

Persuading Voters

Online Appendix

Ricardo Alonso and Odilon Câmara*

May 25, 2016

B Online Appendix

This online Appendix complements “Persuading Voters” by Ricardo Alonso and Odilon Câmara (PV henceforth). It provides additional results and discusses extensions of the model.

Section B.1 describes the general model. Section B.2 presents additional results. Section B.3 presents extensions of the model. Section B.4 discusses relevant applications and examples. Section B.5 describes an alternative interpretation of the model, in which we substitute the politician’s choice of a policy experiment for the choice of an optimal endorser (intermediary).

B.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 policy, and x_1 is the proposal. The collective decision is made following established institutional rules, which we discuss momentarily. Each voter $i \in I \equiv \{1, \dots, n\}$ has preferences over policies that are characterized by

*Alonso: London School of Economics, Houghton Street, London WC2A 2AE, United Kingdom, (e-mail: R.Alonso@lse.ac.uk); Câmara: University of Southern California, 3670 Trousdale Parkway Ste. 308, BRI-308 MC-0804, Los Angeles, CA 90089-0804 (e-mail: ocamara@marshall.usc.edu).

a von Neumann-Morgenstern utility function $u_i(x, \theta)$, $u_i : X \times \Theta \rightarrow \mathbb{R}$, with Θ a finite state space. All players share a common prior belief $p = (p_\theta)_{\theta \in \Theta}$, which has full support on Θ .

Information Controller: One information controller C (called “politician” in PV), who is not a member of the group, has preferences over policies characterized by a von Neumann-Morgenstern utility function u_C . We first consider the case of pure-persuasion in which the controller’s preferences are state-independent, $u_C(x) : X \rightarrow \mathbb{R}$ (in Section B.3.2 we consider a controller with state-dependent payoffs). The controller can influence the decision of the group by designing a public signal (called “policy experiment” in PV) that is correlated with the state (as in Kamenica and Gentzkow 2011, KG henceforth). 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 $\pi(\cdot|\theta)$, $\theta \in \Theta$. Players process information according to Bayes rule. Let $q(s|\pi, p)$ be the updated posterior belief of every voter after observing π and its realization s .

B.1.1 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 $\Gamma = (\Gamma_1, \dots, \Gamma_n, h)$, which defines a strategy set Γ_i for each member i and an outcome function $h : \Gamma_1 \times \dots \times \Gamma_n \rightarrow X$. 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 $g(q) : \Delta(\Theta) \rightarrow X$, 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

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

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*, in which a proposal replaces the status quo if it receives k or more votes. We now formally define these institutional rules.

Delegation: Decision rights are fully delegated to a particular player $d \in I$. Mechanism $\Gamma = \{\Gamma_1, \dots, \Gamma_n, h\}$ has $\Gamma_d = X$, where individual d chooses a policy $\gamma_d(q)$ and this policy is implemented, $h(\gamma_1(q), \dots, \gamma_n(q)) = \gamma_d(q)$. 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 (B.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 $k \in \{1, \dots, n\}$ is the established electoral rule. Mechanism $\Gamma = \{\Gamma_1, \dots, \Gamma_n, h\}$ 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

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

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 equilibria, which guarantees that the controller's expected

payoff v is an upper semicontinuous 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 $g(q) = h(\gamma_1^*(q), \dots, \gamma_n^*(q))$.

As in Alonso and Câmara (forthcoming), 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.

B.1.2 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 (B.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,

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

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

$$\tilde{f}(q) = \sup \{w | (q, w) \in \text{co}(f)\},$$

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

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

$$V = \tilde{v}(p); \tag{B.2}$$

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

For the remaining of our analysis we focus on the case where $u_C(x_0) < u_C(x_1)$ ¹. 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.

B.1.3 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 $\sum_{\theta \in \Theta} q_\theta (u_i(x_1, \theta) - u_i(x_0, \theta)) \geq 0$, 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 ,

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

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

Relevant Sets — Individual Voter: Consider a voter with type δ . Define the *set of approval states* $D(\delta) = \{\theta \in \Theta \mid \delta_\theta \geq 0\}$ and the set of rejection states $D^C(\delta) = \Theta \setminus D(\delta)$. Define the *set of approval beliefs* $A(\delta) = \{q \in \Delta(\Theta) \mid \langle q, \delta \rangle \geq 0\}$ 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

¹The 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.

the *set of strong rejection beliefs* $R(\delta) = \{q \in \Delta(\Theta) \mid \theta \in D(\delta) \Rightarrow q_\theta = 0\}$, that is, the set of beliefs that assign probability zero to every approval state.

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

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

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 \mathcal{B} as the collection of all coalitions of at least $n - k + 1$ voters, with typical element $b \in \mathcal{B}$. Define the *set of strong rejection beliefs*

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

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 $z : \Theta \rightarrow \{1, \dots, \text{card}(\Theta)\}$ that strictly orders the states. Define the class of types

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

That is, class \mathcal{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 $\delta^i \in \mathcal{F}_z$. We say that voters δ^i and δ^j “rank states in the same order” if for some permutation z we have $\delta^i, \delta^j \in \mathcal{F}_z$.

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, 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 $\{\delta^1, \dots, \delta^n\}$. 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 .² Voter δ is a “*weak* representative voter” if the equilibrium expected payoff profile of all players under delegation to voter δ is the same as under the k -voting rule. That is, the probability of approval is the same, $V(\delta) = V(W_k)$, and every voter is indifferent between the k -voting rule and delegation to the weak representative voter δ .

B.2 Additional Results

In the main text of PV we solve for the optimal signal in the cases of delegation to a voter δ (Propositions 1 and 2), and a k -voting rule (Proposition 3). We next present a series of useful additional results.

B.2.1 Weak Representative Voter

We first contrast equilibrium payoffs with a k -voting rule to equilibrium payoffs with an equivalent, but simpler, institutional rule. We perform this analysis in two steps. We first show that, for every k , one can find an electorate, composed of possibly different voters than the original, such that, for any prior belief, a k -voting rule under the original electorate is payoff equivalent to unanimity among this new electorate. That is, while a representative voter may not always exist, we can nevertheless always find a “representative” electorate with a unanimous voting rule. We then turn to the analysis of weak representative voters for a given prior belief. We first show that if all voters in an electorate rank states in the same order, then a weak representative voter exists. Moreover, this weak representative voter can

²A 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.

be chosen such that it ranks states in the same order as the electorate.

