Persuasion and Transparency*

Ronen Gradwohl[†] Timothy Feddersen[‡]

Abstract

An advisory committee with common values and asymmetric information provides a recommendation to a decision maker facing a binary choice. We investigate the effect of a transparency requirement—requiring committee members' actions to be observable—on the committee's ability to influence the decision maker. We show that unless the preferences of the committee and decision maker are sufficiently close, requiring transparency eliminates the committee's ability to provide *any* useful information. In contrast, if preferences are very close or if committee members are able to verifiably reveal their signals then transparency is beneficial.

Keywords: Cheap talk, transparency, persuasion, information aggregation.

*Supplementary material for this article is available in an online appendix, which has conveniently been appended to the end of the document you are reading. Gradwohl gratefully acknowledges support from NSF award #1216006. We would also like to thank the editor and anonymous referees for their insightful comments.

†Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA. E-mail: r-gradwohl@kellogg.northwestern.edu.

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

Transparency requirements are ubiquitous in legislative and regulatory settings, as mandated by the Government in the Sunshine Act.[1] The FDA, for example, extols transparency as a way to "help maximize the integrity, consistency, and utility of advisory committee voting results" (FDA, 2008). While transparency requirements are typically understood to be desirable, it is widely appreciated that such requirements come at a cost. Transparency requirements exacerbate "career concerns", in which politicians may have an incentive to pander to the public, or they may allow third parties to incentivize behavior not in the public interest. In this paper we highlight a different and more fundamental problem with transparency that does not depend upon any possible threat from outside interests distorting agent preferences: namely, transparency inhibits the ability of a committee to persuade a decision maker. When there is no transparency requirement the committee is often able to provide very accurate information to the decision maker. In the case of transparency, however, even an almost-perfectly-informed committee is unable to provide the decision maker with *any* useful information.

Substantively, the most direct application of our model is advisory committees. Such committees deliberate prior to making a non-binding recommendation to a decision maker (Shepsle and Weingast, 1987). Organizations ranging from small groups to large societies routinely set up advisory committees; examples include advisory committees for the U.S. Food and Drug Administration (FDA, 2015), the Intergovernmental Panel on Climate Change (IPCC, 2015), the Federal Advisory Council (FAC, 2015), and the Investor Advisory Committee that advises the U.S. Securities and Exchange Commission (IAC, 2015). While our results extend to non-advisory committees we believe that advisory committees are both understudied and provide the cleanest illustration of our insight.

The basic logic driving our results can be illustrated with a very simple example. Suppose there are two information sources, one providing an accurate signal and the

[1] The goal of the law is to make the government more transparent, for example by requiring that deliberations and votes be recorded and made available.

other an inaccurate signal. A decision maker (DM) will have access to one source, whereas an advisor with different preferences will have access to the other. The basic insight is that if the DM can only be persuaded by observing the information from both sources then, given the choice, she prefers to assign the more accurate information source to the advisor and the less accurate one to herself.

It will help to make this argument just a bit more formal. Suppose that a DM is deciding whether to hire a job candidate. There are two equally likely states of the world, good (G) and bad (B). The DM wants to hire the candidate if she believes the state is G with probability sufficiently higher than 1/2, and prefers not to hire otherwise. Suppose the DM observes the realization of a signal, either g or b. We are interested in the case where neither realization is sufficiently informative to persuade the DM to hire.[2]

Now suppose there is an advisor (A) who observes the realization of a different signal, also either g or b. The advisor would like the DM to hire if the probability of state G is at least 1/2, and to not hire otherwise. The preferences of A and DM are different: the DM prefers to hire only if both signals are g, whereas A prefers to hire if either both are g or the more accurate signal is g and the less accurate signal is b.

The basic question is whether the DM can rely upon A to truthfully reveal his information. The answer is quite intuitive: only when A's signal is more accurate than the DM's signal. The reason is that A must be willing to reveal that he has observed the bad signal conditional on the event that the DM has observed the good signal, i.e., conditional on the event that A's signal is pivotal. If the DM's signal is more accurate than A's, then A would prefer to tell the DM he has observed the good signal even when he has actually observed the bad one. It follows that if the DM can decide who observes the more accurate signal, she will choose A.

To see how this insight applies to advisory committees suppose that the DM does

[2]More formally, suppose the probability of the DM observing signal realization g (b) in state G (B) is $p > 1/2$. On observing a realization g (respectively, b), then, the DM infers that the state is G with probability p (respectively, $1 - p$). She wishes to hire if her posterior on the state being G is at least some β_D, and we are interested in the case where $\beta_D > p$.

not have a signal at all but can empanel an advisory committee consisting of several members, each of whom observes a private signal. If the members are allowed to meet and share their private information and only issue a recommendation to the DM then, in effect, the situation is analogous to the example above in which a single advisor has a very accurate signal while the DM has no private signal at all. Hence, the opaque committee persuasively reveals information to the DM.

But what happens if we introduce transparency? Now the game is changed so that each individual member faces a DM who has observed information gleaned from the other advisors. We show that for a range of parameter values it is necessarily the case that if the committee is to be persuasive, the accuracy of the information the DM gleans from the other committee members must be greater than the accuracy of any individual member's signal. As a consequence, no individual member wishes to reveal his information. Information transmission unwinds, and no persuasive information can be revealed.

The following example illustrates the results in the committee context. A five-member committee advises a DM about whether to hire a potential job candidate. There are two equally likely states of the world, G and B. All members of the committee and the DM agree that in state G (B) the candidate should (not) be hired. Each member of the committee (but not the DM) observes a private, non-verifiable and conditionally independent signal $s \in \{g, b\}$. The probability a member observes signal g (b) in state G (B) is .875. Each member prefers that the DM hire the candidate if the probability the state is G is greater than .5. The DM, on the other hand, prefers to hire if and only if the probability of state G is greater than .98.

Suppose that the committee meets privately and simultaneously votes on whether to recommend hiring the candidate or not. Each member votes to hire if and only if they observe signal g. The decision maker does not observe the profile of votes cast by the committee members but only the outcome of the vote – a recommendation to hire if and only if a majority voted to hire. In this case we say that the committee is using an *opaque* process. Given that the DM anticipates both the voting process and behavior of committee members she will hire the candidate if and only if the committee recommends hiring. This is because upon observing a recommendation to hire the DM believes that the probability of state G is greater than .98. This simple opaque process is incentive compatible for both committee and DM. It guarantees that the committee receives its optimal outcome given the information available and also makes the DM better off since she gets usable information from the committee.

Observe that if we fix the voting behavior of the committee, the DM would be even better off had the committee used a *transparent* process in which she observes the profile of votes cast. This, for example, would permit her to not hire the candidate if only a bare majority vote to hire. The problem with this intuition is that transparency requirements are imposed ex ante rather than ex post.

Under transparency, informative voting is no longer incentive compatible for the committee members. To see this, suppose the DM believes that committee members do vote sincerely according to their private signal. In that case the DM will hire if and only if at least 4 out 5 members vote in favor. But now suppose the committee anticipates this behavior by the DM. Consider a member who has observed the b signal but knows her vote is pivotal, i.e., 3 out of the 4 other committee members have voted in favor and therefore observed the g signal. In this case the member infers that 3 out of 5 signals are good, and so she prefers the candidate to be hired. As is well known in the literature, sincere voting is not an equilibrium in this case (see Austen-Smith and Banks, 1996; McLennan, 1998). But if the committee votes strategically, with individuals sometimes voting to hire even when their private signal is bad, it is no longer incentive compatible for the DM to hire if only 4 out of 5 members vote in favor. Instead she would hire only if all members vote to hire. If the committee members anticipate that the DM will hire only under a unanimous vote they will alter their voting behavior again. However, in any equilibrium under unanimous rule it can be shown that the probability the state is G given a unanimous vote is bounded above by .97...and this is insufficient to persuade the DM to hire. Thus, transparency eliminates even the possibility of information transmission due to a kind of unwinding process.

The first contribution of this paper is to show that the problem in the simple example holds for all possible decision processes that the committee might choose, across a range of parameters. For the purposes of this paper we restrict ourselves to demonstrating this for simultaneous voting games. In an online appendix the results are extended to any possible simultaneous or sequential decision process the committee might choose, including for example straw polls, sequential voting, or arbitrary communication. We show that even for large committees, in which an opaque committee using majority rule provides a very high confidence of a good state conditional upon a positive recommendation, the requirement of transparency leads to no information transmission at all. The bottom line is that, even without considering career concerns or distortions due to third party interventions (e.g., bribery), requiring transparency is often harmful to both the DM and the committee.

Second, we show that there are some parameter values for which this result does not hold and transparency is strictly preferred by the DM. This result depends upon the difference between the preferences of the committee members and the DM being small enough.

Our third contribution is to show that transparency may be beneficial when information is verifiable. Consider the committee example above but assume that members' signals are verifiable. That is, a member who observes the b signal may choose to reveal it or stay silent, but cannot claim to have seen the g signal. Now, suppose that the committee is operating under a transparency requirement using simultaneous revelation of private signals. Each member may choose to reveal their signal or stay silent. Suppose as above that the DM hires the candidate if and only if at least 4 positive signals are revealed. Every member knows that their signal is pivotal only when 3 out of the other 4 have revealed a g signal. In that event each member with a g signal prefers to reveal it. But what about a member who has observed signal b? If her vote is pivotal she prefers to mislead the DM and claim to have seen a g signal, but she is unable to do so. Thus, truthful revelation is incentive compatible for the committee.

At first blush it might seem that the simple example above proves that transparency is optimal for the DM. In the paper we show that strategic incentives do not entirely disappear for such committee members, as they may in fact be better off by strategically withholding g signals. Under the assumption that the committee chooses its most preferred transparent process some additional work is required to show that transparency is optimal.

