Ricardo Alonso and Odilon Câmara
Persuading voters

Working paper

http://eprints.lse.ac.uk

Persuading Voters *

RICARDO ALONSO†   ODILON CÂMARA†

Marshall School of Business

University of Southern California

June 3, 2014

Abstract

In a symmetric information voting model, an individual (information controller) can influence voters' choices by designing the information content of a public signal. We characterize the controller's optimal signal. With a non-unanimous voting rule, she exploits voters' heterogeneity by designing a signal with realizations targeting different winning-coalitions. Consequently, under simple-majority voting rule, a majority of voters might be *strictly* worse off due to the controller's influence. We characterize voters' preferences over electoral rules, and provide conditions for a majority of voters to prefer a supermajority (or unanimity) voting rule, in order to induce the controller to supply a more informative signal.

JEL classification: D72, D83.

Keywords: Information control, persuasion, voting.

1 Introduction

Uncertainty gives rise to persuasion.

— Anthony Downs (1957)

Information is the cornerstone of democracy, as it allows voters to make better choices. In many important cases, however, uninformed voters are not free to launch their own investigations, and must rely on the inquiries of others. For example, in most trials a juror may not choose which tests are performed during the investigation, or which questions are asked to a witness — jurors must rely on the prosecutor's investigation and questions. In politics, the Legislative branch often must rely on information generated by investigative reports produced by the Executive. In firms, shareholders and the Board of Directors typically depend on reports commissioned by the CEO. If the individual choosing the questions and the voters have different preferences, then she may strategically design her investigation to persuade voters to choose her preferred alternative.

In this paper we study how an individual (the "information controller") influences the decision of voters by strategically designing a public signal, that is, by engaging in information control (e.g., Brocas and Carrillo 2007, Duggan and Martinelli 2011, Kamenica and Gentzkow 2011). We first study how the voting rule and the distribution of voters' preferences affect the controller's choice of a signal. We show that with a non-unanimous voting rule and heterogenous voters, the controller can design a signal with realizations targeting different winning coalitions. That is, the signal exploits preference disagreement across voters. Consequently, under simple-majority, a majority of voters might be *strictly* worse off due to the information supplied by the controller. To prevent this negative impact, voters may adopt a supermajority voting rule that induces the controller to supply a more informative signal.

In our model, a group of uninformed voters must choose whether to keep the status quo (or default) policy, or to implement a proposed new policy. This can be interpreted as voters choosing between an incumbent politician and a challenger, shareholders choosing to approve or not a merger, members of a jury choosing between a guilty or not guilty verdict, or members of a legislature choosing to approve or not a new law. Prior to the election,

1

the information controller can sway voters' decision by designing what voters can learn from a public signal, i.e. by specifying the statistical relation of the signal with the underlying state.[1] After observing the signal and its realization, voters apply Bayes' rule and reach a common posterior belief. They then choose an action (vote) and the electoral rule dictates the electoral outcome.

We focus on the case of pure-persuasion, where the information controller wants to maximize the probability of approval of the proposal, independently of the state (in Section 5.2 we consider a controller with state-dependent payoffs). We study two classes of institutional rules. The first is *delegation*, and serves as a benchmark. Under delegation, one member of the group acts as a dictator and chooses his preferred policy given his beliefs. The second is the class of *k-voting rules*, where a proposal replaces the status quo if it receives k or more votes. We focus on k-voting rules because they are important and prevalent in practice, and because they allow us to derive sharp equilibrium characterizations and comparisons. Since the controller's signal is public and voters have no private information, there is no information aggregation problem. Hence, the strategic voting considerations related to the probability of being pivotal are absent in our model.

We start by characterizing the optimal signal when the approval decision is delegated to voter i, who absent further information would reject the proposal. The controller maximizes the probability of approval by designing a signal such that, whenever voter i approves the proposal, he is just indifferent between approval and rejection. This implies that voter i does not benefit from this signal when he is the decision maker. In fact, voter i would prefer to delegate the approval decision to a "tougher" voter j to induce the controller to supply a more informative signal (where voter j is "tougher" than i if, for every belief, approval by j implies approval by i).

We then characterize the optimal signal under a k-voting rule. We establish that a k-voting rule is payoff-equivalent to delegating the decision to a particular "weak representative voter." This result allows us to relate the optimal signal under the k-voting rule to that under delegation. Given a $k-$voting rule, do voters benefit from the controller's signal? We show

---

[1]In our basic setup we consider an information controller who has no private information. In Section 5.1 we consider a controller who learns the state before choosing the signal.

that under a simple majority rule, the controller's influence always makes a majority of voters weakly worse off. We then define conditions such that a majority of voters is *strictly* worse off. This happens when the controller's optimal signal targets different winning coalitions, that is, when the controller exploits voters' preference disagreement to increase the probability of approving the proposal. Interestingly, this can happen even when all voters would agree on their decision if they knew the true state.

Anticipating the controller's influence, which k-voting rule do voters prefer? Voting rules affect outcomes not only by the amount of information voters have of the proposal's relative merits, but also by the consensus required to approve the proposal. In particular, requiring a higher consensus (higher k-voting rule) may lead to an excessive rejection of the proposal. However, it may also induce the controller to provide a more informative signal. We study this trade-off between control and information by posing two questions: (i) under delegation, what are each voter's preferences over decision makers?, and (ii) when decisions are made collectively, how do preferences over decisions makers translate into preferences over k-voting rules?

Consider a set of potential decision makers and a voter who can choose to whom to delegate the approval decision. The decision makers are totally ordered according to their "toughness", and rank states in the same order as the voter, where "higher" states imply that the proposal delivers a higher net payoff. For any such set, we show that the voter has single-peaked preferences over decision makers. Consider now a k-voting rule and an electorate who shares the same ranking of states. We extend the previous result to show that each voter has single-peaked preferences over k-voting rules. Moreover, a majority of voters always prefers a supermajority rule over simple majority. Finally, if voters also agree under full information, then *every voter* prefers unanimity over any other k-voting rule. That is, even heterogenous voters may agree on the optimal electoral rule.

Our paper is related to the recent literature on information control. In Brocas and Carrillo (2007), a leader without private information sways the decision of a follower in her favor by deciding the timing at which a decision must be made. As information arrives sequentially, choosing the timing of the decision is equivalent to shaping (in a particular way) the information available to the follower. Duggan and Martinelli (2011) consider one media outlet that

can affect electoral outcomes by choosing the "slant" of its news reports. The media in Duggan and Martinelli is constrained in the set of signals it can design (it must be a "slant"). In contrast, we consider an information controller who is unconstrained in her choice of signal. Our paper is most closely related to Kamenica and Gentzkow (2011) (KG henceforth). They develop the fundamental methodology to solve this class of unconstrained information control problems, when players have common priors. Alonso and Câmara (2014a) study information control when players have different prior beliefs. As in our paper, Michaeli (2014), Taneva (2014) and Wang (2013) focus on information control when there are multiple receivers.[2]

Our paper also relates to the broad literature on how institutional rules endogenously affect the information available to voters. Following the work of Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1996, 1997, 1998), a large literature has focused on how voting rules affect information aggregation via the strategic behavior of privately-informed voters. The literature has also studied the effects of voting rules on the information available to voters when voters can deliberate (e.g., Austen-Smith and Feddersen 2005, and Gerardi and Yariv 2007), when voters can postpone a decision in order to wait for more information (e.g., Messner and Polborn 2012 and Lizzeri and Yariv 2013), and when voters can acquire costly information (e.g., Li 2001, Persico 2004, Martinelli 2006, and Gerardi and Yariv 2008). Recent papers focus on how voting rules affect information provision by privately-informed experts. Jackson and Tan (2013) consider experts who can reveal verifiable information, while Schnakenberg (2014a,b) considers cheap-talk. In these papers experts are endowed with private information about the state, while in our paper the controller chooses the informational content of a public signal. Moreover, in Section 5.1 we contrast a

---

[2]There are other recent papers that study the strategic design of a public signal. Gill and Sgroi (2008, 2012) consider a privately-informed principal who can subject herself to a test designed to provide public information about her type, and can optimally choose the test's difficulty. Li and Li (2013) study a privately-informed candidate who can choose the accuracy of a costly public signal (campaign) about the qualifications of the politicians competing for office. Rayo and Segal (2010) study optimal advertising when a company can design how to reveal the attributes of its product, but it cannot distort this information. Kolotilin (2014) focus on how the receiver's private information affects the sender's choice of a signal. In a somewhat different setting, Ivanov (2010) studies the benefit to a principal of limiting the information available to a privately informed agent when they both engage in strategic communication.

controller without private information and a controller who privately knows the state.

The rest of the paper is organized as follows. Section 2 presents the basic model. Section 3 solves for the controller's optimal signal. Section 4 studies voters' preferences over institutional rules. Section 5 extends the basic model. Section 6 applies our results to relevant voting models. Section 7 concludes. All proofs are in the Appendix.

2 The Model

2.1 General Setup

Policy and Decision Makers: A group of n voters must choose one alternative from a binary *policy set* $X = \{x_0, x_1\}$, where x_0 is the status quo (or default) policy, and x_1 is the proposal. This can be interpreted as voters choosing between an incumbent politician and a challenger, choosing to approve or not a ballot measure, members of a jury choosing between a guilty or not guilty verdict, or members of a legislature choosing to approve or not a new law. The collective decision is made following established institutional rules, which we discuss momentarily. Each voter $i \in I \equiv \{1, \ldots, n\}$ has preferences over policies that are characterized by a continuous von Neumann-Morgenstern utility function $u_i(x, \theta)$, $u_i : X \times \Theta \to \mathbb{R}$, with Θ a finite state space. State $\theta \in \Theta$ captures the realized value of payoff relevant variables, such as the relative competence of different politicians, or the productivity of different sectors of the economy. All players share a common prior belief $p = \{p_\theta\}_{\theta \in \Theta}$, which has full support in Θ.

Information Controller: One information controller C, who is not a member of the group, has preferences over policies characterized by a continuous von Neumann-Morgenstern utility function u_C. We focus on the case of pure-persuasion were the controller's preferences are independent of the realized state, $u_C(x) : X \to \mathbb{R}$ (in Section 5.2 we consider a controller with state-dependent payoffs). The controller can influence the decision of the group by designing a public signal that is correlated with the state. Before the group selects a policy, the controller chooses a signal π, consisting of a finite realization space S and a family of likelihood functions over S, $\{\pi(\cdot|\theta)\}_{\theta \in \Theta}$, with $\pi(\cdot|\theta) \in \Delta(S)$. Signal π is "commonly

understood": π is observed by all players who agree on the likelihood functions $\pi(\cdot|\theta), \theta \in$ Θ (see Alonso and Câmara 2014a for a discussion of this assumption). Players process information according to Bayes rule. Let $q(s|\pi, p)$ be the updated posterior belief[3] of every voter after observing π and its realization s.

We make two important assumptions regarding the set of signals available to the controller. First, she can choose *any signal* that is correlated with the state. Thus, our setup provides an upper bound on the controller's benefit from information control in settings with more restricted spaces of signals. In particular, he controller will not engage in designing a signal when she faces additional constraints if there is no value of information control in our unrestricted setup. Second, signals are costless to the controller. This is not a serious limitation if each signal is equally costly, and would not affect the choice of signal if the controller decides to influence voters. However, the optimal signal may change if different signals have different costs. Gentzkow and Kamenica (2013) offer an initial exploration of persuasion with costly signals, where the cost of a signal is given by the expected relative entropy of the beliefs that it induces.

2.2   Institutional Rules

After observing the realization of the controller's signal, the group chooses one policy $x \in X$. The institutional rules governing the collective decision process are summarized by a mechanism $\Gamma = (\Gamma_1, \ldots, \Gamma_n, h)$, which defines a strategy set Γ_i for each member i and an outcome function $h : \Gamma_1 \times \ldots \times \Gamma_n \to X$. Given belief q, mechanism Γ and utility functions $\{u_i\}_{i \in I}$ define a Bayesian game $\mathcal{G}$. Let $\gamma^*(q) \equiv \{\gamma_i^*(q)\}_{i \in I}$ be a Perfect Bayesian equilibrium strategy profile played in this game. Together Γ and $\gamma^*(q)$ implement a social choice function $g(q) : \Delta(\Theta) \to X$, which defines the group's equilibrium policy choice as a function of beliefs. Therefore, for any signal π and realization $s \in S$ that yields belief q, the controller's payoff is given by

$$v(q) \;=\; u_C(g(q)). \tag{1}$$

---

[3]We use "posterior belief" to indicate the players' belief about the state after the signal realization but before voting. To simplify notation, we use $q(s)$ or q as a shorthand for $q(s|\pi, p)$.

Our main goal is to study how different institutional rules affect the optimal choice of a signal and the equilibrium payoff of players. We focus on two classes of institutional rules: *delegation*, which serves as a benchmark, and *k-voting rules*, where a proposal replaces the status quo if it receives k or more votes. We focus on *k-voting rules* because they are important and prevalent in practice, and because they allow us to derive sharp equilibrium characterizations and comparisons. We now formally define these institutional rules.

Delegation: Decision rights are fully delegated to a particular player $d \in I$. Mechanism $\Gamma = \{\Gamma_1, \ldots, \Gamma_n, h\}$ has $\Gamma_d = X$, where individual d chooses a policy $\gamma_d(q)$ and this policy is implemented, $h(\gamma_1(q), \ldots, \gamma_n(q)) = \gamma_d(q)$. In equilibrium, player d acts as a dictator and chooses $x \in X$ that maximizes his expected payoff, $\gamma_d^*(q) \in \arg\max_{x \in X} \sum_\Theta q_\theta u_d(x, \theta)$. If there are multiple optimal policies, we assume he chooses the one preferred by the information controller. Delegation implements $g(q) = \gamma_d^*(q)$, and (1) becomes $v(q) = u_C(\gamma_d^*(q))$. It follows from Berge's maximum theorem that v is upper-semicontinuous.

k-voting rule: Proposal x_1 is selected if and only if it receives at least k votes, where $k \in \{1, \ldots, n\}$ is the established electoral rule. Mechanism $\Gamma = \{\Gamma_1, \ldots, \Gamma_n, h\}$ has $\Gamma_i = \{0, 1\}$, where $\gamma_i(q) = 1$ represents voting for proposal x_1, and $\gamma_i(q) = 0$ represents voting for x_0 — we abstract from abstention. The outcome function h is

$$h(\gamma_1(q), \ldots, \gamma_n(q)) = \begin{cases} x_1 & \text{if } \sum_{i \in I} \gamma_i(q) \geq k, \\ x_0 & \text{if } \sum_{i \in I} \gamma_i(q) < k. \end{cases}$$

Given a belief q, we apply the following two equilibrium selection criteria in case of multiple equilibria:

1. If policy x yields voter i a strictly higher expected payoff than policy x', then he votes for x;

2. If the two policies yield voter i the same expected payoff, then he votes for the policy preferred by the information controller.

The first criterion rules out uninteresting equilibria such as, when $k < n$, all voters vote for the status quo independently of expected payoffs. Importantly, in our model voters have

no private information about the state, so there is no information aggregation problem. Hence, the strategic voting considerations related to the probability of being pivotal are not relevant in our setup. From the set of equilibria satisfying the first criterion, we select the subset of controller-preferred equibria, which guarantees that the controller's expected payoff v is an upper semicontinous function of posterior beliefs. Let $\gamma_i^*(q)$ be the equilibrium choice of voter i that satisfies the previous selection criteria. The social choice function is then $g(q) = h(\gamma_1^*(q), \ldots, \gamma_n^*(q))$.

As in Alonso and Câmara (2014a), we focus on *language-invariant Perfect Bayesian equilibrium*: a Perfect Bayesian equilibrium in which individual decisions depend on posterior beliefs, but not on the actual signal or realization — for every signals π and π', and signal realizations s and s' for which individual i has the same posterior belief q, he chooses the same equilibrium strategy $\gamma_i^*(q)$. Note that if game $\mathcal{G}$ has multiple equilibria, then the social choice function g implicitly selects which equilibrium is played.

2.3 Information Controller's Problem

For any signal π and realization $s \in S$ that yields posterior q, the social choice function g determines the implemented policy — the controller's payoff $v(q)$ is then defined by (1). The information controller selects a signal that maximizes $E_\pi[v(q)]$. Upper-semicontinuity of v both with delegation and with k-voting rules ensures the existence of an optimal signal (see KG). Moreover, choosing an optimal signal is equivalent to choosing a probability distribution σ over q, subject to the constraint $E_\sigma[q] = p$. That is,

$$V \;\; = \;\; \max_\sigma E_\sigma[v(q)], \qquad \text{s.t. } E_\sigma[q] = p.$$

For an arbitrary real-valued function f define $\widetilde{f}$ as the concave closure of f,

$$\widetilde{f}(q) = \sup \left\{ w | (q, w) \in co(f) \right\},$$

where $co(f)$ is the convex hull of the graph of f. The following remarks follow immediately from KG:

(R1) An optimal signal exists;

(R2) If the approval decision is delegated to one voter, then there exists an optimal signal

with $card(S) \leq 2$. With a k-voting rule, there exists an optimal signal with $card(S) \leq \min\{\frac{n!}{(n-k)!k!} + 1, card(\Theta)\}$;

(R3) The information controller's expected utility under an optimal signal is

$$V = \widetilde{v}(p); \tag{2}$$

(R4) The value of information control is $V - v(p) = \tilde{v}(p) - v(p)$.

For the remaining of the paper we focus on the case where $u_C(x_0) < u_C(x_1)$[4]. Without loss of generality, set $u_C(x_1) = 1$ and $u_C(x_0) = 0$. Therefore, the controller's expected payoff V is simply the equilibrium approval probability under an optimal signal.

2.4 Definitions and Notation

We next present a series of definitions and notation that will be useful in our analysis.

Notational Conventions: For vectors $q, w \in \mathbb{R}^J$, we denote by $\langle q, w \rangle$ the standard inner product in $\mathbb{R}^J$, i.e. $\langle q, w \rangle = \sum_{j=1}^{J} q_j w_j$, and we denote by qw the component-wise product of vectors q and w, i.e. $(qw)_j = q_j w_j$.

Voter's Type: Define the conditional net payoff for voter i when the state is θ as

$$\delta_\theta^i \equiv u_i(x_1, \theta) - u_i(x_0, \theta).$$

The vector $\delta^i = \{\delta_\theta^i\}_{\theta \in \Theta}$ captures the preferences of the voter, and we call δ^i the *type* of voter i. When voter i holds belief q, he votes for x_1 if and only if $\sum_\Theta q_\theta (u_i(x_1, \theta) - u_i(x_0, \theta)) \geq 0$, that is, if and only if $\langle q, \delta^i \rangle \geq 0$. Hence, equilibrium voting strategies γ_i^* are fully defined by δ^i and q,

$$\gamma_i^*(q) \equiv a(q, \delta^i) = \begin{cases} 1 & \text{if} \quad \langle q, \delta^i \rangle \geq 0, \\ 0 & \text{if} \quad \langle q, \delta^i \rangle < 0. \end{cases}$$

Since a voter's type defines his voting behavior, we use the term "voter δ" to refer to a voter with type δ.

---

[4]The remaining case $u_C(x_0) > u_C(x_1)$ is equivalent to a proposal $\hat{x}_1 = x_0$ and a status quo $\hat{x}_0 = x_1$, with the corresponding relabeling of the collective decision process and social choice function.

