Decentralized Advice

Timothy Feddersen^{*} Ronen Gradwohl[†]

Abstract

We compare the amount of information credibly transmitted by cheap talk when information is centralized to one sender and when it is decentralized, with each of several senders holding a distinct but interdependent piece. Under centralization, full information transmission is typically impossible. Under decentralization, however, the number of receivers is decisive: decentralized communication with one receiver is completely uninformative, but decentralized communication with multiple receivers can be fully informative. We analyze the extent of such fully-informative communication, and apply our results to the issue of transparency in advisory committees.

1 Introduction

The cheap-talk model and its many variants have served as workhorse models of communication across the social sciences, from political science to lingistics.¹ In the standard model a *sender* has private information that is decision

^{*}Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA. E-mail: tfed@kellogg.northwestern.edu.

[†]Department of Economics and Business Administration, Ariel University, Ariel 40700, Israel. E-mail: roneng@ariel.ac.il. Gradwohl gratefully acknowledges the support of United States-Israel Binational Science Foundation and National Science Foundation grant 2016734.

¹See Battaglini (2002) for references in political science, finance, and macroeconomics. See Stalnaker (2006) for a connection to linguistics and philosophy.

relevant for a *receiver*. The sender chooses a message from a set of available messages, and the receiver, who observes this message, makes a decision that determines both agents' payoffs. Communication is called cheap talk if the set of available messages does not depend on the sender's information, and if his choice of message does not directly influence payoffs. The central question in the cheap talk literature is, how much information can be credibly communicated?

One main conclusion is that, unless sender and receiver preferences are sufficiently aligned, it is not possible to credibly communicate all the information. There are two well-known exceptions: when there is one sender communicating with multiple receivers (Farrell and Gibbons, 1989), and when there are multiple senders with different preferences and perfect information about a multidimensional state, communicating with one receiver (Battaglini, 2002).² In this paper we show that a third case exists: when there are multiple senders with common values and multiple receivers with different preferences.

The motivating application of this paper is focused on information transmission from advisory committees to decision makers and, in particular, on whether or not transparency requirements can facilitate information transmission. There is a large literature examining the costs and benefits of transparency in agency relationships,³ and recent studies have demonstrated a clear cost to transparency in terms of its effect on information aggregation (Fehrler and Hughes, 2018; Gradwohl and Feddersen, 2018). Modeling the members of the advisory committee as senders we show conditions under which such transparent committees can fully reveal information even when preferences between committee members and decision makers are quite different. Thus, our formal analysis provides a rationale for mandated transparency in advisory

²There are additional examples in somewhat different settings, e.g., when information is certifiable (Mathis, 2008; Hagenbach et al., 2014) and when communication is dynamic (Renault et al., 2013; Golosov et al., 2014; Margaria and Smolin, 2018).

³This literature typically studies senders with career concerns—see Prat (2005) and Malesky et al. (2012) for extensive reviews on the literatures in political science and economics.

committees.

To understand the mechanism underlying our result, consider first a sender who observes a good or bad signal, and must decide which of two messages to send to a receiver. The receiver observes the sender's message, obtains additional information about the realization of an event, and then decides either yes or no. To make things simple, suppose that there is only one event E in which the receiver might be influenced by the sender's information. If it is the case that the sender and receiver's preferences are aligned in that event then the standard result in the cheap-talk literature is that there exists an *informative equilibrium* in which the receiver learns the sender's information prior to making her decision.

Let's assume instead that, in the event E, the receiver prefers to choose yes if the sender's signal is good and no otherwise. However, the sender, knowing that event E has occurred, prefers that the receiver choose yes regardless of his signal. In this case the cheap-talk model predicts that the sender will be unable to credibly reveal any persuasive information to the receiver. The intuition is that if there were a message the sender could send that would cause the receiver to believe he had observed the good signal (and, as a result, persuade the receiver to choose yes) then the sender would send that message even when he observes the bad signal.⁴

The central intuition underlying the results of the current paper can be understood if we postulate a second event, E', in which the receiver might be influenced by the sender's information. Like event E, in event E' the receiver would like to choose yes if and only if the sender's signal is good. But unlike event E, in event E' the sender would always like the receiver to choose no. If, in addition, the sender's signal is sufficiently correlated with the events E and E', then the sender may prefer to truthfully report his signal. More specifically, if observing the good signal causes the sender to believe event E is

⁴This intuition underlies the results of Wolinsky (2002), Battaglini (2017), and Gradwohl and Feddersen (2018) on the impossibility of information transmission from a committee to a receiver.

more likely than event E', while observing the bad signal induces the opposite inference, then truthful reporting may be incentive compatible.

In this paper we consider an environment in which there are multiple senders with common values, each with a bit of information, and each sending a cheap-talk message to a receiver. From the point of view of each sender, the receiver observes not just that sender's message, but also an event that consists of the other senders' messages. Hence, when the receiver's preferences are known and sufficiently different from the senders', truthful revelation is not incentive compatible: If each sender truthfully reports his information, the event E in which his message matters is one in which he wants the receiver to make the same decision, regardless of his signal. However, when there are multiple receivers with preferences different both from the senders' and from each other, then truthful revelation may produce two different kinds of events: ones in which the senders all want to lie in one direction (E), and ones in which they want to lie in the other (E'). We show that senders' signals are sufficiently correlated with the events, leading to the possibility of truthful reporting in equilibrium.

We will compare a setting in which information is centralized, with a single sender obtaining all information, to one in which it is decentralized, with each of several senders obtaining some information. In terms of the application to advisory committees, centralized information corresponds to an opaque committee, whereas decentralized information corresponds to a transparent committee.

Under centralization, standard cheap talk analysis concludes that, regardless of the number of receivers, one of the following occurs: either no persuasive information can be credibly transmitted, or there is partial but not full information transmission.⁵ Which of the two possibilities is realized depends on the quality of information, as characterized by the number and accuracy of the sender's signals. Under decentralization, however, the amount of in-

⁵See, e.g., the analyses in Battaglini (2017) and in Gradwohl and Feddersen (2018).

	1 receiver	2 receivers		1 receiver	2 receivers
centralized	none	none	centralized	some	some
decentralized	none	all	decentralized	none	all
(-) I lit if-			(1) TT 1 1	· · c	

(a) Low-quality info

(b) High-quality info

Figure 1: Amount of information communicated in equilibrium

formation transmitted depends on whether there is one receiver (and so one pivotal event, E) or multiple receivers (and so multiple pivotal events, E and E'). As discussed above, in the former case *no* persuasive information can be communicated in equilibrium, whereas in the latter case *all* information may be communicated. This comparison is summarized in Figure 1, where the contribution of the current paper consists of the second column of each table.

Our analysis can be applied to a variety of settings beyond that of multiple receivers. For example, when there is one receiver who must decide on one of two actions, there is one pivotal event. When she has more actions from which to choose, or when the senders are uncertain about the receiver's pair of available actions, then there are multiple pivotal events. Similarly, if there are multiple receivers, or when the senders are uncertain about the receiver's preferences, full-information transmission may be possible. We illustrate our result in the following example, contrasting the impossibility of any information transmission with the possibility of full-information transmission.

Example 1 There are two decision makers (DMs), each of whom must decide whether or not to implement a particular gun-control policy in their respective districts. The first DM only prefers the policy if she is sufficiently certain that the policy will reduce crime (say, with probability above 75%). The second DM prefers the policy even if she assigns somewhat low probability to its effectiveness (as long as it is above, say, 25%). In addition, there are several experts with information about the effectiveness of the policy. The experts care about the actions of both receivers, and prefer that both implement the policy whenever it is more likely to be effective than not.

Now, if the experts communicate with each DM separately, then for each

such interaction there is one pivotal event E, and no credible communication could take place in equilibrium. In contrast, if the experts communicate with the DMs simultaneously, where all messages are seen by both DMs, then there are two pivotal events: the event in which an expert's message changes the first DM's decision, and the event in which it alters the second DM's decision. The insight of our paper implies that here, full-information transmission may possible, leading to more-informed decisions about the poilcy.

In our formal analysis we provide conditions under which full-information transmission is possible, and apply this result to argue for the benefit of mandating transparency in advisory committees. In particular, we show that if the committee's information is of low quality—there are few members or their signals are not very accurate—then transparency leads to higher welfare for both the committee and the DMs. In this case, there is no need to mandate transparency, as the committee would voluntarily choose to act transparently. In contrast, if the information is of high quality, then the committee prefers opacity while the DMs prefer transparency. In this case, the DMs benefit from mandating transparency.

Next, we analyze the robustness of such beneficial transparency, and of fully-informative communication in general. We show that the result relies critically on the lack of observability of communication between each sender and the receivers. That is, if communication with the receivers is sequential rather than simultaneous, informative communication is once more unattainable.

Finally, we consider two extensions to the model. First, we study more general preferences, and examine the extent of fully-informative communication under a centralized sender. Second, we consider the possibility and extent of partially-informative communication when senders value the actions of one receiver significantly more than those of the other.

The rest of the paper proceeds as follows. Immediately following is a review of the relevant literature. Section 2 describes our model of receivers, and Sections 3 and 4 contain our model of and main results on centralized and decentralized senders, respectively. These are followed by Section 5 on advisory committees. The extensions to the more general model and to partially-informative communication are in Sections 6 and 7, followed by the conclusion in Section 8. Most of the proofs are deferred to the Appendix.

Literature review This paper fits into the large literature on cheap talk (see Farrell and Rabin, 1996; Sobel, 2013, for excellent surveys). It is most closely related to a model of cheap talk in which there are multiple receivers, introduced by Farrell and Gibbons (1989), in which fully-informative communication may be possible (see also Goltsman and Pavlov, 2011). The driving force behind their possibility result, however, is distinct from that of our paper. In Theorem 5 we apply the insight of Farrell and Gibbons (1989) to our model, and show that it leads to full information transmission only in quite limited circumstances.

Our paper is also related to the large literature on cheap talk with multiple senders. There are two strands of this literature, depending on whether the senders have identical or different preferences. The first case has been extensively studied in various contexts, including legislative politics (Gilligan and Krehbiel, 1987; Austen-Smith, 1993; Li et al., 2001), polling (Morgan and Stocken, 2008), public protests (Battaglini, 2017), expect advice and advisory committees (Wolinsky, 2002; Gradwohl and Feddersen, 2018). The main conclusion from this literature is that preference differences between the senders and the receiver lead to losses in the the informativeness of communication. In Gradwohl and Feddersen (2018) (henceforth GF), for example, we show that regardless of the structure of communication between the senders and receiver, no communication is possible in equilibrium.⁶ The current paper builds on the model of GF, and shows that this conclusion is reversed when there are multiple receivers.

⁶Similar results in different contexts appear also in Wolinsky (2002) and Battaglini (2017).

A second strand of the literature on multiple senders considers the case in which senders have different preferences. Battaglini (2002) shows that full information transmission is possible when the state space is multidimensional and the senders each have perfect information. In subsequent work, Battaglini (2004) studies imperfectly-informed senders and Ambrus and Takahashi (2008) consider a restricted state space, and both show that the possibility of fullyinformative communication in this setting is limited.

Our paper is also related to work on information aggregation in committees, (e.g. Austen-Smith and Banks, 1996; Gerardi et al., 2009; Plott and Llewellyn, 2015) and particularly to the paper of Austen-Smith and Feddersen (2006). They show that when legislators deliberate prior to voting, preference uncertainty increases their ability to reach informed decisions. Our paper expands their insight to a general cheap talk environment.

Finally, the paper fits into the large literature on the costs and benefits of transparency to decision making (e.g. Hansen et al., 2014; Fehrler and Hughes, 2018; Gradwohl and Feddersen, 2018; Paetzel et al., 2018; Shambaugh and Shen, 2018, and many others).

2 Model

There are two possible, equally-likely states of the world, $\Theta = \{G, B\}$, and one or two decision makers (receivers), each of whom must decide between two possible outcomes, $\mathcal{O} = \{y, n\}$. The presence of multiple receivers is amenable to two additional interpretations: First, that there is one receiver but uncertainty about her preferences, and second, that there is one receiver but uncertainty about the actions available to her. We expand on these interpretations in Appendix A, but for the main body of the paper use the language of the initial, multiple-receiver interpretation.