We provide a much more general model in an online appendix, in which the committee may choose any decision making process (and not just simultaneous voting). There we show that restricting the committee to using voting processes only reduces the complexity of the discussion and is without loss of generality when signals are non-verifiable. That is, the unwinding that occurs in this paper with transparent voting processes continues to hold for a wide class of possible processes. The story is more nuanced when preferences are very close or when there are verifiable signals, but we do provide examples in these settings where requiring transparency is strictly beneficial to the DM.

Our paper proceeds as follows. Immediately following is a brief literature review. We then set up the model. The subsequent three sections contain our negative results, our positive results when the preferences of the DM and committee are close but not identical, and our positive results under verifiable signals. The last section concludes.

Literature review

There is a large literature examining the benefits of transparency in agency relationships in both political science and economics—see Prat (2005) and Malesky et al. (2012) for extensive reviews. A growing literature examines whether there is a fundamental connection between the requirement of transparency and democratic institutions (see Berliner (2014) and Hollyer et al. (2011)). The essential idea that underlies both the desirability of transparency as well as its connection to democracy is that a principal delegates a decision to an agent whose action is unobservable to the principal. If principal and agent preferences are not aligned then making the agent's action observable (i.e., transparent) may allow the principal to create incentives that better align preferences. It is well known that transparency also comes with costs. In particular, when agents worry about the impact of third parties observing their behavior or about career and reputation concerns more generally, such alignment of incentives may be hindered. For example, a variety of papers examine incentives for position taking (c.f. Mayhew, 1974) and vote buying, in which policymakers are rewarded by outsiders, such as special-interest groups, for voting a particular way. Such policymakers, in addition to being concerned about the outcome of a vote, are thus also concerned about how their individual votes are perceived. Similarly, papers such as Snyder and Ting (2005), Dal Bó (2007), Felgenhauer and Grüner (2008), and Seidmann (2011) analyze decision-making in committees when members are influenced by outsiders in this way. Additionally, a growing literature (including Sibert, 2003; Fingleton and Raith, 2005; Levy, 2007a,b; Stasavage, 2007; Gersbach and Hahn, 2008; Swank and Visser, 2013; Mattozzi and Nakaguma, 2017; Fehrler and Hughes, forthcoming and the empirical work of Meade and Stasavage, 2008; Hansen et al., 2014) examines the effects of committee members' career concerns on decision making.

It is also possible that reputation concerns create a benefit to transparency for the committee. Moffitt (2010) finds that bureaucracies such as the FDA may desire the work of advisory committees to be public when there is a significant chance the recommended policy fails. While this can explain an internal requirement of transparency imposed by the bureaucracy it does not explain a statutory requirement. In our setting there is nothing preventing an advisory committee from choosing to be transparent. Our interest is in understanding why transparency requirement are imposed upon such committees.

The basic structure of our model follows the literature on voting and information aggregation studied by Austen-Smith and Banks (1996) and Feddersen and Pesendorfer (1998), among others. In those papers an electorate is composed of individuals with private information and common values who face incentives to vote strategically, i.e., not vote according to their private signal, as a function of the voting rule being used. In this paper we adapt a similar information structure but there is no fixed voting rule that determines the outcome. Instead, the outcome of the committee's activities, such as voting, is reported to a decision maker, who then updates her beliefs about the correct decision and makes a final choice. In this respect paper is more closely related to the literature on cheap talk communication initiated by Crawford and Sobel (1982). Crawford and Sobel show that if preferences between a Sender and a Receiver are close enough, then information can be transmitted from the former to the latter. One can view an *opaque* advisory committee as the Sender, and the DM as the Receiver, and then our baseline model fits perfectly into the cheap talk framework. Our interest lies in the additional imposition of a transparency requirement on the committee's process, a requirement that we show destroys any meaningful communication.

Our extended model, in which we allow committee members to verifiably reveal their signals, is part of the large literature on communication of verifiable information – see the recent survey of Milgrom (2008) and the references therein.

Perhaps the most relevant paper is that of Wolinsky (2002). He does not discuss substantive issues related to transparency, but his model and some of his results are a special case of ours. In Wolinsky's paper a set of experts, each of whom observes a private binary signal, can provide a vote to a decision maker. The DM makes a choice between H and L, and prefers the former if and only if the number of high signals observed by the experts is above some threshold. The experts prefer the DM choose H if and only if the number of high signals is above a threshold higher than the

DM's. Wolinsky shows that when the DM queries each expert separately there is no equilibrium in which the DM learns information that causes her to choose H.[3] This is similar to our Lemma 1. Wolinsky considers two possible solutions: commitment by the DM to a particular decision rule, and communication amongst the experts. Wolinsky does not discuss transparency per se, but it is somewhat implicit—when the DM queries each expert separately, that is equivalent to transparent voting, and when he allows communication he implicitly allows them not to reveal their votes.

Our results differ from Wolinsky's in two important respects. First, Wolinsky considers the problem of information transmission from a committee only in an example—namely, in simultaneous voting. Various problems arise in such settings due to voting rules being suboptimal for the structure of information possessed by committee members (Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1998). However, these problems can often be solved by pre-play communication, such as a straw poll (Coughlan, 2000). Our main negative result shows that neither straw polls nor any communication amongst committee members will solve the information transmission problem caused by the transparency requirement. Thus, this problem with transparency is persistent and robust. Second, as our aim is to understand when transparency might be beneficial, we also consider other settings of parameters not analyzed by Wolinsky. In particular, we show that when the preferences of the DM and the committee are close (but not identical), transparency is strictly beneficial. We also show that when information is verifiable transparency is beneficial.[4]

[3]More recently, Battaglini (2015) considers a model of informed protestors who choose to engage in costless protests in order to inform a policy maker. He provides a negative result similar to Wolinsky's in which protestors are unable to provide information to a policy maker with different preferences unless their private signals are sufficiently informative.

[4]Strictly speaking, in Wolinsky's model signals are verifiable in one direction—specifically, a positive signal cannot be imitated by a player who obtained a negative signal. However, as Wolinsky states, this assumption is made only for convenience and has no substantive effect on the result, since in his model no player will ever *want* to imitate a positive signal.

Model and Definitions

There are two possible, equally likely states of the world, $\Theta = \{G, B\}$, and a decision maker (DM) who must decide between two possible outcomes, $\mathcal{O} = \{y, n\}$. Choosing n always yields a utility of 0, whereas choosing y yields a gain in state G and a loss in state B, where the loss incurred from a bad decision is greater than the gain from a good decision. Formally, the DM is equipped with a utility function $u_D : \Theta \times \mathcal{O} \mapsto \mathbb{R}$ that satisfies $u_D(G, y) > 0$, $u_D(B, y) < -u_D(G, y)$, and $u_D(\theta, n) = 0$ for both $\theta \in \Theta$.[5] We assume without loss of generality that $u_D(G, y) = 1$.

Given a belief $\beta = \Pr[\theta = G]$ about the probability that the state is G, the DM's expected utility on choosing outcome o is $U_D(\beta, o) \stackrel{\text{def}}{=} \beta \cdot u_D(G, o) + (1 - \beta) \cdot u_D(B, o)$. A rational DM will choose outcome y if and only if $U_D(\beta, y) \geq 0 = U_D(\beta, n)$.[6] Since $U_D(\beta, y)$ is increasing in β, there exists a threshold β_D such that the DM will choose y if and only if $\beta \geq \beta_D$. By our assumption on u_D it holds that $\beta_D > 1/2$.

Next, there is an advisory committee of N members, also called agents, each of whom receives a conditionally-independent, identically-distributed signal $s_i \in \{g, b\}$ satisfying

$$\Pr[s_i = g | \theta = G] = \Pr[s_i = b | \theta = B] = p,$$

where $p \in (1/2, 1)$ is the *accuracy* of the signal. Each of the N committee members has a common utility function $u : \Theta \times \mathcal{O} \mapsto \mathbb{R}$. The utility function is identical to the DM's, except that the loss they incur from a bad decision is the same as the gain from a good decision, formally $u(B, y) = -u(G, y)$.[7] Observe that the committee members prefer outcome y whenever their belief β_C of the state being G is above a threshold of $1/2$.

In the main body of the paper we will limit the interaction of committee members

[5] The assumption that $u_D(\theta, n) = 0$ is for simplicity only—all our results hold for any utility function in which the threshold $\beta_D > 1/2$ (see next paragraph).

[6] Assume he always chooses y if indifferent.

[7] This symmetry assumption is made solely for simplicity. All our negative results hold as long as the committee's threshold for desiring outcome y is sufficiently lower than the DM's.

to a *poll*, in which each committee member simultaneously chooses a vote $v_i \in \{y, n\}$. We stress that this limitation is for simplicity only: our main negative results hold much more generally, allowing, for example, sequential voting or arbitrary communication amongst committee members.

We are interested in comparing two settings, one with a transparency requirement and one without. The main analysis involves the former, and so most of the definitions here are for transparent polls. The case of no transparency requirement is introduced at the end of this section.

A strategy σ_i of agent i in a poll is a function from his signal to a distribution over $\{y, n\}$. Denote by $\sigma = (\sigma_1, \ldots, \sigma_N)$ a profile of strategies, by $\sigma(s)$ the profile of strategies given signal profile $s = (s_1, \ldots, s_N)$, by $V = \{y, n\}^N$ the set of vote profiles, by $V(\sigma(s))$ the distribution over vote profiles when the strategy profile $\sigma(s)$ is played, and by $v \in V$ a generic vote profile.

After the committee members vote, the DM sees the profile of votes. He then updates his prior over the state, and takes an action that depends on whether the posterior surpasses his threshold β_D or not. Formally, given a strategy profile σ, denote the rational decision rule used by the DM on realized voting profile v as $r(\sigma, v)$, where $r(\sigma, v) = y$ if and only if $\Pr[\theta = G \mid \sigma, v] \geq \beta_D$, and $r(\sigma, v) = n$ otherwise.