Relevant Sets — Individual Voter: Consider a voter with type δ. Define the *set of approval states* $D(\delta) = \{\theta \in \Theta | \delta_\theta \geq 0\}$ and the set of rejection states $D^C(\delta) = \Theta \setminus D(\delta)$. Define the *set of approval beliefs* $A(\delta) = \{q \in \Delta(\Theta) | \langle q, \delta \rangle \geq 0\}$ and the set of rejection beliefs $A^C(\delta) = \Delta(\Theta) \setminus A(\delta)$. Under full information, voter δ approves x_1 if and only if $\theta \in D(\delta)$; while under uncertainty he approves x_1 if and only if $q \in A(\delta)$. Finally, define the *set of strong rejection beliefs* $R(\delta) = \{q \in \Delta(\Theta) | \theta \in D(\delta) \Rightarrow q_\theta = 0\}$, that is, the set of beliefs that assign probability zero to every approval state.

Relevant Sets — Electorate: Consider an electorate $\{\delta^1, \ldots, \delta^n\}$ and a k-voting rule. Define the *win set*

$$W_k = \{q \in \Delta(\Theta) | \sum_{i=1}^n a(q, \delta^i) \geq k\}.$$

That is, voters implement x_1 if and only if $q \in W_k$. Given the k-voting rule, there are $\frac{n!}{(n-k)!(k!)}$ possible minimal winning coalitions of k voters. The win set is then the union of all possible minimal winning coalitions. Under unanimity rule $k = n$, the win set is the intersection of all approval sets, $W_n = \cap_{i \in I} A(\delta^i)$. If $k = 1$, then the win set is the union of all approval sets, $W_1 = \cup_{i \in I} A(\delta^i)$. Note that W_n is convex, but W_k might be a non-convex set when $k < n$. Given the electorate, define $\mathscr{B}$ as the collection of all coalitions of at least $n - k + 1$ voters, with typical element $b \in \mathscr{B}$. Define the *set of strong rejection beliefs*

$$R_k = \cup_{b \in \mathscr{B}} \left(\cap_{\delta \in b} R(\delta) \right).$$

That is, R_k is the set of beliefs such that there exists a "blocking" coalition b, with voters $\delta \in b$ assigning probability zero to every approval state.

Finally, we use $V(\delta)$ and $V(W_k)$ to denote the equilibrium approval probability with delegation to voter δ and with a k-voting rule with win set W_k.

Classes of Voters' Types: It is useful to group voters according to their types. To this end, let z be a permutation $z : \Theta \rightarrow \{1, \ldots, card(\Theta)\}$ that strictly orders the states. Define the class of types

$$\mathscr{F}_z = \{\delta \in \mathbb{R}^{card(\Theta)} | \delta_\theta > \delta_{\theta'} \iff z(\theta) > z(\theta')\}.$$

That is, class $\mathscr{F}_z$ includes all voter types who (strictly) rank states according to the conditional net payoff δ_θ in the order defined by z. We say that voter δ^i "ranks states" according

to z if $\delta^i \in \mathscr{F}_z$.

Ordering Voters: We introduce two orders on the space of voter types. First, we say that voter δ is "tougher" than voter δ' if $A(\delta) \subset A(\delta')$. Second, we say that voter δ is (weakly) "harder-to-persuade" than voter δ' if $V(\delta) \leq V(\delta')$. That is, under the optimal signal, the equilibrium approval probability under an optimal signal with delegation to voter δ is (weakly) lower than with delegation to δ'.

Representative Voter: Fix a k-voting rule and an electorate $\{\delta^1, \ldots, \delta^n\}$. Voter δ is a "representative voter" if $A(\delta) = W_k$, that is, the proposal is approved with a k-voting rule if and only if it would be approved with delegation to voter δ. Voter δ is a "*weak* representative voter" if $V(\delta) = V(W_k)$, that is, if the information controller's equilibrium expected payoff is the same when she only has to convince voter δ and when she has to convince at least k voters.

3 Information Control

In this section we first solve for the controller's optimal signal both with delegation and with a k-voting rule. We then study how the controller's gain from designing the signal varies with the electoral rule and the distribution of voters' preferences.

3.1 Information Control with Delegation

Suppose that the approval decision is delegated to voter δ. Which signal would the controller optimally supply? If $p \in A(\delta)$, then the controller provides a completely uninformative signal, as the voter approves the proposal in the absence of additional information. Now suppose $p \notin A(\delta)$ and $A(\delta) \neq \emptyset$. The characterization of the optimal signal obtains from two observations. First, the controller would never benefit from providing additional information after the realization of an optimal signal. Therefore, after observing a signal realization that induces rejection voter δ must assign zero probability to every approval state $\theta \in D(\delta)$ (see also Proposition 4 in KG), and thus his posterior belief must lie in the strong rejection set $R(\delta)$. Second, any signal realization that leads to approval must induce a posterior belief in $A(\delta)$.

While there is a multiplicity of optimal signals, under delegation one can always construct an optimal signal with only two signal realizations. To do so, starting with an arbitrary optimal signal π^* one can group all approval signal realizations into a single realization associated to a posterior that is the probability-weighted convex combination of all approval posteriors. One can likewise obtain a single signal realization by grouping all rejection signal realizations in a similar fashion. Importantly, $A(\delta)$ and $R(\delta)$ are both convex sets, so the posterior belief obtained from combining all approval realizations still leads the voter to approve, and combining all rejection realizations still leads the voter to reject, so that the probability of approval remains the same.

We now provide a geometric interpretation of the controller's optimal signal. Consider a signal π supported on $\{s^-, s^+\}$, where s^+ induces approval posterior $q^+ \in A(\delta)$ and s^- induces strong rejection posterior $q^- \in R(\delta)$. Holding constant q^-, s^+ becomes more likely as the posterior q^+ becomes closer to the prior p. Conversely, holding constant q^+, s^- becomes less likely as q^- is further away from p. Consequently, the controller would like to resort to both an approval belief q^+ in $A(\delta)$ that is closest to the prior, and a rejection belief q^- in $R(\delta)$ that is farthest from the prior. The martingale property of Bayesian updating requires, however, that q^+, q^- and p must all be collinear. The following lemma shows that the optimal signal balances these two goals as it corresponds to lines through the prior that maximize the ratio of the distances described in (3).

Lemma 1 *Consider delegation to voter δ, with $p \notin A(\delta)$ and $A(\delta) \neq \emptyset$. Let π^* be any controller's optimal signal supported on $\{s^-, s^+\}$, where voter δ approves the proposal if and only if $s = s^+$, with $l^* = q(s^+) - p$. Let $d_l(p, A(\delta))$ and $d_l(p, R(\delta))$ be the (Euclidean) distances from the prior to the sets $A(\delta)$ and $R(\delta)$ along the line l. Then*

$$\frac{d_{l^*}(p, R(\delta))}{d_{l^*}(p, A(\delta))} = \max_l \frac{d_l(p, R(\delta))}{d_l(p, A(\delta))}, \tag{3}$$

and

$$\Pr[Approval] = \frac{d_{l^*}(p, R(\delta))}{d_{l^*}(p, R(\delta)) + d_{l^*}(p, A(\delta))}.$$

The next Proposition shows that the solution to (3) can be understood as the optimal choice of a cutoff state.

Proposition 1 *Consider delegation to voter δ, with $p \notin A(\delta)$ and $A(\delta) \neq \emptyset$. Let π^* be any controller's optimal signal supported on $\{s^-, s^+\}$, where voter δ approves the proposal if and only if $s = s^+$. Letting $\alpha_\theta = \Pr[s^+|\theta]$, there exists $\theta' \in \Theta$ such that*

$$\alpha_\theta = \begin{cases} 0 & \delta_\theta < \delta_{\theta'} \\ 1 & \delta_\theta > \delta_{\theta'} \end{cases} , \quad and \quad \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta = 0. \tag{4}$$

Moreover, while voter δ never gains by making decisions with the signal π^*, the controller's expected utility under π^* is $V(\delta) = \sum_{\theta \in \Theta} \alpha_\theta p_\theta$.

To understand (4), first consider the voter's ideal signal. Voter δ would like to know whether an approval state occurred; thus his best signal would induce s^+ if $\delta_\theta \geq 0$, and s^- if $\delta_\theta < 0$. With this signal, the approval probability is $\sum_{\{\theta:\delta_\theta \geq 0\}} p_\theta$, and the voter's net value from approval is $\sum_{\{\theta:\delta_\theta \geq 0\}} p_\theta \delta_\theta$. If $\delta_\theta > 0$ for at least one θ, then the controller can increase the probability of approval by distorting this signal in a way that rejection states with a small incremental loss (i.e. small $|\delta_\theta|$) still induce approval. The controller can do so until the voter's net value from approval is identically zero, as indicated by (4). This also implies that the voter gains nothing from making decisions with π^*, as he is indifferent between approval and rejection after observing s^+.

Example 1: Consider three states, $\Theta = \{\theta_1, \theta_2, \theta_3\}$, and a single voter δ such that $\delta_{\theta_1} < \delta_{\theta_2} < 0 < \delta_{\theta_3}$. Figure 1 depicts on the simplex the approval set $A(\delta)$ and the strong rejection set $R(\delta)$, which is the bottom line segment. The prior p is such that $p_3\delta_{\theta_3} + p_2\delta_{\theta_2} < 0$. Figure 1(a) shows the approval posterior $q^+ \in A(\delta)$ that is closest to the prior, and the strong rejection posterior $q^- \in R(\delta)$ that is furthest from the prior. However, there is no binary signal that induces these posteriors, since q^+, p and q^- must be collinear. Figure 1(b) shows the posterior beliefs induced by an optimal signal satisfying condition (3). Using (4), θ_2 is the cutoff state: $Pr[s^+|\theta_3] = 1$, $Pr[s^+|\theta_2] = \frac{p_3\delta_{\theta_3}}{p_2|\delta_{\theta_2}|}$, $Pr[s^+|\theta_1] = 0$. The maximum approval probability is then $V(\delta) = 1 \cdot p_3 + \frac{p_3\delta_{\theta_3}}{p_2|\delta_{\theta_2}|} \cdot p_2$. $\square$

3.2   Information Control with a k-Voting Rule

A basic insight of information control under delegation is the existence of a binary optimal signal: one realization leading to approval, and the other leading to rejection. This is possible
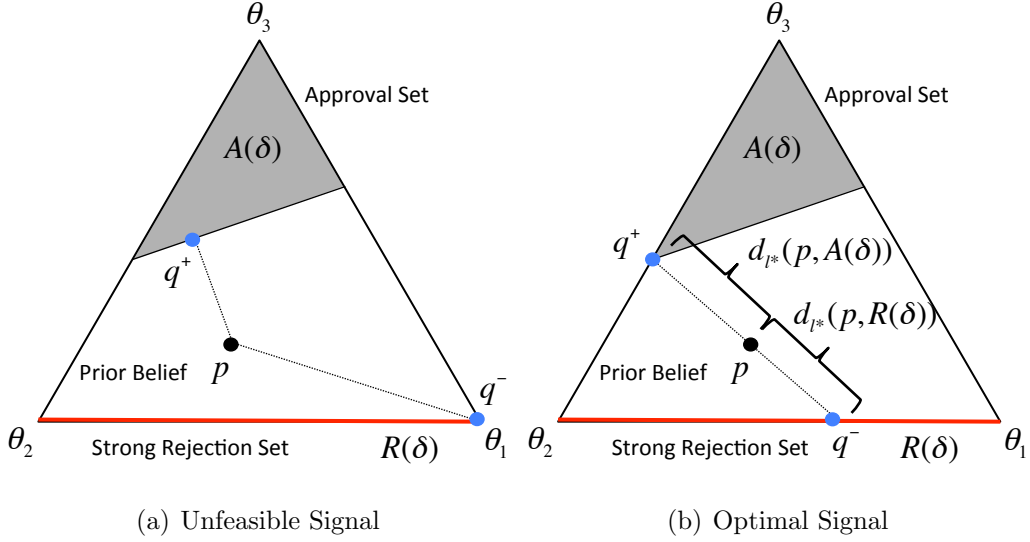
Figure 1: Optimal Signal for Example 1

as the set of approval beliefs of any voter is convex. However, with a k-voting rule the set of approval beliefs W_k is, in general, not convex. Nevertheless, persuading voters with approval beliefs W_k is payoff-equivalent for the controller to persuading voters with approval beliefs equal to the convex hull of W_k. To see this, note that any belief in $co(W_k)$ can be expressed as a convex combination of posterior beliefs each of which ensures approval. Therefore, if $q \in co(W_k)$, then there is a signal that ensures approval with certainty. Lemma 2 shows that the same logic of Lemma 1 holds with a k-voting rule, replacing $A(\delta)$ and $R(\delta)$ with the sets $co(W_k)$ and R_k.

Lemma 2 *Fix a k-voting rule and electorate $\{\delta^1, \ldots, \delta^n\}$. The probability of approval under an optimal signal is*

$$\Pr[Approval] = \frac{d_{l^*}(p, R_k)}{d_{l^*}(p, R_k) + d_{l^*}(p, co(W_k))}, \quad with \quad \frac{d_{l^*}(p, R_k)}{d_{l^*}(p, co(W_k))} = \max_l \frac{d_l(p, R_k)}{d_l(p, co(W_k))}. \quad (5)$$

We now contrast equilibrium payoffs with a k-voting rule to equilibrium payoffs with delegation to a voter δ. Voter δ is a *representative voter* if $A(\delta) = W_k$: in this case, the expected utility of all players with the k-voting rule is equivalent to delegating the decision to voter δ.[5] In many situations, however, a representative voter does not exit. Nevertheless,

---

[5] A representative voter exists for each $k-$voting rule if all voters in the electorate are totally ordered according to toughness. This is the case, for example, if there are only two states and voters have the same ranking of states.

the next proposition establishes that one can always construct a *weak representative voter*.

Proposition 2 *Fix a k-voting rule and electorate $\{\delta^1, \ldots, \delta^n\}$. There exists a weak representative voter $\delta^*(k)$, such that $\{q : \langle q, \delta^*(k) \rangle = 0\}$ is a supporting hyperplane of $co(W_k)$. If $\delta_\theta^*(k) \neq \delta_{\theta'}^*(k)$ for $\theta \neq \theta'$, then the expected utility of all players under a k-voting rule is the same as under delegation to voter $\delta^*(k)$.*

Typically voter $\delta^*(k)$ is not part of the electorate. As $\{q : \langle q, \delta^*(k) \rangle = 0\}$ is a supporting hyperplane of W_k, then $W_k \subset A(\delta^*(k))$. That is, a k-voting rule is equivalent, from the controller's perspective, to delegation to a "less tough" decision maker. Notwithstanding, $\delta^*(k)$ is also the hardest-to-persuade among all voters with approval beliefs containing W_k. To see this, let $G(W_k) = \{\delta : W_k \subset A(\delta)\}$ be the set of voters who are "less tough" than the $k-$voting rule. The proof of the Proposition shows that

$$\Pr[Approval] = \inf_{\delta \in G(W_k)} \Pr[Approval(\delta)] = \Pr[Approval(\delta^*(k))].$$

Moreover, if the weak representative voter $\delta^*(k)$ strictly ranks states, then the ex-ante expected utility of *all voters* is the same with the k-voting rule as with delegation to voter $\delta^*(k)$. This reflects the fact that if $\delta_\theta^*(k) \neq \delta_{\theta'}^*(k)$ for $\theta \neq \theta'$, then the optimal signal in Proposition 1 is unique and provides the same expected utility to all voters as the outcome of a k-voting rule.

An important difference between delegation and k-voting rules is that, except for unanimity, the controller can exploit conflict between voters to increase the chance of approving the proposal. This is always the case if the optimal binary signal for $\delta^*(k)$ would not induce approval under a k-voting rule, i.e. letting $q^+(\delta^*(k))$ be the approval posterior belief associated to (4), then $q^+(\delta^*(k)) \in co(W_k)$ but $q^+(\delta^*(k)) \notin W_k$. Letting $Q^+(\pi^*)$ be the set of approval posterior beliefs induced by an optimal signal, we must have $\langle q, \delta^*(k) \rangle = 0$ for $q \in Q^+(\pi^*)$. Therefore, the representative voter $\delta^*(k)$ describes precisely the direction of voters conflict: if an approval realization doesn't secure the support of one winning coalition, it must secure the support of another. Essentially, the controller exploits the change in the identity of the winning coalition (and hence the identity of the pivotal voter) with the signal realization by tailoring the signal to persuade winning coalitions with opposing interests on the hyperplane $\langle q, \delta^*(k) \rangle = 0$.

In many important applications, voters have different preferences but agree on the ranking of the states — for example, state θ represents the "quality" of the proposal, as in the public good example in Section 6.1. The next lemma characterizes the weak representative voter from Proposition 2 when all voters agree on the ranking of the states.

Lemma 3 *Consider an electorate $\{\delta^1, \ldots, \delta^n\}$ with all $n \geq 2$ voters on the same class, $\delta^i \in \mathscr{F}_z$ for some permutation z. Given a k-voting rule, there exists a weak representative voter $\delta^*(k)$ in the same class as the electorate, $\delta^*(k) \in \mathscr{F}_z$.*

By establishing that $\delta^*(k)$ must be in the same class as the electorate, Lemma 3 allows us to relate the optimal signal with a k-voting rule to the optimal signal with delegation. With delegation to $\delta^*(k)$, there is an optimal signal supported on $\{s^-, s^+\}$, with a cutoff state θ' as described in Proposition 1. As voters rank states in the same order as $\delta^*(k)$, they all agree that s^+ is "good news" about the proposal, while s^- is "bad news". Although signal s^+ is enough to persuade $\delta^*(k)$ to approve, it might not secure k votes from the electorate. As the next example shows, the controller then decomposes s^+ into realizations targeting different winning coalitions.

Example 2: Consider 3 states, $\Theta = \{\theta_1, \theta_2, \theta_3\}$, and two voters, δ^A and δ^B, where prior beliefs and net values are as follows:

State	Prior	δ_θ^A	δ_θ^B
θ_3	0.2	+1	+1
θ_2	0.1	-0.5	-1.5
θ_1	0.7	-6	-2

Without further information voters strictly prefer to reject the proposal. First suppose that the voting decision is delegated to voter δ^A. Figure 2(a) depicts the posterior beliefs induced by an optimal signal π_A^*, using Proposition 1.[6] The equilibrium probability of approval is 0.325. Signal π_A^* is also optimal if the decision is delegated to δ^B. In both cases, the controller's influence does not change voters' expected payoff.

---

[6]Formally, $S = \{s^-, s^+\}$, $Pr(s^+|\theta_3) = Pr(s^+|\theta_2) = 1$, and $Pr(s^+|\theta_1) = \frac{1}{28}$. The possible posterior beliefs are $q^- = (1, 0, 0)$ and $q^+ = (\frac{1}{13}, \frac{4}{13}, \frac{8}{13})$. Probability of approval is $Pr(s = s^+) = \frac{1}{28} \times 0.7 + 1 \times 0.1 + 1 \times 0.2 = 0.325$.

