi_i(L) < 1$. Suppose that for all i , $\pi_i(L) = 1$. Then R_i 's belief of her state being H if she conditions on the approval of all other voters is equal to the prior belief $\sum_{\theta \in \Theta_i^H} f(\theta)$. But given assumption 1, $(\pi_i(H), \pi_i(L)) = (1, 1)$ is not incentive compatible under the prior. Suppose now that $\pi_i(L) < 1$ for some i and IC^a - j is slack for all j . Then increasing $\pi_i(L)$ by a small amount strictly increases the probability that the project is approved without violating any IC^a constraint.

We then show that $\pi_i(L) > 0$ for some i . Suppose $\pi_i(L) = 0$ for all i . The sender fully reveals the payoff state to each voter. Hence, all IC^a constraints are slack. This contradicts the first part of this proof. \square

Proof for proposition 4.2. Suppose there exist R_i and R_j such that $\ell_j > \ell_i$ and $\pi_j(L) > \pi_i(L)$. We can rewrite IC^a - i and IC^a - j as:

$$\ell_i \pi_i(L) \leq \frac{a_1 + a_2 \pi_j(L)}{a_2 + a_3 \pi_j(L)}, \quad \ell_j \pi_j(L) \leq \frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)},$$

where

$$a_1 := \sum_{\Theta_i^H \cap \Theta_j^H} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k), \quad a_3 := \sum_{\Theta_i^L \cap \Theta_j^L} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k), \\ a_2 := \sum_{\Theta_i^H \cap \Theta_j^L} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k) = \sum_{\Theta_i^L \cap \Theta_j^H} f(\theta) \prod_{k \neq i,j} \pi_k(\theta_k).$$

It is easily verified that $a_1, a_2, a_3 \geq 0$. Moreover, using an argument similar to the

proof for lemma 4.1, we can show that $a_2^2 \leq a_1 a_3$. Therefore, $(a_1 + a_2 x)/(a_2 + a_3 x)$ decreases in x .

Multiplying the LHS of IC^{a-j} by $\ell_i/\ell_j \in (0, 1)$, we obtain the following inequality:

$$\ell_i \pi_j(L) < \frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)}. \quad (5)$$

Moreover, it is easy to show that $\frac{a_1 + a_2 x}{a_2 + a_3 x} x$ increases in x . Given the presumption that $\pi_i(L) < \pi_j(L)$, we thus have

$$\frac{\pi_i(L)}{\pi_j(L)} \leq \frac{\frac{a_1 + a_2 \pi_j(L)}{a_2 + a_3 \pi_j(L)}}{\frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)}}.$$

Multiplying the LHS of IC^{a-j} by $\pi_i(L)/\pi_j(L)$ and the RHS of IC^{a-j} by

$$\left(\frac{a_1 + a_2 \pi_j(L)}{a_2 + a_3 \pi_j(L)} \right) / \left(\frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)} \right),$$

we then obtain the following inequality:

$$\ell_j \pi_i(L) \leq \frac{a_1 + a_2 \pi_j(L)}{a_2 + a_3 \pi_j(L)}. \quad (6)$$

Based on (5) and (6), we can assign the lower $\pi_i(L)$ to R_j and the higher $\pi_j(L)$ to R_i without violating their IC^a constraints. Moreover, this switch does not affect the objective or any other voter's IC^a constraint. This shows that the sender is weakly better off after the switch. We have proved that there exists an optimal policy with the monotonicity property.