We first introduce some notation for the different geometric objects that will aid our analysis in this section. First, to a set of voters s we associate the set

$$A_s = \{q \in \Delta(\Theta) : \langle \delta, q \rangle \geq 0, \delta \in s\},$$

which is the win-set under unanimity for s . Note that A_s is a polyhedron (as the intersection of half-spaces) which is also bounded. We will sometimes refer to “phantom voters” (p-voters) to describe voters who may or may not be part of the original electorate. Finally, for a bounded polyhedron K we denote by $\varepsilon(K)$ the set of its extreme points (vertices). A face of K is the intersection of K with any supporting hyperplane, a proper face is a face that is not K , and a facet of K is any proper face that is not contained in some other face. These concepts and some results in polyhedral geometry that we use in the proofs can be found in Ziegler (1995).

Lemma B.1 *Consider two finite groups of p-voters s_1 and s_2 . Then, there exists a finite set of p-voters s , such that*

$$\text{co}(A_{s_1} \cup A_{s_2}) = A_s. \tag{B.3}$$

That is, $\text{co}(A_{s_1} \cup A_{s_2})$ is the win-set with unanimity for s . If all voters in s_1 and s_2 rank states in the same order, then there exists a set s satisfying (B.3) such that every voter in s ranks states in the same order as voters in s_1 and s_2 .

Proof of Lemma B.1 Since both A_{s_1} and A_{s_2} are polyhedra, we can represent its convex hull as the projection of a higher dimensional polyhedron (see, e.g., Balas 1985)

$$\text{co}(A_{s_1} \cup A_{s_2}) = \left\{ q : q = q_1 + q_2, 0 \leq \lambda \leq 1, 0 \leq q_1 \leq \lambda, 0 \leq q_2 \leq 1 - \lambda, \begin{array}{l} \langle \delta_1, q_1 \rangle \geq 0, \delta_1 \in s_1, \\ \langle \delta_2, q_2 \rangle \geq 0, \delta_2 \in s_2 \end{array} \right\}.$$

As the projection of a finite, bounded polyhedron, $\text{co}(A_{s_1} \cup A_{s_2})$ is also finite and bounded.

We can then find vectors $\tilde{\delta}_i$ and scalars b_i , $i = 1, \dots, I$, such that

$$\text{co}(A_{s_1} \cup A_{s_2}) = \left\{ q \in \Delta(\Theta) : \langle \tilde{\delta}_i, q \rangle \geq b_i, i = 1, \dots, I \right\}.$$

Defining $s = \left\{ \delta_i : \delta_i = \tilde{\delta}_i - b_i \vec{1}, i = 1, \dots, I \right\}$ we have

$$\text{co}(A_{s_1} \cup A_{s_2}) = \{q \in \Delta(\Theta) : \langle \delta_i, q \rangle \geq 0, \delta_i \in s\},$$

which implies (B.3).

We now show that ranking of states by voters is preserved under convexified unions, in the sense that one can represent $co(A_{s_1} \cup A_{s_2})$ as the unanimous choice of a collection of voters that rank states in the same order as s_1 and s_2 . We first note that if all voters share the same (strict) ranking, then there are only three possible cases: (i) $A_{s_1} \cup A_{s_2}$ is empty, or (ii) $A_{s_1} \cup A_{s_2}$ consists of a single point, or (iii) $\dim(co(A_{s_1} \cup A_{s_2})) = \dim(\Delta(\Theta)) = \text{card}(\Theta) - 1$. To see this, let θ^* be the state corresponding to the highest payoff for all voters. Then, we have three distinct situations: (i) $\langle \delta_i, 1_{\theta^*} \rangle < 0$, for a pair $\delta_i \in s_i, i = 1, 2$; (ii) $\langle \delta, 1_{\theta^*} \rangle \geq 0, \delta \in s_i, i = 1, 2$ with $\langle \delta_i, 1_{\theta^*} \rangle = 0$ for a pair $\delta_i \in s_i, i = 1, 2$; or (iii) either for $i = 1$ or for $i = 2$ we have $\langle \delta, 1_{\theta^*} \rangle > 0, \delta \in s_i$. This implies that in case (i) each A_{s_i} is empty, in case (ii) $A_{s_1} \cup A_{s_2}$ consists of a single point, and (iii) $A_{s_1} \cup A_{s_2}$ has full dimensionality as either A_{s_1} or A_{s_2} has a non-empty relative interior. It is immediate to find a voter that satisfies the conditions of the Lemma for cases (i) or (ii) -any voter that shares the same ranking as the electorate but never approves would satisfy (B.3) for case (i), while a voter that shares the same ranking with $\delta_{\theta^*} = 0$ would satisfy (B.3) for case (ii). Therefore, in the remainder we will focus on case (iii) and assume full dimensionality of $co(A_{s_1} \cup A_{s_2})$.

First, we have that

$$\varepsilon(co(A_{s_1} \cup A_{s_2})) \subseteq \varepsilon(A_{s_1}) \cup \varepsilon(A_{s_2}),$$

that is, any extreme point of $co(A_{s_1} \cup A_{s_2})$ is either an extreme point of A_{s_1} or it is an extreme point of A_{s_2} (or both)(see Ziegler, 1995). Knowledge of extreme points will be useful as we can represent each facet F of $co(A_{s_1} \cup A_{s_2})$ through its extreme points, in the sense that any hyperplane containing these points is a supporting hyperplane of F . Finally, we will say that F is an ‘‘interior’’ facet of $co(A_{s_1} \cup A_{s_2})$ if it is not contained in a proper facet of the simplex $\Delta(\Theta)$. Our proof strategy will be to find a supporting hyperplane associated to each ‘‘interior’’ facet of $co(A_{s_1} \cup A_{s_2})$ that corresponds to a voter that ranks states in the same order as s_1 and s_2 .