(a) Decision Delegated to δ^A or δ^B (b) Simple Majority   (c) Decision Delegated to δ^*
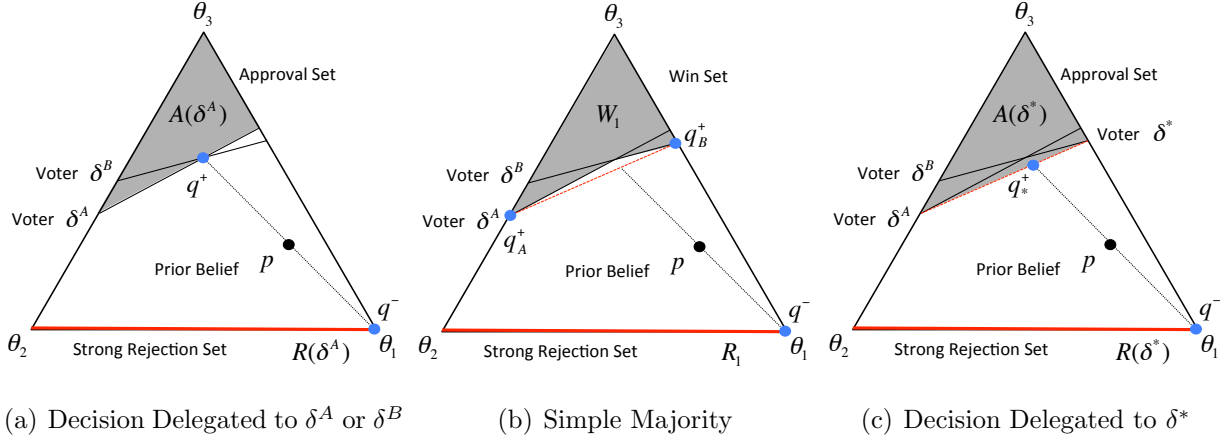
Figure 2: Optimal Signal for Example 2

Now consider simple majority ($k = 1$). Figure 2(b) depicts the posterior beliefs induced by an optimal signal π_1^* supported on $S = \{s^-, s_A^+, s_B^+\}$.[7] Voter δ^A approves if and only if s_A^+ occurs, while δ^B approves if and only if s_B^+ occurs. Realizations s_A^+ and s_B^+ occur with probabilities 0.15 and 0.225, respectively, hence the probability of approval is 0.375.

To understand this optimal signal, we first analyze the weak representative voter δ^* with simple majority. Voter $\delta^* = (-2, -0.5, 1)$ is represented in 2(c) by the dotted line, which delineates the convex hull of the win set W_1. The optimal signal $\pi_{\delta^*}^*$ with delegation to δ^* induces posteriors q_*^+ and q^-.[8] Note that belief q_*^+ is the weighted average of beliefs q_A^+ and q_B^+. With posterior q_*^+, voter δ^* approves, but voters δ^A and δ^B strictly prefer to reject. This implies that the controller's influence *strictly reduces* the expected payoff of both voters δ^A and δ^B.

We now use $\pi_{\delta^*}^*$ to decompose π_1^* into two components: *collective persuasion* and *targeted persuasion*. The straight line connecting q^- and q_*^+ in Figure 2(c) is a direction of *common interest*: all voters agree that moving beliefs from q^- in the direction of q_*^+ represents "good news" about the proposal. The collective persuasion component of π_1^* is the belief change

---

[7]Formally, $S = \{s^-, s_A^+, s_B^+\}$, $Pr(s^-|\theta_3) = 0$, $Pr(s_B^+|\theta_3) = \frac{3}{4}$, $Pr(s_A^+|\theta_3) = \frac{1}{4}$, $Pr(s^-|\theta_2) = Pr(s_B^+|\theta_2) = 0$, $Pr(s_A^+|\theta_2) = 1$, $Pr(s^-|\theta_1) = \frac{25}{28}$, $Pr(s_B^+|\theta_1) = \frac{3}{28}$, $Pr(s_A^+|\theta_1) = 0$. The possible posterior beliefs are $q^- = (1, 0, 0)$, $q_B^+ = (\frac{1}{3}, 0, \frac{2}{3})$, and $q_A^+ = (0, \frac{2}{3}, \frac{1}{3})$. Probability of approval is $Pr(s = s_B^+) + Pr(s = s_A^+) = (\frac{3}{28} \times 0.7 + 0 \times 0.1 + \frac{3}{4} \times 0.2) + (0 \times 0.7 + 1 \times 0.1 + \frac{1}{4} \times 0.2) = 0.225 + 0.15 = 0.375$.

[8]Formally, $S = \{s^-, s_*^+\}$, $Pr(s_*^+|\theta_3) = Pr(s_*^+|\theta_2) = 1$, and $Pr(s_*^+|\theta_1) = \frac{3}{28}$. The possible posterior beliefs are $q^- = (1, 0, 0)$, $q_*^+ = (\frac{3}{15}, \frac{4}{15}, \frac{8}{15})$. Probability of approval is $Pr(s = s_*^+) = \frac{3}{28} \times 0.7 + 1 \times 0.1 + 1 \times 0.2 = 0.375$.

along this common interest direction. However, although belief q_*^+ is good news about the proposal, it is not enough to convince voters δ^A and δ^B. The controller then relies on targeted persuasion. Starting from q_*^+, signal π_1^* moves the belief to either q_A^+ or q_B^+. The straight line connecting q_A^+ and q_B^+ in Figures 2(b) is a direction of *opposing interest*: moving beliefs from q_A^+ in the direction of q_B^+ represents "good news" about the proposal to voter δ^B, but "bad news" to voter δ^A. Importantly, the weak representative voter corresponds precisely to this direction of opposing interest — as Figures 2(b) and (c) illustrate. From belief q_*^+ the controller ensures approval by exploiting the opposing interests of voters.

3.3   Value of information control

The next Corollary provides comparative statics of k-voting rules on the value of information control.

Corollary 1 *Consider an electorate $\{\delta^1, \ldots, \delta^n\}$ and a $k-$voting rule. Then*
(i) Information control is not valuable if and only if the set W_k is empty or $p \in W_k$.
(ii) Information control is most valuable when $p \notin W_k$ and $p \in co(W_k)$.
(iii) Equilibrium probability of approval weakly decreases with k.
(iv) The value of information control is a single-peaked function of k, possibly non-monotone.

Parts (i) to (iii) follow immediately from Lemma 2. To understand (iv), suppose that the value of information control strictly decreases from rule k to rule $k' = k + 1$. This means that the value of information control was strictly positive under k. From (i) this implies $p \notin W_k$, yielding $p \notin W_{\hat{k}}$ for any $\hat{k} > k$. From (iii) we know that the probability of approval decreases in k, hence the value of information control must weakly decrease from rule k on.

3.4   Voter Heterogeneity and Information Control

Proposition 2 showed that the controller can, under non-unanimous voting rules, exploit voter heterogeneity by designing a signal that induces approval from different winning coalitions. In effect, under a k-voting rule the controller designs approval signal realizations along directions of voter disagreement in such a way that there is always a coalition of at least k voters willing to approve the proposal.

A natural question then is: would the controller prefer to persuade a group of voters rather than an individual voter to whom the decision is delegated? To make this statement precise, suppose that voters are ordered according to how "hard" it is for the controller to persuade them, i.e., if $i < i'$ then $V(\delta^i) \geq V(\delta^{i'})$. Thus, voter δ^1 is the easiest voter to persuade, while voter δ^n is the hardest. The following proposition provides a sufficient condition for the controller to prefer a $k-$voting rule to delegation to the k-th hardest voter.

Proposition 3 *Consider an electorate* $\{\delta^1, \ldots, \delta^n\}$, *and index voters according to how hard it is to persuade them individually,* $V(\delta^{i'}) \leq V(\delta^i)$ *for* $i < i'$. *Then*
(i) For any voter δ^i, $V(W_n) \leq V(\delta^i)$ *and* $V(W_1) \geq V(\delta^i)$;
(ii) If voters rank states in the same order, $\delta^i \in \mathscr{F}_z$, $i \in I$, *then* $V(W_n) = V(\delta^n)$ *and*

$$V(W_k) \geq V(\delta^k). \tag{6}$$

Part (i) captures the immediate observation that the controller can do no worse if she only requires one vote, regardless of the voter's identity, rather than the vote of a given voter. Conversely, the controller cannot benefit from securing the approval of all voters simultaneously rather than the approval of a given voter.

Part (ii) states that if voters are sufficiently aligned — i.e., all voters rank states in the same order — then the controller would prefer a decision process where he needs to persuade at least k voters, rather than persuading the $k-$th hardest-to-persuade voter. That is, the controller benefits from some heterogeneity, but requires some alignment between voters. The intuition is that, when voters rank states in the same order, then the approval signal realization under an optimal signal to the k-th hardest-to-persuade voter also induces approval for any voter $i < k$. Therefore, $V(W_k)$ cannot fall below $V(\delta^k)$. Finally, the controller suffers no loss from persuading a collection of voters under a unanimity rule rather than the hardest-to-persuade individual. That is, under unanimity (6) is satisfied with equality.

Inequality (6) holds whenever voters agree on the ranking of states. If voters rank states differently, then the reverse inequality to (6) may hold. The reason is that an optimal signal when facing the k-th hardest-to-persuade voter may not secure approval from all easier-to-persuade voters $i < k$ (see Example 5 in the online Appendix B). Interestingly, sometimes an

optimal signal does not target the easiest-to-persuade voter, even when voters agree under full information and rank states in the same order (see Example 6 in the online Appendix B).

4   Institutional Design

We start this section with a simple question: given a voting rule, do voters benefit from the signal chosen by the controller? We then study how different voting rules affect the payoffs of different voters. Importantly, voting rules affect outcomes not only through the consensus required to approve a proposal, but also by the amount of information that voters endogenously receive. For example, under delegation the controller provides the sole decision maker with some information about the state, although the signal has no value for the latter (cf. Proposition 1). As a result, a single decision maker may prefer to delegate the choice to someone with different preferences than himself, but who will elicit more information about the benefits of the proposal.[9] Voters face a similar trade-off between control and information when evaluating different k-voting rules: a higher consensus (i.e. higher k) may lead to excessive rejection of the proposal, but may induce the controller to provide a more informative signal. We study this trade-off by posing two questions: (i) under delegation, what are each voter's preferences over decision makers?, and (ii) when decisions are made collectively, how do preferences over decision makers translate into preferences over k-voting rules?

4.1   Do voters benefit from the controller's signal?

We start by comparing each voter's ex ante expected payoff under two scenarios: their equilibrium payoff when the controller provides signal π^*, and their equilibrium payoff if there was no controller providing a signal, and voters had to choose a policy solely on the basis of their prior beliefs.

Clearly, if there is a single voter, then he cannot be made worse off be the controller's

---

[9]In the context of organizations, Jensen and Meckling (1976) are among the first to point out that delegation decisions are guided by a fine balance between the loss of control owing to conflict of interest, and the gain of information when delegating to experts (see also, Holmstrom (1982), Dessein (2002), Alonso and Matouschek (2008), and Armstrong and Vickers (2010)).

influence. In fact, the controller's optimal signal does not change his expected payoff. Similarly, if all voters in the electorate have the same type, then the expected payoff of all voters is the same with or without the controller, independently of the k-voting rule. This is not the case when voters have different preferences, as summarized by the next Corollary.

Corollary 2 *Fix a k-voting rule and consider the electorate $\{\delta^1, \ldots, \delta^n\}$. Compare voters' ex ante expected payoff under the controller's optimal signal π^* and under no signal.*
(i) If the voting rule is unanimity, then all voters are weakly better off under the controller's influence, independently of the prior belief;
(ii) If $k < n$ and $p \in W_k$, then the controller's influence does not affect payoffs;
(iii) If $k < n$ and $p \notin W_k$, then at most $k-1$ voters are strictly better off under the controller's influence. Thus, at least $n-k+1$ voters are weakly worse off under the controller's influence. These voters are strictly worse off if there is no optimal signal with a binary realization space.

In particular, with a simple majority voting rule, a majority of voters is weakly worse off because of the controller's influence.

For any given k-rule, if $p \in W_k$, then the controller's optimal signal reveals no relevant information and voters approve the proposal. In this case, the controller's influence has no impact on voters' expected payoffs — which concludes part (ii). Part (i) follows from the same logic when $p \in W_n$, and the veto power of voters when $p \notin W_n$: if the rule is unanimity, then in order to approve the proposal the controller must convince all voters at the same time. However, for any non-unanimous voting rule, the controller can exploit preference disagreement by choosing signal realizations that target different winning coalitions. Part (iii) highlights that it cannot be the case that k voters are strictly better off by the controller's influence. Otherwise, the controller could strictly increase the probability of approval by choosing a less informative signal that leaves the same k voters weakly better off, but at least one of them indifferent. Moreover, whenever the posterior belief $q^+ \in co(W_k)$ obtained from combining all approval realizations of π^* is such that $q^+ \notin W_k$, then $n-k+1$ voters are strictly worse off. This is the case if there is no optimal signal with only two signal realizations, which implies that the controller must be targeting different winning coalitions.[10]

---

[10]There is always an optimal signal with only one realization that leads to rejection — therefore, if every

Finally, with a simple majority rule, a majority of voters can be made strictly worse off by the controller's signal even when all voters agree under full information and rank states in the same order (see Example 4 in Section 6.1) .

4.2   Voter preferences over decision makers

Suppose that the approval decision is made by a single voter δ: the controller only needs to persuade this voter. Now suppose that voter δ can choose whom to delegate the approval decision. How would voter δ rank different decision makers? As mentioned earlier, voter δ faces a well known trade-off between the gain in information and a loss of control: delegating to someone with different preferences can lead to inferior decisions, but may induce the controller to provide a more valuable signal. To study this trade-off, we first characterize voter preferences over decision makers for a suitably-defined restricted domain. We then show that, in these domains, a voter can always resolve the previous trade-off perfectly as a voter's preferred decision maker would (i) induce from the information controller a most valuable signal for voter δ, and (ii) for that signal, there is no loss of control.

The next proposition describes the preferences of a voter over decision makers that belong to the same class $\mathscr{F}_z$, that is, rank states in the same order z.

Proposition 4 *Fix a permutation z and let $\delta^v \in \mathscr{F}_z$. Consider any totally ordered (according to toughness) set of voters $\mathcal{D} \subset \mathscr{F}_z$, and suppose that the approval decision is delegated to a voter in $\mathcal{D}$ prior to the controller supplying a signal π. Then,*
(i) Voter δ^v has single-peaked preferences over decision makers in $\mathcal{D}$. That is, there exist $\bar{\delta} \in \mathcal{D}$ such that for $\delta, \delta' \in \mathcal{D}$, voter δ^v would (weakly) prefer to delegate to voter δ' instead of voter δ if either $A(\bar{\delta}) \subset A(\delta') \subset A(\delta)$ or $A(\delta) \subset A(\delta') \subset A(\bar{\delta})$.
(ii) If all voters in $\mathcal{D}$ agree with δ^v under full information, then voter δ^v has monotone preferences over decision makers in $\mathcal{D}$. That is, for $\delta, \delta' \in \mathcal{D}$, voter δ^v would (weakly) prefer to delegate to voter δ' instead of voter δ if δ' is tougher.
(iii) The maximum expected utility of voter δ^v when delegating to any decision maker in $\mathbb{R}^{|\Theta|}$,

<hr>

optimal signal must have at least three signal realizations, then the controller needs at least two different signals leading to approval, which implies that the controller must be targeting different winning coalitions.

is achieved by any voter $\delta^* \left(\hat{\delta}, \delta^v \right) = \hat{\delta} - \hat{\gamma}(\hat{\delta})\mathbf{1} \in \mathscr{F}_z$, where $\hat{\delta} \in \mathscr{F}_z$ and

$$\hat{\gamma}(\hat{\delta}) = \sum_{\theta \in \left\{ \theta : \delta_\theta^v \geq 0 \right\}} p_\theta \hat{\delta}_\theta. \tag{7}$$

Parts (i) and (ii) of the proposition describe the preferences of voter δ^v over decision makers who share his ranking of states and are ordered according to toughness. This condition on alignment does not guarantee that there is no loss of control under delegation, as these decision makers may not have the same approval set as δ^v. Part (i) shows that a voter has single-peaked preferences over such decision makers. That is, the set inclusion ordering derived from toughness translates naturally to single-peaked preferences when one restricts attention to voters in the same class. Part (ii) shows that the voter's preferences become monotone when the decision makers agree with δ^v under full information.

These results follow from the basic structure of an optimal signal with delegation to a voter in $\mathscr{F}_z$: the controller sets a threshold state and the optimal signal induces approval if a state with a higher net value occurs. Then, switching to a tougher decision maker implies a (weakly) higher threshold state and a (weakly) smaller set of approval states. Importantly, a tougher decision maker induces a signal that discriminates better between states of higher net value and states of lower net value for all voters in $\mathscr{F}_z$. Therefore, switching to a marginally tougher decision maker benefits voter δ^v whenever the current threshold state has a negative net payoff, but it proves detrimental whenever this net payoff is positive. If all decision makers agree with δ^v under full information, then this net payoff is always negative.

Part(iii) identifies in $\mathscr{F}_z$ an ideal decision maker for voter δ^v. If voter δ^v could both choose the signal π and decide whether to approve the proposal, then he only needs to learn whether the realized state corresponds to a positive net value. He can induce the controller to produce such a signal by delegating to a voter $\delta^* \left(\hat{\delta}, \delta^v \right) = \hat{\delta} - \hat{\gamma}(\hat{\delta})\mathbf{1}$, with $\hat{\gamma}(\hat{\delta})$ given by (7). Note however that voter $\delta^* \left(\hat{\delta}, \delta^v \right)$ and voter δ^v disagree under full information: voter $\delta^* \left(\hat{\delta}, \delta^v \right)$ would reject the proposal more often than δ^v if they perfectly learned the state. Nevertheless, they fully agree on the decision given the controller's optimal signal. In this sense, the fact that the signal is not fully revealing eliminates the loss of control when delegating to a tougher voter. Therefore, by delegating to $\delta^* \left(\hat{\delta}, \delta^v \right)$ voter δ^v achieves the same expected value as if he both made decisions and controlled the signal himself.

4.3 Voter preferences over k-voting rules

How does each voter rank different voting rules? Recall that voting rules affect voters' payoffs via two channels: the consensus required to approve a proposal, and the endogenous signal chosen by the controller. The following result follows from Corollary 2.

Corollary 3 *Consider an electorate $\{\delta^1, \ldots, \delta^n\}$ with an odd number $n \geq 3$ of voters. If $p \notin W_{\frac{n+1}{2}}$, then a majority of voters weakly prefer unanimity voting rule over simple majority.*[11]

To draw stronger inferences about voters' equilibrium payoff under different voting rules, we need to consider the nature of voters' preference heterogeneity. The next lemma shows that, if voters belong to the same class, then each voter has single peaked preferences over k.