Let $t \in \{D, L, H\}$ index the possible receivers: If there is one receiver index her by D, and if there are two index them by L and H, as described below. Receiver $t(\mathbf{R}_t)$ taking action $o \in \mathcal{O}$ in state $\theta \in \Theta$ derives utility $u_t(\theta, o)$, where $u_t(G, y) > u_t(G, n)$ and $u_t(B, n) > u_t(B, y)$. Given a belief $\beta = P(\theta = G)$ about the probability that the state is G, \mathbf{R}_t 's expected utility on choosing outcome o is $U_t(\beta, o) \stackrel{\text{def}}{=} \beta \cdot u_t(G, o) + (1 - \beta) \cdot u_t(B, o)$. A rational receiver will choose outcome y if and only if $U_t(\beta, y) \ge U_t(\beta, n)$.⁷ Since $U_t(\beta, y)$ is increasing in β and $U_t(\beta, n)$ is decreasing in β , there exists a threshold β_t such that receiver t will choose y if and only if $\beta \ge \beta_t$.

Now, if there are two receivers, index them so that $\beta_L \leq \beta_H$. For simplicity and tractability we will assume throughout that $\beta_L = 1 - \beta_H$, but our main results do not depend on this (see Appendix G).

Before making a decision, the receivers may obtain information from either a centralized or decentralized source, which we describe in Sections 3 and 4, respectively.

3 Centralized Information

Suppose first that decision-relevant information is centralized, held by a single agent called the sender. To facilitate the comparison with the decentralized setting, suppose the sender has access to an odd number N of conditionally-independent, identically-distributed signals (s_1, \ldots, s_N) of accuracy $p \in (1/2, 1)$, where each signal satisfies

$$P(s_i = g|\theta = G) = P(s_i = b|\theta = B) = p.$$

The sender's utility u is additive in the actions of the receivers who are present. Specifically, the sender obtains utility $u^t(\theta, o_t)$ from choice o_t by R_t in state θ . If there is only one receiver then the sender's total utility is $u(\theta, o_D) = u^D(\theta, o_D)$. If there are two receivers then $u(\theta, o_L, o_H) = u^L(\theta, o_L) + u^H(\theta, o_H)$. As with the receivers, we suppose that $u^t(G, y) > u^t(G, n)$ and $u^t(B, n) > u^t(B, y)$ for each t. The sender prefers outcome y from R_t whenever

⁷Assume she always chooses y if indifferent.

his belief about the probability $P(\theta = G)$ about the state being G is above some threshold γ_t . In most of this paper we will assume that, when there are two receivers, $\gamma_L = \gamma_H$: that is, for any belief the sender may have about the state, he prefers outcome y from one receiver if and only if he prefers it also from the other receiver (but see Section 6 for a more general setting). For simplicity we will also assume that $u^t(G, n) = u^t(B, n) = 0$ and $u^t(G, y) =$ $c_t = -u^t(B, y)$.⁸ Note that the sender's utility from the receivers may be different, since $u^H(G, y)$ need not equal $u^L(G, y)$, but that he prefers outcome y from both receivers whenever his belief about the probability $P(\theta = G)$ about the state being G is above the threshold $\gamma = 1/2$, and outcome n from both otherwise. This means that the sender strictly prefers outcome y if the number of good signals is above N/2, and otherwise strictly prefers outcome n.

After observing the profile of signals, the sender sends a message $m \in M$ to the receivers, where M is some message space with $|M| \ge 2^N$. Denote by $\sigma : \{g, b\}^N \mapsto M$ a strategy of the sender. Upon observing a message, each receiver then updates her prior over the state, and takes an action that depends on whether the posterior surpasses her threshold β_t or not. Formally, given a strategy profile σ , denote the rational decision rule used by \mathbb{R}_t on message mas $r_t(\sigma, m)$, where for all m in the support of σ it holds that $r_t(\sigma, m) = y$ if and only if $P(\theta = G \mid \sigma, m) \ge \beta_t$, and $r_t(\sigma, m) = n$ otherwise.⁹ Denote by $r(\sigma) \equiv (r_L(\sigma, \cdot), r_H(\sigma, \cdot)).$

Observe that without any information, R_H will choose outcome n and R_L will choose outcome y. A profile σ is *persuasive* if there exists a message $m \in \text{supp}(\sigma)$ such that $r_L(\sigma, m) = n$ or $r_H(\sigma, m) = y$. Additionally, since we are interested in the possibility of information transmission in equilibrium, denote a strategy profile σ as *optimal* if it is optimal for the sender given the decision rules $r(\sigma)$ of the receivers that it induces. Optimal profiles, together

⁸This symmetry assumption is made for tractability – the important feature is that the sender's threshold γ lie between those of the different receivers.

⁹Assume each receiver chooses y when indifferent. For messages m that are not in the support of σ the choice of $r_t(\sigma, m)$ does not matter.

with the corresponding r, form a Perfect Bayesian Equilibrium, the standard notion of equilibrium in cheap talk games.¹⁰

Before stating our main result for centralized information we need one more definition: Let β_{maj} be the posterior probability on $(\theta = G)$, given that at least half of the signals are good. Formally,

$$\beta_{\mathrm{maj}}(p,N) \stackrel{\mathrm{def}}{=} \Pr\left[\theta = G \mid \#\{i: s_i = g\} \geq \frac{N}{2}\right].$$

The following theorem characterizes the kind of information transmission possible in equilibrium. The result is straightforward and is similar to results in Battaglini (2017) and Gradwohl and Feddersen (2018), but we find it useful to restate and parametrize to allow for two receivers.

Theorem 1 For any N and p

- there exists an optimal persuasive strategy σ of the sender if and only if $\beta_t \in [1 - \beta_{maj}(p, N), \beta_{maj}(p, N)]$ for all participating receivers R_t ;
- if σ is optimal and persuasive then the participating receivers choose outcome y if and only if $\#\{i : s_i = g\} \geq \frac{N}{2}$, and choose outcome n otherwise.¹¹

An immediate implication of the second bullet is that unless the receivers' utilities are almost identical to the sender's—namely, if they agree on the preferred outcome on every possible realized signal profile—there does not exist an optimal σ in which the receivers learn the realization of all the signals (see Claim 1 for a formal statement of this). That is, fully-informative communication is not possible with a centralized sender. Furthermore, note that Theorem 1 applies to both the case in which there is only one receiver and

¹⁰With appropriately defined off-equilibrium beliefs, for example that on $m \notin \operatorname{supp}(\sigma)$ the posterior is equal to the prior.

¹¹Except for the degenerate case in which $\beta_H = \beta_{\text{maj}}$ and $\beta_L = 1 - \beta_{\text{maj}}$, in which case only R_H chooses these outcomes, whereas R_L always chooses y.

the case in which there are two, and so the amount of information transmission, and particularly whether there is any, does not depend on the number of receivers.

Finally, when $\beta_H > \beta_{\text{maj}}(p, N)$ there is no optimal and persuasive strategy any optimal strategy cannot be persuasive. For fixed β_H and p, then, there is a minimum number of signals for which communication is persuasive. Denote this minimum by $N_C(\beta_H, p) \stackrel{\text{def}}{=} \min\{N \in \mathbb{Z}_+ : \beta_{\text{maj}}(p, N) \ge \beta_H\}$. We will show that under decentralized information fewer signals are necessary.

4 Decentralized Information

Instead of one sender with N signals, suppose now that there are N decentralized senders, numbered $\{1, \ldots, N\}$, each with his own signal. Furthermore, the decentralized senders have a common utility function u that is identical to that of the centralized sender of Section 3.

A strategy σ_i of sender *i* is a function from his signal to a distribution over $\{y, n\}$. Denote by $\sigma = (\sigma_1, \ldots, \sigma_N)$ a profile of strategies, and by $\sigma(s)$ the profile of strategies given signal profile $s = (s_1, \ldots, s_N)$. We will restrict ourselves to symmetric strategy profiles, ones in which $\sigma_i \equiv \sigma_j$ for all senders *i* and *j*. The only relevant aspect of realized profiles is thus the realized number of *y* votes, which we call the vote profile *v*. Also, denote by $\sigma(\theta)$ the distribution over vote profiles under σ is state θ .

After the senders vote, the receivers observe the realized vote profile v. As in the case of a centralized sender, the receivers update their beliefs about the state and take actions that depend on whether the posterior surpasses β_t or not. Formally, given a strategy profile σ , denote the rational decision rule used by \mathbf{R}_t on realized voting profile v as $r_t(\sigma, v)$.

A strategy profile σ of the senders is *informative* if it conveys some information: if there is some vote profile v that occurs with positive probability under σ , and such that $\Pr[\theta = G|\sigma, v] \neq 1/2$. A strategy profile σ is *persuasive* if it sometimes leads some receiver to choose differently: There is some vote profile v that occurs with positive probability under σ , and for which either $\Pr[\theta = G \mid \sigma, v] < \beta_L$ or $\Pr[\theta = G \mid \sigma, v] \ge \beta_H$. Note that, as with a centralized sender, if σ is not persuasive then the receivers base their choices only on the prior distribution over states.

When the receivers update their priors they condition on both the vote profile v and on the strategy profile σ . But what prevents a sender from deviating from σ , unbeknownst to the receivers? In the case of the centralized sender, we required his strategy to be optimal given the receivers' decision rules. For decentralized senders we will require each σ_i to be optimal for sender i conditional on the receivers acting rationally and given the strategies of the other senders. That is, the profile σ must constitute a Nash equilibrium given the receivers' rational decision rule that it induces. In a standard voting game, where senders vote and there is a fixed decision rule mapping vote profiles to outcomes, one may require that the voting strategy be in equilibrium. The difference here is that there is no fixed decision rule: instead, the decision rule is chosen endogenously by the receivers, given σ , t, and v. A profile σ is then an *equilibrium* if it is in equilibrium given the decision rules that it induces. Formally,

Definition 1 (equilibrium) A strategy profile σ is an equilibrium if for each sender *i*, signal s_i , and strategy σ'_i ,

$$\mathbb{E}\left[u(\theta, \overline{r}(\sigma(s))) \mid s_i\right] \ge \mathbb{E}\left[u(\theta, \overline{r}(\sigma'_i, \sigma_{-i}(s))) \mid s_i\right],$$

where $\overline{r}(\cdot) \equiv r_D(\sigma, \cdot)$ when there is one receiver and $\overline{r}(\cdot) \equiv (r_L(\sigma, \cdot), r_H(\sigma, \cdot))$ when there are two receivers, and the expectation is over θ , s_{-i} , and σ .

Example 2 below illustrates the idea.

4.1 One receiver

Is decentralization better than centralization? The case of one receiver was studied by GF, and the following example illustrates the main (negative) result.

Example 2 Let N be odd, and consider the strategy of fully-informative voting, in which each sender i votes $v_i = y$ if and only if $s_i = g$. Such voting is not an equilibrium when there is only one receiver with $\beta_D > p$: To see this, suppose each sender votes according to his signal, and note that on profiles in which only a bare majority voted y (specifically, if exactly $\lceil N/2 \rceil$ voted y) the posterior of the receiver will be p. She will thus choose outcome n on these profiles, and so the induced decision rule is a supermajority rule. But in this case it is well-known that fully-informative voting is not an equilibrium (Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1998).

GF prove a general theorem about the impossibility of any communication between decentralized senders and one receiver. For the theorem, define the threshold $\overline{\beta}(p) \stackrel{\text{def}}{=} p^2/(p^2 + (1-p)^2).^{12}$

Theorem 2 (GF) Fix N and p > 1/2. If $\beta_D \notin [1 - \overline{\beta}(p), \overline{\beta}(p)]$ then there does not exist any persuasive equilibrium strategy profile.¹³

Thus, if there is only a single receiver and preferences are not sufficiently close, centralized information is better than decentralized information.

4.2 Two receivers

We next consider decentralized information when there are two receivers. This setting is the main contribution of our paper.

¹²The interpretation is the following: Starting with a prior $P(\theta = G) = 1/2$, if the receiver observes that sender *i* has a good signal, then she updates to $P(\theta = G \mid s_i = g) = p$. If she then also observes that sender $j \neq i$ has a good signal, she updates to $P(\theta = G \mid s_i = s_j = g) = p^2/(p^2 + (1-p)^2)$, which is precisely $\overline{\beta}(p)$.