Let F be an interior facet of $co(A_{s_1} \cup A_{s_2})$ and enumerate all its extreme points according to $\{\hat{q}_1, \dots, \hat{q}_J\} = \varepsilon(F)$. Note that full dimensionality of $co(A_{s_1} \cup A_{s_2})$ implies that if F does not exist, then $co(A_{s_1} \cup A_{s_2}) = \Delta(\Theta)$ (in which case any voter δ that ranks states as s_1 and s_2 and always approves would satisfy (B.3)). To each extreme point \hat{q} in F such that \hat{q} is

also an extreme point of A_{s_i} we associate the set of indifferent voters $J(\hat{q}) \subseteq s_i$ such that, for $\delta \in s_i$, we have

$$\delta \in J(\hat{q}) \Leftrightarrow \langle \delta, \hat{q} \rangle = 0.$$

Note that $J(\hat{q}) \neq \emptyset$. Indeed, any extreme point of A_{s_i} must satisfy with equality a subset of the inequalities $\{\langle \delta, q \rangle \geq 0, \delta \in s_i, q_\theta \geq 0, \langle 1, q \rangle \geq 1, \langle -1, q \rangle \geq -1\}$. Moreover, if for $\hat{q} \in \varepsilon(A_{s_i})$ we have $\langle \delta, \hat{q} \rangle > 0$ for every $\delta \in s_i$ then \hat{q} corresponds to one of the extreme points of the simplex $\Delta(\Theta)$, which is not in an interior facet of $co(A_{s_1} \cup A_{s_2})$. Finally, we group all these voters in the set $T = \cup_{i=1}^J J(\hat{q}_i)$ and let $\{\delta^1, \dots, \delta^M\} = T$ be an enumeration of such voters.

Let γ be a supporting hyperplane of the facet F in the sense that

$$\begin{aligned} \langle \gamma, q \rangle &= k, q \in F, \\ \langle \gamma, q \rangle &\geq k, q \in co(A_{s_1} \cup A_{s_2}). \end{aligned}$$

for some $k \in \mathbb{R}$. By construction, γ is a separating hyperplane for the set

$$P = \{x \in \mathbb{R}^{card(\Theta)} : \langle \delta^i, x \rangle \geq 0, i \in \{1, \dots, M\}\}, \quad (\text{B.4})$$

in the sense that $\langle \gamma, x \rangle \geq k$ for $x \in P$.

We will now show that there exist a non-negative vector $\nu = \{\nu^1, \dots, \nu^M\}$ such that

$$\gamma = \sum_{i=1}^M \nu^i \delta^i, \sum_{i=1}^m \nu^i = 1, \nu^i \geq 0. \quad (\text{B.5})$$

To see how existence of such ν establishes the last claim in Lemma B.1, suppose all voters in s_1 and s_2 rank states in the same order. As $\nu^i \geq 0$, then $\gamma_F^* = \gamma - k \vec{1}$ corresponds to a voter that ranks states in the same order as voters in s_1 and s_2 . This implies that to each interior facet F of $co(A_{s_1} \cup A_{s_2})$ corresponds a voter γ_F^* who ranks states in the same order as voters in s_1 and s_2 , and

$$co(A_{s_1} \cup A_{s_2}) = \{q \in \Delta(\Theta) : \langle \gamma_F^*, q \rangle \geq 0, F \text{ interior Facet of } co(A_{s_1} \cup A_{s_2})\}.$$

To show that (B.5) must have a solution, we construct a linear system of equations associated to (B.5), and its Farkas alternative (see Ziegler 1995), and show that the Farkas alternative cannot have a solution.

The fact that γ is a supporting hyperplane of the facet F with extreme points $\{\hat{q}_1, \dots, \hat{q}_J\}$, implies that $\langle \gamma, \hat{q}_i \rangle = \langle \gamma, \hat{q}_k \rangle$, $\hat{q}_i, \hat{q}_k \in \varepsilon(F)$. Thus (B.5) has a solution if and only if the following system of linear equations has a solution.

$$\sum_{i=1}^M \nu^i \langle \delta^i, \hat{q}_{j+1} - \hat{q}_j \rangle = 0, j \in \{1, \dots, J-1\}, \sum_{i=1}^M \nu^i = 1, \nu^i \geq 0 \quad (\text{B.6})$$

The Farkas alternative to (B.6) is the system of equations

$$\left\langle \delta^i, \sum_{j=1}^{J-1} y^j (\hat{q}_{j+1} - \hat{q}_j) \right\rangle + y^J \geq 0, i \in \{1, \dots, M\}, y^J < 0 \quad (\text{B.7})$$

Farkas lemma implies that either (B.6) has a solution or (B.7) has a solution, but not both. Suppose that (B.7) has a solution y^* and let

$$\tau = y^{*J-1} \hat{q}_J + \sum_{j=1}^{J-1} (y^{*j-1} - y^{*j}) \hat{q}_j - y^{*1} \hat{q}_1.$$

Note that τ belongs to the linear subspace generated by $\varepsilon(F)$. In particular, τ is on the supporting hyperplane γ . Recall that γ also separates P , as given by (B.4), in the sense that $\langle \gamma, x \rangle \geq k$, $x \in P$. Note, however, that (B.7) and the definition of γ imply that

$$\begin{aligned} \langle \delta^i, \tau \rangle &> 0, i \in \{1, \dots, M\}, \\ \langle \gamma, \tau \rangle &= 0. \end{aligned}$$

That is, γ also intersects the interior of P as $\tau \in \text{int}(P)$. Thus, we reach a contradiction, and (B.7) cannot have a solution. Therefore, there exist a non-negative vector ν that satisfies (B.5). This implies that the facet F admits a hyperplane associated to a voter that ranks states in the same order as s_1 and s_2 . ■

The interest in this lemma is that it allows us to represent the convex-hull of the win-set under a k -voting rule as the unanimous choice of an auxiliary electorate, where ranking of states is preserved in this auxiliary electorate.

Proposition B. 1 *Fix a k -voting rule and electorate $\{\delta^1, \dots, \delta^n\}$. Then, there exist a set of p -voters $s^* = \{\gamma^1, \dots, \gamma^L\}$ such that*

$$\text{co}(W_k) = \{q \in \Delta(\Theta) : \langle \gamma, q \rangle \geq 0, \gamma \in s^*\}.$$

That is, from the point of view of the information controller and each voter $\{\delta^1, \dots, \delta^n\}$, for any prior belief, persuading the electorate $\{\delta^1, \dots, \delta^n\}$ under a k -voting rule is equivalent to persuading the electorate s^* under unanimity. If all voters in $\{\delta^1, \dots, \delta^n\}$ rank states in the same order, then the set s^* can be chosen so that all $\{\gamma^1, \dots, \gamma^L\}$ rank states in the same order as voters in $\{\delta^1, \dots, \delta^n\}$.