We next show that if IC^{a-i} binds, then $\pi_j(L) < \pi_i(L)$ if $\ell_j > \ell_i$. Suppose not: there exists R_j such that $\ell_j > \ell_i$ and $\pi_j(L) \geq \pi_i(L)$. Since IC^{a-i} binds, $\pi_i(L)$ is strictly positive. So is $\pi_j(L)$. Combining IC^{a-j} and the binding IC^{a-i} , we have

$$\frac{\ell_i \pi_i(L)}{\ell_j \pi_j(L)} \geq \frac{\frac{a_1 + a_2 \pi_j(L)}{a_2 + a_3 \pi_j(L)}}{\frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)}}.$$

On the other hand, given the presumption that $\pi_j(L) \geq \pi_i(L)$, it is easy to show that

$$\frac{\pi_i(L)}{\pi_j(L)} \leq \left(\frac{a_1 + a_2 \pi_j(L)}{a_2 + a_3 \pi_j(L)} \right) / \left(\frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)} \right).$$

This leads to the relation that $\pi_i(L)/\pi_j(L) \leq \ell_i \pi_i(L)/(\ell_j \pi_j(L))$, which contradicts the presumption that $\ell_j > \ell_i$. \square

Proof for proposition 4.3. We first prove that among those voters who have an interior $\pi_i(L)$, at most one voter has a slack IC^a constraint. Suppose not. Suppose that $\ell_j > \ell_i$ and $0 < \pi_j(L) \leq \pi_i(L) < 1$. (From proposition 4.2, it is without loss to assume that the more lenient voter R_i has a higher $\pi_i(L)$.) Suppose that both IC^{a-i}

and $IC^{\alpha-j}$ are slack. Then we can find a pair of small positive numbers $(\varepsilon_1, \varepsilon_2)$ such that replacing $(\pi_i(L), \pi_j(L))$ with $(\pi_i(L) + \varepsilon_1, \pi_j(L) - \varepsilon_2)$ makes the sender strictly better off without violating any IC^{α} constraint. First, when $\varepsilon_1, \varepsilon_2$ are sufficiently small, $IC^{\alpha-i}$ and $IC^{\alpha-j}$ are satisfied since they were slack.

In order to leave $IC^{\alpha-k}$ intact for any $k \neq i, j$, we need the following ratio to be weakly higher:

$$\frac{\mathcal{S}(HHH) + \mathcal{S}(HLH)\pi_j(L) + \mathcal{S}(LHH)\pi_i(L) + \mathcal{S}(LLH)\pi_i(L)\pi_j(L)}{\mathcal{S}(HHL) + \mathcal{S}(HLL)\pi_j(L) + \mathcal{S}(LHL)\pi_i(L) + \mathcal{S}(LLL)\pi_i(L)\pi_j(L)}, \quad (7)$$

where $\mathcal{S}(\theta_i\theta_j\theta_k) := \sum_{\theta_{-ijk} \in \Theta_{-ijk}} f(\theta_i\theta_j\theta_k, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk})$. Given the exchangeability assumption, we have $\mathcal{S}(HHL) = \mathcal{S}(HLH) = \mathcal{S}(LHH)$ and $\mathcal{S}(HLL) = \mathcal{S}(LHL) = \mathcal{S}(LLH)$.

We next show that $\mathcal{S}(HHH)\mathcal{S}(LLL) > \mathcal{S}(HHL)\mathcal{S}(HLL)$ by applying the strict Ahlswede-Daykin inequality introduced and proved in part B.6 of the online appendix. We define f_1, f_2, f_3, f_4 as follows:

$$\begin{aligned} f_1(\theta_{-ijk}) &= f(HHL, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk}), & f_2(\theta_{-ijk}) &= f(HLL, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk}), \\ f_3(\theta_{-ijk}) &= f(HHH, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk}), & f_4(\theta_{-ijk}) &= f(LLL, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk}). \end{aligned}$$