¹³In fact, GF show that this impossibility extends beyond voting.

We begin with some definitions. For a given strategy profile σ , let $k_L(\sigma)$ be the number such that, if $k_L(\sigma)$ senders vote y then \mathbf{R}_L prefers outcome o = n, but if $k_L(\sigma) + 1$ senders vote y then she prefers outcome o = y. Similarly, Let $k_H(\sigma)$ be the same but for \mathbf{R}_H . We will omit the dependence on σ when clear from context. Formally, under strategy profile σ , for $t \in T$,

$$P(\theta = G|v = k_t) < \beta_t \le P(\theta = G|v = k_t + 1).$$

With some abuse of notation, we will also denote by k_t the event that $(v_{-i} = k_t)$.

Let $\operatorname{piv}_i(\sigma)$ be the event that sender *i* is pivotal, namely that his vote will change the chosen outcome of some receiver, when senders play strategy profile σ . Formally, $\operatorname{piv}_i(\sigma) \stackrel{\text{def}}{=} (v_{-i} = k_L(\sigma)) \cup (v_{-i} = k_H(\sigma))$. Again, we will omit the dependence on σ when clear from context.

Finally, recall that the senders' utility is such that $u^t(G, y) = c_t$, $u^t(B, y) = -c_t$, and $u^t(\theta, n) = 0$. In what follows it will be useful to denote by $h \stackrel{\text{def}}{=} c_H/(c_H + c_L)$ the weight the senders put on the decision of R_H relative to that of R_L . We will also refer to $\ell = 1 - h$.

Fully-informative equilibrium Let τ be the fully-informative strategy profile. We are interested in the question of when τ is an equilibrium—that is, when does a fully-informative equilibrium (FIE) exist. In order for τ to be an equilibrium it must be the case that each sender prefers to vote informatively. Since senders only affect the outcome when they are pivotal, this is equivalent to each sender preferring to vote informatively, conditional on being pivotal. Let r be the decision rules of the receivers, with thresholds k_L and k_H , under τ . On signal $s_i = g$, then, sender i should prefer to vote y, which requires

$$E[u(\theta, r(\tau(s)))|s_i = g, \operatorname{piv}_i] \ge E[u(\theta, r(\sigma_i, \tau_{-i}(s)))|s_i = g, \operatorname{piv}_i],$$

where σ_i is the deviation of sender *i* to voting *n* on signal *g*. This is equivalent to

$$P(\theta = G \cap k_L | s_i = g)c_L + P(\theta = G \cap k_H | s_i = g)c_H$$

$$\geq P(\theta = B \cap k_L | s_i = g)c_L + P(\theta = B \cap k_H | s_i = g)c_H,$$

which is equivalent to

$$\frac{\ell P(\theta = G \cap k_L | s_i = g) + hP(\theta = G \cap k_H | s_i = g)}{\ell P(k_L | s_i = g) + hP(k_H | s_i = g)} \ge \frac{1}{2}$$

An analogous inequality must hold for $s_i = b$, and it is straightforward to see that both hold if and only if

$$\frac{\ell P(\theta = G \cap k_L) + hP(\theta = G \cap k_H)}{\ell P(k_L) + hP(k_H)} \in [1 - p, p].$$

$$\tag{1}$$

The LHS of (1) can intuitively be understood as $P(\theta = G | \text{piv}_i)$, except that the elements referring to k_L and k_H are weighted by h and ℓ , respectively.

Equation (1) can be further simplified under our assumption that $\beta_L = 1 - \beta_H$, as in this case $k_L = N - 1 - k_H$, which implies that $P(v_{-i} = k_L) = P(v_{-i} = k_H)$. Thus, in this case there exists a FIE if and only if $\ell P(\theta = G|k_L) + hP(\theta = G|k_H) \in [1 - p, p]$, where $P(\theta = G|k_L)$ is close to β_L and $P(\theta = G|k_H)$ is close to β_H . Thus, there is a FIE if and only if the weighted average of the posteriors on $(\theta = G)$ at the pivotal events, weighted according to h and 1 - h, is close to the senders' threshold.

The intuition for the possibility of fully-informative equilibria builds on the impossibility of such equilibria when there is only one receiver. Consider first this latter case, in which only R_D is present, as in Example 2. Suppose all decentralized senders play the fully-informative strategy, and consider one sender's reasoning. On either signal, he conditions on being pivotal, as this is the only case in which his vote matters. If he is pivotal, this means that the posterior on ($\theta = G$) must be close to β_D . On a good signal his posterior is even higher, and so he certainly wishes to vote y, and on a bad signal the posterior is a bit below β_D . But if β_D is sufficiently higher than 1/2 then his posterior on a bad signal is still above 1/2, and so he wishes to vote y here as well. Thus, fully-informative voting is not an equilibrium.

Now consider the case in which there are two receivers. When a given sender is pivotal for R_H , his posterior is close to β_H , whereas if he is pivotal for R_L his posterior is close to β_L . The sender must then weigh the relative weights of being pivotal for each of the two receivers, namely the probability $(v_{-i} = k_H)$ weighted by h versus the probability $(v_{-i} = k_L)$ weighted by ℓ . When $\beta_H = 1 - \beta_L$ the probabilities of $(v_{-i} = k_H)$ and $(v_{-i} = k_L)$ are the same, and so only h is relevant. When h is not too far from ℓ the two pivotal events are roughly equally-weighted, and so the sender places roughly equal weight on the posterior close to β_L and the posterior close to β_H . The average is close to 1/2, and so the sender's own signal is the determining factor in assessing which state is more likely. Thus, he votes informatively.

The following example formalizes this logic:

Example 3 Suppose N = 3, $\beta_H < p^3/(p^3 + (1-p)^3)$, and $h = \ell = 1/2$. The bound on β_H implies that if all senders have the good signal (respectively, the bad signal), then R_H (respectively, R_L) would choose outcome y (respectively, n). Then under fully-informative voting

$$\frac{\ell P(\theta = G \cap k_L) + hP(\theta = G \cap k_H)}{\ell P(k_L) + hP(k_H)} = \frac{\ell \cdot P(k_L|G) + h \cdot P(k_H|G)}{\ell \cdot P(k_L|G) + h \cdot P(k_H|G) + \ell \cdot P(k_L|B) + h \cdot P(k_H|B)} = \frac{P(k_L|G) + P(k_H|G)}{(P(k_L|G) + P(k_H|B)) + (P(k_H|G) + P(k_L|B))} = \frac{1}{2},$$

since $\beta_H = 1 - \beta_L$ implies that $k_H = 2 - k_L$, and so under fully-informative voting we have that $P(k_L|G) = P(k_H|B)$ and $P(k_L|B) = P(k_H|G)$. Thus, there is a FIE. Furthermore, by the assumption on β_H , this FIE is persuasive.

Now suppose everything is as in Example 3, except that $\beta_H = p^4/(p^4 + p^4)$

 $(1-p)^4$). That is, R_H requires at least 4 good signals in order to choose y, and R_L requires at least 4 bad signals in order to choose n. If there are only 3 senders, however, there is never enough information for either receiver, and so no equilibrium will be persuasive.

Increasing the number of senders will help: If N = 5, for example, then an analysis similar to that of Example 3 will imply that there is a FIE. This FIE is persuasive, since \mathbb{R}_H will choose y if all 5 voters vote y (and will choose n otherwise), and \mathbb{R}_L will choose n if all 5 voters vote n (and will choose yotherwise). In fact, at least 5 senders are required to persuade both receivers in this example. To generalize, denote by $N(\beta_H, p)$ the size of the smallest number of senders that are able to persuade both receivers with thresholds β_H and $\beta_L = 1 - \beta_H$, namely $N(\beta_H, p) \stackrel{\text{def}}{=} \min\{N \in \mathbb{Z}_+ : \beta_H \leq p^N/(p^N + (1-p)^N)\}$ observe that, in general, $N(\beta_H, p) < N_C(\beta_H, p)$, and so fewer signals are necessary for a persuasive FIE under decentralization than any persuasive profile under centralization.

Note that $h = \ell$ is not necessary for the existence of a FIE, and instead there is an interval of h's for which they exist. Furthermore, as the following theorem states, this interval is independent of N:

Theorem 3 For any β_H and p there exists an interval $H^{\text{FIE}} = [h_1, h_2]$ with $h_1 < h_2$ such that there is a persuasive FIE for every odd $N \ge N(\beta_H, p)$ if and only if $h \in H^{\text{FIE}}$.¹⁴

5 Advisory Committees

In this section we view the decentralized senders as forming an advisory committee, and examine the potential benefit of transparency from the point of view of the receivers, in light of Theorem 3. We then analyze the robustness of this benefit, as well as the existence of a FIE, to variations in the structure of the committee's communication.

¹⁴An identical theorem holds for even N, but the interval H^{FIE} will be slightly different.

5.1 Transparency vs. Opacity

Suppose that the senders, who now comprise a committee, observe their respective signals and vote. The information observed by the receivers before choosing y or n is then one of the following:

- Under transparency, the receivers observe the entire profile of votes.
- Under opacity, only the committee observes the profile of votes, whereas the receivers observe a message $m \in M$ subsequently sent by a specific member of the committee called the committee chair.

Observe that transparency is analogous to the decentralized setting of Section 4, whereas opacity is analogous to the centralized setting of Section 3.

A first question is, when senders' strategies are an equilibrium, do the receivers prefer transparency or opacity? A second question is, when is there a benefit to *mandating* transparency? Note that if all parties prefer transparency, then there is no reason to require it – the committee will conduct itself transparently by choice. Mandated transparency will be beneficial if the receivers prefer transparency whereas the senders prefer opacity.

The senders' and receivers' preferences partly depend on the strategy profile played by the committee. For example, an uninformative (babbling) profile always exists under both transparency and opacity, rendering all parties indifferent. In the following, then, we suppose that the committee plays a profile that is Pareto optimal for the receivers, out of all equilibrium profiles. We note that under both opacity and transparency, whenever there exists a persuasive equilibrium profile that is Pareto optimal, it is unique.

In their study of transparency, GF use Theorem 2 to show that mandating transparency is harmful.

Proposition 1 If there is one receiver with $\beta_D > \overline{\beta}(p)$ then the committee and the receiver prefer opacity. The intuition is straightforward: under transparency, Theorem 2 shows that there is no persuasive communication. Under opacity, however, persuasive communication is possible when $\beta_D \leq \beta_{\text{maj}}$ (by Theorem 1), in which case committee members and receiver are strictly better off.

When there are multiple receivers, however, transparency can be beneficial:

Proposition 2 If there are two receivers and $h \in H^{\text{FIE}}$ then

- both receivers prefer transparency;
- the committee prefers transparency if $\beta_H > \beta_{maj}$, and opacity otherwise.

When $\beta_H > \beta_{\text{maj}}$ all parties prefer transparency. When $\beta_H \leq \beta_{\text{maj}}$, however, there is a benefit to *mandating* transparency: the receivers prefer it, but the committee would not voluntarily choose it, as they prefer opacity.

The intuition for Proposition 2 is also straightforward. By Theorem 3, if $h \in H^{\text{FIE}}$ then under transparency there is a persuasive FIE. This is bestpossible for the receivers, and so they always prefer it. For opacity, in contrast, Theorem 1 states that there is either no persuasive equilibrium (when $\beta_H > \beta_{\text{maj}}$) or a partially-informative persuasive equilibrium in which the senders obtain their optimal outcomes (when $\beta_H \leq \beta_{\text{maj}}$). The senders prefer the latter most and the former least, with the FIE in the middle.

5.2 Sequential Voting

We now argue that the existence of a FIE, and hence also the benefit of transparency, relies crucially on the structure of communication between the senders and the receivers. Theorem 3 shows that fully-informative voting is an equilibrium when senders vote simultaneously. But when senders vote sequentially, no information transmission is possible in equilibrium:

Theorem 4 Suppose senders vote sequentially. Then there is no persuasive equilibrium profile for any h, p, N, and $\beta_H > \overline{\beta} = p^2/(p^2 + (1-p)^2)$.