Proof of Proposition B.1 For an electorate $\{\delta^1, \dots, \delta^n\}$ and k -voting rule, let \mathcal{S}_k be the set of k -coalitions of voters in $\{\delta^1, \dots, \delta^n\}$, and let $\{s_1, \dots, s_C\} = \mathcal{S}_k$ be an enumeration of all k -coalitions. Note that $W_k = \cup_{i=1}^C A_{s_i}$. Define recursively \tilde{A}_i , $i \in \{1, \dots, C\}$, with $\tilde{A}_1 = A_{s_1}$ and

$$\tilde{A}_i = co(A_{s_i} \cup \tilde{A}_{i-1}).$$

If we let $B_i = \cup_{j=1}^i A_{s_j}$ then we have

$$co(B_i) = \tilde{A}_i.$$

To see this, note that it is true for $i = 1$. Suppose it is true for $i > 1$. Then

$$co(B_{i+1}) = co(A_{s_{i+1}} \cup B_i) = co(A_{s_{i+1}} \cup co(B_i)) = co(A_{s_{i+1}} \cup \tilde{A}_i) = \tilde{A}_{i+1}.$$

By induction, we then have $\tilde{A}_C = co(\cup_{i=1}^C A_{s_i}) = co(W_k)$. By Lemma B.1, each \tilde{A}_i corresponds to a unanimous decision among a set of voters s_i . Thus, $co(W_k)$ corresponds to a unanimous decision among an electorate s_C . By Lemma B.1 we can choose at each step the set s_i so that it is composed of voters that rank states in the same order as those in $\{\delta^1, \dots, \delta^n\}$. Setting $s^* = s_C$ in such case would satisfy the conditions of the Proposition. ■

The previous proposition showed that for a given k -voting rule, one can find a set of p -voters such that their unanimous choice replicates the decisions under a k -voting rule for *any* prior belief. We now show that, given prior belief $p \in \Delta(\Theta)$, if all voters in the electorate rank states in the same order, then a weak representative voter exists. Furthermore, one can find a weak representative voter that shares the ranking of the states of the electorate.

Proposition B. 2 Consider an electorate $\{\delta^1, \dots, \delta^n\}$ and suppose all voters rank states in the same order, $\delta^i \in \mathcal{F}_z$, $i = 1, \dots, n$, for some permutation z of states. Given prior belief

$p \in \Delta(\Theta)$, for each k -voting rule there exists a weak representative voter $\delta^*(k)$ who ranks states in the same order as the electorate, $\delta^*(k) \in \mathcal{F}_z$. Furthermore, if $p \notin \text{co}(W_k)$, then one can select $\delta^*(k)$ such that: (i) $\{q : \langle q, \delta^*(k) \rangle = 0\}$ is a supporting hyperplane of $\text{co}(W_k)$, and (ii) $\delta^*(k) \in \mathcal{F}_z$.

Proof of Proposition B.2: Let W_k be the win set under a k -voting rule. If $p \in \text{co}(W_k)$, then the controller's optimal signal guarantees approval with certainty. Hence, it is immediate to construct a weak representative voter who ranks the state in the same order as the voters, $\delta^*(k) \in \mathcal{F}_z$, and approves the proposal without further information, $\langle p, \delta^*(k) \rangle > 0$.

Now consider the remaining case, $p \notin \text{co}(W_k)$. By Proposition B.1, there exist a set $s^* = \{\tilde{\delta}^1, \dots, \tilde{\delta}^M\}$, with $\tilde{\delta}^i \in \mathcal{F}_z$, $i = 1, \dots, M$, and such that

$$\text{co}(W_k) = \left\{ q \in \Delta(\Theta) : \langle q, \tilde{\delta}^i \rangle \geq 0, i = 1, \dots, M \right\}.$$

As p has full support on Θ , we can equivalently represent $\text{co}(W_k)$ as

$$q \in \text{co}(W_k) \Leftrightarrow q = \frac{\hat{\alpha}p}{\langle \hat{\alpha}, p \rangle} \text{ with } \hat{\alpha} \in \left\{ \alpha : 0 \leq \alpha_\theta \leq 1, \langle \alpha, \tilde{\delta}^i p \rangle \geq 0, i = 1, \dots, M \right\}.$$

We use this second representation to describe and solve the controller's problem. Fix an optimal signal π^* . Let s^+ be the event corresponding to approval of the proposal under this optimal signal, and let $\alpha_\theta^* = \Pr[s^+ | \theta]$ so that $\Pr[\text{Approval}] = \sum_{\theta \in \Theta} \alpha_\theta^* p_\theta$. Since the expected approval posterior must be in $\text{co}(W_k)$, then the controller's optimal signal must satisfy

$$\sum_{\theta \in \Theta} \alpha_\theta^* p_\theta = \max_{\alpha} \sum_{\theta \in \Theta} \alpha_\theta p_\theta, \text{ s.t. } 0 \leq \alpha_\theta \leq 1, \langle \alpha, \tilde{\delta}^i p \rangle \geq 0, i = 1, \dots, M. \quad (\text{B.8})$$

Program (B.8) is linear. Consider the Lagrangian L associated to (B.8)

$$L = \langle \alpha, p \rangle + \sum_{\theta} \nu_\theta \langle \alpha, 1_\theta \rangle + \sum_{\theta} \mu_\theta \langle \vec{1} - \alpha, 1_\theta \rangle + \sum_{i=1}^M \kappa_i \langle \alpha, \tilde{\delta}^i p \rangle,$$

with $\nu_\theta, \mu_\theta, \kappa_i \geq 0$, and 1_θ is the unitary vector whose θ -component equals 1. As shown in the proof of Lemma B.1, if voters rank states in the same order then W_k has full dimensionality when it 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 $p \notin \text{co}(W_k)$, α^* is an optimal solution if and only if there exist $\lambda^*, \nu_\theta^*, \mu_\theta^*, \kappa_i^* > 0$, $\theta \in \Theta, i = 1, \dots, M$, such that

$$\lambda^* p + \sum_{\theta} \nu_\theta^* 1_\theta - \sum_{\theta} \mu_\theta^* 1_\theta + \sum_{i=1}^M \kappa_i^* \tilde{\delta}^i p = 0, \quad (\text{B.9})$$

with complementary slackness conditions

$$\nu_\theta \alpha_\theta = 0, \mu_\theta (1 - \alpha_\theta) = 0, \kappa_i^* \langle \alpha^*, \tilde{\delta}^i p \rangle = 0, i = 1, \dots, M. \quad (\text{B.10})$$