Observe that without any information, the DM will choose outcome n. A strategy profile σ of the committee is *persuasive* if it sometimes leads the DM to choose y. This means that there is some vote profile v that occurs with positive probability under σ, and for which $\Pr[\theta = G \mid \sigma, v] \geq \beta_D$. Note that if σ is not persuasive, the DM makes his choice based only on the prior distribution over states, and so will always choose action n. This yields him and the committee members utility 0.

When the DM updates his prior he conditions on both the vote profile v and on the strategy profile σ. But what prevents a committee member from deviating from σ, unbeknownst to the DM? In order to prevent this, we will additionally restrict the committee to strategy profiles that are *feasible*: ones that constitute a Nash equilibrium for the agents conditional on the DM acting rationally. In a standard voting game, where agents 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 DM, given σ. A profile σ is then feasible if it is in equilibrium given the decision rule that it induces. Formally,

Definition 1 (feasibility) *A strategy profile σ is* feasible *if for each agent i, signal s_i, and strategy σ_i',*

$$\mathrm{E}\left[u(\theta,\bar{r}(\sigma(s))) \mid s_i\right] \geq \mathrm{E}\left[u(\theta,\bar{r}(\sigma_i',\sigma_{-i}(s))) \mid s_i\right],$$

where $\bar{r}(\cdot) \equiv r(\sigma,\cdot)$ and the expectation is over θ, s, and σ.

As an example, suppose N is odd and consider the strategy of *sincere* voting, in which each agent i votes $v_i = y$ if and only if $s_i = g$. Sincere voting is not feasible when $\beta_D > p$: To see this, suppose each agent votes sincerely, 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 DM will be p. He 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 sincere voting is not an equilibrium (Austen-Smith and Banks, 1996; Feddersen and Pesendorfer, 1998).

Now, for each p, N, and β_D there may be multiple feasible strategy profiles. We assume that the committee chooses which profile to play, and so will consider the *committee-optimal* feasible strategy profile – one that yields the highest expected utility to the committee, out of all feasible profiles. This modeling choice is consistent with our motivating examples of federal advisory committees, in which behavior is determined by the committees themselves, while the law (namely, the Federal Advisory Committee Act) only requires them to be transparent.

No transparency requirement We will assume that when there is no transparency requirement, the committee engages in *opaque, sincere majority*: first, each committee member votes sincerely. Second, if at least $N/2$ members vote y then the committee recommends to the DM to choose outcome y; otherwise, they recommend that he choose n. Importantly, the DM observes *only* the recommendation, and not the vote profile. This assumption—that without a transparency requirement, the committee engages in opaque, sincere majority—is nearly without loss of generality for our results.

Consider a signal profile s, and note that if there are more good signals than bad the committee members prefer outcome y, since in that case the posterior on $\theta = G$ is greater than $1/2$. Similarly, if there are more bad signals than good they prefer outcome n. If N is even and the committee is split, they are indifferent. Under opaque, sincere majority, if the DM acts on the committee's recommendation (that is, if it is persuasive), the committee obtains its preferred outcome on each signal profile, and so this is committee-optimal.

To be persuasive, the DM's posterior on recommendation y must surpass his threshold β_D. We will denote the posterior on recommendation y by

$$\beta_{\mathrm{maj}}(p, N) \stackrel{\mathrm{def}}{=} \Pr\left[\theta = G \mid \#\{i : s_i = g\} \geq \frac{N}{2}\right].$$ [8]

It will be easiest to focus the case in which $\beta_{\mathrm{maj}} \geq \beta_D$ (and this is indeed always the case for large enough N – see Claim 1). When $\beta_{\mathrm{maj}} < \beta_D$ the situation is a bit more subtle, and depends on whether or not the committee can commit to some feasible process. Regardless, this will not alter our main result, Theorem 1.

Finally, our only result that is slightly affected by the assumption of sincere, opaque majority is Theorem 2 – see footnote 11 for details.

[8] We omit the dependence on p and N when clear from context.

Transparency is harmful

We begin by defining the threshold $\overline{\beta}(p) \stackrel{\text{def}}{=} p^2/(p^2+(1-p)^2)$, which will play a vital role in our analysis.[9] The interpretation of $\overline{\beta}$ is as follows: Starting with a prior $\Pr[\theta = G] = 1/2$, if the DM observes that agent i has a good signal, then he updates to $\Pr[\theta = G \mid s_i = g] = p$. If he then also observes that agent $j \neq i$ has a good signal, he updates to $\Pr[\theta = G \mid s_i = s_j = g] = p^2/(p^2+(1-p)^2)$, which is precisely $\overline{\beta}(p)$.

We can now state our main theorem.

Theorem 1 *For any N and $p > 1/2$, if $\beta_D > \overline{\beta}(p)$ then the DM (weakly) prefers* ***not*** *to require transparency.*

The theorem relies on the following lemma:

Lemma 1 *Under transparency, there does not exist any feasible and persuasive strategy profile for any N, $p > 1/2$, and $\beta_D > \overline{\beta}(p)$.*

Theorem 1 follows directly from Lemma 1. The latter implies that any feasible strategy profile cannot be persuasive, and so it yields utility 0 to the DM. Without requiring transparency—under opaque, sincere majority—the expected utility of the DM is at least 0.

The intuition for Lemma 1 is quite simple and can be understood by thinking of the problem from the perspective of the DM and a single member i of the committee. Because a single signal is not persuasive to the DM (i.e., $\beta_D > p$) the only time the agent's action can be pivotal is if some event E (i.e., a profile of votes by other committee members) has occurred that makes the agent's signal pivotal. Let $\Pr(\theta = G \mid E) = \rho$ be the probability the state is good given that the DM observes an event E. Notice that, as in the example in the introduction, ρ amounts to the accuracy of the DM's information. By the logic of the example we know that truthful revelation of the agent's signal requires that the agent's signal be as least as accurate

[9] We omit the dependence on p when clear from context.

as the DM's, namely $p \geq \rho$. But if the agent is to be persuasive it must be the case that a single positive signal plus the event E produce beliefs in the DM that exceed her threshold. That is,

$$\Pr\left[\theta = G \mid E, s_i = g\right] = \frac{\rho p}{\rho p + (1-\rho)(1-p)} \geq \beta_D.$$

This threshold is maximized when the event E is as informative as possible, i.e., $\rho = p$. So this gives us the requirement that

$$\overline{\beta}(p) = \frac{p^2}{p^2 + (1-p)^2} \geq \beta_D.$$

In words, if there is any information transmission under transparency then it must be the case that the DM is persuadable by two good signals.

Lemma 1 states that when $\beta_D > \overline{\beta}(p)$ the only feasible strategy profiles are not persuasive, with the DM always choosing the same outcome (which must thus be n).

Theorem 1 states that the DM weakly prefers not to require transparency. This preference is strict whenever the DM obtains positive expected utility when not requiring transparency—that is, when sincere, opaque majority yields positive utility to the DM. This is the case, for example, if the DM's posterior on a y recommendation makes him strictly prefer outcome y, namely when $\beta_{\text{maj}} > \beta_D$. Now, for any signal accuracy p, as the number of agents grows, the probability that the majority of agents correctly determines the state grows. This implies that β_{maj} grows, and for a sufficiently large committee surpasses the DM's threshold. When this occurs, the DM strictly prefers opaque, sincere majority to the utility 0 he obtains under transparency:

Claim 1 *For any p and $\beta_D > \overline{\beta}(p)$ there exists N_0 such that for all $N \geq N_0$, transparency is* strictly *worse for the DM than opaque, sincere majority.*

Claim 1 follows from Lemma 1 together with the following lemma, which stands in sharp contrast with Lemma 1:

Lemma 2 *For any $\beta_D > 1/2$ and p there exists N_0 such that for all $N \geq N_0$ it holds that $\beta_{\text{maj}}(p, N) > \beta_D$.*

Small β_D

Theorem 1 states that when the preferences of the DM and committee are not sufficiently close, when $\beta_D > \overline{\beta}$, transparency is harmful to the DM. In contrast, if the preferences of the DM and the committee are identical, there is no need for transparency. In this section we show that when the preferences are close enough, but not identical, transparency may be *strictly* beneficial to the DM. In fact, we show that the benefit of transparency is not only non-monotone in β_D, but also discontinuous. For large N, Claim 1 states that transparency is strictly harmful to the DM when $\beta_D > \overline{\beta}$. In contrast, Theorems 2 and 3 below state that when $\beta_D = \overline{\beta}$, transparency is strictly beneficial to the DM.

To see why transparency might be beneficial, consider a simple example with 3 committee members and with $p < \beta_D \leq \overline{\beta}(p)$. Also suppose that $\beta_D \leq \beta_{\text{maj}}(p, 3)$, and so opaque, sincere majority is persuasive. Ideally, the DM would like outcome y only when all 3 agents have good signals, but with opaque majority the outcome is y also when only 2 out of 3 agents have good signals. What is the effect of imposing a transparency requirement? Recall that sincerity is not feasible. However, unlike the case of $\beta_D > \overline{\beta}$, here there *do* exist feasible, persuasive strategy profiles. In fact, there are several candidates for such profiles: One is the asymmetric, pure profile in which two agents vote sincerely and the third takes a signal-independent action (à la McLennan, 1998). Another is a symmetric, mixed profile, in which all agents vote y on a good signal but randomize on a bad signal. The asymmetric profile is always feasible, whereas the symmetric one is feasible only if β_D is small enough. But even in the latter case, it turns out that the *committee-optimal* strategy profile under transparency is the asymmetric one. Observe that under this asymmetric profile, the

DM is strictly better off than under opaque, sincere majority: he still obtains outcome y when all 3 agents have good signals, but also avoids outcome y on two undesirable signal profiles, namely the ones in which the signal-independent agent has a good signal and the others are split. Thus, the committee-optimal feasible profile under transparency is strictly better for the DM than opaque, sincere majority.