Lemma 4 *Consider an electorate $\{\delta^1, \ldots, \delta^n\}$, with $\delta^i \in \mathscr{F}_z$, for some permutation z. Then each voter δ^i has single peaked preferences over k, in the sense that there exists $k^*(\delta^i)$ such that his expected utility is non-decreasing in k for $k < k^*(\delta^i)$, and it is not increasing for $k > k^*(\delta^i)$.*

Proposition 2 shows that voters' expected utilities with a $k-$voting rule are the same as with delegation to the weak-representative voter $\delta^*(k)$, as long as $\delta^*(k)$ strictly ranks states. Lemma 3 established that if all voters are in the same class, then the weak representative voter also belongs to that class. The intuition behind Lemma 4 is that since the weak-representative voter $\delta^*(k)$ also belongs to the same class $\mathscr{F}_z$, then a voting rule requiring a higher consensus is equivalent to delegating to a tougher voter. As a result, the collection of representative voters $\delta^*(k)$ describes a totally ordered set of voters in $\mathscr{F}_z$, and Proposition 4(i) implies that each voter has single-peaked preferences over these decision makers, and hence, over k-voting rules.[12]

An important implication of Lemma 4 is that a majority of voters prefer a supermajority voting rule over a simple majority voting rule.

---

[11]If $p \in W_{\frac{n+1}{2}}$, then a majority of voters might prefer simple majority over unanimity when unanimity makes approving the project too unlikely, e.g., if the win set W_n is empty.

[12]If voters do not agree on the ranking of the states, then preferences might not be single-peaked even when voters agree under full information and are totally ordered according to toughness. See Corollary 4 and Example 7 in the online Appendix B.

Lemma 5 *Consider and electorate $\{\delta^1, \cdots, \delta^n\}$ with an odd number $n \geq 3$ of voters in the same class $\delta^i \in \mathscr{F}_z$, and $p \notin W_{\frac{n+1}{2}}$. Then a majority of voters:*
(i) weakly prefer any supermajority voting rule $k' > \frac{n+1}{2}$ over simple majority $k = \frac{n+1}{2}$; and
(ii) strictly prefer supermajority k' over simple majority if it leads to a lower (but positive) equilibrium probability of approval, $0 < V(W_{k'}) < V(W_{\frac{n+1}{2}})$.

The next Proposition provides sufficient conditions for all voters to have the same preferences over k-voting rules.

Proposition 5 *Suppose that all voters are in $\mathscr{F}_z$ and they agree under full information. Then every voter weakly prefers a $(k+1)$-voting rule to a k-voting rule, for $k \in \{1, ..., n-1\}$.*

This proposition implies that even heterogenous voters may have the same preferences over electoral rules. In fact, as long as there is agreement under full information and voters rank states in the same order, then they all prefer a unanimity rule to any other k-voting rule. Essentially, sufficient alignment among voters can induce perfect agreement over electoral rules if information is endogenous to the electoral rule. Indeed, while voters may disagree under uncertainty, if they agree under full information, then they also agree on the signal they would choose if they were in control of decisions and could design the signal themselves. The intuition is that the weak representative voters $\{\delta^*(1), \ldots, \delta^*(n)\}$ (i) belong to the same class, (ii) agree under full information, and (iii) are totally ordered according to toughness. Therefore, the conditions of Proposition 4(ii) apply and every voter has monotone preferences: as a higher k corresponds to a tougher weak-representative voter, voters prefer rules that require more consensus only because they induce the controller to supply a more valuable signal.

5 Extensions

5.1 Controller knows the State

In our basic setup the controller has no private information. Suppose instead that the controller privately observes the true state θ before choosing signal π. In this case, the

choice of π by the informed controller may itself convey information to voters. We ask two questions: does the controller benefit from her private information? and what is the signal that maximizes the expected payoff of the informed controller, when expectation over controller's types is taking according to the prior p? In the online Appendix B we first apply the results from Alonso and Câmara (2014b) to show that the controller cannot benefit from privately observing the state. We then show that the maximum expected payoff is achieved in pooling equilibria where: (i) all controller's types choose the same signal π^*, and (ii) π^* is also an optimal signal in the case of an uninformed controller. Together these two results imply that the equilibrium probability of approving the proposal is unaffected by the controller privately learning the state.

5.2 Controller's Payoff Depends on the State

In our basic setup we focus on the case of pure-persuasion. We now consider a controller with a state-dependent payoff $u_C(x, \theta) : X \times \Theta \to \mathbb{R}$. Let $\delta_\theta^C = u_C(x_1, \theta) - u_C(x_0, \theta)$ and define the controller's type $\delta^C = \{\delta_\theta^C\}_{\theta \in \Theta}$. To simplify presentation, suppose $\delta_\theta^C \neq 0$.

First suppose that the approval decision is delegated to voter δ. Proposition 6 in the online Appendix B generalizes Proposition 1. It shows that the optimal signal always induces an approval realization for states such that both controller and voter agree on approval, and it always induces a rejection realization for states such that both agree on rejection. In the set of states where there is disagreement, the optimal signal again defines a cutoff state. However, in the case of pure persuasion the cutoff state was defined by ordering the states solely according to the voter's net payoff δ_θ; now the cutoff is defined by ordering the disagreement states according to the absolute value of the ratio of players preferences, $\left| \frac{\delta_\theta^C}{\delta_\theta} \right|$.

In many important cases the controller ranks states in the same order as the voter. For example, the controller receives the same payoff as the voter, plus some private benefit from approving the proposal (see also our application in Section 6.1). Proposition 7 in the online Appendix B shows that if the information controller and the electorate rank states in the same order, then each voter has single-peaked preferences over k-voting rules. Moreover, if voters also agree under full information, then the payoff of every voter is weakly increasing in k.

To understand the result, consider delegation to the weak representative voter $\delta^*(k)$. If

all players agree on the ranking of the states, then $\delta^*(k)$ also agrees. The controller's cutoff state is then defined according to this common ranking. A higher k-voting rule implies a tougher $\delta^*(k)$ and a weakly higher cutoff state, without changing the ranking. Consequently, all the results of Lemmas 4 and 5, and Proposition 5 continue to hold. Moreover, the proof of Proposition 7 shows that if all players rank states in the same order and the controller is sufficiently biased towards approval, then π^* is an optimal signal for a controller with type δ^C if and only if π^* is an optimal signal in the case of pure-persuasion.

5.3   Preference Shocks

In our basic setup the controller knows the preference profile of the electorate. However, in some instances voters are subject to idiosyncratic preference shocks, in which case the controller faces a probability distribution over preference profiles. To study the effects of preference shocks, assume that voter i's preferences are given by $u_i(x, \theta, \mu_i)$ and the conditional net payoff from approval with state θ and private shock μ_i is

$$u_i(x_1, \theta, \mu_i) - u_i(x_0, \theta, \mu_i) = \delta_\theta^i - \mu_i.$$

Shocks μ_i are i.i.d. and jointly independent with θ, with each shock distributed according to $F(\mu)$ with support in $[-\bar{\mu}, \bar{\mu}]$. The controller chooses the signal before shocks are realized.

Proposition 8 in the online Appendix B shows that if preference shocks are small and high shocks are sufficiently likely, then with unanimity the controller behaves as if she is facing the toughest electorate. That is, the controller benefits from choosing a more informative signal that persuades even the electorate profile where each voter received the worst possible shock $\bar{\mu}$. Under this optimal signal, voters have almost surely a strict preference between approval and rejection of the proposal, so that voters obtain a strictly positive gain with probability 1 when they approve the proposal. Therefore, unlike the case of non-probabilistic voting, with unanimity all voters can strictly benefit from the controller's influence. For non-unanimous voting rules, the results from Corollary 2 carry over to cases with a small $\bar{\mu}$. In particular, with a simple majority voting rule, a majority of voters can be made strictly worse off by the controller's signal.

The same is true with delegation: the controller's signal targets to persuade the voter with the worst possible shock. Hence, the voter can strictly benefit from the signal.

5.4   Heterogenous Prior Beliefs

In our base model, players share a common prior belief about the consequences of different policies. As argued by Alonso and Câmara (2014a), however, heterogeneous priors provide a powerful motive for persuasion, as a controller typically gains from shaping the learning of decision makers in the face of open disagreement. We can extend our main analysis to the case of heterogenous priors as follows. Suppose players hold different prior beliefs $p^l \in int(\Delta(\Theta))$, with $l \in \{C, 1, \ldots, n\}$. Suppose that the controller's signal is *commonly understood* in that all players agree on the conditional probabilities generating each realization. Then we can use the results from Alonso and Câmara (2014a) to characterize the controller's optimal signal and her gain from information control. We now briefly discuss how heterogenous priors affects our insights from Sections 3 and 4.

With delegation to voter δ and common priors, the controller's optimal signal defines a cutoff state where states are ordered solely according to the voter's net payoff δ_θ. In the case of heterogeneous priors, the optimal signal continues to define a cutoff state. However, the ordering of states might change depending on prior beliefs. Formally, the controller ranks states according to $\delta_\theta^i \frac{p_\theta^i}{p_\theta^C}$ and induces rejection only for the negative states with the lowest $\delta_\theta^i \frac{p_\theta^i}{p_\theta^C}$. For example, the controller might now find it optimal to have a state θ with a very negative δ_θ inducing an approval signal simply because the controller assigns a very high prior belief to θ, while the voter believes that θ is very unlikely. In other words, the controller favors approval realizations for states with negative payoffs whose likelihood he believes the voter underestimates.

Proposition 5 showed that if voters share the same ranking of states and agree under full information, then they all have the same preferences over voting rules. In particular, unanimity is preferred to any other k-voting rule. This does not hold, however, if voters have heterogenous prior beliefs. Note that open disagreement does not *per se* induce disagreement over the public signal. Indeed, under the conditions of Proposition 5 all voters have the same preferences over the class of binary "approve-reject" signals that preserve the ranking

28

of states — i.e., signals with a cutoff state with higher ranked states always inducing the approval realization. The fact that Proposition 5 no longer holds with heterogenous priors owes to the fact that the controller's signal no longer follows a cutoff on the ranking of states given by δ_θ^i, but rather in the ranking according to $\delta_\theta^i \frac{p_\theta^i}{p_\theta^C}$. Nevertheless, if the two rankings coincide, then the results of Proposition 5 still hold with heterogenous priors.

6 Applications

6.1 Voting on a Public Good

Consider a one-period k-voting model where an odd number $n \geq 3$ of voters must choose whether to approve $(x = x_1)$ or not $(x = x_0)$ the investment on a new public good, e.g., construction of a new highway overpass to improve traffic. If implemented, the cost c of the project is paid through a proportional tax t. Each voter i has a pre-tax income w_i and the government budget must balance. For simplicity, suppose there are no other government expenditures. Hence, the status quo tax is $t_0 \equiv 0$, and it increases to $t_1 \equiv \frac{c}{\sum_{i \in I} w_i}$ if the project is implemented. Voters' payoff from the project depends on state $\theta \in \Theta \subset \mathbb{R}$. This represents the uncertainty about how the overpass will affect the overall traffic flow. A voter-specific payoff $y_i : \Theta \to \mathbb{R}$ captures how each voter is affected by traffic flow changes, depending on factors such as where the voter lives and works. Let y_i be strictly increasing, so that a higher "quality" θ means a better traffic outcome. The utility function of each voter is then

$$u_i(x, \theta) = \begin{cases} (1 - t_1)w_i + y_i(\theta) & \text{if} \quad x = x_1, \\ w_i & \text{if} \quad x = x_0. \end{cases}$$

For each voter i compute the net payoff from approval

$$\delta_\theta^i = (1 - t_1)w_i + y_i(\theta) - w_i = y_i(\theta) - t_1 w_i. \tag{8}$$

All voters belong to the same class $\mathscr{F}_z$ since δ_θ^i strictly increases in θ. Voter i with posterior belief q votes to approve the project if and only if the expected payoff from the traffic outcome is greater than how much he has to pay in taxes to implement it, $E[y_i(\theta)|q] \geq t_1 w_i$.

Consider an information controller who has vested interests on the project — e.g., the controller is the Governor who proposed the project, but she needs voters to approve the ballot measure. Suppose that the Governor ranks states in the same order as voters. For example, her net payoff is proportional to the change in her "political capital," which is increasing in the quality of the project. Moreover, suppose she receives additional private benefits (e.g., ego rents) from approving the project.[13]

Lemma 3 imply that for each $k-$voting rule there is a weak representative voter $\delta^*(k) \in \mathscr{F}_z$, and from the point of view of all players the k-voting rule is payoff-equivalent to delegating the approval decision to $\delta^*(k)$. Moreover, the controller's optimal signal π^* defines a cutoff quality θ_k^* such that the project is always rejected if the quality is below the cutoff, $\theta < \theta_k^*$, and the project is approved with certainty if the quality is above the cutoff, $\theta > \theta_k^*$. If it is optimal to target different winning coalitions, then π^* contains multiple signal realizations that lead to approval. Cutoff θ_k^* weakly increases with k. Importantly, if the controller is more biased towards approval than voters, that is, $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta^*(k) < 0$, then a signal is optimal for controller δ^C if and only if it is optimal under the pure-persuasion benchmark (see the proof of Proposition 7 in the online Appendix B).

Next we present two examples based on this general setup. Example 3 considers voters with homogenous preferences for the public good but different incomes, which affects their tax burden. It shows that the voter with the median income can benefit from delegating the approval decision to a richer voter. Example 4 considers voters with heterogeneous preferences. It shows that under a simple majority voting rule a majority of voters can be made strictly worse off by the controller's influence, even when voters have the same income, agree under full information, and rank states in the same order.

Example 3: Suppose voters have homogeneous quality preferences $y_i = y, i \in I$. Voter i approves the project if and only if $E[y(\theta)|q] \geq t_1 w_i$. Therefore, voters are totally ordered — voters with higher income are both harder-to-persuade and tougher, $w_i < w_j$ implies $V(\delta^i) \geq V(\delta^j)$ and $A(\delta^i) \supset A(\delta^j)$. Let δ^k be the voter with the $k-$th lowest income. Voter δ^k is then a representative voter and a k-voting rule is equivalent to delegating the

---

[13]Note that the controller's ranking of the states does not change if her private benefit from approving the project is either constant or strictly increasing with the project's quality.

decision to him. Increasing the k-voting rule implies that the controller must target a richer voter. Suppose that the controller is more biased towards approval than the median voter δ^m, $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta^m < 0$. Lemma 5 implies that a majority of voters (the median and richer voters) weakly prefer any supermajority voting rule over simple majority. Moreover, this preference relation is *strict* if voter δ^k is strictly richer than the median voter and his approval set is not empty. By delegating the approval decision to a richer voter, who pays more to implement the project, the electorate induces the controller to supply a more informative signal. This result does not require the median voter to agree with δ^k under full information. □

Example 4: Suppose $y_i = \theta^{\beta_i}$, and consider three voters with $\beta_1 = 0.1$, $\beta_2 = 0.5$, $\beta_3 = 0.9$. Voters have the same income $w_i = 5$. If implemented, the project costs 1.5, so the proposed tax $t_1 = 0.1$ runs against the status quo $t_0 = 0$. There are three possible quality levels for the project: it does not improve traffic ($\theta = 0$), it moderately improves traffic ($\theta = 0.7$), or it greatly improves traffic ($\theta = 1.4$), so that $\Theta = \{0, 0.7, 1.4\}$. From (8) we have $\delta_\theta^i = \theta^{\beta_i} - 0.5$, so $\delta^1 \approx \{-0.5, 0.46, 0.53\}$, $\delta^2 \approx \{-0.5, 0.34, 0.68\}$, $\delta^3 \approx \{-0.5, 0.23, 0.85\}$. Voters would like to reject the project if it does not improve traffic, and approve if it has a moderate or great impact on traffic. Figure 3 depicts the prior belief p, the approval set of each voter, and the win set with simple majority. Note that there is no representative voter. The win set is not convex and the dotted lines delineate the convex hull of W_2. Consider a controller who prefers to approve the project in every state, which implies that she is more biased towards approval than voters. There is no optimal signal with only two signal realizations, but there is a π^* with three signal realizations. One realization induces posterior q^- and all voters reject the project. Another induces posterior q_1^+: voters 1 and 3 approve the project, while voter 2 strictly prefers to reject. The remaining realization induces posterior q_2^+: voters 2 and 3 approve the project, while voter 1 strictly prefers to reject. Note that the weighted average of the two approval posterior beliefs is a belief on the dotted line connecting q_1^+ and q_2^+. This average approval belief belongs to the convex hull of W_2, but it does not belong to W_2. Consequently, a majority of voters (voters 1 and 2) are made *strictly* worse off by the controller's influence. They strictly prefer the controller not to release the signal π^*, so that voters keep their prior and vote to reject the proposal. Even though all voters

agree under full information and rank states in the same order, they sometimes disagree under uncertainty because of the differences in the curvature of their utility functions. The information controller exploits this disagreement by designing a partially informative signal that targets different winning coalitions. Finally, all voters strictly prefer unanimity over simple majority, to induce the controller to provide a more informative signal. □
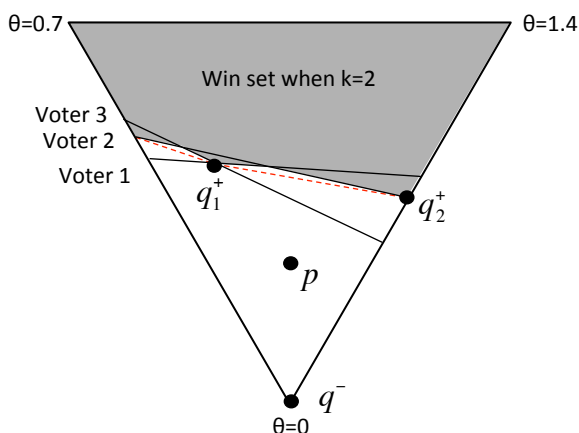

Figure 3: Win Set and Optimal Signal from Example 4

6.2 Spatial Model of Elections

Consider an election where a left-wing incumbent politician is running for re-election against an untried challenger from the opposing right-wing party. Let $X = \{L, R\}$, where L represents re-electing the incumbent and R electing the challenger. Each voter has a utility function $u_i(x, \theta) = -(y_x - y_i)^2$, where y_i captures the ideology of voter i and y_x is the policy implemented by politician x. There is an odd number $n \geq 3$ of voters with ideologies symmetrically distributed around the median voter $y_{\text{median}} = 0$. Voters know more about the incumbent than the challenger. Formally, voters know that the incumbent is committed to a policy $y_L < 0$, but they are uncertain about the policy $y_R > 0$ that the challenger would implement if elected. Voters' prior belief is that $y_R = \theta$ with probability p_θ, where $\theta \in \Theta \equiv \{\theta_1, \ldots, \theta_M\}$, $0 \leq \theta_1 < \ldots < \theta_M$, and $\theta_1 < |y_L| < \theta_M$. Suppose that without further information the median voter strictly prefers the incumbent.

For each voter i the net payoff from electing the challenger is

$$\delta_\theta^i = -(y_R - y_i)^2 + (y_L - y_i)^2 = -(\theta^2 - y_L^2) + 2(\theta - y_L)y_i. \tag{9}$$

Voter i strictly prefers the challenger if $(\theta^2 - y_L^2) + 2(\theta - y_L) > 0$, that is, if $y_i > \frac{\theta + y_L}{2}$, and he strictly prefers the incumbent if $y_i < \frac{\theta + y_L}{2}$. Therefore, the voter with the k-th lowest ideology y_i is the representative voter. Consider a simple majority rule, so that the median voter is decisive. Under full information, the median voter strictly prefers the challengers if $y_R < |y_L|$, and strictly prefers the incumbent if $y_R > |y_L|$. From the median voter's perspective, the relevant information for his decision is: who is more moderate, the challenger or the incumbent?