The intuition is that when senders vote sequentially, senders near the end of the sequence already know which receiver they have a chance of persuading. For example, a sender who observes more y votes than n votes knows he will not be able to persuade R_L , but may be able to persuade R_H . Thus, from the point of view of this sender he is only facing one receiver, in which case he will not vote informatively (by Theorem 2). In this manner, information transmission completely unwinds.

6 Centralized Information in a General Model

In this section we consider a more general setting than the rest of the paper, and examine the possibility of fully-informative communication with a centralized sender. Theorem 5 below will point to the limited applicability of the Farrell and Gibbons (1989) insight to full-information transmission in our model.

Specifically, we drop the assumptions that $\beta_L = 1 - \beta_H$ and that $\gamma_L = \gamma_H$. Recall that the centralized sender's utility from action $o \in \{y, n\}$ of \mathbb{R}_t is $u^t(\theta, o)$, and where the total utility of the sender is $u(\theta, o_L, o_H) = u^L(\theta, o_L) + u^H(\theta, o_H)$. As above, for each t there is a threshold γ_t such that the sender prefers action y by \mathbb{R}_t if and only if his posterior on $(\theta = G)$ is at least γ_t . However, we now allow γ_L to be distinct from γ_H , and both to be distinct from β_L and β_H . For example, the sender may prefer outcome y from \mathbb{R}_L whenever the posterior is above 0.25, and from \mathbb{R}_H whenever it is above 0.4. Also, denote by γ the threshold of the sender relative to the aggregate receiver: Namely, it is the threshold such that the sender prefers both receivers to choose y over both receivers to choose n if and only if the posterior is at least γ .

For any N and p, two thresholds $\beta, \beta' \in [0, 1]$ are (N, p)-equivalent, denoted by $\beta \approx \beta'$ when N and p are clear from context, if for every profile $s \in \{g, b\}^N$ of signals with accuracy p it holds that $P(\theta = G|s) < \beta \Leftrightarrow P(\theta = G|s) < \beta'$. This means that although the thresholds β and β' may be not be exactly equal, every realized signal profile lies on the same side of both.

In this section, assume (without loss of generality) that when there is a centralized sender the message space is equal to the set of possible signal realizations, $M = \{g, b\}^N$. We begin with a claim. Suppose there is only one receiver, with preferences captured by the threshold β_D . Furthermore, let the sender's preferences over the receiver's actions be captured by the threshold γ_D . Then:

Claim 1 When $N \ge N_0(\beta_D, p)$ the sender has an optimal, fully-informative strategy if and only if $\beta_D \approx \gamma_D$.

This is an implication of Theorem 1. Of course, even if $\beta_D \not\approx \gamma_D$, there may be some informative communication between the sender and receiver in equilibrium. The point here, however, is that not all information is disclosed.

Next, suppose there are two receivers with $\beta_L \leq \beta_H$ and arbitrary γ_L and γ_H . Then there may be an optimal, fully-informative strategy under centralized information, but only in two cases:

Theorem 5 When $N \ge N_0(\max\{\beta_H, 1 - \beta_L\}, p)$ the sender has an optimal, fully-informative strategy if and only if at least one of the following holds:

- $\gamma_L \approx \beta_L$ and $\gamma_H \approx \beta_H$;
- $\beta_L \approx \beta_H \approx \gamma$.

In the first case, the addition of a second receiver does not facilitate fullyinformative communication, as such communication would be possible with only one receiver as well, by Claim 1. The second case is essentially the insight of Farrell and Gibbons (1989) applied to our setting. That is, fully-informative communication is possible by the mechanism described by Farrell and Gibbons (1989) only in the restricted case in which $\beta_L \approx \beta_H \approx \gamma$. When neither case of Theorem 5 is satisfied then there is no fully-informative communication under centralized information. In particular, the main setting studied in most of this paper, with $\beta_L = 1 - \beta_H$ and $\gamma_L = \gamma_H = 1/2$ is such a case whenever $\beta_H \not\approx 1/2$, and so here decentralization can be strictly beneficial to the receivers.

A natural question is whether or not decentralization can be harmful—that is, are there situations in which there is fully-informative communication under centralization but *not* under decentralization? The following claim answers negatively:

Claim 2 Fix a centralized sender with N signals of accuracy p as well as one or two receivers, and suppose that the sender has an optimal, fully-informative strategy. Then under decentralization there is a FIE.

7 Partially-Informative Equilibria

Under decentralization, when h is too large to admit a FIE, there may still be a partially-informative equilibrium (PIE), in which senders play a mixed strategy. Such an equilibrium is more informative than any equilibrium under centralized information if $\beta_H \notin [1 - \beta_{maj}(p, N), \beta_{maj}(p, N)]$, as in that case there is no information transmission in any equilibrium. Does there always exist a PIE? The immediate answer is no: if h = 0 or h = 1 then there is effectively only one receiver, and so we know from Theorem 2 that there is no persuasive equilibrium. But what if $h \in (0, 1)$? We will argue that if h is sufficiently large (or small), then there is no PIE, regardless of the number of senders.

Observe first that every symmetric PIE σ must have one-sided mixing: senders mix either on $s_i = g$ or on $s_i = b$, but never on both. Notice also that each strategy profile σ implies unique pivotal thresholds $k_H(\sigma)$ and $k_L(\sigma)$. An equilibrium with mixing on signal $s_i = w$ is then a profile σ such that $P(\theta = G|\text{piv}_i(\sigma), s_i = w) = 1/2$. If mixing on $s_i = g$ then this is equivalent to $P(\theta = G|\text{piv}_i(\sigma)) = 1 - p$, and if mixing on $s_i = b$ then this is equivalent to $P(\theta = G|\text{piv}_i(\sigma)) = p$. Suppose senders mix on signal $s_i = b$, so that they vote y on with probability $\sigma_i(b) > 0$ when $s_i = b$. Then for each $k \in \{1, \ldots, N\}$, the posterior on $(\theta = G)$ given exactly k votes for y decreases as $\sigma_i(b)$ increases. This is because the more senders mix on a bad signal, the less informative a y vote becomes. Thus, as $\sigma_i(b)$ increases, $k_H(\sigma)$ and $k_L(\sigma)$ also increase. Similarly, if senders mix on signal $s_i = g$, then as $\sigma_i(g)$ increases, $k_H(\sigma)$ and $k_L(\sigma)$ decrease.

Is there always a PIE? Suppose $h \notin H^{\text{FIE}}$, and consider a pair of thresholds $m_H, m_L \in \{0, \ldots, N-1\}$ with $m_H > m_L$. For each such pair, it is possible that there is a profile σ with one-sided mixing such that $m_H = k_H(\sigma)$ and $m_L = k_L(\sigma)$. However, for each such pair m_H and m_L there is a maximum amount of mixing in equilibrium, subject to these being the thresholds—if senders were to mix more, then the posteriors on the thresholds would be too high or too low, and the thresholds would change. Now, given m_H and m_L , as well as this maximal amount of mixing, it holds that if h is too large then there is no equilibrium with these thresholds. This follows from the observation that increasing h increases the posterior on ($\theta = G$) conditional on sender i being pivotal, at some point surpassing p. Thus, for any pair of thresholds there is a maximal h for which the thresholds potentially correspond to a PIE. For a fixed number of players, if h surpasses the maximum of all these (over all pairs of thresholds), there will be no PIE.

However, if the number of players increases, then so does the set of possible maximal h's. One might then conjecture that for every $h \in (0, 1)$ there is a PIE if there are sufficiently many players. Theorem 6, however, disproves this conjecture.

Theorem 6 For every p and $\beta_H > p^2/(p^2 + (1-p)^2)$ there exists a nonempty set $H^{\text{NP}} = [0, \overline{h}_1) \cup (\overline{h}_2, 1]$ for which the following holds: if $h \in H^{\text{NP}}$ then there is no persuasive PIE for any N.

In words, if h is too high or too low then increasing the number of voters will not help. To see the intuition, observe that although changing the level of mixing may alter the pivotal thresholds $k_L(\sigma)$ and $k_H(\sigma)$, the posterior on $(\theta = G)$ at each such threshold stays roughly the same: around β_L at $k_L(\sigma)$ and around β_H at $k_H(\sigma)$. The main effect of mixing is to thus vary the probabilities of the pivotal events k_L and k_H . The main idea of the proof is to show that for any level of mixing the ratio of these latter probabilities cannot be either too large or too small, and in particular that it is bounded above and below independently of the number of senders.

8 Conclusion

In this paper, we developed a model of cheap talk communication with multiple senders and multiple receivers, and showed that fully-informative communication may be possible. The possibility applies beyond this specific setting, to ones in which there is one receiver with multiple options, and to the presence of uncertainty about the preferences or available options of the receiver. Drawing an analogy between decentralized senders and a transparent committee, our analysis also provides a rationale for mandating transparency in advisory committees.

There are several interesting questions left open by this paper. Two main questions are to characterize the information structures under which fullyinformative communication is possible, and under which the decentralized setting dominates the centralized setting.

Another interesting direction is to consider a different model of preferences for the senders in which their utilities depend not on the chosen outcome and the realized state, but rather on their individual recommendation and the state, similarly to the career concerns literature. In subsequent and ongoing work, we show that many of the insights of this paper persist under such preferences.

Appendix

A Alternative Interpretations

We now discuss two analogous interpretations of multiple receivers. The first is that instead of two receivers there is just one receiver, but with uncertainty about her bias: Namely, she can be one of two types, high or low, where the high type is realized with probability h and the low type with probability ℓ . The utilities of the high and low type of receiver are equal to the utilities of R_H and R_L , respectively. Furthermore, the common utility function of the senders is $\overline{u} : \Theta \times \mathcal{O} \mapsto \mathbb{R}$, where $\overline{u}(G,n) = \overline{u}(B,n) = 0$, $\overline{u}(G,y) = 1$, and $\overline{u}(B,y) = -1$, regardless of the realized receiver choosing the outcome. This model with bias uncertainty is analogous to the multiple receivers model.

Claim 3 The two-receivers model is identical to the bias uncertainty model with $h = c_H/(c_H + c_L)$, modulo innocuous scaling of the senders' utilities.

The proof is at the end of this section.

The second alternative interpretation, easily seen to be analogous to the first, is that instead of uncertainty about the type of receiver, there is only one receiver but exogenous uncertainty about the options available to her: with probability h she must choose between outcomes y_H and n_H , and with probability ℓ she must choose between outcomes y_L and n_L . The senders' preferences are as above, with $\overline{u}(G,n) = \overline{u}(B,n) = 0$, $\overline{u}(G,y) = 1$, and $\overline{u}(B,y) = -1$, and where $y \in \{y_L, y_H\}$ and $n \in \{n_L, n_H\}$. The receiver's preferences for y_L and n_L (respectively, y_H and n_H) are like those of R_L (respectively, R_H) for y and n.

Proof of Claim 3: Denote the actions of \mathbb{R}_H and \mathbb{R}_L by o_H and o_L . The utility of the senders is $u(\theta, o_H, o_L) = u^H(\theta, o_H) + u^L(\theta, o_L)$, where $u^t(\theta, o_t)$ is the senders' utility from the action o_t of \mathbb{R}_t . Finally, recall that for each t it holds that $u^t(\theta, y) = c_t$ if $\theta = G$, $u^t(\theta, y) = -c_t$ if $\theta = B$, and $u_t(\theta, n) = 0$.

Fix a strategy profile σ for the senders, and let r_H and r_L be the corresponding decision rules of the receivers. In the two-receiver model, for each $\theta \in \Theta$ it holds that

$$E[u(\sigma, r_H(\sigma, \sigma(\theta)), r_L(\sigma, \sigma(\theta)))] = E[u^H(\theta, r_H(\sigma, \sigma(\theta))] + E[u^L(\theta, r_L(\sigma, \sigma(\theta)))]$$

= $c_H \cdot E[\overline{u}(\theta, r_H(\sigma, \sigma(\theta))] + c_L \cdot E[\overline{u}(\theta, r_L(\sigma, \sigma(\theta)))]$
= $(c_H + c_L) \left(\frac{c_H}{c_H + c_L} \cdot E[\overline{u}(\theta, r_H(\sigma, \sigma(\theta)))] + \frac{c_L}{c_H + c_L} \cdot E[\overline{u}(\theta, r_L(\sigma, \sigma(\theta)))] \right)$
= $(c_H + c_L) E[\overline{u}(\sigma, r_t(\sigma, \sigma(\theta)))],$

where the expectation is over θ , σ , and t, where the type t = H with probability $c_H/(c_H + c_L)$ and t = L otherwise. Thus, for every strategy profile the senders' utility is identical in the two-receiver model and in the bias uncertainty model (except that the latter is scaled by $c_H + c_L$). This implies that the incentive compatibility constraints are identical, as are thus the equilibria and utility comparisons.