The argument above generalizes to larger N. Although the information aggregated by opaque, sincere majority becomes increasingly accurate as N grows, there always remains a marginal benefit to the DM from imposing transparency—it allows him to avoid outcome y on some of the signal profiles in which the number of good signals is close to the number of bad ones.

We have two theorems, one for odd N and one for even N. The proofs of both follow similar lines, by arguing that under transparency, the committee-optimal feasible profile consists of at most one agent playing a signal-independent strategy, and all others sincerely revealing their signals. For even N we are able to show that this profile is committee-optimal even when the symmetric, mixed profile is feasible. For the case of odd N our result is slightly more restrictive, and applies only for large enough β_D. This is because here we do not have a general proof that the committee prefers the asymmetric profile to the symmetric one, so instead we show that when β_D is close enough the $\overline{\beta}$ the latter profile is not feasible.[10]

Theorem 2 *Fix any $p > 1/2$, even N, and any $\beta_D \in (1/2, \overline{\beta}(p)]$. Then the DM strictly prefers any committee-optimal feasible strategy profile under transparency to opaque, sincere majority.*[11]

Theorem 3 *Fix any $p > 1/2$ and odd $N \geq 3$, and let $\beta_D = \overline{\beta}(p)$. Then the DM strictly prefers any committee-optimal feasible strategy profile under transparency to opaque, sincere majority.*

[10] For the specific case of $N = 3$, as in the example at the beginning of the section, we *are* able to show that the asymmetric profile is committee-optimal even when the symmetric one is feasible.

[11] This uses the assumption when N is even and the committee is split, members recommend y. See Theorem 6 in the online appendix for a version of Theorem 2 without this assumption.

As noted above, the reason Theorem 3 is stated for $\beta_D = \overline{\beta}(p)$ and not other $\beta_D \leq \overline{\beta}(p)$ is that when N is odd, showing that the asymmetric equilibrium is committee-optimal is more challenging. When $\beta_D = \overline{\beta}(p)$ (and, in fact, as long as β_D is close enough to $\overline{\beta}$) then any mixed equilibrium is not feasible, but if β_D is much smaller than $\overline{\beta}(p)$ there is such a feasible profile. Proving the claim in that case then requires showing that this mixed equilibrium is not committee-optimal. While numerical calculations indicate that this is true, we currently do not have a general proof.

Verifiable signals

As we have seen above, when preferences are not sufficiently close a transparency requirement effectively eliminates the committee as an independent aggregator of information. In this section we show that when committee members possess verifiable messages, useful information may be conveyed to the DM. Indeed, we show that if members can vote by verifiably revealing their signals the DM may benefit from a transparency requirement. It is worth noting that if committee members can be compelled to disclose their private information (as would seem possible if such information is verifiable) then advisory committees might be disposed of entirely.

Formally, suppose that for each agent i there exist two messages, v_i^g and v_i^b, that can be used if and only if $s_i = g$ and $s_i = b$, respectively. Additionally, agents may choose not to use either verifiable vote, and simply abstain (which we denote by a vote $v_i = a$). That is, each agent either truthfully reveals his signal, or abstains. Then Lemma 1 no longer holds: there do exist feasible, persuasive strategy profiles under transparency, even when $\beta_D > \overline{\beta}$.

Example 1 Fix some N and $\beta_D \leq p^N/(p^N + (1-p)^N)$, guaranteeing that when all N agents have good signals the DM prefers outcome y. Consider the strategy profile in which all committee members verifiably disclose their signals, and the DM chooses outcome y if sufficiently many verifiably disclosed v_i^g. This is both persuasive and feasible.

However, Example 1 does not suffice as motivation for requiring transparency: although there is a feasible, persuasive strategy profile under transparency, it may not be committee-optimal. And the committee-optimal strategy profile might not be more beneficial to the DM than opaque, sincere majority.

We will prove the following theorem, a sharp contrast with Theorem 1, which states that with verifiable messages, the committee-optimal feasible process is strictly preferred by the DM.

Theorem 4 *For any N and $p > 1/2$, if $\beta_D > \overline{\beta}(p)$ and $\beta_D < p^N/(p^N + (1-p)^N)$ then when there are verifiable messages the DM strictly prefers any committee-optimal, feasible strategy profile under transparency to opaque, sincere majority.*[12]

Theorem 4 relies on a lemma that states that when there are verifiable messages, in any feasible strategy profile every agent who is sometimes pivotal will always vote verifiably on a good signal. Note that if, in a committee-optimal equilibrium, all agents were sometimes pivotal, the lemma would immediate imply Theorem 4. But this is actually not the case.[13] Nonetheless, in the proof of Theorem 4 we show that in every committee-optimal equilibrium at most one agent is never pivotal, and that this implies that transparency is strictly preferred by the DM.

Conclusion

In this paper we developed a model of transparency in advisory committees. The advantage of an advisory committee is that it allows a DM to aggregate information from multiple sources in order to make better decisions. The requirement of transparency has previously been seen as a way to enhance the ability of such a committee

[12] If $\beta_D = p^N/(p^N + (1-p)^N)$ then the DM's utility is 0 even if all agents have good signals, and if $\beta_D > p^N/(p^N + (1-p)^N)$ then the committee is never persuasive.

[13] For example, if $N = 4$ and $\beta_D = p^3/(p^3 + (1-p)^3)$, then the committee-optimal equilibrium is the asymmetric one in which the first three agents reveal their signals and the fourth abstains.

to provide persuasive information. Clearly, if there were no potential for disagreement between the committee members and the DM requiring transparency would not be necessary – the committee itself could aggregate the members' information and provide a simple recommendation to the ultimate DM. If transparency is to be imposed then two things have to be true. First, the ultimate decision maker must benefit from the imposition of transparency. This requires that committee and DM preferences must be different. Second, it must be that the committee itself would prefer not to use a transparent process. When the committee members' information is not verifiable we have shown that these conditions are never met when the preferences between the committee members and the DM are not sufficiently close.

In contrast, we have shown that when preferences are very close or when signals are verifiable then there is a much larger range of situations in which transparency benefits the DM.

The model we presented in the paper is stylized and considers a committee restricted to using simultaneous voting. In an online appendix we show that the results are very robust to a wide variety of different mechanisms committees might use during deliberations. Our results on the problems of transparency complement existing results that consider the ways in which transparency requirements may distort committee member preferences via outside influences and career concerns. Notwithstanding all the results showing the problems with transparency requirements, such requirements remain popular. This raises the question, why? As we show in our paper, perhaps the explanation lies in the availability of verifiable signals.

Appendix

Proof of Lemma 1

Fix N and $p > 1/2$, as well as a strategy profile that is feasible and persuasive under transparency. We will show that this implies $\beta_D \leq \overline{\beta}(p)$.

In the following, we slightly abuse notation and denote by $v \in V$, a vote profile, also the event $(V(\sigma(s)) = v)$. We also denote by $v_{-i} \in \{y, n\}^{N-1}$ and $v_i \in \{y, n\}$ the events $(V(\sigma(s))_{-i} = v_{-i})$ and $(V(\sigma(s))_i = v_i)$, respectively.

Next, let $\text{piv}_i \subseteq \{y, n\}^{N-1}$ be the set of pivotal vote profiles for agent i:

$$\text{piv}_i = \left\{ v_{-i} \in \{y, n\}^{N-1} : \text{ s.t. } r(y, v_{-i}) \neq r(n, v_{-i}) \right\}.$$

Finally, we can assume without loss of generality that $\sigma(\cdot)$ has full support over V, for otherwise some agents always play the same action and can thus be disregarded.

We begin with the following lemma:

Lemma 3 *If σ is feasible and persuasive then there exists an agent i and a $v_{-i} \in \text{piv}_i$ for which*

$$\Pr\left[\theta = G \mid v_{-i} \cap s_i = b\right] \leq \frac{1}{2}.$$

To prove Lemma 3 we need a definition and a claim.

Definition 2 (monotone decision rule) *A decision rule $r : V \mapsto \mathcal{O}$ of the DM is* monotone *if:*

- *There is a complete ordering $\succeq$ on elements V.*
- *For any $v, v' \in V$, if $v \succeq v'$ and $r(v') = y$ then $r(v) = y$.*

Consider the following claim:

Claim 2 *The decision rule r of the DM is monotone.*

Proof: For each agent i define $\succeq$ so that, for all $v, v' \in V$,

$$v \succeq v' \;\Leftrightarrow\; \Pr\left[\theta = G | v\right] \geq \Pr\left[\theta = G | v'\right].$$

In words, $v \succeq v'$ whenever the posterior on state G is higher after observing v than after observing v'.

Now fix any $v, v' \in V$ with $v \succeq v'$. Then $r(v') = y$ implies that $\Pr[\theta = G|v'] \geq \overline{\beta}(p)$. By the above, $\Pr[\theta = G|v] \geq \Pr[\theta = G|v'] \geq \overline{\beta}(p)$ and so $r(v) = y$. Thus, r is monotone. ■

We can now prove Lemma 3.