Suppose that among all voters other than i, j, k , only voters in $C \subseteq \{1, \dots, n\} \setminus \{i, j, k\}$ have a fully revealing individual policy, i.e. $\pi_m^L(\theta_{-ijk}) = 0$ for any $m \in C$. Then, $\pi_{-ijk}(\theta_{-ijk}) \neq 0$ only for $\{\theta_{-ijk} \in \Theta_{-ijk} : \theta_m = H \text{ for all } m \in C\} \equiv \mathcal{L}$. The set \mathcal{L} forms a sublattice because: 1) it is nonempty, as $\theta_{-ijk}^H \in \mathcal{L}$, 2) for any $\theta_{-ijk}, \theta'_{-ijk} \in \mathcal{L}$, it is easily verified that $\theta_{-ijk} \vee \theta'_{-ijk} \in \mathcal{L}$ and $\theta_{-ijk} \wedge \theta'_{-ijk} \in \mathcal{L}$. Given that for any $\theta_{-ijk}, \theta'_{-ijk} \in \mathcal{L}$, the premise of the Ahlswede-Daykin inequality is satisfied strictly due to our assumption of strict affiliation for any three voters' states, i.e. $f_1(\theta_{-ijk})f_2(\theta'_{-ijk}) < f_3(\theta_{-ijk} \wedge \theta'_{-ijk})f_4(\theta_{-ijk} \vee \theta'_{-ijk})$, for all $\theta_{-ijk}, \theta'_{-ijk} \in \mathcal{L}$, we can apply the strict Ahlswede-Daykin inequality with $A = B = \mathcal{L}$ to conclude that $f_1(\mathcal{L})f_2(\mathcal{L}) < f_3(\mathcal{L})f_4(\mathcal{L})$. But, because $\pi_{-ijk}(\theta_{-ijk}) = 0$ for all $\theta_{-ijk} \in \Theta_{-ijk} \setminus \mathcal{L}$,

$$\begin{aligned} \mathcal{S}(\theta_i\theta_j\theta_k) &= \sum_{\theta_{-ijk} \in \mathcal{L}} f(\theta_i\theta_j\theta_k, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk}) \\ &+ \sum_{\theta_{-ijk} \in \Theta_{-ijk} \setminus \mathcal{L}} f(\theta_i\theta_j\theta_k, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk}) \\ &= \sum_{\theta_{-ijk} \in \mathcal{L}} f(\theta_i\theta_j\theta_k, \theta_{-ijk})\pi_{-ijk}(\theta_{-ijk}). \end{aligned}$$

Therefore, $\mathcal{S}(HHH)\mathcal{S}(LLL) > \mathcal{S}(HHL)\mathcal{S}(HLL)$. Using the same method, we can prove that $\mathcal{S}(HHL)\mathcal{S}(LLL) > \mathcal{S}(HLL)^2$ and $\mathcal{S}(HHH)\mathcal{S}(HLL) > \mathcal{S}(HHL)^2$. We omit the details here. These inequalities imply that the ratio (7) strictly decreases in $\pi_i(L), \pi_j(L)$.

To make sure that the ratio (7) is higher after we replace $(\pi_i(L), \pi_j(L))$ with

$(\pi_i(L) + \varepsilon_1, \pi_j(L) - \varepsilon_2)$, we have to put a lower bound on ε_2 in terms of ε_1 . After we substitute $\mathcal{S}(HLH)$, $\mathcal{S}(LHH)$ with $\mathcal{S}(HHL)$ and $\mathcal{S}(LHL)$, $\mathcal{S}(LLH)$ with $\mathcal{S}(HLL)$, the lower bound on ε_2 in terms of ε_1 can be written as:

$$\varepsilon_2 \geq \frac{\pi_j(L)(s_1s_4 - s_2s_3) + s_1s_3 - s_2^2 + (\pi_j(L))^2(s_2s_4 - s_3^2)}{\pi_i(L)(s_1s_4 - s_2s_3) + s_1s_3 - s_2^2 + (\pi_i(L))^2(s_2s_4 - s_3^2)}\varepsilon_1, \quad (8)$$

where $s_1 = \mathcal{S}(HHH)$, $s_2 = \mathcal{S}(HHL)$, $s_3 = \mathcal{S}(HLL)$, $s_4 = \mathcal{S}(LLL)$.

Note that the lower bound on the RHS of (8) varies as we vary $k \neq i, j$. There exists a maximum lower bound. For the rest of the proof, we let R_k be the voter corresponding to this maximum lower bound.