B Proof of Theorem 1

Proof of Theorem 1: Suppose h > 0, and so there are either two receivers, or if h = 1 then just the high receiver (a symmetric proof holds if h = 0).

Consider first the case in which $\beta_H > \beta_{\text{maj}}(p, N)$ (a symmetric case holds for $\beta_H < 1 - \beta_{\text{maj}}$ in the case of one receiver). Suppose towards a contradiction that there is some optimal persuasive σ . Without loss of generality, suppose R_H is persuaded. Let r_t be the corresponding decision rule of R_H , where $r_t :$ $M \mapsto \{y, n\}$. If r_H is such that R_H always chooses y when $\#\{i : s_i = g\} \geq \frac{N}{2}$, then $P(\theta = G | r_t = y) \leq \beta_{\text{maj}}$. But since $\beta_H > \beta_{\text{maj}}(p, N)$ this cannot be an optimal decision rule for R_H , a contradiction. Suppose then that the outcome is not always y when $\#\{i : s_i = g\} \geq \frac{N}{2}$. Persuasiveness implies that there is some message $m_y \in M$ such that $r_H(m_y) = y$. Thus, a profitable deviation for the sender is to send message m whenever $\#\{i : s_i = g\} \geq \frac{N}{2}$, contradicting optimality.

Next, consider the case in which $\beta_H < \beta_{\text{maj}}(p, N)$ (and $\beta_H > 1 - \beta_{\text{maj}}$ in the case of one receiver). One optimal persuasive strategy sends a message m_y whenever $\#\{i : s_i = g\} \ge \frac{N}{2}$, leading to outcome y for both receivers, and a message m_n otherwise, leading to outcome n for both receivers. Suppose that there is some other optimal persuasive strategy σ in which the receivers do not choose y if and only if $\#\{i : s_i = g\} \ge \frac{N}{2}$.

We first claim that since σ is persuasive, both receivers must be persuaded. Suppose not, and only one is persuaded, say \mathbb{R}_H . This means that there is some message m_y such that both receivers choose outcome y on message m_y . Optimality implies that \mathbb{R}_H chooses y if and only if $\#\{i : s_i = g\} \geq \frac{N}{2}$ (otherwise the sender will have a profitable deviation). But this implies that the posterior on $(\theta = G|r_H = n) = 1 - \beta_{\text{maj}}$, as a majority of the signals must have been bad. This further implies that there is some message sent by the sender, say m_n , such that $(\theta = G|m_n) \leq 1 - \beta_{\text{maj}}$. This message persuades \mathbb{R}_L , as claimed.

Thus, persuasiveness implies that there are two messages m_y and m_n such that both receivers choose y on message m_y and n on message m_n . Now, suppose that under σ there is some signal profile with fewer than N/2 good realizations, on which one of the receivers chooses y with positive probability. Then the sender has a profitable deviation from σ – namely, to send m_n on this signal profile. Similarly, if there is a signal profile with more than N/2 good realizations on which one of the receivers chooses n with positive probability, a profitable deviation of the sender would be to send m_y on this signal profile. Either case contradicts feasibility. This contradiction implies that if σ is a persuasive equilibrium then the receivers choose outcome y if and only if $\#\{i: s_i = g\} \geq \frac{N}{2}$.

C Proof of Theorem 3

Proof of Theorem 3: For this proof, let us assume an alternative interpretation of the proof, in which there is one receiver with threshold β_H with probability h and with threshold β_L with probability $\ell = 1 - h$, and in which the senders' utilities are \overline{u} (see Section A). For simplicity, denote by h (respectively, ℓ) also the event that the realized type of receiver is high (respectively, low).

When is sincere voting an equilibrium? It must be the case that, conditional on being pivotal, each voter weakly prefers to vote sincerely for both possible signals. Formally, it must be the case that $P(\theta = G|\text{piv}_i) \in [1 - p, p]$. For ease of notation, denote the number of senders by N + 1. For a fixed sender *i*, the pivotalness probability is calculated with respect to the remaining *N* senders. Now,

$$P(\theta = G|\text{piv}_i) = \frac{P(G \cap \text{piv}_i)}{P(\text{piv}_i)}$$
$$= \frac{P(G \cap k_L \cap \ell) + P(G \cap k_H \cap h)}{P(k_L \cap \ell) + P(k_H \cap h)}$$
$$= \frac{\ell \cdot \binom{N}{k_L} p^{k_L} (1-p)^{N-k_L} + h \cdot \binom{N}{k_H} p^{k_H} (1-p)^{N-k_H}}{Z},$$

where

$$Z = \ell \cdot \binom{N}{k_L} p^{k_L} (1-p)^{N-k_L} + h \cdot \binom{N}{k_H} p^{k_H} (1-p)^{N-k_H} + \ell \cdot \binom{N}{k_L} (1-p)^{k_L} p^{N-k_L} + h \cdot \binom{N}{k_H} (1-p)^{k_H} p^{N-k_H}$$

and where k_L and k_H are the pivotal events given the low and high types of receiver, respectively, when the senders play the fully-informative profile.

The assumption that the receiver is symmetric, namely that $\beta_H = 1 - \beta_L$,

implies that $k_H = N - k_L$, and so we get that

$$\begin{split} P(\theta &= G | \text{piv}_i) \\ &= \frac{\ell \cdot p^{k_L} (1-p)^{N-k_L} + h \cdot p^{k_H} (1-p)^{N-k_H}}{\ell \cdot p^{k_L} (1-p)^{N-k_L} + h \cdot p^{k_H} (1-p)^{N-k_H} + \ell \cdot (1-p)^{k_L} p^{N-k_L} + h \cdot (1-p)^{k_H} p^{N-k_H}} \\ &= \frac{\ell \cdot p^{k_L} (1-p)^{N-k_L} + h \cdot p^{k_H} (1-p)^{N-k_H} + \ell \cdot (1-p)^{N-k_H} p^{k_H} + h \cdot (1-p)^{N-k_L} p^{k_L}}{\ell + h) \cdot p^{k_L} (1-p)^{N-k_L} + h \cdot p^{k_H} (1-p)^{N-k_H}} \\ &= \frac{\ell \cdot p^{k_L} (1-p)^{N-k_L} + h \cdot p^{N-k_L} (1-p)^{N-k_H}}{(\ell + h) \cdot p^{k_L} (1-p)^{N-k_L} + \ell (\ell + h) \cdot p^{N-k_L} (1-p)^{k_L}} \\ &= \frac{\ell \cdot p^{k_L} (1-p)^{N-k_L} + \ell (\ell + h) \cdot p^{N-k_L} (1-p)^{k_L}}{(\ell + h) \cdot p^{k_L} (1-p)^{N-k_L} + \ell (\ell + h) \cdot p^{N-k_L} (1-p)^{k_L}} \\ &= \frac{(p^{k_L} (1-p)^{k_L}) \left(\ell \cdot (1-p)^{N-2k_L} + h \cdot p^{N-2k_L} \right)}{(p^{k_L} (1-p)^{k_L} + h \cdot p^{N-2k_L} + p^{N-2k_L})} \\ &= \frac{\ell \cdot (1-p)^{N-2k_L} + h \cdot p^{N-2k_L}}{(1-p)^{N-2k_L} + p^{N-2k_L}}. \end{split}$$

Observing that $N - 2k_L = N - k_L - (N - k_H) = k_H - k_L$ yields

$$P(\theta = G|\text{piv}_i) = \frac{\ell \cdot (1-p)^{k_H - k_L} + h \cdot p^{k_H - k_L}}{(1-p)^{k_H - k_L} + p^{k_H - k_L}}.$$

Finally, since $k_H - k_L$ depends only on β_H and β_L , and is independent of N (but note that it depends on the parity of N), it holds that whether or not $P(\theta = G|\text{piv}_i) \in [1 - p, p]$ depends only on h, β_H , and β_L , and not on N. In particular, it holds whenever

$$h \in \left[\frac{(1-p) \cdot p^{k_H - k_L} - p \cdot (1-p)^{k_H - k_L}}{p^{k_H - k_L} - (1-p)^{k_H - k_L}}, \frac{p^{k_H - k_L + 1} - (1-p)^{k_H - k_L + 1}}{p^{k_H - k_L} - (1-p)^{k_H - k_L}}\right].$$

D Proof of Theorem 4

Proof of Theorem 4: Consider a strategy profile σ of the senders, where each sender's strategy now depends on both his signal and the realized votes

of senders who precede him in the sequence. Suppose towards a contradiction that σ is a persuasive equilibrium.

Let us view the sequence of voting as a tree, where in each level of the tree a different sender votes. Consider the last level of the tree, after which the receiver makes a decision. If for every sender at this last level, the actions of both types of receiver are unchanged by this sender's vote, then we can delete the last level and consider the tree with one fewer level. So suppose that some sender's vote on the last level is pivotal for some receiver, and denote the sender by i and the sequence of votes leading up to this sender's pivotal vote as w.

Now, if sender *i* is pivotal for the high-type receiver, then $P(\theta = G|w \cap m_i = y) \geq \beta_H$ and $P(\theta = G|w \cap m_i = n) < \beta_H$. Similarly, if sender *i* is pivotal for the low-type receiver, then $P(\theta = G|w \cap m_i = y) \geq \beta_L$ and $P(\theta = G|w \cap m_i = n) < \beta_L$. Because $\beta_H > p^2/(p^2 + (1-p)^2)$, sender *i* can be pivotal for at most one type of receiver. Without loss of generality, suppose he is pivotal for the high type. But then $P(\theta = G|w \cap m_i = y) \geq \beta_H$ implies that $P(\theta = G|w) > p$. Thus, even if sender *i* gets the low signal he will vote *y*. This implies that for σ to be an equilibrium, sender *i* always votes *y* on history *w*. Thus, we can delete sender *i*'s action at history *w* from the tree.

Since we can repeat the argument above for any pivotal sender at the last level, feasibility implies that we can remove the entire last level of the tree. Iterating this argument leads to an empty tree, and so feasibility implies that σ is not informative and so also not persuasive, a contradiction.

E Proofs from Section 6

Proof of Claim 1: It is clear that there is a persuasive FIE when $\beta_D \approx \gamma_D$. Suppose then that $\beta_D \not\approx \gamma_D$, and without loss suppose $\gamma_D < \beta_D$. Then there is a signal profile *s* for which the sender prefers outcome *y* while the receiver prefers outcome *n*. Thus, there cannot be a FIE: on realization *s*, the sender strictly benefits from deviating and reporting the profile $s' = (b \dots b)$, where $r_D(s') = n$.

Proof of Theorem 5: It is clear that under the first bullet there is a persuasive FIE. Now suppose $\beta_L \approx \beta_H \approx \gamma$, and that the sender plays the fully-informative strategy. Since $\beta_L \approx \beta_H$ both receivers always choose the same action. Any deviation by the sender will thus either leave the outcomes unchanged, or will lead outcomes (n, n) to (y, y) or (y, y) to (n, n). However, since $\gamma \approx \beta_H$ none of these deviations will be strictly beneficial to the sender. Thus, there is a FIE.

For the "only if" direction, suppose that neither of the bullets in the theorem hold, and that the sender plays the fully-informative strategy. If $\beta_L \approx \beta_H$ then $\gamma \not\approx \beta_H$. Thus, there is some signal profile *s* such that the receivers prefer outcomes (y, y) whereas the sender prefers outcome (n, n), or vice versa. A profitable deviation for the sender is thus to send message $s' = (b \dots b)$ on realization *s* (or $s' = (g \dots g)$ in the vice versa case), leading to outcome (n, n)(or (y, y) in the vice versa case). Thus, there is no FIE.