Proof of Lemma 3: Suppose towards a contradiction that there is no i and v_{-i} as desired. This implies that whenever i is pivotal his posterior on both signals $s_i = g$ and $s_i = b$ is $\Pr[\theta = G|v_{-i} \cap s_i] > 1/2$, and so whenever he is pivotal he will prefer outcome y. Thus, by the full support assumption, for both $v_i \in \{y, n\}$ the probability of getting outcome y conditional on being pivotal must be the same:

$$\Pr_{v_{-i} \in \mathrm{piv}_i}[r(y, v_{-i}) = y] = \Pr_{v_{-i} \in \mathrm{piv}_i}[r(n, v_{-i}) = y]. \tag{1}$$

Furthermore, for any $v_{-i} \in \mathrm{piv}_i$ and any $v_i \neq v_i'$, either $(v_i, v_{-i}) \succeq (v_i', v_{-i})$, or vice versa. Suppose without loss of generality that the former holds. Then by Claim 2, $r(v_i', v_{-i}) = y$ implies that $r(v_i, v_{-i}) = y$. Combining this with (1) implies that $r(y, v_{-i}) = r(n, v_{-i})$, and so $v_{-i} \notin \mathrm{piv}_i$. But this implies that i is never pivotal, i.e., $\mathrm{piv}_i = \emptyset$. Since this holds for every agent i, no agent is ever pivotal. This implies that $r(\cdot)$ is constant, contradicting the assumption that the strategy profile is persuasive. ■

Next, we prove the following lemma:

Lemma 4 *For any agent i and $v_{-i} \in \{y, n\}^{N-1}$ the following holds: If*

$$\Pr[\theta = G \mid v_{-i} \cap s_i = b] \leq \frac{1}{2}$$

then

$$\Pr[\theta = G \mid v_{-i} \cap s_i = g] \leq \frac{p^2}{p^2 + (1-p)^2}.$$

The intuition behind the lemma is simple. A g signal is equivalent to two g signals and one b signal. So going from the premise of the lemma to the conclusion is analogous to receiving two additional g signals. But two g signals move a prior of $1/2$ to a posterior of $p^2/(p^2+(1-p)^2)$.

Proof of Lemma 4: First, observe that

$$\begin{aligned}
&\Pr\left[s_i = g \cap v_{-i}\right] \\
&\quad = \Pr\left[s_i = g \cap v_{-i} \cap \theta = G\right] + \Pr\left[s_i = g \cap v_{-i} \cap \theta = B\right] \\
&\quad = \Pr\left[s_i = g \mid v_{-i} \cap \theta = G\right] \cdot \Pr\left[v_{-i} \cap \theta = G\right] \\
&\qquad + \Pr\left[s_i = g \mid v_{-i} \cap \theta = B\right] \cdot \Pr\left[v_{-i} \cap \theta = B\right] \\
&\quad = p \cdot \Pr\left[v_{-i} \cap \theta = G\right] + (1-p) \cdot \Pr\left[v_{-i} \cap \theta = B\right]
\end{aligned}$$

since the probability of obtaining signal $s_i = g$ or $s_i = b$ depends only on the state θ, and on v_{-i} also only through θ.

An analogous calculation with b replacing g yields

$$\Pr\left[s_i = b \cap v_{-i}\right] = (1-p) \cdot \Pr\left[v_{-i} \cap \theta = G\right] + p \cdot \Pr\left[v_{-i} \cap \theta = B\right].$$

Next,

$$\begin{aligned}
&\Pr\left[\theta = G \mid v_{-i} \cap s_i = g\right] \\
&\quad = \frac{\Pr\left[s_i = g \mid \theta = G \cap v_{-i}\right] \cdot \Pr\left[\theta = G \cap v_{-i}\right]}{\Pr\left[s_i = g \cap v_{-i}\right]} \\
&\quad = \frac{p \cdot \Pr\left[\theta = G \cap v_{-i}\right]}{\Pr\left[s_i = g \cap v_{-i}\right]} \\
&\quad = \frac{p \cdot \Pr\left[\theta = G \cap v_{-i}\right]}{p \cdot \Pr\left[v_{-i} \cap \theta = G\right] + (1-p) \cdot \Pr\left[v_{-i} \cap \theta = B\right]}. \qquad (2)
\end{aligned}$$

Again, an analogous calculation with b replacing g yields

$$\begin{aligned}
&\Pr\left[\theta = G \mid v_{-i} \cap s_i = b\right] \\
&\quad = \frac{(1-p) \cdot \Pr\left[\theta = G \cap v_{-i}\right]}{(1-p) \cdot \Pr\left[v_{-i} \cap \theta = G\right] + p \cdot \Pr\left[v_{-i} \cap \theta = B\right]}.
\end{aligned}$$

The premise of the lemma,

$$\Pr\left[\theta = G \mid v_{-i} \cap s_i = b\right] \leq \frac{1}{2},$$

thus implies that

$$\frac{(1-p) \cdot \Pr\left[\theta = G \cap v_{-i}\right]}{(1-p) \cdot \Pr\left[v_{-i} \cap \theta = G\right] + p \cdot \Pr\left[v_{-i} \cap \theta = B\right]} \leq \frac{1}{2},$$

and so

$$\Pr\left[\theta = G \cap v_{-i}\right] \leq \frac{p}{1-p} \cdot \Pr\left[\theta = B \cap v_{-i}\right].$$

Plugging this into equation (2) yields

$$\begin{aligned}
&\Pr\left[\theta = G \mid v_{-i} \cap s_i = g\right] \\
&\quad = \frac{p \cdot \Pr\left[\theta = G \cap v_{-i}\right]}{p \cdot \Pr\left[v_{-i} \cap \theta = G\right] + (1-p) \cdot \Pr\left[v_{-i} \cap \theta = B\right]} \\
&\quad \leq \frac{p \cdot \frac{p}{1-p} \cdot \Pr\left[\theta = B \cap v_{-i}\right]}{p \cdot \frac{p}{1-p} \cdot \Pr\left[\theta = B \cap v_{-i}\right] + (1-p) \cdot \Pr\left[v_{-i} \cap \theta = B\right]} \\
&\quad = \frac{p^2}{p^2 + (1-p)^2},
\end{aligned}$$

as claimed. ■

Finally, we prove the following lemma:

Lemma 5 *If a strategy profile feasible and persuasive under transparency then* $\beta_D \leq \overline{\beta}(p)$.

Proof of Lemma 5: By Lemma 3, there exists an agent i and an $v_{-i} \in \mathrm{piv}_i$ such that

$$\Pr\left[\theta = G \mid v_{-i} \cap s_i = b\right] \leq \frac{1}{2}.$$

By Lemma 4, it thus holds that

$$\Pr\left[\theta = G \mid v_{-i} \cap s_i = g\right] \leq \frac{p^2}{p^2 + (1-p)^2}.$$

Additionally, observe that

$$\Pr\left[\theta = G \mid v_{-i} \cap s_i = s\right] = \Pr\left[\theta = G \mid v_{-i} \cap s_i \cap v_i\right]$$

for any $s_i \in \{g, b\}$ and $v_i \in \{y, n\}$, since, given the signal s_i, the posterior on θ does not depend on v_i.

Since $v_{-i} \in \mathrm{piv}_i$, there exists some $v_i \in \{y, n\}$ for which $r(v_i, v_{-i}) = y$. We will show that conditional on observing the vote profile $v = (v_i, v_{-i})$, the DM's posterior on $\theta = G$ is at most $\overline{\beta}$:

$$\begin{aligned}
&\Pr\left[\theta = G \mid v\right] = \Pr\left[\theta = G \mid v_{-i} \cap v_i\right] \\
&= \Pr\left[\theta = G \cap s_i = g \mid v_{-i} \cap v_i\right] + \Pr\left[\theta = G \cap s_i = b \mid v_{-i} \cap v_i\right] \\
&= \Pr\left[\theta = G \mid v_{-i} \cap v_i \cap s_i = g\right] \cdot \Pr\left[s_i = g \mid v_{-i} \cap v_i\right] \\
&\quad + \Pr\left[\theta = G \mid v_{-i} \cap h_i \cap s_i = b\right] \cdot \Pr\left[s_i = b \mid v_{-i} \cap v_i\right] \\
&\leq \frac{p^2}{p^2 + (1-p)^2} \cdot \Pr\left[s_i = g \mid v_{-i} \cap v_i\right] + \frac{1}{2} \cdot \Pr\left[s_i = b \mid v_{-i} \cap v_i\right] \\
&\leq \frac{p^2}{p^2 + (1-p)^2} = \overline{\beta}.
\end{aligned}$$

Thus, $\Pr\left[\theta = G \mid v\right] \leq \overline{\beta}$. However, $r(v)$ should be an optimal decision for the DM on vote profile v, which implies that $\beta_D \leq \overline{\beta}$ as claimed. ■

Lemma 5 immediately implies Lemma 1.

Proofs for small β_D

Agent i is called *non-pivotal* if $\text{piv}_i = \emptyset$. This means his action is signal-independent, and so $r(\cdot)$ is independent of his action. We begin with a lemma that we will use both for the case of small β_D and for the case of verifiable signals.

Lemma 6 *If $\beta_D \geq p$ then in any committee-optimal persuasive pure strategy equilibrium under transparency at most one agent is non-pivotal.*

Proof: Suppose towards a contradiction that this is not the case, and that in the committee-optimal persuasive pure strategy equilibrium two or more agents are not pivotal. Without loss of generality, denote the pivotal agents by $L = \{1, \ldots, \ell\}$, and two of the non-pivotal agents by u and w. Without loss of generality, we can assume that agents in L vote y (or play action v_i^g under verifiable signals) on signal $s_i = g$, and vote n (or play action v_i^b under verifiable signals) on signal $s_i = b$. If there are no verifiable signals, in the remainder of this proof identify vote y with action v_i^g and vote n with action v_i^b.

We will show that the equilibrium in which agents in $L \cup \{u, w\}$ are pivotal (where the latter also vote sincerely) is strictly better for the committee than the equilibrium in which only the agents in L are pivotal. Denote by $X \in \{v_i^g, v_i^b\}^\ell$ the set of action profiles x that satisfy

$$\Pr\left[\theta = G | x\right] \in \left[\frac{(1-p)^2 \beta_D}{(1-p)^2 \beta_D + p^2 (1-\beta_D)}, \frac{p^2 \beta_D}{p^2 \beta_D + (1-p)^2 (1-\beta_D)}\right).$$

Note that $X \neq \emptyset$ since the strategy profile is persuasive. X consists of the action profiles for which the DM will change his choice of outcome if he learns that $s_u = s_w = g$ or if he learns that $s_u = s_w = b$. For all other action profiles of agents in L, learning the signals of agents u and w will have no effect. Furthermore, even for profiles in $x \in X$, learning the $s_u \neq s_w$ will also not have an effect, since these signals cancel each other and so will not change the DM's posterior.