Consider the voter

$$\delta^* = \sum_{i=1}^M \kappa_i^* \tilde{\delta}^i. \quad (\text{B.11})$$

By construction, $\langle \delta^*, q \rangle = 0$ is a supporting hyperplane of $\text{co}(W_k)$. Moreover, as $\kappa_i \geq 0$ and $\tilde{\delta}^i \in \mathcal{F}_z, i = 1, \dots, M$, then $\delta^* \in \mathcal{F}_z$. Now consider the optimal signal α' under delegation to voter δ^* which must satisfy

$$\sum_{\theta \in \Theta} \alpha'_\theta p_\theta = \max_{\alpha} \sum_{\theta \in \Theta} \alpha_\theta p_\theta, \text{ s.t. } 0 \leq \alpha_\theta \leq 1, \langle \alpha, \delta^* p \rangle \geq 0. \quad (\text{B.12})$$

Again, this is a linear program with non-empty relative interior as $\text{co}(W_k)$ has a non-empty relative interior. Therefore α' is optimal if and only if there exist $\tilde{\lambda}, \tilde{\nu}_\theta, \tilde{\mu}_\theta, \tilde{\kappa} > 0, \theta \in \Theta$, such that

$$\tilde{\lambda} p + \sum_{\theta} \tilde{\nu}_\theta 1_\theta - \sum_{\theta} \tilde{\mu}_\theta 1_\theta + \tilde{\kappa} \delta^* p = 0. \quad (\text{B.13})$$

with complementary slackness conditions

$$\tilde{\nu}_\theta \alpha'_\theta = 0, \tilde{\mu}_\theta (1 - \alpha'_\theta) = 0, \tilde{\kappa} \langle \alpha', \delta^* p \rangle = 0, i = 1, \dots, M \quad (\text{B.14})$$

In particular, as α^* and $\lambda^*, \nu_\theta^*, \mu_\theta^*, \kappa_i^* > 0, \theta \in \Theta, i = 1, \dots, M$, satisfy (B.9) and (B.10), and given (B.11), then they also satisfy (B.13) and (B.14) by setting $\tilde{\kappa} = 1$. Thus, α^* also provides an optimal signal when delegating to voter δ^* .

As δ^* strictly ranks states, then the optimal signal is unique (in the sense that the optimal $\alpha'_\theta = \Pr[s^+|\theta]$ is unique) (see the proof of Proposition 2 in PV). With this observation, we now show that the optimal experiment when persuading the electorate $\{\delta^1, \dots, \delta^n\}$ under a

k -voting rule must also be unique. This also means that δ^* is a weak representative voter, as delegation to δ^* generates a single payoff profile for all players.

Suppose that program (B.8) admits two solutions α^1 and α^2 . Then $\sum_{\theta \in \Theta} \alpha_{\theta}^* p_{\theta} = \sum_{\theta \in \Theta} \alpha_{\theta}^j p_{\theta}$, $j = 1, 2$, and $\langle \alpha^j, \tilde{\delta}^i p \rangle \geq 0$, $i = 1, \dots, M, j = 1, 2$. But this means that (i) α^1 and α^2 are feasible for program (B.12), and (ii) α^1 and α^2 are also optimal solutions. Given the uniqueness of solution to (B.12), we must then have $\alpha^1 = \alpha^2 = \alpha^*$. ■

By establishing that $\delta^*(k)$ must be in the same class as the electorate, Proposition B.2 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 by Proposition 2 of PV. 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”.

B.2.2 Voter Heterogeneity and Information Control

Proposition 3 in PV shows 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 B. 3 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$, 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, $\delta^i \in \mathcal{F}_z$, $i \in I$, then $V(W_n) = V(\delta^n)$ and*

$$V(W_k) \geq V(\delta^k). \tag{B.15}$$

Proof of Proposition B.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 $\delta^i \in \mathcal{F}_z$, then Proposition 2 from PV shows that the structure of the optimal signal is the same for all voters: if $\alpha_\theta(\delta^i) = \Pr[\text{approval}|\theta]$ represents the optimal signal under delegation to voter δ^i , where $\alpha_\theta(\delta^i)$ is given by the proof of Proposition 2, 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 $V(W_k) \geq V(\delta^k)$. ■

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 (B.15) is satisfied with equality.

Inequality (B.15) holds whenever voters agree on the ranking of states. If voters rank states differently, then the reverse inequality to (B.15) 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 B.1 below). 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 B.2 below).

Example B.1: 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 B.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 ($\delta_{\theta_3}^{Ai} > 0 > \delta_{\theta_2}^{Ai} > \delta_{\theta_1}^{Ai}$). Group B has voters $\{\delta^{B1}, \delta^{B2}, \delta^{B3}\}$, who rank state θ_1 higher than θ_2 ($\delta_{\theta_3}^{Bi} > 0 > \delta_{\theta_1}^{Bi} > \delta_{\theta_2}^{Bi}$). Figures B.1(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 B.1(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 B.1(b) depicts the posterior beliefs induced by an optimal signal with unanimity rule, which is also optimal given any $k > 3$. \square



(a) Example B.1, with approval delegated to δ^{A3} (b) Example B.1, with unanimity (c) Example B.2, with $k = 1$

Figure B.1: Optimal Signals for Examples B.1 and B.2

Example B.2: We now 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 voters δ^A , δ^B and δ^C with payoffs described by Table B.1:

| 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 |

Table B.1: Payoffs from Example B.2

If the decision is delegated to 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 approval³ is 34%. Therefore, δ^C is the easiest-to-persuade voter. However, if the voting rule is $k = 1$, then the optimal signal is only supported on three posterior beliefs $\{q_A^+, q_B^+, q^-\}$ and it does not target voter δ^C (see Figure B.1(c)): belief q_A^+ convinces voter δ^A to approve the proposal but it does not convince voters δ^B and δ^C ; belief q_B^+ convinces voter δ^B but not voters δ^A and δ^C ; belief q^- 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 B.1(c) 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

B.2.3 Voter preferences over decision makers

Suppose that the approval decision is made by a single voter (dictator) δ : the controller only needs to persuade this voter. Now suppose that voter δ can choose whom to delegate his approval decision. How would voter δ rank different decision makers? Voter δ faces

³Formally, $S = \{s^-, s^+\}$, $Pr(s^+|\theta_3) = Pr(s^+|\theta_2) = 1$, 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 $Pr(s = s^+) = \frac{2}{35} \times 0.7 + 1 \times 0.1 + 1 \times 0.2 = 0.34$.