If $\beta_L \not\approx \beta_H$ then there is some profile \overline{s} such that given this realization, R_L prefers outcome y whereas R_H prefers outcome n. Furthermore, either $\gamma_L \not\approx \beta_L$ or $\gamma_H \not\approx \beta_H$. Suppose $\gamma_L \not\approx \beta_L$ (the other case is analogous). If $\gamma_L < \beta_L$ then there is a signal profile s such that the sender prefers outcome y from R_L , but on which R_L prefers outcome n (which implies that R_H also prefers outcome n). There are now two cases: If on realized profile s the sender prefers outcome y also from R_H , then she can deviate to the message $(g \dots g)$. This is a strict improvement, since it leads to her most preferred outcomes, (y, y). If on realized profile s the sender prefers outcome n from R_H , then she can deviate to message \overline{s} , leading to her most preferred outcomes.

If, on the other hand, $\gamma_L > \beta_L$, then there is a signal profile *s* such that the sender prefers outcome *n* from \mathbf{R}_L , but on which \mathbf{R}_L prefer outcome *y*. Furthermore, we can assume that \mathbf{R}_H prefers outcome *n* on realization *s*. Again there are two cases: If on realized profile *s* the sender prefers outcome *n* also from \mathbb{R}_H , then she can deviate to the message $(b \dots b)$. This is a strict improvement, since it leads to her most preferred outcomes, (n, n). If on realized profile *s* the sender prefers outcome *y* from \mathbb{R}_H , then she can deviate to message $(g \dots g)$, which, while not leading to her most preferred outcome, is still a strict improvement: it leads from outcomes (y, n) to outcomes (y, y), which she prefers.

Proof of Claim 2: If there is a FIE under centralization, then the fullyinformative strategy τ is optimal for the receiver given the induced decision rules $r(\tau)$. By the main observation of McLennan (1998), a strategy profile that is optimal amongst all strategy profiles is an equilibrium in a common value game. Thus, τ is an equilibrium profile in the decentralized setting, given $r(\tau)$. Thus, it is an equilibrium, and constitutes a FIE.

F Proof of Theorem 6

To simplify notation, in this section we let N + 1 be the number of voters. Furthermore, in this section we will use the bias uncertainty interpretation of the model: there are two types of receiver R, namely R_H and R_L , where the former is realized with probability h and the latter with probability $\ell = 1 - h$. Furthermore, with some abuse of notation also denote by h the event ($R = R_H$) and by ℓ the event ($R = R_L$). This has the benefit of simplifying notation, as then we can denote by $piv_i(\sigma) \stackrel{\text{def}}{=} (\ell \cap v_{-i} = k_L(\sigma)) \cup (h \cap v_{-i} = k_H(\sigma))$.

F.1 When is there no mixed equilibrium?

Fix some strategy profile σ . Let $k_L = k_L(\sigma)$ and $k_H = k_H(\sigma)$, when there are N + 1 senders. Fix an arbitrary voter i, and again denote by k_L (respectively, k_H) the event that, out of the remaining N voters, k_L (respectively, k_H) vote y. Then there is some real $c \geq 0$ for which $P(k_L) = c \cdot P(k_H)$. Thus, we can write

$$P(\theta = G|\text{piv}_i) = \frac{P(G \cap k_L \cap \ell) + P(G \cap k_H \cap h)}{P(k_L \cap \ell) + P(k_H \cap h)}$$
$$= \frac{\ell \cdot P(G|k_L) \cdot P(k_L) + h \cdot P(G|k_H) \cdot P(k_H)}{\ell \cdot P(k_L) + h \cdot P(k_H)}$$
$$= \frac{c\ell P(G|k_L) + hP(G|k_H)}{c\ell + h}.$$

Now, if under σ the voters mix on signal $s_i = g$, then σ is not an equilibrium when

$$P(\theta = G|\text{piv}_i) = \frac{c\ell P(G|k_L) + hP(G|k_H)}{c\ell + h} > 1 - p$$

$$\Leftrightarrow c\ell \left(P(G|k_L) - 1 + p \right) > h \left(1 - p - P(G|k_H)\right)$$

$$\Leftrightarrow c < \frac{h \left(P(G|k_H) - (1 - p) \right)}{\ell \left(1 - p - P(G|k_L)\right)}.$$

If, on the other hand, voters mix on signal $s_i = b$, then σ is not an equilibrium when

$$P(\theta = G|\text{piv}_i) = \frac{c\ell P(G|k_L) + hP(G|k_H)}{c\ell + h} > p$$

$$\Leftrightarrow c < \frac{h\left(P(G|k_H) - p\right)}{\ell\left(p - P(G|k_L)\right)}.$$

Thus, if $c = P(k_L)/P(k_H)$ is bounded above by a constant independent of N then Theorem 6 will follow by choosing a sufficiently large h (and, by symmetry, a sufficiently small h).

We will proceed by consider two potential equilibrium profiles σ : the first are ones in which voters mix on signal $s_i = g$, and the second are ones in which the mix on signal $s_i = b$. Note that mixing on both signals cannot be an equilibrium.

We begin with a claim.

Claim 4 Let \overline{k}_H and \overline{k}_L be the thresholds for the high and low type of receiver,

respectively, when senders vote fully-informatively: $\overline{k}_H = k_H(\tau)$ and $\overline{k}_L = k_L(\tau)$. Then for any strategy profile σ , the respective thresholds $k_H = k_H(\sigma)$ and $k_L = k_L(\sigma)$ satisfy $k_H - k_L \leq 2(\overline{k}_H - \overline{k}_L) + 2$.

Proof: We first make two preliminary claims. First, since k_L and \overline{k}_L are pivotal for the low-type receiver under σ and τ , respectively, it must be the case that $P(G|v = k_L + 1, \sigma) \ge \beta_L$ and $P(G|v = \overline{k}_L, \tau) < \beta_L$, and so $P(G|v = k_L + 1, \sigma) > P(G|v = \overline{k}_L, \tau)$. Similarly, $P(G|v = \overline{k}_H + 1, \tau) \ge \beta_H > P(G|v = k_H, \sigma)$.

Second, we argue that the informational value of two y votes under σ is higher than the value of one y vote under τ . More formally, fix some $\beta \in (0, 1)$. Consider two senders, i and j, playing according σ , fix some profile of votes $v_{-(i,j)}$ of the other voters, and suppose $P(\theta = G|v_i = v_j = n, v_{-(i,j)}, \sigma) \geq \beta$. Consider also one sender, k, playing according to the fully-informative strategy τ_i , fix some profile of votes v'_{-k} of the other voters, and suppose $P(\theta = G|v_k = n, v_{-k}, \tau) = \beta$. Then we claim that $P(\theta = G|v_i = v_j = y, v_{-(i,j)}, \sigma) \geq P(\theta = G|v_k = y, v_{-k}, \tau)$.

To see this, consider first the case in which senders mix on signal $s_i = g$ under σ . This means that when $v_i = v_j = y$, it must be the case that $s_i = s_j = g$. What about $v_i = v_j = n$? In the limit, when senders mix with probability 1, the event $v_i = v_j = n$ yields no information. Thus, $P(\theta = G|v_i = v_j = n, v_{-(i,j)}, \sigma) \leq P(G|v_{-(i,j)}, \sigma)$. Thus,

$$P(\theta = G | v_k = n, v_{-k}, \tau) = \beta \le P(\theta = G | v_i = v_j = n, v_{-(i,j)}, \sigma) \le P(G | v_{-(i,j)}, \sigma)$$

Note that $P(G|v_k = n, v_{-k}, \tau) = P(G|s_k = n, v_{-k}, \tau)$. Adding two good signals is equivalent to changing $s_k = b$ to $s_k = g$, yielding

$$P(G|v_k = y, v_{-k}, \tau) = P(G|s_k = g, v_{-k}, \tau)$$

$$\leq P(G|s_i = s_j = g, v_{-(i,j)}, \sigma) = P(G|v_i = v_j = y, v_{-(i,j)}, \sigma),$$

as claimed. Similarly, if σ is such that senders mix on signal $s_i = b$, then when

 $v_i = v_j = n$, it must be the case that $s_i = s_j = b$. Additionally, in the worst case of mixing with probability 1, the event $v_i = v_j = y$ yields no information, and so $P(G|v_i = v_j = y, v_{-(i,j)}, \sigma) \ge P(G|v_{-(i,j)}, \sigma)$. Thus,

$$P(G|v_k = n, v_{-k}, \tau) = \beta \le P(G|v_i = v_j = n, v_{-(i,j)}, \sigma) = P(G|s_i = s_j = b, v_{-(i,j)}, \sigma).$$

Again, note that $P(G|v_k = n, v_{-k}, \tau) = P(G|s_k = b, v_{-k}, \tau)$. Adding two good signals is equivalent to changing $s_k = b$ to $s_k = g$, and canceling the signals $s_i = s_j = b$, yielding

$$P(G|v_k = y, v_{-k}, \tau) = P(G|s_k = g, v_{-k}, \tau)$$

$$\leq P(G|v_{-(i,j)}, \sigma) \leq P(G|v_i = v_j = y, v_{-(i,j)}, \sigma),$$

as claimed.

Given these two preliminary claims, we can now argue that $k_H - k_L \leq 2(\overline{k}_H - \overline{k}_L) + 2$. Suppose towards a contradiction that $k_H - k_L > 2(\overline{k}_H - \overline{k}_L) + 2$. We begin with a profile $v = k_L$, and ask how many n votes must be changed to y votes in order to yield a profile with $v' = k_H$. Changing one n vote to a y vote leads to the profile v^1 , and using our first observation from above we know that $P(G|v^1, \sigma) > P(G|\overline{k}_L, \tau)$. Now, change 2 more n votes to y votes in v^1 leading to v^3 , and note that the second observation above implies that $P(G|v^3, \sigma) \geq P(G|\overline{k}_L + 1, \tau)$. Iteratively keep changing 2 more n vote to y votes, each time changing v^m to v^{m+2} , and do this $\overline{k}_H - \overline{k}_L$ more times. This leads to a profile $v^{2(\overline{k}_H - \overline{k}_L)+2}$, with the property that $P(G|v^{2(\overline{k}_H - \overline{k}_L)+2}, \sigma) \geq P(G|\overline{k}_H + 1, \tau)$. Recall the first observation above, that $P(G|\overline{k}_H, \sigma)$, it must be the case that $v^{2(\overline{k}_H - \overline{k}_L)+2} \geq k_H$ and so $k_L + 2(\overline{k}_H - \overline{k}_L) + 2 \geq k_H$. This contradicts the assumption that $k_H - k_L > 2(\overline{k}_H - \overline{k}_L) + 2$, completing the proof.

F.2 Bound on $P(k_L)/P(k_H)$ when mixing on $s_i = g$

Let $q = \sigma(g)$, the probability of voting $m_i = y$ on signal $s_i = g$. For each $t \in \{h, \ell\}$ define $P(h \mid C) P(C)$

$$\hat{\beta}_t \stackrel{\text{def}}{=} P(G|k_t) = \frac{P(k_t|G)P(G)}{P(k_t)},$$

the posterior on $(\theta = G)$ when k_t out of N voters vote y. Recall that if k_t voters out of N + 1 vote y then this is insufficient to persuade the receiver of type t, whereas if $k_t + 1$ out of N + 1 vote y then this is sufficient. From this, it follows that

$$\frac{1-p}{p} \cdot \beta_t \le \hat{\beta}_t < \frac{p}{1-p} \cdot \beta_t.$$

Now,

$$c = \frac{P(k_L)}{P(k_H)} = \frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \frac{P(k_L|G)}{P(k_H|G)}$$

= $\frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \frac{\binom{N}{k_L}(pq)^{k_L}(1-pq)^{N-k_L}}{\binom{N}{k_H}(pq)^{k_H}(1-pq)^{N-k_H}}$
= $\frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \frac{k_H!(N-k_H)!}{k_L!(N-k_L)!} \cdot \frac{(1-pq)^{k_H-k_L}}{(pq)^{k_H-k_L}}$
< $\frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \frac{k_H!(N-k_H)!}{k_L!(N-k_L)!} \cdot \left(\frac{1}{pq}-1\right)^{k_H-k_L}$. (2)

Observe that

$$\frac{k_H!(N-k_H)!}{k_L!(N-k_L)!} < \left(\frac{k_H}{N-k_H}\right)^{k_H-k_L}.$$
(3)

We will now bound $\left(\frac{1}{pq}-1\right)^{k_H-k_L}$ from above. We know the posterior on $(\theta = G)$ given a profile of k_L votes for y out of a total of N+1 votes must be

at most β_L . Furthermore, $P(G|k_L) = \frac{P(k_L|G)}{P(k_L|G) + P(k_L|B)}$. Thus,

$$\begin{aligned} \frac{1}{P(G|k_L)} &= 1 + \frac{P(k_L|B)}{P(k_L|G)} = 1 + \frac{((1-p)q)^{k_L}(1-(1-p)q)^{N-k_L+1}}{(pq)^{k_L}(1-pq)^{N-k_L+1}} \ge \frac{1}{\beta_L} \\ &\Leftrightarrow \left(\frac{1-p}{p}\right)^{k_L} \left(\frac{1-q+pq}{1-pq}\right)^{N-k_L+1} \ge \frac{1}{\beta_L} - 1 \\ &\Leftrightarrow \frac{1-q+pq}{1-pq} \ge \left(\left(\frac{1}{\beta_L}-1\right)\left(\frac{p}{1-p}\right)^{k_L}\right)^{\frac{1}{N-k_L+1}} \\ &\Leftrightarrow q \ge \frac{R-1}{pR+p-1}, \end{aligned}$$

where

$$R \stackrel{\text{def}}{=} \left(\left(\frac{1}{\beta_L} - 1 \right) \left(\frac{p}{1-p} \right)^{k_L} \right)^{\frac{1}{N-k_L+1}}$$

The above inequalities hold if and only if

$$\frac{1}{q} \le p + \frac{2p - 1}{R - 1} = \frac{pR - (1 - p)}{R - 1}$$

This holds if and only if

$$\frac{1}{pq}-1\leq \frac{2p-1}{p(R-1)}.$$

Let us now bound R-1 from below, which will provide the desired upper bound on 1/(pq) - 1.