Now, observe that for $x \in X$ such that $\Pr[\theta = G \mid x] < \beta_D$, learning the signals of u and w is always beneficial to the committee: on such an x the DM will choose $r(x) = n$, even though the committee desires the other outcome (since $\Pr[\theta = G \mid x] \geq 1/2$). Thus, while learning $s_u = s_w = b$ will have no effect, learning that $s_u = s_w = g$ will lead to outcome y. In contrast, for $x \in X$ such that $\Pr[\theta = G \mid x] \geq \beta_D$, learning the signals of u and w is always harmful to the committee: on such an x the DM will choose $r(x) = y$, so while $s_u = s_w = g$ will have no effect, learning that $s_u = s_w = b$ will lead to outcome b.

We need to show that the benefits dominate the harms. Denote by $S = \{x \in X : \Pr[\theta = G \mid x] \geq \beta_D\}$, and by $T = \{x \in X : \Pr[\theta = G \mid x] < \beta_D\}$. Observe that profiles in S and those in T differ in the number of agents with a positive signal by exactly 1. In the former, the number of agents with a good signal is some k, whereas in the latter it is $k - 1$. In the former, the addition of the signals of u and w can only be harmful, whereas in the latter it can only by beneficial. Consider a matching from elements of S to elements of T, where each element $x_S \in S$ is matched to a unique element $x_T \in T$, and with the additional constraint that x_S and its match x_T differ in the signal of exactly one agent from L. While not obvious, the fact that this is possible follows from Theorem 3 of Bollobás and Leader (1997). Note that all elements of S are matched, but not all elements of T, since the cardinality of the former is smaller than that of the latter. But for those profiles that remain unmatched, the addition of u and w can only be beneficial.

Fix some $x_S \in S$ and its match $x_T \in T$, and suppose they differ on the signal of agent $z \in L$. Let $\overline{x}$ be the profile of signals in x_S, but excluding that of agent z, and note that $\Pr[\theta = G \mid \overline{x}] \geq 1/2$ (since $\beta_D \geq p$). Condition on the profile $\overline{x}$, and consider now the addition of the signals of agents z, u, and w. The addition of u's and w's signals is beneficial to the committee when $s_z = b$ but $s_u = s_w = g$, and harmful when $s_z = g$ but $s_u = s_w = b$. In all other cases, the addition has no effect. Since $\Pr[\theta = G \mid \overline{x}] \geq 1/2$, the former case is more likely than the latter (2 good signals and 1 bad is more likely than 2 bad signals and 1 good since the probability of the good state is greater than 1/2). Furthermore, the expected utility gain to the committee in the former case is equal to the expected utility loss in the latter. Thus, in total, the committee gains from including the signals of agents u and w. This is a contradiction to the assumption that the committee-optimal equilibrium has two or more non-pivotal agents. ■

Next, consider a signal profile $s \in \{g, b\}^N$, and denote by $\text{diff}(s) = |\{i : s_i = g\}| - |\{i : s_i = b\}|$. Note that if $\text{diff}(s) > 0$, then the committee obtains positive expected utility from outcome y, if $\text{diff}(s) < 0$ the committee obtains negative expected utility, and if $\text{diff}(s) = 0$ they obtain expected utility 0 and are thus indifferent between a y and a n outcome. Thus, a committee-optimal process will lead to one in which the outcome is y whenever s satisfies $\text{diff}(s) > 0$ and to outcome n when s satisfies $\text{diff}(s) < 0$.

Proof of Theorem 3: Observe that under opaque, sincere majority the DM obtains negative expected utility when the signal profile is some s with $\text{diff}(s) = 1$, as these profiles also lead to y recommendation and hence outcome y.

Consider now the case of transparency. We first show that no feasible strategy profile can involve a mixed strategy by an agent who is sometimes pivotal. To see this, consider some agent i with $\text{piv}_i \neq \emptyset$ who mixes in equilibrium. Then he must mix when his signal is $s_i = b$. Lemma 3 implies that there exists some $v_{-i} \in \text{piv}_i$ such that $\Pr[\theta = G \mid v_{-i}, s_i = b] \leq 1/2$. This implies that on signal $s_i = g$, the posterior is $\Pr[\theta = G \mid h_{-i}, s_i = g] \leq \overline{\beta}$. Since $v_{-i} \in \text{piv}_i$, there is some vote $\overline{v}$ for which $r(\overline{v}, v_{-i}) = y$. Because agent i mixes, however, his vote $v_i = \overline{v}$ does not fully reveal his signal, and so the posterior on $v_i = \overline{v}$ must be strictly less than $\overline{\beta}$; formally, $\Pr[\theta = G \mid v_{-i}, v_i = \overline{v}] < \overline{\beta}$. But since $\beta_D = \overline{\beta}$, this contradicts the assumption that $v_{-i} \in \text{piv}_i$.

Since feasible strategy profiles cannot involve mixing, the only candidate strategies are pure, but in which some of the agents may play a signal-independent action (i.e., they are non-pivotal). If $\beta_D < p$ then the only committee-optimal feasible profile is the sincere one, since then the preferences of the DM and committee are perfectly aligned (i.e., on every signal profile s all agree about the preferred outcome). If $\beta_D \geq p$ then by Lemma 6, in such strategy profiles that are committee-optimal, there is at most one non-pivotal agent. If there are no such agents, and so each agent votes sincerely, the DM clearly benefits with respect to opaque, sincere majority. The process in which there is one non-pivotal agent is also strictly preferred by the DM: this is because he avoids outcome y on some signal profiles s with $\text{diff}(s) = 1$, namely the ones where the non-pivotal agent's signal is g. For all other signal profiles, the outcome is the same as in opaque, sincere majority. Thus, in all committee-optimal feasible strategy profiles under transparency, the DM obtains strictly higher expected utility than in opaque, sincere majority. ■

Proof of Theorem 2: Under opaque, sincere majority the outcome is y whenever s satisfies $\text{diff}(s) \geq 0$, and n otherwise.

Consider now the case of transparency. First, note that the sincere strategy profile is feasible and persuasive here. Furthermore, this strategy profile is committee-optimal, as the committee obtains their desired outcome on every signal profile (outcome y if $\text{diff}(s) > 0$ and outcome n if $\text{diff}(s) < 0$). This implies that the symmetric, mixed strategy profile cannot be committee-optimal, since that profile inevitably leads to some incorrect outcome with positive probability.

Other candidates for committee-optimal feasible strategy profiles are the asymmetric profiles à la McLennan (1998), in which one agent is non-pivotal and the rest vote sincerely (note that if $\beta_D > p$ then only the sincere profile is committee-optimal). In all these committee-optimal profiles, the outcome is y whenever $\text{diff}(s) > 0$, as well as some but not all profiles in which $\text{diff}(s) = 0$. In particular, the outcome is n when $\text{diff}(s) = 0$ and $s_i = b$, where i is the non-pivotal agent.

Finally, the only remaining candidates for committee-optimal feasible strategy profiles under transparency are ones where $N-1$ agents vote sincerely, and one agent mixes (voting y also on signal $s_i = b$ with some probability). This is feasible and committee-optimal as long as he does not mix too much, so that the posterior when i as well as $N/2$ agents other than i vote y is at least β_D. But similarly to the asymmetric pure profiles, the outcome is y whenever $\text{diff}(s) > 0$, as well as some but not all profiles in which $\text{diff}(s) = 0$. In particular, the outcome is n with positive probability when $\text{diff}(s) = 0$ and $s_i = b$, where i is the non-pivotal agent.

Since under opaque, sincere majority the outcome is y for all profiles s with $\text{diff}(s) = 0$, the DM strictly prefers transparency. ■

Proofs for verifiable signals

Denote by Z_i the set $\{v_i^g, v_i^b, a\}$. We begin with a lemma:

Lemma 7 *If $\beta_D > \overline{\beta}(p)$ then in every feasible, persuasive strategy profile it holds that* $\text{supp}\,(\sigma_i(g)) = \{v_i^g\}$ *for every agent i that satisfies* $\text{piv}_i \neq \emptyset$.

That is, in any such strategy profile, every agent i who is sometimes pivotal plays the verifiable action v_i^g with probability 1 when $s_i = g$.

Proof of Lemma 7: The proof is similar to that of Lemma 1. Fix a feasible and persuasive strategy profile, and suppose towards a contradiction that there exists a pivotal agent i such that $\text{supp}\,(\sigma_i(g)) \neq \{v_i^g\}$. That is, on signal $s_i = g$, agent i plays the non-verifiable action a with positive probability. Denote by $\text{piv}_i \subseteq \times_{j\neq i} Z_j$ agent i's pivotal action profiles: $\text{piv}_i = \{z_{-i} \in \times_{j\neq i} Z_j : \exists z_i, z_i' \in Z_i \text{ s.t. } r(z_i, z_{-i}) \neq r(z_i', z_{-i})\}$.

We first claim that there exists some $z_{-i} \in \text{piv}_i$ for which $\Pr\left[\theta = G \mid z_{-i} \cap s_i = b\right] \leq 1/2$. The proof of this claim essentially follows that of Lemma 9: If the implication of the claim does not hold, then whenever i is pivotal his posterior on both signals $s_i = g$ and $s_i = b$ is $\Pr\left[\theta = G | z_{-i} \cap s_i\right] > 1/2$, and so whenever he is pivotal he will prefer outcome y. Thus, for any pair of actions $z_i, z_i' \in \text{supp}(\sigma_i(g))$, the probability of getting outcome y conditional on being pivotal must be the same. But $a \in \text{supp}(\sigma_i(g))$, and this implies that in fact for any pair of actions in $\mathrm{supp}(\sigma_i(g)) \cup \mathrm{supp}(\sigma_i(b))$, the probability of getting outcome y conditional on being pivotal must be the same. This, together with Claim 2, implies that $\mathrm{piv}_i = \emptyset$, a contradiction.