The sender's payoff can be written as follows:

$$\{\mathcal{S}(HHH), \mathcal{S}(HHL), \mathcal{S}(HLH), \mathcal{S}(LHH), \mathcal{S}(HLL), \mathcal{S}(LHL), \mathcal{S}(LLH), \mathcal{S}(LLL)\} \\ \{1, \pi_k(L), \pi_j(L), \pi_i(L), \pi_j(L)\pi_k(L), \pi_i(L)\pi_k(L), \pi_i(L)\pi_j(L), \pi_i(L)\pi_j(L)\pi_k(L)\}$$

To make sure that the sender's payoff is weakly higher after we replace $(\pi_i(L), \pi_j(L))$ with $(\pi_i(L) + \varepsilon_1, \pi_j(L) - \varepsilon_2)$, we have to impose an upper bound on ε_2 in terms of ε_1 :

$$\varepsilon_2 \leq \frac{\pi_j(L)(\pi_k(L)s_4 + s_3) + \pi_k(L)s_3 + s_2}{\pi_i(L)(\pi_k(L)s_4 + s_3) + \pi_k(L)s_3 + s_2}\varepsilon_1.$$

In this expression, we have substituted $\mathcal{S}(HLH)$ with $\mathcal{S}(HHL)$, $\mathcal{S}(LHH)$ with $\mathcal{S}(LHL)$, and $\mathcal{S}(LLH)$ with $\mathcal{S}(HLL)$. The RHS decreases in $\pi_k(L)$, so this upper bound is the tightest when $\pi_k(L) = 1$:

$$\varepsilon_2 \leq \frac{(s_3 + s_4)\pi_j(L) + s_2 + s_3}{(s_3 + s_4)\pi_i(L) + s_2 + s_3}\varepsilon_1. \quad (9)$$

The upper bound in (9) is higher than the lower bound in (8), given that $s_1s_4 > s_2s_3$, $s_2s_4 > s_3^2$, and $s_1s_3 > s_2^2$. Therefore, we can find a pair $(\varepsilon_1, \varepsilon_2)$ such that the sender is better off after the change, and all the IC^a constraints are satisfied.

We then prove that among those voters who have interior $\pi_i(L)$, only the strictest voter might have a slack IC^a constraint. Suppose that $\ell_j > \ell_i$ and $0 < \pi_j(L) \leq \pi_i(L) < 1$. Suppose further that IC^a-i is slack and IC^a-j is binding. Then we can replace $(\pi_i(L), \pi_j(L))$ with $(\pi_i(L) + \varepsilon_1, \pi_j(L) - \varepsilon_2)$ where $(\varepsilon_1, \varepsilon_2)$ satisfy the two inequalities (8) and (9). Based on the argument above, this change makes the sender strictly better off without violating the IC^a constraints of the voters other than R_i and R_j . Moreover, since IC^a-i was slack, IC^a-i is satisfied after the change. IC^a-j is satisfied as well since we replace $\pi_j(L)$ with a more informative policy $\pi_j(L) - \varepsilon_2$ and the voters other than R_i and R_j are willing to be obedient if they condition on R_i 's and R_j 's approvals. It is easily verified that the decrease ε_2 required for IC^a-j to be satisfied is smaller than the decrease ε_2 that keeps the sender's payoff

intact. □

Proof for proposition 4.4. Consider the auxiliary problem in which the sender needs only R_1 's approval. The optimal policy is to set $\pi_1(L)$ as high as R_1 's IC^a constraint allows: $\pi_1(L) = \sum_{\theta \in \Theta_i^H} f(\theta) / \left(\sum_{\theta \in \Theta_i^L} f(\theta) \ell_1 \right)$. Conditioning on receiving an approval recommendation, R_1 's posterior belief that her state is H is $\ell_1 / (1 + \ell_1)$. Let f' denote the state distribution such that

$$f'(\theta^H) = \sum_{\theta \in \Theta_i^H} f(\theta), \quad f'(\theta^L) = 1 - f'(\theta^H).$$

Given the single-voter policy above, the belief of any other voter $R_j \neq R_1$ about $\theta_j = H$ approaches R_1 's posterior belief $\ell_1 / (1 + \ell_1)$, if $\|f - f'\|$ is small enough.³⁶ On the other hand, any $R_j \neq R_1$ is willing to approve the project if her belief of being H is above $\ell_j / (1 + \ell_j)$ which is strictly smaller than $\ell_1 / (1 + \ell_1)$. This means that there exists $\varepsilon > 0$ such that every other $R_j \neq R_1$ is willing to rubber-stamp R_1 's approval decision if $\|f - f'\| \leq \varepsilon$.