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 \mathcal{F}_z , that is, rank states in the same order z .

Proposition B. 4 *Fix a permutation z and let $\delta^v \in \mathcal{F}_z$. Consider any totally ordered (according to toughness) set of voters $\mathcal{D} \subset \mathcal{F}_z$, 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 $\mathbb{R}^{|\Theta|}$, is achieved by any voter $\delta^* \left(\hat{\delta}, \delta^v \right) = \hat{\delta} - \hat{\gamma}(\hat{\delta})\mathbf{1} \in \mathcal{F}_z$, where $\hat{\delta} \in \mathcal{F}_z$ and*

$$\hat{\gamma}(\hat{\delta}) = \sum_{\theta \in \{\theta: \delta_{\theta}^v \geq 0\}} p_{\theta} \hat{\delta}_{\theta}. \quad (\text{B.16})$$

Proof of Proposition B.4: Without loss of generality, suppose that $z(i) = i$ so that for $\delta \in \mathcal{F}_z$ we have $\delta_{\theta_i} < \delta_{\theta_{i+1}}$, $i \in \{1, \dots, \text{card}(\Theta) - 1\}$. From Proposition 2 in PV, the controller's optimal signal when the decision is made by voter $\delta \in \mathcal{F}_z$ is characterized by the approval conditional probabilities $\alpha_{\theta}(\delta) = \Pr[\text{approval}|\theta]$ 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 $\delta \in \mathcal{F}_z$ 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 2, 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 $\mathcal{F}(D)_z$ be the set of voters that rank states according to z and who share the same set of approval states D . For any $\delta, \delta' \in \mathcal{F}(D)_z$ 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 B.4 implies that all voters in $\mathcal{F}(D)_z$ have monotone preferences over totally ordered chains in $\mathcal{F}(D)_z$.

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 . ■

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 \mathcal{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 \mathcal{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 \mathcal{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 (B.16). 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.

B.2.4 Voter preferences over k -voting rules

How does each voter rank different voting rules? The next lemma shows that, if voters belong to the same class, then each voter has single peaked preferences over k .

Lemma B.2 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$, with $\delta^i \in \mathcal{F}_z$, 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 $k > k^*(\delta^i)$.*

Proof of Lemma B.2: Proposition B.2 implies that if all voters are in the same class \mathcal{F}_z , then for each k there exists a weak representative voter $\delta^*(k) \in \mathcal{F}_z$ and, furthermore, $A(\delta^*(k')) \subset A(\delta^*(k))$ for $k' > k$. Therefore, $D = \{\delta^*(k) : k \in \{1, \dots, n\}\}$ forms a totally ordered chain, and Proposition B.4 implies that each voter in the electorate has single-peaked preferences in D . This implies that each voter has single peaked preferences over k .

■

Proposition B.2 shows that, if voters rank states in the same order, then voters' expected utilities with a k -voting rule are the same as with delegation to a weak-representative voter $\delta^*(k)$, who also ranks states in the same order as the voters. The intuition behind Lemma B.2 is that since the weak-representative voter $\delta^*(k)$ also belongs to the same class \mathcal{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 \mathcal{F}_z , and Proposition B.4(i) implies that each voter has single-peaked preferences over these decision makers, and hence, over k -voting rules.

An important implication of Lemma B.2 is that a majority of voters prefer a supermajority voting rule over a simple majority voting rule.

Lemma B.3 *Consider an electorate $\{\delta^1, \dots, \delta^n\}$ with an odd number $n \geq 3$ of voters in the same class $\delta^i \in \mathcal{F}_z$, 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, $0 < V(W_{k'}) < V(W_{\frac{n+1}{2}})$.*

Proof of Lemma B.3: 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, $\text{card}(M) \geq \frac{n+1}{2}$, such that for each $\delta \in M$ we have $\langle q_{\frac{n+1}{2}}^+, \delta \rangle \leq 0$. 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 $\langle q_n^+, \delta \rangle \geq 0$. Therefore all voters in M weakly prefer unanimity over simple majority. Using Lemma B.2, 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 $\langle q_{\frac{n+1}{2}}^+, \delta \rangle > 0$. 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. ■

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

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

Proof of Proposition B.5: Together Propositions B.2 and B.4(ii) imply that all voters have monotone preferences over a chain of weak-representative voters $D = \{\delta^*(k) : k \in \{1, \dots, n\}\}$, when they all agree under full information. This implies that each voter has monotone preferences over k . ■

We conclude this section with an example to show that, in order to obtain the results in Lemmas B.2 and B.3, one may not replace the same class assumption with an agreement under full information assumption.

Example B.3: 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 $\delta_{\theta_2} > \delta_{\theta_1}$, while voters δ^D and δ^E (who have the same type) rank $\delta_{\theta_2} < \delta_{\theta_1}$. Approval sets and win sets for different voting rules are depicted in Figure B.2. 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 B.2, 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 B.3, a majority of voters (voters δ^A , δ^B and δ^C) strictly prefer simple majority over the supermajority $k = 4$. Nevertheless, a majority of voters (voters δ^C , δ^D and δ^E) weakly prefer unanimity over simple majority. □




Figure B.2: Optimal Signal for Example B.3

B.3 Extensions

B.3.1 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 ? We next 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.

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 (2016) 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 $\{\delta^1, \dots, \delta^n\}$ 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 believe 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.3.2 Controller's Payoff Depends on the State

In this section we study the case in which 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, $\delta_\theta^C \neq 0$ for all $\theta \in \Theta$.

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 2 in PV.