Given a $z_{-i} \in \mathrm{piv}_i$ for which $\Pr[\theta = G \mid z_{-i} \cap s_i = b] \leq 1/2$, Lemma 4 implies that $\Pr[\theta = G \mid z_{-i} \cap s_i = g] \leq \overline{\beta}(p)$.[14] The argument in the proof of Lemma 5 can then be directly applied to i and z_{-i}, leading to the contradiction that for some $z_i \in Z_i$ for which $r(z_i, z_{-i}) = y$ it holds that $\Pr[\theta = G \mid (a, z_{-i})] \leq \overline{\beta}(p) < \beta_D$. Thus, there is no agent i such that $\mathrm{supp}(\sigma_i(g)) \neq \{v_i^g\}$. ■

Proof of Theorem 4: By Lemma 7, for each agent $i \in L$ it holds that $\mathrm{supp}(\sigma_i(g)) = \{v_i^g\}$. Since this action is verifiable, any other action by agent i is taken to imply that $s_i = b$. Without loss of generality, then, we simply assume that all actions played in equilibrium are verifiable ones, either v_i^g or v_i^b. Thus, the equilibrium consists of pure strategies. By Lemma 6, in the committee-optimal pure equilibrium there is at most one agent who is never pivotal.

We now show that the equilibrium with at most one non-pivotal agent, where all others reveal their signals, is strictly better for the DM than opaque majority. First note that this is clearly true if all agents are pivotal. Suppose then that one agent is not pivotal, say agent i. The DM will choose outcome y on action profiles z_{-i} for which $\Pr[\theta = G \mid z_{-i}] \geq \beta_D$. Importantly, the outcome is y only on some signal profiles s with $\mathrm{diff}(s) \in \{1, 2\}$, and in particular not on those with $s_i = g$ (it is also possible that the outcome is y on no such profiles), and on no profiles with $\mathrm{diff}(s) = 0$. Under opaque majority, however, the outcome will be y on the same profiles as under transparency, but also on additional profiles s: namely, all the profiles in which $\mathrm{diff}(s) \in \{0, 1, 2\}$. Thus, the DM is strictly better off under transparency. ■

[14]Strictly speaking, this follows from a generalization of Lemma 4, namely Lemma 10, in the online appendix.

References

Austen-Smith, D. and Banks, J. S. (1996). Information aggregation, rationality, and the condorcet jury theorem. *American Political Science Review*, **90** 34–45.

Battaglini, M. (2015). Public protests, social media and policy making. *Princeton University William S. Dietrich II Economic Theory Center Research Paper*.

Berliner, D. (2014). The political origins of transparency. *The Journal of Politics*, **76** 479–491.

Bollobás, B. and Leader, I. (1997). Matchings and paths in the cube. *Discrete applied mathematics*, **75** 1–8.

Coughlan, P. J. (2000). In defense of unanimous jury verdicts: Mistrials, communication, and strategic voting. *American Political science review*, **94** 375–393.

Crawford, V. P. and Sobel, J. (1982). Strategic information transmission. *Econometrica*, **50** 1431–1451.

Dal Bó, E. (2007). Bribing voters. *American Journal of Political Science*, **51** 789–803.

FAC (2015). Federal Advisory Committee. URL `http://www.federalreserve.gov/aboutthefed/fac.htm`.

FDA (2008). Guidance for FDA Advisory Committee members and FDA staff : Voting procedures for advisory committee meetings. URL `http://www.fda.gov/downloads/RegulatoryInformation/Guidances/UCM125641.pdf`.

FDA (2015). URL `http://www.fda.gov/AdvisoryCommittees/`.

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. (forthcoming). How transparency kills information aggregation: theory and experiment. *American Economic Journal: Microeconomics*.

Felgenhauer, M. and Grüner, H. P. (2008). Committees and special interests. *Journal of Public Economic Theory*, **10** 219–243.

Fingleton, J. and Raith, M. (2005). Career concerns of bargainers. *Journal of Law, Economics, and Organization*, **21** 179–204.

Gersbach, H. and Hahn, V. (2008). Should the individual voting records of central bankers be published? *Social Choice and Welfare*, **30** 655–683.

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`.

Hollyer, J. R., Rosendorff, B. P. and Vreeland, J. R. (2011). Democracy and transparency. *Journal of Politics*, **73** 1191–1205.

IAC (2015). Investor Advisory Committee. URL `http://www.sec.gov/spotlight/investor-advisory-committee-2012.shtml`.

IPCC (2015). Intergovernmental panel on climate change. URL `http://www.ipcc.ch/`.

Levy, G. (2007a). Decision making in committees: Transparency, reputation, and voting rules. *American Economic Review*, **97** 150–168.

Levy, G. (2007b). Decision-making procedures for committees of careerist experts. *American Economic Review*, **97** 306–310.

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.

Mattozzi, A. and Nakaguma, M. Y. (2017). Public versus secret voting in committees.

Mayhew, D. R. (1974). *Congress: The Electoral Connection*. Yale University Press.

McLennan, A. (1998). Consequences of the condorcet jury theorem for beneficial information aggregation by rational agents. *American Political Science Review*, **92** 413–418.

Meade, E. E. and Stasavage, D. (2008). Publicity of debate and the incentive to dissent: evidence from the US Federal Reserve. *The Economic Journal*, **118** 695–717.

Milgrom, P. (2008). What the seller won't tell you: Persuasion and disclosure in markets. *The Journal of Economic Perspectives*, **22** 115–131.

Moffitt, S. L. (2010). Promoting agency reputation through public advice: Advisory committee use in the FDA. *The Journal of Politics*, **72** 880–893.

Osborne, M. J. and Rubinstein, A. (1994). *A course in game theory*. MIT press.

Prat, A. (2005). The wrong kind of transparency. *American Economic Review*, **95** 862–877.

Seidmann, D. J. (2011). A theory of voting patterns and performance in private and public committees. *Social Choice and Welfare*, **36** 49–74.

Shepsle, K. A. and Weingast, B. R. (1987). The institutional foundations of committee power. *The American Political Science Review*, **81** 85–104.

Sibert, A. (2003). Monetary policy committees: Individual and collective reputations. *The Review of Economic Studies*, **70** 649–665.

Snyder, J. M. and Ting, M. M. (2005). Why roll calls? A model of position-taking in legislative voting and elections. *Journal of Law, Economics, and Organization*, **21** 153–178.

Stasavage, D. (2007). Polarization and publicity: Rethinking the benefits of deliberative democracy. *Journal of Politics*, **69** 59–72.

Swank, O. H. and Visser, B. (2013). Is transparency to no avail? *The Scandinavian Journal of Economics*, **115** 967–994.

Wolinsky, A. (2002). Eliciting information from multiple experts. *Games and Economic Behavior*, **41** 141–160.

Online Appendix

General model

In addition to the N members of the committee, there is a committee chair whose utility function is identical to the committee's, but who does not obtain any information.

To describe the interaction of committee members, first recall the following definition of an extensive game with perfect information and simultaneous moves (see for example Osborne and Rubinstein, 1994):

Definition 3 (game) *An N-agent* game Γ *is a pair* (H, P) *where:*

- H *is a set of finite* history *sequences such that the empty string* $\epsilon \in H$*. A history* $h \in H$ *is* terminal *if* $\{a : (h, a) \in H\} = \emptyset$*. The set of terminal histories is denoted* $Z \subseteq H$*.*

- P *is a function that maps nonterminal histories* h *to sets of agents, namely subsets of* $\{1, \ldots, N\}$*. For an agent* i *such that* $i \in P(h)$*, denote by* $A_i(h)$ *the set of* actions *available to agent* i *at history* h*, such that* $\times_{i \in P(h)} A_i(h) = \{a : (h, a) \in H\}$*.*

A strategy σ_i of agent i in a game Γ is a function that, for every h and $i \in P(h)$, maps every (s_i, h) to a distribution over $A_i(h)$. Denote by $\sigma = (\sigma_1, \ldots, \sigma_N)$ a profile of strategies, by $\sigma(s)$ the profile of strategies given signal profile $s = (s_1, \ldots, s_N)$, and by $Z(\sigma(s))$ the distribution over terminal histories reached when the profile $\sigma(s)$ is played.

While the committee members play an extensive game, the DM only observes a message about the path of play. This message is specified by the *disclosure rule*: a (possibly random) function $d : Z \mapsto D$, where D is some set of messages. A *mechanism* is then a game Γ coupled with a disclosure rule d.

Definition 4 (process) *A* process *is a tuple* (Γ, d, σ) *that consists of a mechanism and a strategy profile. A process is* transparent *if* $d(h) = h$ *for every* $h \in Z$*, and is* non-transparent *otherwise. A process is* opaque *if the set* D *in the range of* d *has cardinality 2.*

Note that if a process is opaque, then one can think of each $d \in D$ as a recommendation, one for y and one for n. In this case we will use the convention $D = \{m_y, m_n\}$.

Example 2 A *voting process* is one in which Γ is a simultaneous-move game with each member voting y or n. Formally, $P(\epsilon) = \{1, \ldots, N\}$, and $H = \{\epsilon\} \cup \{y, n\}^N$. That is, every agent i has an action that is a vote $A_i(\epsilon) = \{y, n\}$, actions are simultaneous and so $A_i(h) = \emptyset$ for all $h \neq \epsilon$, and terminal histories consist of all profiles of votes $\{y, n\}^N$.

If a voting process is transparent, then the DM learns the profile of votes. An example of an opaque voting process is one that "discloses the majority", in which $d : \{y, n\}^N \mapsto \{m_y, m_n\}$, with $d(z) = m_y$ if and only if $|\{i : z_i = y\}| \geq N/2$.

Of course, processes can be more complex, and may involve sequential voting, communication amongst committee members, straw polls, and more.