Let the information controller be a right-wing Interest Group (IG) with ideology $y_C > 0$ and payoff $u_C(x, \theta) = -(y_x - y_C)^2$. In this spatial model, the controller and voters rank states in different orders.[14] Nevertheless, in the online Appendix B we show that the optimal signal π^* defines a cutoff state $\theta_R^* > |y_L|$: the challenger losses for sure if he is "too radical", $\theta_R > \theta_R^*$, and he wins for sure if $\theta_R < \theta_R^*$. Importantly, we also show that there exists an ideology cutoff $\bar{y} > 0$ such all radical IG's with ideology $y_C > \bar{y}$ behave as in the pure-persuasion benchmark. That is, π^* is an optimal signal for these policy-motivated IGs if and only if π^* is an optimal signal for a purely office-motivated IG.

Suppose that the IG is radical, $y_C > \bar{y}$. Proposition 1 implies that the IG's influence does not affect the payoff of the decisive median voter. However, voters to the left of the median are hurt by the IG's influence, while voters to the right are better off. Although players do not rank states in the same order, the results from Lemma 5 continue to hold: a strict majority of voters (the median voter and all left-wing voters) prefer any supermajority voting rule over simple majority. Supermajority implies that a voter to the left of the median becomes decisive, which induces the IG to provide a more informative signal.

If the IG is moderate, $y_C < \bar{y}$, then the IG's preferences are sufficiently aligned with the median voter. In this case, the more informative signal provided by the IG strictly benefits the median voter.

---

[14]Note that $\frac{\partial \delta_\theta^i}{\partial \theta} = 2(y_i - \theta)$. Therefore, a strict majority of voters — voters with ideology $y_i < \theta_1$ — rank all states in the same decreasing order. Voters $y_i > \theta_M$ rank all states in the same increasing order.

6.3 Winners and Losers

We now study an application inspired by the model of Fernandez and Rodrik (1991), who highlight the role of individual-specific uncertainty when voters must decide whether or not to engage in an economic reform.[15]

There are three sectors in the economy, L, M and R. The population of workers, who are also voters, is distributed uniformly across the sectors. Voters must decide whether to implement an economic reform x_1 (e.g., sign a trade agreement with other countries) that increases the productivity of one sector, but decreases the productivity of the other sectors. Players have a uniform prior believe over which sector $\theta \in \Theta = \{L, M, R\}$ will benefit from the reform. The reform increases the payoff of workers in sector θ by $+1$, and decreases the payoff of all other workers by -1.

Consider a simple majority voting rule. Without further information, each worker believes that he is more likely to be a loser than a winner. Therefore, the proposal delivers a negative expected payoff and all voters reject the proposal. With full information about the state, voters in the winning sector θ vote to approve, but voters in the two losing sectors form a majority and reject the proposal.

Consider an information controller who wants to maximize the probability of approval. The controller can design a partially informative signal that guarantees the approval of the proposal. The optimal signal does not reveal the identity of the winning sector. Instead, it reveals the identity of one losing sector.[16] Upon learning this information, the losing sector votes to reject, but the two other sectors vote to approve. They now believe that there is an equal chance of being a winner or a loser.

With the controller's influence and a simple majority rule, the proposal is approved independently of the state. Consequently, the controller's strategic information provision strictly lowers the expected payoff of *all* voters. All voters would strictly prefer a unanimity voting rule to block the influence of the controller. With unanimity, the win set is empty

---

[15]We are also grateful for suggestions by Navin Kartik.

[16]Formally, let $s \in S = \{L, M, R\}$, $Pr[s|\theta] = 0$ if $s = \theta$, and $Pr[s|\theta] = 0.5$ if $s \neq \theta$. Therefore, upon observing s, all players know that sector s is not the winner θ, and the two remaining sectors are equally likely to be the winner.

and the reform cannot be implemented.

7   Conclusion

In important cases, acquiring information is infeasible or prohibitively expensive for individual voters. Voters must then rely on the information generated by certain individuals, who control the design of a public signal (e.g., jurors and prosecutor, voters and media, shareholders and CEO). Obviously, if the controller and voters share the same preferences, then the controller's signal always benefits voters, as it allows them to make better decisions. However, this is not true if there is a conflict of interest between the controller and voters. We show that, with a simple majority rule, a majority of voters is *always* weakly worse off by observing the information provided. In fact, all voters can be *strictly* worse off, even when they would agree on their decision if they knew the true state. This is so because the controller strategically designs a signal with realizations targeting different winning coalitions. To prevent this negative impact, voters may adopt a supermajority voting rule that induces the controller to supply a more informative signal. We also provide conditions for unanimity to be the rule preferred by all voters.

We extend our analysis in a number of ways. We show that the controller cannot benefit from privately observing the state prior to choosing her signal. We study situations in which the controller also cares about the state and situations where voters are subject to idiosyncratic preference shocks. In these cases, a voter may now strictly benefit from the controller's signal if he is the sole decision maker, although a majority of voters can still be worse off under a simple majority rule. We also extend the analysis to allow voters to have heterogeneous prior beliefs, so that they openly disagree about the likelihood of the state. Importantly, even if they all share the same ranking over states and agree under full information, belief disagreement can translate into disagreement over the optimal electoral rule.

Two interesting extensions are to allow for voters' to privately acquire information and then deliberate prior to voting, and to allow voters to choose among multiple policy options. We see these extensions as promising and leave them for future work.

A Appendix

Proof of Lemma 1: Under the assumptions, the set $A(\delta)$ is non empty and the voter rejects the proposal if he has no additional information. Let π' be an arbitrary binary signal that induces posterior beliefs $\{q^-(\pi'), q^+(\pi')\}$ such that $q^-(\pi') \in R(\delta)$ and $q^+(\pi') \in A(\delta)$ with $\sum_{\theta \in \Theta} q_\theta^+(\pi')\delta_\theta = 0$. Define the vector l as $l = q^+(\pi') - p$. Then, Bayesian rationality implies that average posteriors must equal the prior so that

$$\Pr[Approval] \langle q^+(\pi') - p, l \rangle + (1 - \Pr[Approval]) \langle q^-(\pi') - p, l \rangle = 0,$$

and $q^-(\pi') - p$ and $q^+(\pi') - p$ are collinear so

$$\langle q^-(\pi') - p, q^+(\pi') - p \rangle = - \left\| (q^+(\pi') - p) \right\| \left\| (q^-(\pi') - p) \right\|.$$

Therefore,

$$\Pr[Approval] = \frac{\langle p - q^-(\pi'), l \rangle}{\langle q^+(\pi') - p, l \rangle + \langle p - q^-(\pi'), l \rangle} = \frac{\| (q^-(\pi') - p) \|}{\| (q^+(\pi') - p) \| + \| (q^-(\pi') - p) \|},$$

where, by construction, $\| (q^+(\pi') - p) \| = d_l(p, A(\delta))$ and $\| (q^-(\pi') - p) \| = d_l(p, R(\delta))$. As the optimal signal maximizes $\Pr[Approval]$, it must be that the optimal signal corresponds to a vector l^* that maximizes the ratio $\| (q^-(\pi') - p) \| / \| (q^+(\pi') - p) \|$. ∎

Proof of Proposition 1: The existence of an optimal binary signal is established in KG (Proposition 1, p. 2595). Let π be an optimal binary signal supported on $S = \{s^-, s^+\}$ where the voter approves the proposal if and only if he observes s^+, and let $\alpha_\theta = \Pr[s^+|\theta]$ so that $\Pr[Approval] = \sum_{\theta \in \Theta} \alpha_\theta p_\theta$. A Bayesian voter δ will approve after observing s^+ if and only if

$$E[\delta|s^+] = \sum_{\theta \in \Theta} q_\theta^+ \delta_\theta = \sum_{\theta \in \Theta} \frac{\alpha_\theta p_\theta \delta_\theta}{\Pr[Approval]} \geq 0$$

Therefore, α must solve the following linear program

$$\sum_{\theta \in \Theta} \alpha_\theta p_\theta = \max \sum_{\theta \in \Theta} \alpha_\theta' p_\theta, \quad s.t. \quad 0 \leq \alpha_\theta' \leq 1, \sum_{\theta \in \Theta} \alpha_\theta' p_\theta \delta_\theta \geq 0. \tag{10}$$

Note that for any θ' such that $\delta_{\theta'} \geq 0$ we must then have $\alpha_{\theta'} = 1$, as increasing $\alpha_{\theta'}$ whenever $\alpha_{\theta'} < 1$ relaxes the approval constraint and increases the approval probability. Therefore,

suppose that $\delta_\theta < \delta_{\theta'} < 0$ for $\theta, \theta' \in \Theta$. If $\alpha_\theta > 0$ but $\alpha_{\theta'} < 1$, then increasing $\alpha_{\theta'}$ by $\varepsilon \left(|\delta_\theta| p_\theta / |\delta_{\theta'}| p_{\theta'}\right)$ while reducing α_θ by ε leaves unchanged the approval constraint but increases the probability of approval by $\varepsilon p_\theta \left(|\delta_\theta| / |\delta_{\theta'}|\right) - \varepsilon p_\theta > 0$, thus leading to a contradiction. Therefore, if $\alpha_\theta > 0$, then $\alpha_{\theta'} = 1$ for any $\delta_{\theta'} > \delta_\theta$.

Now suppose that voter δ strictly ranks states so that $\delta_\theta \neq \delta_{\theta'}$ for $\theta \neq \theta'$. Then there can only be one "cutoff" state that satisfies (4), and given that the approval constraint is met with equality, the optimal binary signal must be unique.

If $p \in A(\delta)$ then $\alpha_\theta = 1$, $\theta \in \Theta$, and the voter receives a completely uninformative signal. If $p \notin A(\delta)$ then the solution to (10) must have a binding approval constraint, implying that the voter is indifferent between approval and rejection after observing s^+. That is, his expected gain from making decisions with π is again zero. ∎

Proof of Lemma 2: Follows immediately by replacing $A(\delta)$ with $co(W_k)$, and $R(\delta)$ with R_k, in Lemma 1 and then applying the same reasoning as in the proof of Lemma 1. ∎

Proof of Proposition 2: Let W_k be the win set under a k-voting rule and suppose that $p \notin co(W_k)$. Define $G(W_k)$ as

$$G(W_k) = \{\delta : q \in W_k \Rightarrow \langle q, \delta \rangle \geq 0\}$$

Note that each $\delta \in G(W_k)$ corresponds to a "less tough" voter than the k-voting rule, in the sense that any voter in $G(W_k)$ would approve the proposal if the electorate does so under a k-voting rule. Moreover, as $G(W_k)$ describes all hyperplanes that contain W_k, then we have (i) $co(W_k) = \cap_{\delta \in G(W_k)} A(\delta)$, and (ii) $V(W_k) \leq \inf_{\delta \in G(W_k)} V(\delta)$. We now show that there is a $\delta^* \in G(W_k)$ with $V(co(W_k)) = V(\delta^*)$ so that δ^* is a weak representative voter when the win set is $co(W_k)$. As explained in the text, for any belief in $co(W_k)$ the controller can find a signal that induces approval with probability 1. Therefore $V(W_k) = V(co(W_k))$, and δ^* is also a representative voter for W_k.

Define the function $f(\alpha) = \inf_{\delta \in G(W_k)} \langle \alpha, \delta p \rangle$, $\mathbf{0 \leq \alpha \leq 1}$, which is concave as it is the infimum of affine functions. The function $f(\alpha)$ provides a representation of $co(W_k)$, since $q = \frac{\alpha p}{\langle \alpha, p \rangle} \in co(W_k)$ if and only if $f(\alpha) \geq 0$.[17] Let s^+ be the event corresponding to approval of the

---

[17]For every $q \in \Delta(\Theta)$, the existence of a corresponding $\alpha \leq 1$ is guaranteed by simply choosing $\alpha_\theta =$

proposal under an optimal signal, and let $\alpha_\theta^* = \Pr[s^+|\theta]$ so that $\Pr[Approval] = \sum_{\theta \in \Theta} \alpha_\theta^* p_\theta$. Since the expected approval posterior must be in $co(W_k)$, then we must have $f(\alpha^*) \geq 0$. Thus, the controller's optimal signal must maximize $\Pr[Approval]$, i.e.

$$\sum_{\theta \in \Theta} \alpha_\theta^* p_\theta = \max_\alpha \sum_{\theta \in \Theta} \alpha_\theta p_\theta, \text{s.t. } 0 \leq \alpha_\theta \leq 1, f(\alpha) \geq 0. \tag{11}$$

Program (11) is concave (as it maximizes a concave function over a convex set). Consider the Lagrangian L associated to (11)

$$L = <\alpha, p> - \sum_\theta \nu_\theta < \alpha, 1_\theta > + \sum_\theta \mu_\theta < \alpha - \mathbf{1}, 1_\theta > -\kappa f(\alpha),$$

with $\nu_\theta, \mu_\theta, \kappa \geq 0$, and 1_θ is the unitary vector whose θ-component equals 1. Suppose that W_k is non-empty and has at least two different elements. This implies that W_k has a non-empty relative interior, so that the constraint qualification is satisfied and the Karush-Kuhn-Tucker conditions are both necessary and sufficient for optimality (Boyd and Vandenberghe 2004). In particular, when $p \notin co(W_k)$, α^* is an optimal solution if and only if there exist $\lambda^*, \nu_\theta^*, \mu_\theta^* > 0$, $\theta \in \Theta$, such that

$$\tilde{\delta} \equiv -\lambda^* p - \sum_\theta \nu_\theta^* < \alpha^*, 1_\theta > + \sum_\theta \mu_\theta^* < \alpha^* - \mathbf{1}, 1_\theta > \in \partial f(\alpha^*) \text{ and } f(\alpha^*) = 0, \tag{12}$$

where $\partial f(\alpha^*)$ is the set of subgradients of f at the point α^*, Define

$$\gamma^* = \left\langle \tilde{\delta}, \frac{\alpha^* p}{< \alpha^*, p >} \right\rangle,$$

and consider the voter $\delta^* = \tilde{\delta} - \gamma^* \mathbf{1}$. By construction, $\langle \delta^*, q \rangle = 0$ is a supporting hyperplane of $co(W_k)$. Now consider the optimal signal α' under delegation to voter δ^* which must satisfy

$$\sum_{\theta \in \Theta} \alpha_\theta' p_\theta = \max_\alpha \sum_{\theta \in \Theta} \alpha_\theta p_\theta, \text{s.t. } 0 \leq \alpha_\theta \leq 1, \langle \alpha, \delta^* p \rangle \geq 0.$$

Again, this is a concave program (in fact a linear program) with non-empty relative interior if $co(W_k)$ has a non-empty relative interior. Therefore α' is optimal if and only if there exist $\lambda^*, \nu_\theta^*, \mu_\theta^* > 0$, $\theta \in \Theta$, such that

$$-\lambda^* p - \sum_\theta \nu_\theta^* < \alpha', 1_\theta > + \sum_\theta \mu_\theta^* < \alpha' - \mathbf{1}, 1_\theta >= \tilde{\delta} \tag{13}$$

$\left(\frac{q_\theta}{p_\theta}\right) / \max \frac{q_{\theta'}}{p_{\theta'}}.$

In particular, as α^* satisfies (12) it also satisfies (13) and thus provides an optimal signal when delegating to voter δ^*. Therefore δ^* is a weak representative voter. Finally, suppose that the representative voter $\delta^*(k)$ strictly ranks states. Then, following Proposition 1 the binary optimal signal is unique and thus the expected utility of every player is the same under delegation to δ^* or under a k-voting rule. ∎

Proof of Lemma 3: We will prove the lemma by showing that if all voters in the electorate belong to $\mathscr{F}_z$ for some permutation z, then if $\{q \in \Delta(\Theta) : \langle q, v \rangle = 0\}$ is a supporting hyperplane of $co(W_k)$ then $v \in \mathscr{F}_z$.

Let S_k be the set of all $k-$coalitions of voters with generic element s. Then $\cap_{\delta \in s} A(\delta)$ describes the win set associated with a unanimous decision when the electorate is restricted to the coalition s, and $W_k = \cup_{s \in S_k} \cap_{\delta \in s} A(\delta)$. As $\cap_{\delta \in s} A(\delta)$ is the finite intersection of half-spaces $\{q \in \Delta(\Theta) : \langle q, \delta \rangle \geq 0\}$, then any supporting hyperplane of $\cap_{\delta \in s} A(\delta)$ at $q \in int(\Delta(\Theta))$ can be represented as a convex combination of $\{\delta : \delta \in s\}$. Moreover, $co(s) \subset \mathscr{F}_z$, as ranking of states is preserved under convex combinations. Thus any supporting hyperplane of $\cap_{\delta \in s} A(\delta)$ at an interior belief corresponds to a voter in $\mathscr{F}_z$.

Turning to $co(W_k)$, consider any point $q' \in int(\Delta(\Theta))$ with $q' \in \partial(co(W_k))$ and a supporting hyperplane $\{q : \langle q, v(q') \rangle = 0\}$ of $co(W_k)$ at q'. Since $q' \in co(W_k)$, Caratheodory's theorem guarantees the existence of $J \leq card(\Theta)+1$ points, $\{q^i\}$, $i = \{1, ..., J\}$, with $q^i \in W_k$ and $q' \in co(q^i, i = \{1, ..., J\})$. We consider two possibilities: (i) at least one of the points q^i is in $int(\Delta(\Theta))$, (ii) any representation of q' as a convex combination of points $\{q^i\}$ with $q^i \in W_k$ must correspond to points on the faces of the simplex $\Delta(\Theta)$, i.e. for each $i = \{1, ..., J'\}$ there exists θ_i such that $q^i_{\theta_i} = 0$.

Consider first case (i) in which $q^i \in int(\Delta(\Theta))$ for some i. Then $\{q : \langle q, v(q') \rangle = 0\}$ must also be a supporting hyperplane of W_k at q^i. Furthermore, since $q^i \in W_k$ there exists a $k-$coalition s such that $q^i \in \cap_{\delta \in s} A(\delta)$ and thus $\{q : \langle q, v(q') \rangle = 0\}$ is a supporting hyperplane of $\cap_{\delta \in s} A(\delta)$. However, as all supporting hyperplanes to $\cap_{\delta \in s} A(\delta)$ in $int(\Delta(\Theta))$ can be associated to a voter in $\mathscr{F}_z$, then $v(q') \in \mathscr{F}_z$.

Consider now case (ii) where every representation of q' as a convex combination of points in W_k involves points that lie on (possibly different) faces of the simplex $\Delta(\Theta)$, and let $\{q^i\}$,

$i = \{1, ..., J\}$ be such a collection of points. As each $q^i \in W_k$, and $\{q : \langle q, v(q') \rangle = 0\}$ is a supporting hyperplane of $co(W_k)$ then there exists a coalition s_i such that $\min_{\delta \in s_i} \langle q^i, \delta \rangle = 0$, and let δ^i be a voter at which this minimum is achieved. That is, coalition s_i would support approval for belief q^i with at least one voter being indifferent between approval and rejection.