We then argue that when the above policy is incentive compatible, it is never optimal to fully reveal her payoff state to some voter. To the contrary, suppose that in the optimal policy $R_i \neq R_1$ learns about her state fully: $\pi_i(L) = 0$. Then the project is never approved when R_i 's state is L . The sender's payoff is at most $\sum_{\theta \in \Theta_i^H} f(\theta)$, which is strictly below the payoff under the single-voter policy specified in the first paragraph. □

Proof for proposition 4.5. Let f' denote the state distribution such that voters' states are independent and each voter's state is H with probability $\sum_{\theta \in \Theta_i^H} f(\theta)$. We want to argue that if $\|f - f'\|$ is sufficiently small, all IC^a constraints bind. Suppose not. Suppose there exists R_i for which IC^a - i is slack. Let R_j be another voter whose IC^a - j constraint binds. Then,

$$\ell_i \pi_i(L) < \frac{a_1 + a_2 \pi_j(L)}{a_2 + a_3 \pi_j(L)}, \quad \text{and} \quad \ell_j \pi_j(L) = \frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)}.$$

The coefficients a_1, a_2, a_3 are as defined in the proof of proposition 4.2. When $\|f - f'\|$ is sufficiently small, the LHS of each voters' IC^a constraints becomes independent of the information of the other voters:

$$\lim_{\|f - f'\| \rightarrow 0} \frac{a_1 + a_2 \pi_i(L)}{a_2 + a_3 \pi_i(L)} = \frac{\sum_{\theta \in \Theta_i^H} f'(\theta)}{\sum_{\theta \in \Theta_i^L} f'(\theta)}.$$

If the sender increases $\pi_i(L)$, the LHS of IC^a - j will be only slightly affected if $\|f - f'\|$ is sufficiently small. Hence, $\pi_j(L)$ has to decrease only slightly so that IC^a - j continues to bind. For sufficiently independent states, i.e. for $\|f(\cdot) - f'(\cdot)\| < \varepsilon$

³⁶Here, $\|f - f'\|$ denotes the Euclidean distance between f and f' .

for some small ε , setting all IC^a constraints binding strictly improves the payoff of the sender. \square

References

- Ahlsvede, Rudolf, and David E. Daykin.** 1978. "An Inequality for the Weights of Two Families of Sets, Their Unions and Intersections." *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 43: 183–185.
- Alonso, Ricardo, and Odilon Câmara.** 2016a. "Bayesian Persuasion with Heterogeneous Priors." *Journal of Economic Theory*, 165: 672–706.
- Alonso, Ricardo, and Odilon Câmara.** 2016b. "Persuading Voters." *American Economic Review*, 106(11): 3590–3605.
- Arieli, Itai, and Yakov Babichenko.** 2016. "Private Bayesian Persuasion."
- Austen-Smith, David, and Jeffrey S. Banks.** 1996. "Information Aggregation, Rationality, and the Condorcet Jury Theorem." *American Political Science Review*, 90(1): 34–45.
- Austen-Smith, David, and Timothy J. Feddersen.** 2006. "Deliberation, Preference Uncertainty, and Voting Rules." *American Political Science Review*, 100(2): 209–217.
- Bergemann, Dirk, and Stephen Morris.** 2016a. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics*, 11(2): 487–522.
- Bergemann, Dirk, and Stephen Morris.** 2016b. "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium." *American Economic Review: Papers and Proceedings*, 106(5): 586–591.
- Bergemann, Dirk, Tibor Heumann, and Stephen Morris.** 2015. "Information and volatility." *Journal of Economic Theory*, 158: 427–465.
- Caillaud, Bernard, and Jean Tirole.** 2007. "Consensus Building: How to Persuade a Group." *American Economic Review*, 97(5): 1877–1900.
- Callander, Steven.** 2007. "Bandwagons and Momentum in Sequential Voting." *Review of Economic Studies*, 74(3): 653–684.
- Chan, Jimmy, Alessandro Lizzeri, Wing Suen, and Leat Yariv.** 2016a. "Deliberating Collective Decisions."
- Chan, Jimmy, Seher Gupta, Fei Li, and Yun Wang.** 2016b. "Pivotal Persuasion."
- Coughlan, Peter J.** 2000. "In Defense of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting." *American Political Science Review*, 94(2): 375–393.
- Crawford, Vincent P., and Joel Sobel.** 1982. "Strategic Information Transmission." *Econometrica*, 50(6): 1431–1451.