Claim 5 There is a positive number D, independent of N, for which $R-1 > D \cdot \frac{k_L}{N-k_L+1}$.

Proof: Suppose not. Then for any positive D

$$\left(\left(\frac{1}{\beta_L} - 1\right) \left(\frac{p}{1-p}\right)^{k_L} \right)^{\frac{1}{N-k_L+1}} - 1 \le D \cdot \frac{k_L}{N-k_L}$$

$$\Rightarrow \left(\left(\frac{1}{\beta_L} - 1\right) \left(\frac{p}{1-p}\right)^{k_L} \right)^{\frac{1}{N-k_L+1}} \le 1 + D \cdot \frac{k_L}{N-k_L+1}$$

$$\Rightarrow \left(\frac{1}{\beta_L} - 1\right) \left(\frac{p}{1-p}\right)^{k_L} \le \left(1 + D \cdot \frac{k_L}{N-k_L+1}\right)^{N-k_L+1} \le e^{Dk_L}$$

$$\Rightarrow \left(\frac{1}{\beta_L} - 1\right)^{\frac{1}{k_L}} \cdot \frac{p}{1-p} \le e^D.$$

Note that

$$\frac{1}{\beta_L} - 1 > \frac{1}{1 - \overline{\beta}} - 1 = \frac{p^2}{(1 - p)^2} > 1,$$

and so the LHS above, $\left(\frac{1}{\beta_L} - 1\right)^{\frac{1}{k_L}} \cdot \frac{p}{1-p}$, is strictly greater than p/(1-p) > 1. In contrast, the RHS, e^D , approaches 1 from above as D decreases. This is thus a contradiction for small enough D > 0.

Plugging in the conclusion of Claim 5 we get that

$$\frac{1}{pq} - 1 \le \frac{(2p-1)(N-k_L+1)}{pDk_L},$$

and so

$$\left(\frac{1}{pq}-1\right)^{k_H-k_L} < \left(\frac{2p+D-1}{pD}\right)^{k_H-k_L} \left(\frac{N-k_L+1}{k_L}\right)^{k_H-k_L}$$

Combining this with (3) into (2) yields the bound

$$\frac{P(k_L)}{P(k_H)} < \frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \left(\frac{k_H}{N - k_H}\right)^{k_H - k_L} \cdot \left(\frac{2p + D - 1}{pD}\right)^{k_H - k_L} \cdot \left(\frac{N - k_L + 1}{k_L}\right)^{k_H - k_L} \\ = \frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \left(\frac{2p + D - 1}{pD}\right)^{k_H - k_L} \cdot \left(\frac{k_H(N - k_L + 1)}{k_L(N - k_H)}\right)^{k_H - k_L}.$$

Observing that

$$\left(\frac{k_H}{k_L}\right)^{k_H - k_L} = \left(1 + \frac{k_H - k_L}{k_L}\right)^{k_H - k_L} < e^{(k_H - k_L)^2}$$

and that

$$\left(\frac{N-k_L+1}{N-k_H}\right)^{k_H-k_L} = \left(1+\frac{k_H-k_L+1}{N-k_H}\right)^{k_H-k_L} < e^{(k_H-k_L+1)^2}$$

yields

$$c = \frac{P(k_L)}{P(k_H)} < \frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \left(\frac{2p + D - 1}{pD}\right)^{k_H - k_L} \cdot e^{2(k_H - k_L + 1)^2}.$$

To see that this bound is independent of N, notice first that $\hat{\beta}_H/\hat{\beta}_L$ is bounded above independently of N. Furthermore, the difference $k_H - k_L$ is bounded above independently of N (for fixed β_H and β_L) by Claim 4 and the observation that $\overline{k}_H - \overline{k}_L$ is independent of N.

F.3 Bound on $P(k_L)/P(k_H)$ when mixing on $s_i = b$

Let $q = \sigma_i(b)$ be the probability of voting $m_i = n$ on signal $s_i = b$.

Recall that

$$\hat{\beta}_t \stackrel{\text{def}}{=} P(G|k_t) = \frac{P(k_t|G)P(G)}{P(k_t)}.$$

Thus,

$$c = \frac{P(k_L)}{P(k_H)} = \frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \frac{P(k_L|G)}{P(k_H|G)}$$

= $\frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \frac{\binom{N}{k_L}(1 - (1 - p)q)^{k_L}((1 - p)q)^{N - k_L}}{\binom{N}{k_H}(1 - (1 - p)q)^{k_H}((1 - p)q)^{N - k_H}}$
= $\frac{\hat{\beta}_H}{\hat{\beta}_L} \cdot \frac{k_H!(N - k_H)!}{k_L!(N - k_L)!} \cdot \left(\frac{(1 - p)q}{1 - (1 - p)q}\right)^{k_H - k_L}.$ (4)

Now,

$$\frac{k_H!(N-k_H)!}{k_L!(N-k_L)!} < \left(\frac{k_H}{N-k_H}\right)^{k_H-k_L}.$$

If $N - k_H > k_H/C$ for some positive constant C then the above inequality is at most $C^{k_H-k_L}$. Furthermore, as $\frac{(1-p)q}{1-(1-p)q} \leq \frac{1-p}{p}$, this bounds $\frac{P(k_L)}{P(k_H)}$ from above by some constant (that depends on C). We will choose C below, in the proof of Claim 6.

When $N-k_H \leq k_H/C$ we need a tighter bound. To bound $\left(\frac{(1-p)q}{1-(1-p)q}\right)^{k_H-k_L} \leq \left(\frac{(1-p)q}{p}\right)^{k_H-k_L}$ from above in that case, let us bound q from above. We know the posterior on $(\theta = G)$ given a profile of $k_H + 1$ votes for y out of a total of N+1 votes must be at least β_H , and that $P(G|k_H) = \frac{P(k_H|G)}{P(k_H|G)+P(k_H|B)}$. Thus,

$$\begin{aligned} \frac{1}{P(G|k_H)} &= 1 + \frac{P(k_H|B)}{P(k_H|G)} = 1 + \frac{(1-pq)^{k_H+1}(pq)^{N-k_H}}{(1-(1-p)q)^{k_H+1}((1-p)q)^{N-k_H}} \leq \frac{1}{\beta_H} \\ &\Leftrightarrow \left(\frac{p}{1-p}\right)^{N-k_H} \left(\frac{1-pq}{1-(1-p)q}\right)^{k_H+1} \leq \frac{1}{\beta_H} - 1 \\ &\Leftrightarrow \frac{1-pq}{1-(1-p)q} \leq \left(\left(\frac{1}{\beta_H} - 1\right) \left(\frac{1-p}{p}\right)^{N-k_H}\right)^{\frac{1}{k_H+1}} \\ &\Leftrightarrow q \leq \frac{R-1}{(1-p)R-p}, \end{aligned}$$

where

$$R \stackrel{\text{def}}{=} \left(\left(\frac{1}{\beta_H} - 1 \right) \left(\frac{1-p}{p} \right)^{N-k_H} \right)^{\frac{1}{k_H+1}}$$

Now, $\beta_H > p^2/(p^2 + (1-p)^2)$, and so $1/\beta_H - 1 < (1-p)^2/p^2$. Plugging this into the definition of R we get that

$$R < \left(\frac{1-p}{p}\right)^{\frac{N-k_H+2}{k_H+1}} < 1.$$

Thus, we have

$$q \le \frac{R-1}{(1-p)R-p} = \frac{1-R}{p-(1-p)R} \le \frac{1-R}{2p-1}.$$

We will now bound R from below, thus bounding 1-R from above, leading to an upper bound on q.

Claim 6 There is a number $D < \frac{k_H+1}{N-k_H}$, independent of N, for which $R > 1 - D \cdot \frac{N-k_H}{k_H+1}$.

Proof: Suppose not. Then for any $D < \frac{k_H + 1}{N - k_H}$

$$\left(\left(\frac{1}{\beta_H}-1\right)\left(\frac{1-p}{p}\right)^{N-k_H}\right)^{\frac{1}{k_H+1}} \le 1-D\cdot\frac{N-k_H}{k_H+1}$$
$$\Rightarrow \left(\frac{1}{\beta_H}-1\right)^{\frac{1}{N-k_H}}\cdot\frac{1-p}{p} \le \left(1-D\cdot\frac{N-k_H}{k_H+1}\right)^{\frac{k_H+1}{N-k_H}} \le e^{-D}$$

Recall that $\beta_H > \overline{\beta}$, and so

$$\frac{1}{\beta_H} - 1 < \frac{1}{\overline{\beta}} - 1 = \frac{(1-p)^2}{p^2} < 1.$$

Thus,

$$\left(\frac{1}{\beta_H} - 1\right)^{\frac{1}{N-k_H}} > \frac{1}{\beta_H} - 1,$$

and so it must be that

$$\left(\frac{1}{\beta_H} - 1\right) \cdot \frac{1 - p}{p} \le e^{-D}.$$

The is a contradiction for large enough D, since the LHS is a positive constant, whereas the RHS approaches 0 from above as D increases.

The remaining detail is to confirm that one can indeed make D large enough, while still maintaining the inequality $D < \frac{k_H+1}{N-k_H}$. Recall that $N-k_H \leq k_H/C$, and note that $k_H/C < (k_H+1)/C$. Thus, $\frac{k_H+1}{N-k_H} > C$, so $D < \frac{k_H+1}{N-k_H}$ whenever $D \leq C$. So as long as C is chosen to be large enough, we can choose D = C and simultaneously satisfy

$$\left(\frac{1}{\beta_H} - 1\right) \cdot \frac{1 - p}{p} > e^{-D}.$$

Plugging in the previous claim we get that

$$q \le \frac{1-R}{2p-1} \le D \cdot \frac{N-k_H}{(2p-1)(k_H+1)}$$

and so

$$\left(\frac{(1-p)q}{1-(1-p)q}\right)^{k_H-k_L} \le \left(D \cdot \frac{(1-p)(N-k_H)}{p(2p-1)k_H}\right)^{k_H-k_L}$$

Combining this with (4) yields

$$\begin{aligned} \frac{P(k_L)}{P(k_H)} &< \frac{\hat{\beta_H}}{\hat{\beta_H}} \cdot \left(\frac{k_H}{N - k_H}\right)^{k_H - k_L} \cdot \left(\frac{D(1-p)}{p(2p-1)}\right)^{k_H - k_L} \cdot \left(\frac{N - k_H}{k_H}\right)^{k_H - k_L} \\ &= \left(\frac{D(1-p)}{p(2p-1)}\right)^{k_H - k_L}, \end{aligned}$$

which is independent of N since the difference $k_H - k_L$ is bounded above independently of N, for fixed β_H and β_L (by Claim 4 and the observation that $\overline{k}_H - \overline{k}_L$ is independent of N).