Now suppose by way of contradiction that $v(q') \notin \mathscr{F}_z$. This means that there exist two states θ and θ' with $(\delta^i_{\theta'} - \delta^i_\theta)(v_{\theta'}(q') - v_\theta(q')) < 0$ for all $i \in \{1, ..., J\}$. Now consider the edge of beliefs $\Psi(\theta, \theta')$ that put positive probability only on θ and θ', i.e. $\Psi(\theta, \theta') = \{q : q = \alpha 1_{\theta'} + (1 - \alpha)1_\theta\}$ and let $\tilde{q} \in \Psi(\theta, \theta')$ be such that $\langle \tilde{q}, v(q') \rangle = 0$. As $v(q')$ is a support hyperplane, we must have $\langle \tilde{q}, \delta^i \rangle = 0$ for some $i = \{1, ..., J\}$. The fact that δ^i and $v(q')$ rank θ and θ' differently implies that either $v_{\theta'}(q') < 0 < \delta^i_{\theta'}$ or $v_\theta(q') < 0 < \delta^i_\theta$. In either case, it implies that there is one state that belongs to W_k (as it is approved by the coalition represented by) but does not lead to approval by voter $v(q')$. Therefore, $v(q')$ cannot be a supporting hyperplane of $co(W_k)$ and we reach a contradiction. ∎

Proof of Corollary 1: In the text.

Proof of Proposition 3: Part (i)- Follows immediately as any optimal signal under unanimity must induce approval of every voter, while an optimal signal for a voter δ^i would also induce approval if $k = 1$.

Part (ii)- Note that if all $\delta^i \in \mathscr{F}_z$, then Proposition 1 shows that the structure of the optimal signal is the same for all voters: if $\alpha_\theta(\delta^i) = \Pr[approval|\theta]$ represents the optimal signal under delegation to voter δ^i, where $\alpha_\theta(\delta^i)$ is given by (4), then $\alpha_\theta(\delta^{i'}) - \alpha_\theta(\delta^i) \leq 0$, $\theta \in \Theta$ if $V(\delta^{i'}) \leq V(\delta^i)$. This implies that signal $\alpha(\delta^k)$ would induce approval for any $i < k$ such that $V(\delta^k) \leq V(\delta^i)$. Therefore, the optimal signal to persuade voter δ^k has an approval signal realization that would induce the approval vote of at least k voters. Therefore $V(W_k) \geq V(\delta^k)$. ∎

Proof of Corollary 2: In the text.

Proof of Proposition 4: Without loss of generality, suppose that $z(i) = i$ so that for $\delta \in \mathscr{F}_z$ we have $\delta_{\theta_i} < \delta_{\theta_{i+1}}$, $i \in \{1, ..., card(\Theta) - 1\}$. From Proposition 1, the controller's optimal signal when the decision is made by voter $\delta \in \mathscr{F}_z$ is characterized by the approval

conditional probabilites $\alpha_\theta(\delta) = \Pr[approval|\theta]$ such that there exists $i^\alpha(\delta)$ with (i) $\alpha_{\theta_i}(\delta) = 0$ if $i < i^\alpha(\delta)$, (ii) $\alpha_{\theta_i}(\delta) = 1$ if $i > i^\alpha(\delta)$, and (iii) $\sum \alpha_\theta(\delta) p_\theta \delta_\theta = 0$. Also, for $\delta \in \mathscr{F}_z$ let $\underline{i}(\delta) = \min\{i : \delta_{\theta_i} \geq 0\}$. In words, if the realized state is θ_i then voter δ would approve the proposal under full information as long as $i \geq \underline{i}(\delta)$, while the optimal signal induces approval by voter δ only if $i \geq i^\alpha(\delta)$.

Part (i)- The increment in the expected utility of voter δ^v under delegation to δ rather than choosing always the status quo is

$$\Delta U = E[u_i(x(\delta), \theta)] - E[u_i(x_0, \theta)] = P(q^+(\delta)) \langle q^+(\delta), \delta^v \rangle = \sum \alpha_\theta(\delta) p_\theta \delta_\theta^v.$$

We now show that voter δ^v has single peaked preferences among voters in D. Select two voters $\delta, \delta' \in D$ with $A(\delta') \subset A(\delta)$. From Proposition 1, this implies that $\alpha_\theta(\delta') - \alpha_\theta(\delta) \leq 0$, $\theta \in \Theta$. First, suppose that $i^\alpha(\delta), i^\alpha(\delta') < \underline{i}(\delta^v)$. Then, $\alpha_{\theta_i}(\delta) = \alpha_{\theta_i}(\delta') = 1$ if $i \geq i'(\delta^v)$, and thus

$$\Delta U(\delta') - \Delta U(\delta) = \sum_{i < i'(\delta^v)} (\alpha_{\theta_i}(\delta') - \alpha_{\theta_i}(\delta)) p_{\theta_i} \delta_{\theta_i}^v \geq 0,$$

where the inequality follows from $\delta_{\theta_i}^v < 0$ if $i < i'(\delta^v)$. Second, suppose that $i^\alpha(\delta), i^\alpha(\delta') \geq \underline{i}(\delta^v)$. Then, $\alpha_{\theta_i}(\delta) = \alpha_{\theta_i}(\delta') = 0$ if $i < i'(\delta^v)$, and thus

$$\Delta U(\delta') - \Delta U(\delta) = \sum_{i \geq i'(\delta^v)} (\alpha_{\theta_i}(\delta') - \alpha_{\theta_i}(\delta)) p_{\theta_i} \delta_{\theta_i}^v \leq 0,$$

where the inequality follows from $\delta_{\theta_i}^v \geq 0$ if $i \geq i'(\delta^v)$.

Finally, divide voters in D into two groups $D^+(\delta^v) = \{\delta \in D : i^\alpha(\delta) \geq \underline{i}(\delta^v)\}$ and $D^-(\delta^v) = \{\delta \in D : i^\alpha(\delta) < \underline{i}(\delta^v)\}$. Then, for any $\delta, \delta' \in D^-(\delta^v)$, voter δ^v preferences over decision makers are given by their toughness, while if $\delta, \delta' \in D^+(\delta^v)$, voter δ^v prefers decision makers that are less tough. Therefore, δ^v has single peaked preferences over voters in any totally ordered chain (ordered according to toughness).

Part (ii)- Let $\mathscr{F}(D)_z$ be the set of voters that rank states according to z and who share the same set of approval states D. For any $\delta, \delta' \in \mathscr{F}(D)_z$ we have that $i^\alpha(\delta) < \underline{i}(\delta')$ (equality is ruled out as voters strictly rank states). In words, if voters both agree on the ranking of states and on decisions under full information, then the controller would provide voter δ with a signal that always induces approval in states for which voter δ' would want to approve.

Therefore, Proposition 4 implies that all voters in $\mathscr{F}(D)_z$ have monotone preferences over totally ordered chains in $\mathscr{F}(D)_z$.

Part (iii)- The maximum expected gain to voter δ^v (with respect to always selecting the status quo) if he can design the signal himself is

$$\Delta U^* = \sum_{i \geq i'(\delta^v)} p_\theta \delta^v(\theta).$$

This corresponds to (i) a signal that reveals whether or not a state with a non-negative net value occurred, i.e. if a state θ_i with $i \geq i'(\delta^v)$ occurred, and (ii) the proposal is selected in that case. But, this is precisely the signal that the controller provides to a voter $\delta^*\left(\hat{\delta}, \delta^v\right)$, as to induce approval the controller would need to supply a signal such that

$$E[\delta^*\left(\hat{\delta}, \delta^v\right) | s^+] = \sum_{\theta \in \Theta} q_\theta^+ \delta_\theta - \hat{\gamma}(\hat{\delta}) = \sum_{\theta \in \Theta} \left(\alpha_\theta - 1_{\{\theta : \delta_\theta^v \geq 0\}}\right) p_\theta \delta_\theta = 0,$$

which implies that $\alpha_\theta = 1$ only if $\delta_\theta^v \geq 0$, which corresponds to the optimal signal to voter δ^v. ■

Proof of Corollary 3: Corollary 2(i) implies that all voters weakly prefer unanimity with the controller's influence than rejecting the proposal without further information; assumption $p \notin W_{\frac{n+1}{2}}$ and Corollary 2(iii) imply that a majority of voters prefers to reject the proposal without further information than having simple majority under the controller's influence. Corollary 3 then follows immediately. ■

Proof of Lemma 4: Lemma 3 implies that if all voters are in the same class $\mathscr{F}_z$, then for each k there exists a representative voter $\delta^*(k) \in \mathscr{F}_z$ and, furthermore, $A\left(\delta^*(k')\right) \subset A\left(\delta^*(k)\right)$ for $k' > k$. Therefore $D = \{\delta^*(k) : k \in \{1, ..., n\}\}$ forms a totally ordered chain, and Proposition 4 implies that each voter in the electorate has single-peaked preferences in D. Finally, since $\delta^*(k) \in \mathscr{F}_z$, Proposition 2 implies that the expected utility of each voter under a k-voting rule is the same as delegating to $\delta^*(k)$. This implies that each voter has single peaked preferences over k. ■

Proof of Lemma 5: Consider the optimal binary signal targeting the weak representative voter $\delta^*(W_k)$. Let q_k^+ be the posterior belief after the approval signal. Under simple majority

rule there is a set M of voters, $card(M) \geq \frac{n+1}{2}$, such that for each $\delta \in M$ we have $\left\langle q^+_{\frac{n+1}{2}}, \delta \right\rangle \leq 0$. Hence the expected payoff of those voters under simple majority is weakly lower than their expected payoff from always rejecting the proposal. Moreover, since $W_n \neq \emptyset$, under unanimity the payoff of all voters $\delta \in M$ is weakly higher than their payoff from always rejecting the proposal, since unanimity implies $< q^+_n, \delta > \geq 0$. Therefore all voters in M weakly prefer unanimity over simple majority. Using Lemma 4, single-peaked preferences over k implies that all voters in M weakly prefer k' over simple majority, concluding the proof of part (i). Part (ii) follows from $0 < V(W_{k'}) < V(W_{\frac{n+1}{2}})$ because it implies that the optimal signal under k' is not the same as the signal under simple majority.

To see that the set M must exist, suppose by contradiction that it does not exist. Then there are at least $n - \frac{n+1}{2} + 1 = \frac{n+1}{2}$ voters such that $\left\langle q^+_{\frac{n+1}{2}}, \delta \right\rangle > 0$. Therefore, after observing $q^+_{\frac{n+1}{2}}$ a majority of voters strictly prefer to approve the proposal, a contradiction to the optimality of the signal. ∎

Proof of Proposition 5: Proposition 4(ii) applied to Lemma 3 implies that all voters have monotone preferences over voters in $D = \{\delta^*(k) : k \in \{1, ..., n\}\}$ when they all agree under full information. Since $\delta^*(k) \in \mathscr{F}_z$, Proposition 2 implies that the expected utility of each voter under a k-voting rule is the same as delegating to $\delta^*(k)$. This implies that each voter has monotone preferences over k. ∎

References

[1] ALONSO, R., AND O. CÂMARA (2014a): "Persuading Skeptics and Reaffirming Believers," *mimeo*.

[2] ALONSO, R., AND O. CÂMARA (2014b): "On the Value of Persuasion by Experts," *mimeo*.

[3] ALONSO, R., AND N. MATOUSCHEK (2008): "Optimal Delegation," *Review of Economic Studies*, 75(1), 259-293.

[4] ARMSTRONG, M., AND J. VICKERS (2010): "A Model of Delegated Project Choice," *Econometrica*, 78(1), 213-244.

[5] AUSTEN-SMITH, D., AND J. S. BANKS (1996): "Information Aggregation, Rationality, and the Condorcet Jury Theorem," *American Political Science Review*, 90(1), 34-45.

[6] AUSTEN-SMITH, D., AND T. FEDDERSEN (2005): "Deliberation and Voting Rules",*In Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks*, edited by David Austen-Smith and John Duggan, Springer-Verlag.

[7] BOYD, S., AND VANDENBERGHE, L. (2004): *Convex Optimization*, Cambridge University Press.

[8] BROCAS, I., AND J. CARRILLO (2007): "Influence through Ignorance," *Rand Journal of Economics*, 38(4), 931-947.

[9] DESSEIN, W. (2002): "Authority and Communication in Organizations," *Review of Economic Studies*, 69(4), 811-838.

[10] DOWNS, A. (1957): *An Economic Theory of Democracy*, Harper and Row, New York.

[11] DUGGAN, J., AND C. MARTINELLI (2011): "A Spatial Theory of Media Slant and Voter Choice," *Review of Economic Studies*, 78(2), 640-666.

[12] FEDDERSEN, T. J., AND W. PESENDORFER (1996): "The Swing Voter's Curse," *American Economic Review*, 86(3), 408-424.

[13] FEDDERSEN, T. J., AND W. PESENDORFER (1997): "Voting Behavior and Information Aggregation in Elections with Private Information," *Econometrica*, 65(5), 1029-1058.

[14] FEDDERSEN, T. J., AND W. PESENDORFER (1998): "Convicting the Innocent: The Inferiority of Unanimous Jury Verdicts Under Strategic Voting," *American Political Science Review*, 92(1), 23-35.

[15] FERNANDEZ, R., AND D. RODRIK (1991): "Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty," *American Economic Review*, 81(5), 1146-1155.

[16] GENTZKOW, M., AND E. KAMENICA (2013): "Costly Persuasion," *mimeo*, University of Chicago.

[17] GERARDI, D., AND L. YARIV (2008): "Information Acquisition in Committees," *Games and Economic Behavior*, 62(2), 436-459.

[18] GILL, D., AND D. SGROI (2008): "Sequential Decisions with Tests," *Games and Economic Behavior*, 63(2), 663-678.

[19] GILL, D., AND D. SGROI (2012): "The Optimal Choice of Pre-Launch Reviewer," *Journal of Economic Theory*, 147(3), 1247-1260.

[20] Ivanov, M. (2010): "Informational Control and Organizational Design," *Journal of Economic Theory*, 145(2), 721-751.

[21] Jackson, M. O., and X. Tan (2013): "Deliberation, Disclosure of Information, and Voting," *Journal of Economic Theory*, 148(1), 2-30.

[22] Jensen, M. C. and W. H. Meckling (1976): "Theory of the firm: Managerial Behavior, Agency Costs and Ownership Structure,", *Journal of Financial Economics*, 3(4), 305-360.

[23] Kamenica, E., and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590-2615.

[24] Kolotilin, A. (2014): "Optimal Information Disclosure: Quantity vs. Quality," *mimeo*.

[25] Li, H. (2001): "A Theory of Conservatism," *Journal of Political Economy*, 109(3), 617-636.

[26] Li, H. and Li, W. (2013): "Misinformation," *International Economic Review*, 54(1), 253-277.

[27] Lizzeri, A. and L. Yariv (2013): "Sequential Deliberation," mimeo.

[28] Martinelli, C. (2006): "Would Rational Voters Acquire Costly Information?," *Journal of Economic Theory*, 129(1), 225-251.

[29] Messner, M., and M. K. Polborn (2012): "The Option to Wait in Collective Decisions and Optimal Majority Rules," *Journal of Public Economics*, 96(5), 524-540.

[30] Michaeli, B. (2014): "Divide and Inform: Rationing Information to Facilitate Persuasion," *mimeo*.

[31] Persico, N. (2004): "Committee Design with Endogenous Information," *Review of Economic Studies*, 71(1), 165-191.

[32] Rayo, L., and I. Segal (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118(5), 949-987.

[33] Schnakenberg, K. E. (2014a): "Expert Advice to a Voting Body," *mimeo*.

[34] Schnakenberg, K. E. (2014b): "Informational Lobbying and Legislative Voting," *mimeo*.

[35] Taneva, I. (2014): "Information Design," *mimeo*.

[36] Wang, Y. (2013): "Bayesian Persuasion with Multiple Receivers," *mimeo*.

B   Online Supplemental Material

This online Appendix complements "Persuading Voters," by Alonso and Câmara (2014).

B.1   Additional Examples

Example 5:   This example illustrates that when voters disagree on the ranking of states, a signal realization that convinces one voter to support the proposal sometimes does not guarantee the support of all easier-to-persuade voters (see Proposition 3). Consider 3 states, $\Theta = \{\theta_1, \theta_2, \theta_3\}$, and $n = 6$ voters. All voters agree that θ_3 is the only approval state, but they disagree on the ranking of rejection states θ_2 and θ_1. Specifically, voters are divided into two groups. Group A has voters $\{\delta^{A1}, \delta^{A2}, \delta^{A3}\}$, who rank state θ_2 higher than θ_1 ($\delta^{Ai}_{\theta_3} > \delta^{Ai}_{\theta_2} > \delta^{Ai}_{\theta_1}$). Group B has voters $\{\delta^{B1}, \delta^{B2}, \delta^{B3}\}$, who rank state θ_1 higher than θ_2 ($\delta^{Bi}_{\theta_3} > \delta^{Bi}_{\theta_1} > \delta^{Bi}_{\theta_2}$). Figures 4(a) and (b) depict voters' approval sets, which are rotations around the same point. Voters with a lower numerical index are easier to persuade — voter δ^{Ai} is easier to persuade than voter $\delta^{A(i+1)}$, while voter δ^{Ai} is as easy to persuade as voter δ^{Bi}. Importantly, voters from both groups are so misaligned that an optimal signal to persuade a voter from one group never induces approval from *any* voter from the other group. In Figure 4(a), the optimal signal when the decision is delegated to voter δ^{A3} induces approval posterior q^+. Belief q^+ convinces voter δ^{A3} and the easier-to-persuade voters δ^{A1} and δ^{A2} that the proposal is better than the status quo. However, it does not convince any voter in group B, including the easier-to-persuade voters δ^{B1} and δ^{B2}. Moreover, no single group has enough votes to approve the proposal when the voting rule requires a strict majority $k > \frac{n}{2}$. Consequently, any $k > 3$ is payoff-equivalent to requiring a unanimous vote. Figure 4(b) depicts the posterior beliefs induced by an optimal signal with unanimity rule, which is also optimal given any $k > 3$. □
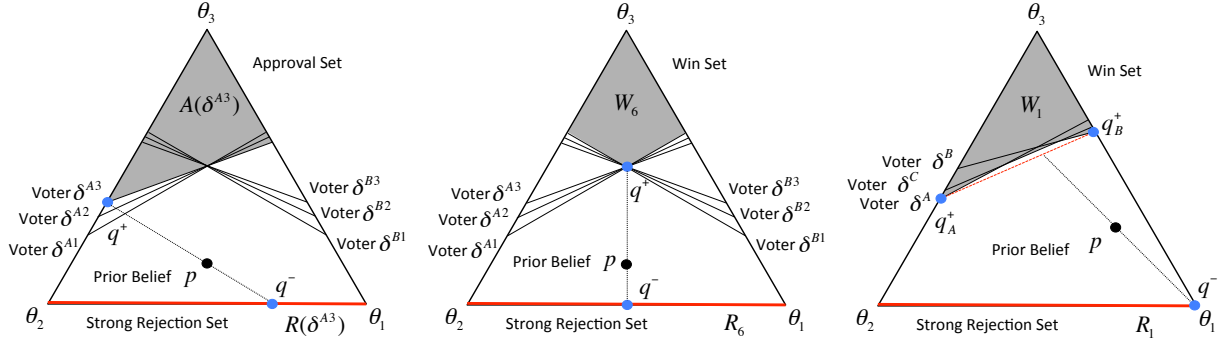
(a) Example 5, approval delegated   (b) Example 5, with unanimity   (c) Example 6, with $k = 1$
to δ^{A3}

Figure 4: Optimal Signals for Examples 5 and 6

Example 6:   We now extend Example 2 to illustrate that sometimes an optimal signal does not target the easiest-to-persuade voter, even when voters agree under full information and rank states in the same order. Consider 3 states, $\Theta = \{\theta_1, \theta_2, \theta_3\}$, and in addition to voters δ^A and δ^B from Example 2, consider a third voter δ^C as follows:

State	Prior	δ_θ^A	δ_θ^B	δ_θ^C
θ_3	0.2	+1	+1	+1
θ_2	0.1	-0.5	-1.5	-0.8
θ_1	0.7	-6	-2	-3

We know from Example 2 that if the decision is delegated do either δ^A or δ^B, then the equilibrium probability of approval is 32.5%. If the decision is delegated to δ^C, then the equilibrium probability of approval[18] is 34%. Therefore, δ^C is the easiest-to-persuade voter. However, if the voting rule is $k = 1$, then the signal described in Example 2 remains optimal: realization s_A^+ convinces voter δ^A to approve the proposal but it does not convince voters δ^B and δ^C; realization s_B^+ convinces voter δ^B but not voters δ^A and δ^C; realization s^- does not convince any voter; the probability of approval is then 37.5%. Therefore, the optimal signal does not target the approval of the easiest-to-persuade voter — it is optimal for the controller to instead exploit the disagreement between voters δ^A and δ^B. As Figure 4(c)

---

[18]Formally, $S = \{s^-, s^+\}$, $Pr(s^+|\theta_3) = Pr(s^+|\theta_2) = 1$, and $Pr(s^+|\theta_1) = \frac{2}{35}$. The possible posterior beliefs are $q^- = (1, 0, 0)$ and $q^+ = (\frac{2}{17}, \frac{5}{17}, \frac{10}{17})$. Probability of approval is $Pr(s = s^+) = \frac{2}{35} \times 0.7 + 1 \times 0.1 + 1 \times 0.2 = 0.34$.