- Feddersen, Timothy J., and Wolfgang Pesendorfer.** 1996. "The Swing Voter's Curse." *American Economic Review*, 86(3): 408–424.
- Feddersen, Timothy J., and Wolfgang Pesendorfer.** 1997. "Voting Behavior and Information Aggregation in Elections with Private Information." *Econometrica*, 65(5): 1029–1058.
- Gerardi, Dino, and Leat Yariv.** 2008. "Information Acquisition in Committees." *Games and Economic Behavior*, 62(2): 436–459.
- Gershkov, Alex, and Balazs Szentes.** 2009. "Optimal Voting Schemes with Costly Information Acquisition." *Journal of Economic Theory*, 144: 36–68.
- Glassman, Matthew Eric, and Amber Hope Wilhelm.** 2015. "Congressional Careers: Service Tenure and Patterns of Member Service, 1789-2015." Congressional Research Service CRS Report for Congress.
- Glazer, Jacob, and Ariel Rubinstein.** 2004. "On Optimal Rules of Persuasion." *Econometrica*, 72(6): 1715–1736.
- Harris, Charles E., and Julie H. Firestone.** 2014. "When Multiple Regulators Are at the Door: Strategies For Managing Parallel Proceedings by Government Agencies in Today's Competitive Regulatory Environment." *Section Annual Conference*.
- Kamenica, Emir, and Matthew Gentzkow.** 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.
- Li, Hao.** 2001. "A Theory of Conservatism." *Journal of Political Economy*, 109(3): 617–636.
- Mathevet, Laurent, Jacopo Perego, and Ina Taneva.** 2017. "On Information Design in Games."
- McCarty, Nolan.** 2015. "Regulation and Self-regulation of a Complex Industry."
- McCarty, Nolan, Keith T. Poole, and Howard Rosenthal.** 2013. *Political Bubbles: Financial Crises and the Failure of American Democracy*. Princeton, NJ: Princeton University Press.
- Milgrom, Paul.** 2008. "What the Seller Won't Tell You: Persuasion and Disclosure in Markets." *Journal of Economic Perspectives*, 22(2): 115–131.
- Omarova, Saule T.** 2011. "Wall Street as Community of Fate: Toward Financial Industry Self-regulation." *Cornell Law Faculty Publications*.
- Persico, Nicola.** 2004. "Committee Design with Endogeneous Information." *Review of Economic Studies*, 71: 165–191.
- Rayo, Luis, and Ilya Segal.** 2010. "Optimal Information Disclosure." *Journal of Political Economy*, 118(5): 949–987.

- Schnakenberg, Keith E.** 2015. "Expert Advice to a Voting Body." *Journal of Economic Theory*, 160: 102–113.
- Taneva, Ina.** 2016. "Information Design."
- Wang, Yun.** 2015. "Bayesian Persuasion with Multiple Receivers."
- Woodage, Trevor.** 2012. "Gatekeepers and Goalposts: The Need for a New Regulatory Paradigm for Whole Genome Sequence Results." *Northwestern Journal of Technology and Intellectual Property*, 11(1): 1–20.