Proposition B. 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

$$\alpha_\theta = \begin{cases} 1 & \text{if } \delta_\theta^C > 0, \left| \frac{\delta_\theta^C}{\delta_\theta} \right| > \left| \frac{\delta_{\theta'}^C}{\delta_{\theta'}} \right| \text{ or } \delta_\theta^C < 0, \left| \frac{\delta_\theta^C}{\delta_\theta} \right| < \left| \frac{\delta_{\theta'}^C}{\delta_{\theta'}} \right| \\ 0 & \text{if } \delta_\theta^C > 0, \left| \frac{\delta_\theta^C}{\delta_\theta} \right| < \left| \frac{\delta_{\theta'}^C}{\delta_{\theta'}} \right| \text{ or } \delta_\theta^C < 0, \left| \frac{\delta_\theta^C}{\delta_\theta} \right| > \left| \frac{\delta_{\theta'}^C}{\delta_{\theta'}} \right| \end{cases}, \text{ and } \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta = 0. \quad (\text{B.17})$$

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 B.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 $\Pr[\text{Approval}] = \sum_{\theta \in \Theta} \alpha_\theta p_\theta$. Voter δ will approve after observing s^+ if and only if

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

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:

$$\max \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta^C, \quad \text{s.t. } 0 \leq \alpha_\theta \leq 1, \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta \geq 0. \quad (\text{B.18})$$

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

$$\begin{aligned} \mathcal{L} &= \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta^C + \lambda \sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta + \sum_{\theta \in \Theta} \nu_\theta^+ \alpha_\theta + \sum_{\theta \in \Theta} \nu_\theta^- (1 - \alpha_\theta), \\ \frac{\partial \mathcal{L}}{\partial \alpha_\theta} &= p_\theta \delta_\theta^C + \lambda p_\theta \delta_\theta + \nu_\theta^+ - \nu_\theta^- = 0. \end{aligned} \quad (\text{B.19})$$

It must be the case that $\sum_{\theta \in \Theta} \alpha_\theta p_\theta \delta_\theta = 0$ and $\lambda > 0^4$. If players agree on approval, $\delta_\theta^C > 0$ and $\delta_\theta \geq 0$, then (B.19) implies $\nu_\theta^- > 0$ and $\alpha_\theta = 1$. If players agree on rejection, $\delta_\theta^C < 0$ and

⁴By contradiction, suppose $\lambda = 0$. From (B.19) 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 $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta < 0$.

$\delta_\theta \leq 0$, then (B.19) 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 (B.19) 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 (B.19) 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 (B.19) 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 (B.19) 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 . ■

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 B.4.1). We find that if all players rank states in the same order, then our main results on voters' preferences over k-voting rules continue to hold.

Proposition B. 7 *Consider an information controller δ^C and an electorate $\{\delta^1, \dots, \delta^n\}$, with $\delta^C, \delta^i \in \mathcal{F}_z$. 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 B.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 $\delta_\theta \geq 0$.⁵ 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

⁵By contradiction, suppose there is a state θ such that $\delta_\theta^C < 0$ and $\delta_\theta \geq 0$. Then every state θ' such that

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 B.6 becomes equivalent to the problem in Proposition 2 in PV 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 B.4 as follows. Take any totally ordered (according to toughness) set of voters $\mathcal{D} \in \mathcal{F}_z$. Then any voter $\delta \in \mathcal{F}_z$ has single-peaked preferences over decision makers in \mathcal{D} . Now consider a k -voting rule. Proposition B.2 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 B.2 and B.3 and Proposition B.5 continue to hold.

B.3.3 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 $z(\theta') > z(\theta)$ 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, $\sum_{\theta \in D(\delta^C)} p_\theta \delta_\theta < 0$.

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. $P_i(q) = \Pr[\mu_i \leq \langle q, \delta^i \rangle]$. 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 B. 8 *Suppose that for every voter i we have $\delta_\theta^i \notin [-\bar{\mu}, \bar{\mu}]$ with at least one state with $\delta_\theta^i > \bar{\mu}$, 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 $\{\check{\delta}^i\}_{i \in I}$, 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.

Proof of Proposition B.8 : Consider an electorate $\{\delta^i\}_{i \in I}$ and let $V_i^{PV}(q) = P_i(q)$ and $V_U^{PV}(q) = \prod_{i \in I} P_i(q)$ 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 $A(\check{\delta}^i) = \{q \in \Delta(\Theta) : \bar{\mu} \leq \langle q, \check{\delta}^i \rangle\} = A^P(\delta^i)$, 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 $j = \{1, \dots, N, U\}$, 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 $j = \{1, \dots, N, U\}$. 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 $j = \{1, \dots, N, U\}$, 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 $A(\tilde{\delta}^i) = A^P(\delta^i)$ 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 its lowest value for any belief $q \in R_i^P(\delta^i) \equiv \text{co}(D_i^C(\delta^i))$, 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

$$F_U(\langle q, \delta^i \rangle) = 1 - \frac{\langle q^+ - q, \delta^i \rangle}{2\bar{\mu}} = 1 - \frac{(1 - \lambda) \langle q^+ - q^-, \delta^i \rangle}{2\bar{\mu}} < \lambda, \quad (\text{B.20})$$

which leads to

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

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 $0 \leq \prod_{i \in I} P_i(q) < 1$. 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 (B.20) and $F(\mu) \leq F_U(\mu)$, we have $F(\langle q, \delta^{i'} \rangle) < \lambda$ and

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

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 $j = \{1, \dots, N\}$. 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 $V_U^{PV}(q) = \tilde{V}_U^{PV}(q)$. ■

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 B.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.⁶ Proposition B.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 $\check{\delta}^i = \delta^i - \mathbf{1}\bar{\mu}$.

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'

⁶This 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, i.e. $P_i(q_s) \notin \{0, 1\}$.

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 2). If all voters would reject the proposal under the prior belief, the same would be true under a unanimity rule. Proposition B.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 1 in PV 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.

B.3.4 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 (forthcoming), 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 $p^l \in \text{int}(\Delta(\Theta))$, with $l \in \{C, 1, \dots, n\}$. 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 (forthcoming) to characterize the controller’s optimal signal and her gain from information control. We now briefly discuss how heterogenous priors affects our insights.