2

illustrates, when $k = 1$ the weak-representative voter $\delta^* = (-2, -0.5, 1)$ is supported on the approval sets of voters δ^A and δ^B, but not on the approval set of δ^C. $\square$

Example 7: This example highlights that in order to obtain the results in Lemmas 4 and 5, one may not replace the same class assumption with an agreement under full information assumption. Consider an electorate with five voters, $\{\delta^A, \delta^B, \delta^C, \delta^D, \delta^E\}$, and three states $\Theta = \{\theta_1, \theta_2, \theta_3\}$. All voters agree under full information — they want to approve the proposal if the state is θ_3, and to reject if the state is θ_2 or θ_1. However, they do not have the same ranking of states. Voters δ^A and δ^B (who have the same type) and voter δ^C rank $\delta_{\theta_2} > \delta_{\theta_3}$, while voters δ^D and δ^E (who have the same type) rank $\delta_{\theta_2} < \delta_{\theta_3}$. Approval sets and win sets for different voting rules are depicted in Figure 5. The optimal signal under simple majority $k = 3$ induces posteriors $\{q_3^-, q_3^+\}$, the optimal signal under supermajority $k = 4$ induces posteriors $\{q_4^-, q_4^+\}$, and the optimal signal under unanimity $k = 5$ induces posteriors $\{q_5^-, q_5^+\}$. Differently than Lemma 4, voters δ^A, δ^B and δ^C have non-single-peaked preferences over voting rules: they consider $k = 4$ strictly worse than $k = 3$ and $k = 5$. To see this, note that voters δ^A and δ^B strictly prefer to approve the proposal under beliefs q_3^+ and q_5^+, but they are indifferent under q_4^+. Voter δ^C weakly prefers to approve under beliefs q_3^+ and q_5^+, and he strictly prefers to reject under q_4^+. Differently than Lemma 5, a majority of voters (voters δ^A, δ^B and δ^C) strictly prefer simple majority over the supermajority $k = 4$. Nevertheless, in accordance with Corollary 3, a majority of voters (voters δ^C, δ^D and δ^E) weakly prefer unanimity over simple majority. $\square$
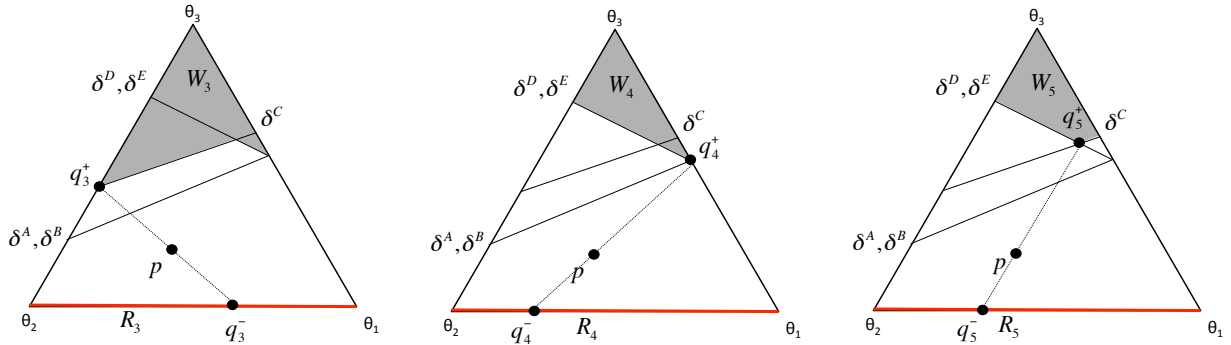

Figure 5: Optimal Signal for Example 7

3

B.2 Voter preferences over k-voting rules

Corollary 4 *Consider an electorate $\{\delta^1, \ldots, \delta^n\}$.*
(i) Suppose that $\delta^i \in \mathscr{F}_z$, $i \in I$, and there is a toughest voter, i.e. there exists $j \in \{1, ..., n\}$ such that $A(\delta^j) \subset A(\delta^i)$, $i \in I$. If unanimity is not a Condorcet winner, then δ^j disagrees with a majority of voters under full information.
(ii) Suppose that $A(\delta^{i+1}) \subset A(\delta^i)$ and all voters agree under full information. If unanimity is not a Condorcet winner, then δ^n must rank states differently than a majority of voters.

Proof of Corollary 4: Part (i)- If there is a toughest voter, say δ^j, then this voter is a representative voter under unanimity. If a majority of voters does not prefer unanimity then the conditions of Proposition 5 are violated and we must have that a majority of voters must have a set of approval states different from the set of approval states of voter δ^j.

Part (ii)- If all voters are totally ordered then a representative voter exists for any $k-$voting rule. If a majority of voters does not prefer unanimity then the conditions of Proposition 5 are violated and we must have that the toughest voter δ^n must rank states differently than the majority of voters. ∎

B.3 Controller Knows the State

Consider our basic model of pure-persuasion where all players have a common prior p, and suppose that before choosing signal π the controller privately observes the true state $\theta \in \Theta$. We can apply the results from Alonso and Câmara (2014b) to conclude that the controller cannot benefit from observing θ. Moreover, our voting model allows us to derive a second result: if π^* is an optimal signal when the controller has no private information, then there is a pooling equilibrium where all controller's types supply the same π^*. Therefore, a pooling equilibrium with π^* maximizes the controller's expected payoff in the case when the controller can observe the state.

To understand the second result, let π^* be an optimal signal in the case when the controller has no private information. Consider an electorate $\{\delta^1, \ldots, \delta^n\}$ and a k-voting rule. We can partition the realization space S into two sets: the set of approval signals that induce a

posterior belief in the win set W_k, and the set of rejection signals that induce a posterior belief in the strong rejection set R_k. Consequently, signal π^* also partitions the state space Θ into three sets: the set of states Θ_1 that induce an approval realization with probability one, the set of states Θ_2 that induce a rejection realization with probability one, and the set of states Θ_3 that induce an approval signal with probability strictly between zero and one. Importantly, since any state θ in the sets Θ_2 and Θ_3 can induce a posterior in the strong rejection set, it must be the case that a blocking coalition of at least $n - k + 1$ voters would reject the proposal if they knew that $\theta \in \Theta_2 \cup \Theta_3$.

Now suppose that the controller knows the state prior to choosing the signal. We construct the following equilibrium strategies and beliefs. The controller's equilibrium strategy is to choose the original signal π^*, independently of her private information. Along the equilibrium path, upon observing the choice of π^* voters do not update their priors (since all controllers' types choose the same π^* in equilibrium), so they use π^* and the realized signal s to vote as in the original equilibrium. Before we define the out-of-equilibrium-path behavior, note that if the controller knows that the state is in Θ_1, then she knows that the proposal will be approved for sure with the signal π^*. Therefore, these types have no incentive to deviate from the original signal. The only types who could possibly benefit are the ones who know that $\theta \in \Theta_2 \cup \Theta_3$. Therefore, out of the equilibrium path, when voters observe a choice of signal π' different than π^*, we let them update their beliefs as follows: they assign probability zero to the controller's types that know that $\theta \in \Theta_1$. Therefore, upon observing any choice of signal π' different than π^*, voters belief that $\theta \in \Theta_2 \cup \Theta_3$ and reject the proposal. Hence, no controller benefits from deviating, and pooling is indeed an equilibrium.

The case of delegation follows since it is equivalent to unanimity with a homogeneous electorate.

B.4   Controller's Payoff Depends on the State

In this section we study the case when the controller's payoff depends on the realized state. Consider the controller's payoff $u_C(x, \theta) : X \times \Theta \rightarrow \mathbb{R}$. Let $\delta_\theta^C = u_C(x_1, \theta) - u_C(x_0, \theta)$ and define the controller's type $\delta^C = \{\delta_\theta^C\}_{\theta \in \Theta}$. The controller prefers the proposal to be approved in states $\theta \in D(\delta^C)$, and rejected in states $\theta \notin D(\delta^C)$. To simplify presentation, suppose the

controller is never indifferent, $\delta_\theta^C \neq 0$.

First suppose that the approval decision is delegated to voter δ, with $p \notin A(\delta)$ and $A(\delta) \cap A(\delta^C) \neq \emptyset$. As in the case of pure-persuasion with delegation, there is a multiplicity of optimal signals. However, one can always construct an optimal signal with only two signal realizations: the voter approves if s^+ is realized, and rejects if s^-. There are two possible cases: preference alignment and preference misalignment. We say that there is preference alignment if the controller suffers no loss of control by the fact that the voter makes the approval decision. This is the case if and only if $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta \geq 0$. In this case, the controller's optimal signal induces approval realization s^+ for every $\delta_\theta^C > 0$, and rejection realization s^- for every $\delta_\theta^C < 0$. Consequently, the controller's preferred policy is implemented in each state. Moreover, the voter benefits from this signal. The more interesting case is preference misalignment, $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta < 0$, when the information controller can no longer guarantee implementation of her preferred policy in each state. The following Proposition generalizes the optimal signal from Proposition 1.

Proposition 6 *Consider controller δ^C and suppose that the approval decision is delegated to voter δ, $p \notin A(\delta)$, $A(\delta) \cap A(\delta^C) \neq \emptyset$, and $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta < 0$. Let π^* be any controller's optimal signal supported on two realizations $\{s^-, s^+\}$, where voter δ approves the proposal if and only if $s = s^+$. Letting $\alpha_\theta = Pr[s^+|\theta]$, for each state θ we have:*

(i) If players agree on approval, $\theta \in \Theta_A \equiv \{\theta \in \Theta | \delta_\theta^C > 0, \delta_\theta \geq 0\}$, then $\alpha_\theta = 1$;

(ii) If players agree on rejection, $\theta \in \Theta_R \equiv \{\theta \in \Theta | \delta_\theta^C < 0, \delta_\theta \leq 0\}$, then $\alpha_\theta = 0$;

(iii) If players disagree, $\theta \in \Theta_D \equiv \{\theta \in \Theta | \delta_\theta^C > 0, \delta_\theta < 0 \text{ or } \delta_\theta^C < 0, \delta_\theta > 0\}$, then there exists $\theta' \in \Theta_D$ such that

$$\alpha_\theta = \begin{cases} 1 & \text{if } \delta_\theta^C > 0, \left|\frac{\delta_\theta^C}{\delta_\theta}\right| > \left|\frac{\delta_{\theta'}^C}{\delta_{\theta'}}\right| \text{ or } \delta_\theta^C < 0, \left|\frac{\delta_\theta^C}{\delta_\theta}\right| < \left|\frac{\delta_{\theta'}^C}{\delta_{\theta'}}\right| \\ 0 & \text{if } \delta_\theta^C > 0, \left|\frac{\delta_\theta^C}{\delta_\theta}\right| < \left|\frac{\delta_{\theta'}^C}{\delta_{\theta'}}\right| \text{ or } \delta_\theta^C < 0, \left|\frac{\delta_\theta^C}{\delta_\theta}\right| > \left|\frac{\delta_{\theta'}^C}{\delta_{\theta'}}\right| \end{cases} , \text{ and } \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta = 0. \quad (14)$$

Moreover, while voter δ never gains by making decisions with the signal π^*, the controller's value of information control is $V - v(p) = \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta^C$.

Proof of Proposition 6: The existence of an optimal binary signal is established in KG (Proposition 1, p. 2595). Let π be an optimal binary signal with $S = \{s^-, s^+\}$ where the

voter approves the proposal if and only if he observes s^+, and let $\alpha_\theta = \Pr[s^+|\theta]$ so that $\Pr[Approval] = \sum_{\theta \in \Theta} \alpha_\theta p_\theta$. A Bayesian voter δ will approve after observing s^+ if and only if

$$E[\delta|s^+] = \sum_{\theta \in \Theta} \frac{\alpha_\theta p_\theta \delta_\theta}{\Pr[Approval]} \geq 0.$$

The information controller's payoff with signal π is

$$\sum_{\theta \in \Theta} \left[\alpha_\theta p_\theta u_C(x_1, \theta) + (1 - \alpha_\theta) p_\theta u_C(x_0, \theta) \right] = \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta^C + \sum_{\theta \in \Theta} p_\theta u_C(x_0, \theta).$$

The last term $\sum_{\theta \in \Theta} p_\theta u_C(x_0, \theta)$ is not a function of α_θ, hence the optimal signal must solve the following linear program:

$$\max \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta^C, \quad s.t. \quad 0 \leq \alpha_\theta \leq 1, \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta \geq 0. \tag{15}$$

We can write the Lagrangian and the first order condition with respect to α_θ,

$$\mathcal{L} = \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta^C + \lambda \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta + \sum_{\theta \in \Theta} \nu_\theta^+ \alpha_\theta + \sum_{\theta \in \Theta} \nu_\theta^- (1 - \alpha_\theta),$$

$$\frac{\partial \mathcal{L}}{\partial \alpha_\theta} = p_\theta \delta_\theta^C + \lambda p_\theta \delta_\theta + \nu_\theta^+ - \nu_\theta^- = 0. \tag{16}$$

It must be the case that $\sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta = 0$ and $\lambda > 0$[19]. If players agree on approval, $\delta_\theta^C > 0$ and $\delta_\theta \geq 0$, then (16) implies $\nu_\theta^- > 0$ and $\alpha_\theta = 1$. If players agree on rejection, $\delta_\theta^C < 0$ and $\delta_\theta \leq 0$, then (16) implies $\nu_\theta^+ > 0$ and $\alpha_\theta = 0$. Now consider the set of disagreement states such that the controller prefers approval, $\delta_\theta^C > 0$ and $\delta_\theta < 0$. If $p_\theta \delta_\theta^C + \lambda p_\theta \delta_\theta > 0$ (hence $\frac{\delta_\theta^C}{-\delta_\theta} > \lambda$), then (16) implies $\nu_\theta^- > 0$ and $\alpha_\theta = 1$. If $p_\theta \delta_\theta^C + \lambda p_\theta \delta_\theta < 0$ (hence $\frac{\delta_\theta^C}{-\delta_\theta} < \lambda$), then (16) implies $\nu_\theta^+ > 0$ and $\alpha_\theta = 0$. Now consider the set of disagreement states such that the controller prefers rejection, $\delta_\theta^C < 0$ and $\delta_\theta > 0$. If $p_\theta \delta_\theta^C + \lambda p_\theta \delta_\theta > 0$ (hence $\frac{-\delta_\theta^C}{\delta_\theta} < \lambda$), then (16) implies $\nu_\theta^- > 0$ and $\alpha_\theta = 1$. If $p_\theta \delta_\theta^C + \lambda p_\theta \delta_\theta < 0$ (hence $\frac{-\delta_\theta^C}{\delta_\theta} > \lambda$), then (16) implies $\nu_\theta^+ > 0$ and $\alpha_\theta = 0$.

Therefore, the cutoff state $\theta' \in \Theta_D$ defined by the Proposition is the state θ' with the absolute value of the ratio $\left| \frac{\delta_\theta^C}{\delta_\theta} \right|$ closest to λ. ■

---

[19]By contradiction, suppose $\lambda = 0$. From (16) we have: if $\delta_\theta^C > 0$, then $\nu_\theta^- > 0$ and $\alpha_\theta = 1$; if $\delta_\theta^C < 0$, then $\nu_\theta^+ > 0$ and $\alpha_\theta = 0$. In this case, signal s^+ violates the approval constraint $\sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta \geq 0$, since the proposition considers the misalignment case $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta < 0$.

In many important cases the controller ranks states in the same order as the voter. For example, the controller receives the same payoff as the voter, plus some private benefit from approving the proposal (see also our application in Section 6.1). We find that if all players rank states in the same order, then our results on voters' preferences over k-voting rules continue to hold.

Proposition 7 *Consider an information controller δ^C and an electorate $\{\delta^1, \ldots, \delta^n\}$, with $\delta^C, \delta^i \in \mathscr{F}_z$. Then each voter has single-peaked preferences over k-voting rules. If voters also agree under full information, then the payoff of every voter is weakly increasing in k.*

Proof of Proposition 7: To prove this Proposition, we start by solving the delegation benchmark. Suppose that the approval decision is delegated to a voter δ who ranks the state in the same order as the controller. If there is no preference misalignment, $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta \geq 0$, then the controller suffers no loss of control: her optimal signal leads to approval in every approval state $\delta_\theta^C > 0$, and rejection in every rejection state $\delta_\theta^C < 0$. If there is preference misalignment, then π^* is an optimal signal for controller δ^C if and only if π^* is an optimal signal in the case of pure-persuasion. To see this, first note that since both players rank states in the same order z, misalignment implies that there is no state θ such that $\delta_\theta^C < 0$ and $\delta_\theta \geq 0$.[20] Therefore, in every disagreement state $\theta \in \Theta_D$ it must be that $\delta_\theta^C > 0$ and $\delta_\theta < 0$. Consequently, for every $\theta \in \Theta_D$, the absolute value of the ratio $\left|\frac{\delta_\theta^C}{\delta_\theta}\right|$ is increasing in δ_θ (the positive term δ_θ^C increases, while the negative term δ_θ goes to zero). In this case, the controller's problem in Proposition 6 becomes equivalent to the problem in Proposition 1 with pure-persuasion: the controller seeks a cutoff state $\theta' \in \Theta$ such that the proposal is approved in every state $\delta_\theta > \delta_{\theta'}$, rejected in every state $\delta_\theta < \delta_{\theta'}$, and voter's expected payoff conditional on approval equals the expected payoff of rejection. Also note that if controller δ^C and voter δ are misaligned, then the controller is also misaligned with any tougher voter δ', $A(\delta') \subset A(\delta)$. Therefore, we can extend the results from Proposition 4 as follows. Take any totally ordered (according to toughness) set of voters $\mathcal{D} \in \mathscr{F}_z$. Then any voter $\delta \in \mathscr{F}_z$ has

---

[20]By contradiction, suppose there is a state θ such that $\delta_\theta^C < 0$ and $\delta_\theta \geq 0$. Then every state θ' such that $z(\theta') > z(\theta)$ is also an approval state for the voter. Moreover, for the controller there is no approval state θ' such that $z(\theta') < z(\theta)$. Hence, $D(\delta^C) \subset D(\delta)$, which contradicts misalignment, $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta < 0$.

single-peaked preferences over decision makers in $\mathcal{D}$. Now consider a k-voting rule. Lemma 3 continues to hold — if all voters rank states in the same order, then there is a weak representative voter $\delta^*(k)$ who also rank states in the same order. Moreover, recall that in this case the weak representative voters are ordered by toughness, $A(\delta^*(k+1)) \subset A(\delta^*(k))$. Therefore, all the results of Lemmas 4 and 5 and Proposition 5 continue to hold.