G Asymmetric receivers

The situation is a bit more complicated when the receivers are not symmetric. If there is no FIE for some N, then $P(\theta = G|\text{piv}_i) \notin [1 - p, p]$. Whether or not increasing the number of senders leads to the existence of a FIE depends on whether $P(\theta = G|\text{piv}_i)$ is greater than p or less than 1 - p.

Proposition 3 Suppose $\beta_H \ge 1 - \beta_L$ and that $P(\theta = G|\text{piv}_i) \notin [1 - p, p]$. Then there is a number $\overline{p} > p$ for which the following holds:

- If P(θ = G|piv_i) ∉ [1 − p̄, p] then increasing the number of senders will not lead to the existence of a FIE.
- If $P(\theta = G | \text{piv}_i) \in (1 \overline{p}, 1 p)$ then sufficiently increasing the number of senders will lead to the existence of a FIE.

A symmetric proposition holds for the case in which $\beta_H \leq 1 - \beta_L$.

Proof of Proposition 3: The assumption $\beta_H \ge 1 - \beta_L$ implies that $k_H \ge N - k_L$. We have

$$\begin{split} P(\theta = G|\text{piv}_{i}) &= \frac{P(G \cap \text{piv}_{i})}{P(\text{piv}_{i})} \\ &= \frac{P(G \cap k_{L} \cap \ell) + P(G \cap k_{H} \cap h)}{P(k_{L} \cap \ell) + P(k_{H} \cap h)} \\ &= \frac{\ell \cdot \binom{N}{k_{L}} p^{k_{L}} (1-p)^{N-k_{L}} + h \cdot \binom{N}{k_{H}} p^{k_{H}} (1-p)^{N-k_{H}} + \ell \cdot \binom{N}{k_{L}} p^{k_{H}} (1-p)^{N-k_{H}} \\ &= \frac{\ell \cdot \binom{N}{k_{L}} p^{k_{L}} (1-p)^{N-k_{L}} + h \cdot \binom{N}{k_{H}} p^{k_{H}} (1-p)^{N-k_{H}} + \ell \cdot \binom{N}{k_{L}} (1-p)^{k_{L}} p^{N-k_{L}} + h \cdot \binom{N}{k_{H}} (1-p)^{k_{H}-k_{L}} \\ &= \frac{\ell \cdot \binom{N}{k_{L}} p^{k_{H}+k_{L}-N} (1-p)^{k_{H}-k_{L}} + h \cdot \binom{N}{k_{H}} p^{2k_{H}-N} \\ &= \frac{\ell C(N, k_{H}, k_{L}) p^{k_{H}+k_{L}-N} (1-p)^{k_{H}-k_{L}} + hp^{2k_{H}-N} + \ell C(N, k_{H}, k_{L}) (1-p)^{k_{H}+k_{L}-N} p^{k_{H}-k_{L}} + h(1-p)^{2k_{H}-N} \\ &= \frac{\ell C(N, k_{H}, k_{L}) p^{k_{H}+k_{L}-N} (1-p)^{k_{H}-k_{L}} + hp^{2k_{H}-N} + \ell C(N, k_{H}, k_{L}) (1-p)^{k_{H}-k_{L}} + h(1-p)^{2k_{H}-N} \\ &= \frac{\ell C(N, k_{H}, k_{L}) p^{k_{H}+k_{L}-N} (1-p)^{k_{H}-k_{L}} + (1-p)^{k_{H}-k_{L}} + hp^{2k_{H}-N} \\ &= \frac{\ell C(N, k_{H}, k_{L}) p^{k_{H}+k_{L}-N} (1-p)^{k_{H}-k_{L}} + (1-p)^{k_{H}-k_{L}} + hp^{2k_{H}-N} \\ &= \frac{\ell C(N, k_{H}, k_{L}) (p^{k_{H}+k_{L}-N} (1-p)^{k_{H}-k_{L}} + (1-p)^{k_{H}-k_{L}} + hp^{2k_{H}-N} + \ell C(N, k_{H}, k_{L}) (1-p)^{2k_{H}-N} + \ell C(N, k_{H}, k_{L}) + \ell C(N, k_{H}, k_{L}) (1-p)^{2k_{H}-N} + \ell C(N, k_{H}, k_{L}) (1-p)^{2k_{H}-N} + \ell C(N, k_{H}, k_{L}) + \ell C(N, k_{H}, k_{L}) (1-p)^{2k_{H}-N}$$

where $C(N, k_L, k_H) \stackrel{\text{def}}{=} {N \choose k_L} / {N \choose k_H}$. Let $k_H = N/2 + \hat{k}_H$ and $k_L = N/2 - \hat{k}_L$. The assumption that $\beta_H \ge N - \beta_L$ implies that $\hat{k}_H \ge \hat{k}_L$. With this change of variables, we get that $k_H + k_L - N = \hat{k}_H - \hat{k}_L$, that $2k_H - N = 2\hat{k}_H$, and that $k_H - k_L = \hat{k}_H + \hat{k}_L$. Thus, all three kinds of exponents above depend only on \hat{k}_H and \hat{k}_L , and in particular are independent of N. It remains to examine the dependence of $C(N, k_L, k_H)$ on N.

We will show that $C(N, k_L, k_H)$ decreases as N increases. Note that

$$C(N, k_L, k_H) = \frac{\binom{N}{k_L}}{\binom{N}{k_H}} = \frac{k_H!(N - k_H)!}{k_L!(N - k_L)!},$$

and so

$$\frac{C(N+1, k_L, k_H)}{C(N, k_L, k_H)} = \frac{N+1-k_H}{N+1-k_L} \le 1,$$

since $k_H \ge k_L$. Thus, increasing N has the same effect as decreasing ℓ relative to h. This has the effect of increasing $P(\theta = G|\text{piv}_i)$. Thus, if the posterior $P(\theta = G|\text{piv}_i) > p$, then this same inequality will also hold for larger number of voters.

If the posterior $P(\theta = G|\text{piv}_i) < 1 - p$, increasing the number of voters will lead to a slightly higher posterior, and so may render sincerity an equilibrium. However, this is not possible for all parameters. Observe first that $\lim_{N\to\infty} C(N, k_L, k_H) = 1$. Let $1 - \overline{p}$ be equal to

$$\frac{\ell p^{k_H+k_L-N}(1-p)^{k_H-k_L}+hp^{2k_H-N}}{\ell \left(p^{k_H+k_L-N}(1-p)^{k_H-k_L}+(1-p)^{k_H+k_L-N}p^{k_H-k_L}\right)+h \left(p^{2k_H-N}+(1-p)^{2k_H-N}\right)}$$

This is the final value of the posterior from above, but setting $C(N, k_L, k_H) = 1$. Note that it is independent of N. If $1 - \overline{p} \leq 1 - p$ then increasing the number of voters will not lead to a posterior that is greater than 1 - p, and for any finite N the posterior will be strictly less than 1 - p. Thus, there will be no FIE for any number of senders.

If $1 - \overline{p} > 1 - p$, however, then for sufficiently many senders the posterior will be sufficiently close to $1 - \overline{p}$, and so strictly greater to 1 - p. At this point there will be a FIE.

References

- AMBRUS, A. and TAKAHASHI, S. (2008). Multi-sender cheap talk with restricted state spaces. *Theoretical Economics*, **3** 1–27.
- AUSTEN-SMITH, D. (1993). Interested experts and policy advice: Multiple referrals under open rule. *Games and Economic Behavior*, **5** 3–43.

AUSTEN-SMITH, D. and BANKS, J. S. (1996). Information aggregation, ratio-

nality, and the condorcet jury theorem. *American Political Science Review*, **90** 34–45.

- AUSTEN-SMITH, D. and FEDDERSEN, T. J. (2006). Deliberation, preference uncertainty, and voting rules. *American political science review*, **100** 209– 217.
- BATTAGLINI, M. (2002). Multiple referrals and multidimensional cheap talk. *Econometrica*, **70** 1379–1401.
- BATTAGLINI, M. (2004). Policy advice with imperfectly informed experts. Advances in theoretical Economics, 4.
- BATTAGLINI, M. (2017). Public protests and policy making. *The Quarterly Journal of Economics*, **132** 485–549.
- FARRELL, J. and GIBBONS, R. (1989). Cheap talk with two audiences. The American Economic Review, 79 1214–1223.
- FARRELL, J. and RABIN, M. (1996). Cheap talk. Journal of Economic perspectives, 10 103–118.
- FEDDERSEN, T. and PESENDORFER, W. (1998). Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting. American Political Science Review, 92 23–35.
- FEHRLER, S. and HUGHES, N. (2018). How transparency kills information aggregation: theory and experiment. American Economic Journal: Microeconomics, 10 181–209.
- GERARDI, D., MCLEAN, R. and POSTLEWAITE, A. (2009). Aggregation of expert opinions. *Games and Economic Behavior*, **65** 339–371.
- GILLIGAN, T. W. and KREHBIEL, K. (1987). Collective decisionmaking and standing committees: An informational rationale for restrictive amendment procedures. *Journal of Law, Economics, & Organization*, **3** 287–335.

- GOLOSOV, M., SKRETA, V., TSYVINSKI, A. and WILSON, A. (2014). Dynamic strategic information transmission. *Journal of Economic Theory*, **151** 304–341.
- GOLTSMAN, M. and PAVLOV, G. (2011). How to talk to multiple audiences. Games and Economic Behavior, **72** 100–122.
- GRADWOHL, R. and FEDDERSEN, T. (2018). Persuasion and transparency. The Journal of Politics, 80 903–915.
- HAGENBACH, J., KOESSLER, F. and PEREZ-RICHET, E. (2014). Certifiable pre-play communication: Full disclosure. *Econometrica*, **82** 1093–1131.
- HANSEN, S., MCMAHON, M. and PRAT, A. (2014). Transparency and deliberation within the FOMC: a computational linguistics approach. URL www.cepr.org/pubs/dps/DP9994.php.
- LI, H., ROSEN, S. and SUEN, W. (2001). Conflicts and common interests in committees. American Economic Review, 91 1478–1497.
- MALESKY, E., SCHULER, P. and TRAN, A. (2012). The adverse effects of sunshine: a field experiment on legislative transparency in an authoritarian assembly. *American Political Science Review*, **106** 762–786.
- MARGARIA, C. and SMOLIN, A. (2018). Dynamic communication with biased senders. *Games and Economic Behavior*, **110** 330–339.
- MATHIS, J. (2008). Full revelation of information in sender–receiver games of persuasion. *Journal of Economic Theory*, **143** 571–584.
- MCLENNAN, A. (1998). Consequences of the condorcet jury theorem for beneficial information aggregation by rational agents. *American Political Science Review*, **92** 413–418.
- MORGAN, J. and STOCKEN, P. C. (2008). Information aggregation in polls. American Economic Review, **98** 864–96.

- PAETZEL, F., LORENZ, J. and TEPE, M. (2018). Transparency diminishes framing-effects in voting on redistribution: Some experimental evidence. *European Journal of Political Economy*, 55 169–184.
- PLOTT, C. R. and LLEWELLYN, M. (2015). Information transfer and aggregation in an uninformed committee: A model for the selection and use of biased expert advice. *European Journal of Political Economy*, 40 208–223.
- PRAT, A. (2005). The wrong kind of transparency. American Economic Review, 95 862–877.
- RENAULT, J., SOLAN, E. and VIEILLE, N. (2013). Dynamic sender-receiver games. *Journal of Economic Theory*, **148** 502–534.
- SHAMBAUGH, G. E. and SHEN, E. B. (2018). A clear advantage: The benefits of transparency to crisis recovery. *European Journal of Political Economy*, 55 391–416.
- SOBEL, J. (2013). Giving and receiving advice. Advances in economics and econometrics, 1 305–341.
- STALNAKER, R. (2006). Saying and meaning, cheap talk and credibility. In Game theory and pragmatics. Springer, 83–100.
- WOLINSKY, A. (2002). Eliciting information from multiple experts. *Games* and *Economic Behavior*, **41** 141–160.