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 $\delta_\theta^i \frac{p_\theta^i}{p_C^i}$ and induces rejection only for the negative states with the lowest $\delta_\theta^i \frac{p_\theta^i}{p_C^i}$. 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 B.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 B.5 all voters have the same preferences over the class of binary “approve-reject” signals that preserve the ranking of states — i.e., signals with a cutoff state with higher ranked states always inducing the approval realization. The fact that Proposition B.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 B.5 still hold with heterogenous priors.

B.4 Applications

B.4.1 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 $y_i : \Theta \rightarrow \mathbb{R}$ 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

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

For each voter i compute the net payoff from approval

$$\delta_\theta^i = (1 - t_1)w_i + y_i(\theta) - w_i = y_i(\theta) - t_1w_i. \quad (\text{B.21})$$

All voters belong to the same class \mathcal{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, $E[y_i(\theta)|q] \geq t_1w_i$.

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.⁷

Proposition B.2 imply that for each k -voting rule there is a weak representative voter $\delta^*(k) \in \mathcal{F}_z$, 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 B.7).

Next we present two examples based on this general setup. Example B.4 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

⁷Note 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.

the approval decision to a richer voter. Example B.5 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 B.4: 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 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 B.3 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. \square

Example B.5: 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 (B.21) 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 B.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. \square




Figure B.3: Win Set and Optimal Signal from Example B.5

B.4.2 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.⁸

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.⁹ 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 and the reform cannot be implemented.

⁸We are also grateful for suggestions by Navin Kartik.

⁹Formally, 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.

B.5 Optimal Endorsement

We now consider an alternative interpretation of the model, in which we substitute the politician’s choice of a policy experiment for the choice of an optimal endorser (intermediary).

As before, we consider one politician (controller) C and a group $i \in I = \{1, \dots, n\}$ of voters who must approve/reject a proposal, according to a k -voting rule. Suppose the incumbent politician privately draws one “idea” (proposal) θ from the set $\Theta = \{\theta_1, \dots, \theta_T\}$ according to the prior probability $p = (p_\theta)_{\theta \in \Theta}$, with $\theta_1 < \dots < \theta_T$. The novel feature is that the politician has access to a diverse set of potential endorsers: established individuals (other politicians, legislators or bureaucrats) whose policy preferences are publicly known. The politician can either directly show the realized θ to voters, or can privately show θ to an “endorser” (intermediary). If the politician privately shows the state to the endorser, then the endorser can then send a cheap-talk message to voters, in order influence voters and maximize his own payoff.¹⁰ We ask: from the point of view of the politician, who is the optimal endorser?

To gain some intuition, consider the pure-persuasion model in which the politician simply wants to maximize the probability of approval, and suppose voters rank states in the same order. To easy exposition, suppose δ_θ^i strictly increases in θ for each voter, $p \notin W_k$ and $W_k \neq \emptyset$. From Proposition 2 in PV, if the politician is unconstrained in his choice of a policy experiment, her optimal experiment defines a unique cutoff state θ^* such that: the proposal is approved with probability one for all states $\theta > \theta^*$, and rejected for sure if $\theta < \theta^*$. Therefore, the optimal endorser is simply an individual with preference δ^{e*} such that $\delta_\theta^{e*} > 0$ for all $\theta > \theta^*$, $\delta_\theta^{e*} < 0$ for all $\theta < \theta^*$, and $\delta_{\theta^*}^{e*} = 0$. When the politician privately shows θ to this endorser, the endorser can credibly send a cheap talk message to voters stating whether the state is above or below θ^* . Importantly, the equilibrium cheap-talk message is, in general, not a simple “support” or “no support” statement. Isomorphic to the optimal policy experiment in our benchmark model, the equilibrium message is a “targeted endorsement,” where the endorser specifies which coalition of voters should approve the proposal. Moreover, note that

¹⁰Alternatively, we can assume that the politician does not know θ , but she can assign the endorser to investigate and privately learn the realized θ (e.g., the politician commissions the endorser as the head of an investigative committee or government agency).

this optimal endorser is not the weak representative voter.

Interestingly, suppose that there exists a group E of potential endorsers, so that every endorser e ranks states in the same order as voters: δ_θ^e strictly increases in θ . For each endorser e let $\underline{\theta}^e$ be the lowest $\theta \in \Theta$ such that $\delta_\theta^e \geq 0$. Suppose the politician privately shows θ to endorser e , who then sends a cheap talk message to voters to maximize his own payoff. Then the proposal's probability of approval weakly increases in $\underline{\theta}^e$ if $\underline{\theta}^e < \theta^*$, and weakly decreases in $\underline{\theta}^e$ if $\underline{\theta}^e > \theta^*$. Therefore, the politician has single-peaked preferences over endorsers according to $\underline{\theta}^e$.

Finally, suppose the politician's payoff depends on the state. In particular, suppose the politician ranks states in the same order as voters, but she is more biased towards the proposal than voters. We can then apply the same logic to the optimal signal from Proposition B.7. The optimal endorser is someone less biased towards the proposal than the politician, but more biased than voters. Figure B.4 illustrates this point, showing the set of approval states for the politician, endorser and voters.




Figure B.4: Example of Optimal Endorser

References

- [1] **Alonso, Ricardo, and Odilon Câmara.** 2016. “On the Value of Persuasion by Experts.” *Unpublished manuscript*.
- [2] **Alonso, Ricardo, and Odilon Câmara.** Forthcoming. “Bayesian Persuasion with Heterogeneous Priors.” *Journal of Economic Theory*.
- [3] **Balas, Egon.** 1985. “Disjunctive Programming and a Hierarchy of Relaxations for Discrete Optimization Problems.” *SIAM Journal on Algebraic and Discrete Methods* 6(3): 466-486.
- [4] **Boyd, Stephen, and Lieven Vandenberghe.** 2004. *Convex Optimization*. Cambridge University Press.
- [5] **Fernandez, Raquel, and Dani Rodrik.** 1991. “Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty.” *American Economic Review* 81(5): 1146-1155.
- [6] **Kamenica, Emir, and Matthew Gentzkow.** 2011. “Bayesian Persuasion.” *American Economic Review* 101(6): 2590-2615.
- [7] **Ziegler, Günter M.** 1995. *Lectures on Polytopes (graduate texts in mathematics)*. Berlin: Springer.