B.5 Preference Shocks

In our setup, voters' preferences depend on a common state θ and the controller can influence voters' behavior by providing a signal correlated with θ. The controller's ability to predict each vote leads her to resort to signals with realizations that just guarantee approval by the required number of votes. In many instances, however, voters' behavior may not be completely pinned down by the realization of the controller's signal. This is the case when voters are subject to idiosyncratic shocks that affect the relative benefits of the proposal and the status quo.

To study the effect of preference shocks, we assume that voter i's preferences are given by $u_i(x, \theta, \mu_i)$ and the conditional net payoff from approval with state θ and private shock μ_i is

$$u_i(x_1, \theta, \mu_i) - u_i(x_0, \theta, \mu_i) = \delta_\theta^i - \mu_i.$$

Then, when voter i observes μ_i and holds belief q regarding the realization of θ, he votes for x_1 if and only if $\langle q, \delta^i \rangle \geq \mu_i$. To simplify exposition, we assume that each voter i observes the realization of μ_i after observing the realization of the controller's signal. Shocks μ_i are i.i.d. and jointly independent with θ, with each shock distributed according to $F(\mu)$ with support in $[-\bar{\mu}, \bar{\mu}]$.

When designing the signal, the controller must now consider how the joint distribution of private shocks affects voters' willingness to side with the proposal. Given independence, we summarize this information in the form of a function $P_i(q)$ that determines the probability that voter i approves the proposal if he has belief q, i.e. $P_i(q) = \Pr\left[\mu_i \leq \langle q, \delta^i \rangle\right]$. Define the set of guaranteed-approval beliefs $A^P(\delta^i) = \{q \in \Delta(\Theta) : P_i(q) = 1\}$.

The next proposition shows that if preference shocks are small and high shocks are

sufficiently likely, then under both delegation and unanimity the controller behaves as if she is facing non-probabilistic voters with approval sets $A^P(\delta^i)$.

Proposition 8 *Suppose that for every voter i we have $\delta^i_\theta \notin [-\bar{\mu}, \bar{\mu}]$ with at least one state with $\delta^i_\theta > \bar{\mu}$, and $F(\mu) \leq F_U(\mu)$ where U is uniformly distributed in $[-\bar{\mu}, \bar{\mu}]$. Then,*
(i) Suppose that decisions are delegated to voter i. Then the controller's expected utility and optimal signal are the same as if the decision was delegated to a voter with no private shock and type $\check{\delta}^i = \delta^i - \mathbf{1}\bar{\mu}$. Furthermore, if $p \notin A^P(\delta^i)$, then the value to voter i of the controller's optimal signal is strictly positive.
(ii) Suppose that decisions are made under unanimity rule. Then the controller's expected utility and optimal signal are the same as if decisions are made under unanimity by an electorate of voters $\left\{\check{\delta}^i\right\}_{i \in I}$, with $\check{\delta}^i = \delta^i - \mathbf{1}\bar{\mu}$, for $i \in I$. Furthermore, if for all $i \in I$, $p \notin A^P(\delta^i)$, then every voter is strictly better off with the controller's optimal signal.

Part (i) of the proposition considers delegation to voter i. For any belief q such that $|\langle q, \delta^i \rangle| < \bar{\mu}$, we have $0 < P_i(q) < 1$, and voter i takes into account the realization of his private shock when deciding whether to approve the proposal. In designing approval signal realizations, the controller must trade-off a higher probability of approval of any belief that increases $\langle q, \delta^i \rangle$ with a lower probability that the optimal signal induces that belief. Proposition 8(i) states that the controller resolves this trade-off with a signal such that voter $i's$ behavior does not depend on his private shock; in particular, the optimal signal induces beliefs q_s such that $P_i(q_s) \in \{0, 1\}$. That is, the controller behaves as if she is facing a voter that, with certainty, either approves or rejects the proposal, i.e. a voter not subject to a preference shock. Since $P_i(q) = 1$ iff $\bar{\mu} \leq \langle q, \delta^i \rangle$, the controller acts as if she needs to persuade the non-probabilistic voter of type $\check{\delta}^i = \delta^i - \mathbf{1}\bar{\mu}$, i.e., the voter that always experiences the most adverse possible shock. This is true as long as $F(\mu) \leq F_U(\mu)$, which is satisfied for any convex distribution of support contained in $[-\bar{\mu}, \bar{\mu}]$ and implies that high values of the private shock are more likely.[21] Proposition 8(ii) shows that these insights carry over to the case of a unanimous voting rule. With unanimity, the controller behaves as if she faces a non-probabilistic electorate with types $\check{\delta}^i = \delta^i - \mathbf{1}\bar{\mu}$.

---

[21]This condition is also necessary in the sense that if it is violated, then there exists a prior belief p such that the optimal signal induces beliefs that do not guarantee either approval or rejection, ie. $P_i(q_s) \notin \{0, 1\}$.

Under our distributional assumptions, the presence of small private shocks changes neither the controller's behavior nor her expected utility. However, it does impact voters' welfare. To begin with, voters have almost surely a strict preference between approval and rejection of the proposal, so that voters obtain a strictly positive gain with probability 1 when they side with the proposal. Therefore, unlike the case of non-probabilistic voting, voter i strictly benefits from the controller's signal with delegation (cf. Proposition 1). If all voters would reject the proposal under the prior belief, the same would be true under a unanimity rule. Proposition 8(ii) shows that in this case all voters strictly benefit from the signal. In summary, if decisions need the approval of all decision makers, then all decision makers benefit from the controller's influence.

For non-unanimous voting rules, the results from Corollary 2 carry over to cases with a small $\bar{\mu}$. In particular, with a simple majority voting rule, a majority of voters can be made strictly worse off by the controller's signal.

Proof of Proposition 8 : Consider an electorate $\{\delta^i\}_{i \in I}$ and let $V_i^{PV}(q) = P_i(q)$ and $V_U^{PV}(q) = \Pi_{i \in I} P_i(q)$ be the indirect utility of the controller when decisions are delegated to voter i and when decisions are made under unanimity. Similarly, consider the electorate of voters $\{\check{\delta}^i\}_{i \in I}$ that are not subject to a preference shock, with $\check{\delta}_\theta^i = \delta_\theta^i - \bar{\mu}\mathbf{1}$ so that the set of approval beliefs of each voter is $A(\check{\delta}^i) = \{q \in \Delta(\Theta) : \bar{\mu} \leq \langle q, \delta^i \rangle\} = A^P(\delta^i)$, and let $V_i^{NPV}(q) = 1_{\{q \in A(\check{\delta}^i)\}}$ and $V_U^{NPV}(q) = 1_{\{q \in \cap_{i \in I} A(\check{\delta}^i)\}}$ be the indirect utility of the controller when decisions are delegated to voter i and when decisions are made under unanimity. Recall that $\tilde{f}$ denotes the concave closure of function f. KG show that if $V(q)$ is the controller's indirect utility then the controller's expected utility under an optimal signal when the prior belief is p is $\tilde{V}(p)$. We will show that (i) $\tilde{V}_j^{PV}(q) = \tilde{V}_j^{NPV}(q)$, for $j = \{1, .., N, U\}$, and (ii) $V_j^{NPV}(q) = \tilde{V}_j^{NPV}(q)$ if and only if $V_j^{PV}(q) = \tilde{V}_j^{PV}(q)$, for $j = \{1, .., N, U\}$. Identity (i) implies that the controller's expected utility is the same if he faces probabilistic or if he faces non-probabilistic voters, while (ii) and (i) imply that the set of optimal signals for probabilistic and non-probabilistic voters coincides.

Because approval is more likely for the case of preference shocks, then $V_j^{PV}(q) \geq V_j^{NPV}(q)$ and the monotonicity of the concave closure operator implies that $\tilde{V}_j^{PV}(q) \geq \tilde{V}_j^{NPV}(q)$. We

now show that $\tilde{V}_j^{NPV}(q) \geq V_j^{PV}(q)$ for $j = \{1, .., N, U\}$, which implies $\tilde{V}_j^{NPV}(q) \geq \tilde{V}_j^{PV}(q)$. These inequalities together establish that $\tilde{V}_j^{PV}(q) = \tilde{V}_j^{NPV}(q)$.

Consider first the case of delegation to voter δ^i. Suppose that $q \in A^P(\delta^i)$. Because $A(\check{\delta}^i) = A^P(\delta^i)$ and the maximum indirect utility is achieved in these sets, we have that $V_i^{NPV}(q) = \tilde{V}_i^{NPV}(q) = V_i^{PV}(q)$. Now suppose that $q \notin A^P(\delta^i)$. Since shocks are small so that $\delta_\theta^i \notin [-\bar{\mu}, \bar{\mu}]$, then $V_i^{PV}(1_\theta) = 0$ for any state $\theta \in D_i^C$ where $D_i^C(\delta^i) = \{\theta \in \Theta : \delta_\theta^i < -\bar{\mu}\}$ is the set of rejection states. As the indirect utility achieves is lowest value for any belief $q \in R_i^P(\delta^i) \equiv co(D_i^C(\delta^i))$, then $V_i^{NPV}(q) = \tilde{V}_i^{NPV}(q) = V_i^{PV}(q) = 0$. Finally, suppose that $0 \leq P_i(q) < 1$. Note that $\tilde{V}_i^{NPV}(q) = \lambda$ for some λ such that $q = \lambda q^+ + (1-\lambda)q^-$, with $q^+ \in A^P(\delta^i)$, $\langle q^+, \delta^i \rangle = \bar{\mu}$, and $q^- \in R_i^P(\delta^i)$. For any state $\theta \in D_i^C$, we have $\delta_\theta^i < -\bar{\mu}$ implying that $\langle q^-, \delta^i \rangle < -\bar{\mu}$. Since $\langle q^+, \delta^i \rangle = \bar{\mu}$, then $\langle q^+ - q^-, \delta^i \rangle > 2\bar{\mu}$. Overall, because $F(\mu) \leq F_U(\mu)$ we obtain the following inequality

$$F_U(\langle q, \delta^i \rangle) = 1 - \frac{\langle q^+ - q, \delta^i \rangle}{2\bar{\mu}} = 1 - \frac{(1-\lambda)\langle q^+ - q^-, \delta^i \rangle}{2\bar{\mu}} < \lambda, \tag{17}$$

which leads to

$$V_i^{PV}(q) = P_i(q) = F\left(\langle q, \delta^i \rangle\right) \leq F_U(\mu) < \lambda = \tilde{V}_i^{NPV}(q).$$

Consider now the case of unanimity. By the same reasoning as before we have that $V_U^{NPV}(q) = \tilde{V}_U^{NPV}(q) = V_U^{PV}(q)$ for $q \in \cap_{i \in I} A^P(\delta^i)$, and for any belief , then $V_i^{NPV}(q) = \tilde{V}_i^{NPV}(q) = V_i^{PV}(q) = 0$ for $q \in \cup_{i \in I} R_i^P(\delta^i)$. Finally, suppose that $0 \leq \Pi_{i \in I} P_i(q) < 1$. Again, we can write $\tilde{V}_i^{NPV}(q) = \lambda$ with $q = \lambda q^+ + (1-\lambda)q^-$, and $q^+ \in \cap_{i \in I} A^P(\delta^i)$, with $\langle q^+, \delta^i \rangle \geq \bar{\mu}$ for $i \in I$ with at least one strict equality, and $q^- \in \cup_{i \in I} R_i^P(\delta^i)$. Let i' be a voter that always rejects for q^-, i.e. $\langle q^-, \delta^{i'} \rangle < -\bar{\mu}$. Since for this voter $\langle q^+, \delta^{i'} \rangle \geq \bar{\mu}$ then $\langle q^+ - q^-, \delta^{i'} \rangle > 2\bar{\mu}$. From (17) and $F(\mu) \leq F_U(\mu)$, we have $F(\langle q, \delta^{i'} \rangle) < \lambda$ and

$$V_i^{PV}(q) = \Pi_{i \in I} P_i(q) = \Pi_{i \in I} F\left(\langle q, \delta^i \rangle\right) \leq F(\langle q, \delta^{i'} \rangle) < \lambda = \tilde{V}_i^{NPV}(q).$$

Finally, note that under delegation, in all the cases that $q \notin A^P(\delta^i)$ and $q \notin R_i^P(\delta^i)$, we had $\tilde{V}_j^{NPV}(q) \neq V_j^{PV}(q)$. Since in all these cases we also had $\tilde{V}_j^{NPV}(q) = \lambda \neq V_j^{NPV}(q)$, then $V_j^{NPV}(q) = \tilde{V}_j^{NPV}(q)$ if and only if $V_j^{PV}(q) = \tilde{V}_j^{PV}(q)$, for $j = \{1, .., N\}$. The same reasoning can be translated to the case of unanimity to prove that $V_U^{NPV}(q) = \tilde{V}_U^{NPV}(q)$ if and only if $V_U^{PV}(q) = \tilde{V}_U^{PV}(q)$. ∎

B.6 Spatial Model of Elections

We now solve for the optimal signal in the spatial electoral model described in Section 6.2.

Consider a simple majority rule, so that the median voter is the representative voter. A player with ideology y prefers a right-wing challenger with policy $\theta \geq 0$ over the left-wing incumbent with policy $y_L < 0$ if the challenger is no too radical, $\theta \leq |y_L| + 2y_C$. Therefore, a player with ideology y and corresponding type δ (defined by (9)) has a set of approval states $D(\delta) = \{\theta \in \Theta | \theta \leq |y_L| + 2y_C\}$. If player δ' has an ideology to the right of δ, then $D(\delta) \subset D(\delta')$. For any ideology $y \geq 0$, define $T(y) = \{\max \theta \in \Theta | \theta \leq |y_L| + 2y\}$, that is, the most right-wing policy of the challenger that a player y would prefer over the incumbent. Since the example assumes $\theta_1 < |y_L|$, this set is not empty. Moreover, T weakly increases in y.

Next we define a cutoff $\bar{y}$ on the ideology of the controller, and then characterize the optimal signal of moderate ($y_C < \bar{y}$) and radical ($y_C > \bar{y}$) controllers.

Step 1 - Ideology cutoff: Let θ^* to be the smallest $\theta \in \Theta$ such that $\sum_{\{\theta \in \Theta | \theta \leq \theta^*\}} p_\theta \delta_\theta^{median} < 0$. The example assumes $\theta_1 < |y_L|$, therefore $\sum_{\{\theta \in \Theta | \theta \leq |y_L|\}} p_\theta \delta_\theta^{median} > 0$ and $\theta^* > |y_L|$. Since without further information the median voter keeps the incumbent, $\sum_{\theta \in \Theta} p_\theta \delta_\theta^{median} < 0$, then $\theta^* \leq \theta_M$. Define the ideology cutoff $\bar{y} = \frac{\theta^* - |y_L|}{2}$, and note that a player with ideology $\bar{y}$ is just indifferent between electing challenger θ^* or the incumbent. Moreover, $T(\bar{y}) = \{\max \theta \in \Theta | \theta \leq |y_L| + 2\bar{y}\} = \theta^*$. Since $|y_L| < \theta^* \leq \theta_M$, it must be the case that $0 < \bar{y} \leq \frac{\theta_M - |y_L|}{2}$.

Step 2 - Moderate Controller: Suppose that the controller is moderate, $y_C < \bar{y}$. This implies that the controller strictly prefers not to elect challenger θ^* (and every challenger to the right). In this case, the controller is able to persuade the median to elect the challenger by simply revealing that $\theta \in D(\delta^C)$. Therefore, any optimal signal π^* induces an approval realization for every state $\theta \leq T(y_C)$, and a rejection signal for every state $\theta > T(y_C)$. Note that even if the controller does not agree with the median under full information, $T(0) < T(y_C)$, the controller still does not suffer loss of control. That is, she achieves her maximum payoff as if she could both choose the signal and decide whether to approve the proposal (elect the challenger). The median strictly benefits from this signal.

Step 3 - Radical Controller: Suppose that the controller is radical, $y_C > \bar{y}$. Then by simply revealing that $\theta \in D(\delta^C)$ the controller cannot persuade the median voter. We can then apply the results from Proposition 6 as follows. For any state $\theta \leq T(0)$, both the median and the controller agree on approval, while for any $\theta > T(y_C)$ both agree on rejection. Hence, the only disagreement states are $\theta \in (T(0), T(y_C)]$. Using δ_θ as defined by (9), in all these state we have $\delta_\theta^C > 0$, $\delta_\theta^{median} < 0$ and

$$\frac{\delta_\theta^C}{|\delta_\theta^{median}|} = \frac{-(\theta^2 - y_L^2) + 2(\theta - y_L)y_C}{(\theta^2 - y_L^2)} = \frac{2y_C}{\theta + y_L} - 1.$$

Since $y_C > 0$, the ratio $\frac{\delta_\theta^C}{|\delta_\theta^{median}|}$ strictly decreases with θ. Using Proposition 6(iii), there is a cutoff state θ' such that every state $\theta > \theta'$ (hence a lower ratio) induces an rejection signal and every state $\theta < \theta'$ (hence a higher ratio) induces an approval signal. Moreover, 14 requires the median voter to be indifferent when electing the challenger (hence he does not benefit from the signal). In other words, the disagreement does not change the ordering of the states, and the optimal cutoff defined by (14) is the same as (4) in the case of pure-persuasion.

Finally, note that a higher k-voting rule is equivalent to delegating the decision to a representative voter with ideology $y \leq 0$. Every voter $y \leq 0$ ranks states in the same order (a higher θ is bad news about the challenger), and the controller's signal continues to take the form of a cutoff rule on the state. A higher k-voting rule implies that, to persuade the tougher decision maker, the controller's approval signals must show that the challenger is more moderate. Similar argument used on Proposition 7 reveals that the median and all left voters then have single-peaked preferences over k-voting rules above simple majority. Moreover, the median and all left voters weakly prefer unanimity over simple majority. Together these two results imply that any supermajority is weakly better than simple majority, from the point of view of the median and left voters.