After the mechanism is played, the DM obtains a message specified by the disclosure rule. The DM updates his prior over the state, and takes an action that depends on whether the posterior surpasses his threshold β_D or not. Formally, given a process (Γ, d, σ), denote the rational decision rule used by the DM on message $m \in D$ as $r((\Gamma, d, \sigma), m)$, where $r((\Gamma, d, \sigma), m) = y$ if $\Pr[\theta = G \mid m, (\Gamma, d, \sigma)] \geq \beta_D$, and $r((\Gamma, d, \sigma), m) = n$ otherwise.

When the DM updates his prior he conditions on both the message $d(z)$ and on the process (Γ, d, σ) chosen by the chair. But what prevents the committee from deviating from the strategy chosen by the chair, unbeknownst to the DM?[15]

[15]While we assume that the DM observes the committee's choice of process, this assumption is not important for our results—in particular, Theorem 5 goes through also for a variant of the model in which the DM observes only $d(z)$.

In order to prevent this, we will additionally restrict the chair to choose a process that is feasible: one in which the strategy profile constitutes a Nash equilibrium for the agents conditional on the DM acting rationally. Formally,

Definition 5 (feasibility) *A process* (Γ, d, σ) *is* feasible *if for each agent* i*, signal* s_i*, and strategy* σ_i'*,*

$$\mathrm{E}\left[u(\theta, \overline{r}(d(H(\sigma(s))))) \mid s_i\right] \geq \mathrm{E}\left[u(\theta, \overline{r}(d(H(\sigma_i', \sigma_{-i}(s))))) \mid s_i\right],$$

where $\overline{r}(\cdot) \equiv r((\Gamma, d, \sigma), \cdot)$ *and the expectation is over* θ*,* s*, and* σ*.*

Figure 1 outlines the game we will analyze, parametrized by the signal accuracy p and denoted $M(p)$. Note that in this game, the DM's initial choice consists only of whether or not to require transparency, whereas the choice of process belongs to the committee. This modeling choice is consistent with our motivating examples of federal advisory committees, whose processes are determined by the committees themselves, while the law (the Federal Advisory Committee Act) only requires them to be transparent. However, we note that a different modeling choice, in which the DM also chooses the transparent process (subject to feasibility), would lead to the same negative result.

1. The DM chooses to either require transparency, or not.
2. The committee chair chooses a feasible process (Γ, d, σ). If required by the DM, the process must be transparent. Committee members then observe their signals s and play game Γ with strategy profile $\sigma(s)$, leading to some terminal history $z \in Z$.
3. The DM observes the choice of process as well as the realized element $d(z)$, and chooses an outcome from $\{y, n\}$.

Figure 1: The game $M(p)$.

As we are interested in the possible benefit of transparency, we ask the following question: Is there a subgame perfect equilibrium (SPE) of $M(p)$ in which the DM chooses to require transparency? And is this choice strictly preferred?

We define one final property of processes—persuasiveness. A process is persuasive if there are some instances in which the DM is influenced by it—that is, if there exist two messages, sent with positive probability, such that a DM who acts rationally will make different decisions after observing these messages. Formally,

Definition 6 (persuasiveness) *A process* (Γ, d, σ) *is* persuasive *if there exist two messages* $m, m' \in D$ *such that*

- $\Pr\left[d(H(\sigma(s))) = m\right] > 0$ *and* $\Pr\left[d(H(\sigma(s'))) = m'\right] > 0$ *for some signal profiles* $s, s' \in \{g, b\}^N$*, and*
- $r((\Gamma, d, \sigma), m) \neq r((\Gamma, d, \sigma), m')$.

Observe that if a process is not persuasive then the DM makes his choice based only on the prior distribution over states, and so will always choose action n. This yields him and the committee members utility 0.

Transparency is harmful – general model

Theorem 5 *For any N and $p > 1/2$, if $\beta_D > \overline{\beta}(p)$ then in any SPE of $M(p)$ the DM (weakly) prefers **not** to require transparency.*

The theorem relies on the following lemma:

Lemma 8 *There does not exist any feasible, persuasive, and transparent process for any N, $p > 1/2$, and $\beta_D > \overline{\beta}(p)$.*

Theorem 5 follows directly from Lemma 8. The latter implies that any feasible, transparent process cannot be persuasive, and so it yields utility 0 to the DM. Without requiring transparency, the expected utility of the DM is at least 0.

Proof of Lemma 8

Fix N and $p > 1/2$, as well as an N-agent feasible, persuasive, and transparent process (Γ, d, σ). We will show that this implies $\beta_D \leq \overline{\beta}(p)$.

Normal-form mechanisms, arbitrary prior and signals

We first prove Lemma 8 for normal-form mechanisms—ones in which all histories $h \in A_1(\epsilon) \times \ldots \times A_N(\epsilon)$ are terminal—and in a more general setting of parameters: First, fix any prior $\pi_0 \stackrel{\text{def}}{=} \Pr[\theta = G] \in (0, 1)$. Second, it will be useful for the proof with extensive-form mechanisms (in the next section) to allow for more general signal accuracies, and not just p. To that effect, for each agent i define two signal accuracies, denoted by p_i^g and p_i^b, where

$$p_i^g \stackrel{\text{def}}{=} \Pr[s_i = g | \theta = G] \quad \text{and} \quad p_i^b \stackrel{\text{def}}{=} \Pr[s_i = b | \theta = B].$$

We will place one restriction on each p_i^g and p_i^b, whose purpose will become apparent in the proof here and in the section on extensive-form games. Namely, for each agent i, let p_i^g and p_i^b satisfy

$$\frac{p_i^g p_i^b}{p_i^g p_i^b + (1 - p_i^g)(1 - p_i^b)} \leq \frac{p^2}{p^2 + (1 - p)^2}. \tag{3}$$

In particular, this is satisfied whenever $p_i^g = p_i^b = p$.

For simplicity of notation, denote by $A_i \stackrel{\text{def}}{=} A_i(\epsilon)$. In the following, we slightly abuse notation and denote by $h \in Z$, a terminal history, also the event $(d(H(\sigma(s))) = h)$. We also denote by $h_{-i} \in \times_{j \neq i} A_j$ and $h_i \in A_i$ the events $(d(H(\sigma(s)))_{-i} = h_{-i})$ and $(d(H(\sigma(s)))_i = h_i)$, respectively.

Next, let $\text{piv}_i \subseteq \times_{j \neq i} A_j$ be the set of pivotal histories for agent i:

$$\text{piv}_i = \{h_{-i} \in \times_{j \neq i} A_j : \exists h_i, h_i' \in A_i \text{ s.t. } r(h_i, h_{-i}) \neq r(h_i', h_{-i})\}.$$

Finally, we can assume without loss of generality that $\sigma(\cdot)$ has full support over $\times_i A_i$, for otherwise we can just delete those actions that have probability 0.

We begin with the following lemma:

Lemma 9 *If (Γ, d, σ) is feasible, persuasive, and transparent then there exists an agent i and an $h_{-i} \in \mathrm{piv}_i$ for which*

$$\Pr\left[\theta = G \mid h_{-i} \cap s_i = b\right] \leq \frac{1}{2}.$$

To prove Lemma 9 we need a definition and a claim.

Definition 7 (monotone decision rule) *A decision rule $r : Z \mapsto \mathcal{O}$ of the DM is* monotone *if:*

- *There is a complete ordering $\succeq$ on elements Z.*
- *For any $h, h' \in Z$, if $h \succeq h'$ and $r(h') = y$ then $r(h) = y$.*

Consider the following claim:

Claim 3 *In any transparent process, the decision rule r of the DM is monotone.*

Proof: For each agent i define $\succeq$ so that, for all $h, h' \in A$,

$$h \succeq h' \;\Leftrightarrow\; \Pr\left[\theta = G | h\right] \geq \Pr\left[\theta = G | h'\right].$$

In words, $h \succeq h'$ whenever the posterior on state G is higher after observing h than after observing h'.

Now fix any $h, h' \in Z$ with $h \succeq h'$. Then $r(h') = y$ implies that $\Pr\left[\theta = G | h'\right] \geq \overline{\beta}(p)$. By the above, $\Pr\left[\theta = G | h\right] \geq \Pr\left[\theta = G | h'\right] \geq \overline{\beta}(p)$ and so $r(h) = y$. Thus, r is monotone. ∎

We can now prove Lemma 9.

Proof of Lemma 9: Suppose towards a contradiction that there is no i and h_{-i} as desired. This implies that whenever i is pivotal his posterior on both signals $s_i = g$ and $s_i = b$ is $\Pr[\theta = G | h_{-i} \cap s_i] > 1/2$, and so whenever he is pivotal he will prefer outcome y. Thus, for any pair of actions $h_i, h_i' \in H_i$, the probability of getting outcome y conditional on being pivotal must be the same:

$$\Pr_{h_{-i} \in \mathrm{piv}_i} [r(h_i, h_{-i}) = y] = \Pr_{h_{-i} \in \mathrm{piv}_i} [r(h_i', h_{-i}) = y] \tag{4}$$

(by the full support assumption).

Furthermore, for any $h_{-i} \in \mathrm{piv}_i$ and any $h_i \neq h_i'$, either $(h_i, h_{-i}) \succeq (h_i', h_{-i})$, or vice versa. Suppose without loss of generality that the former holds. Then by Claim 3, $r(h_i', h_{-i}) = y$ implies that $r(h_i, h_{-i}) = y$. Combining this with (4) implies that $r(h_i', h_{-i}) = r(h_i, h_{-i})$ for all $h_{-i} \in \mathrm{piv}_i$ and all $h_i, h_i' \in H_i$, and so $h_{-i} \notin \mathrm{piv}_i$. But this implies that i is never pivotal, i.e., $\mathrm{piv}_i = \emptyset$. Since this holds for every agent i, no agent is ever pivotal. This implies that $r(\cdot)$ is constant, contradicting the assumption that the process is persuasive. ■

Next, we prove the following lemma:

Lemma 10 *For any agent i and $h_{-i} \in \times_{j \neq i} A_j$ the following holds: If*

$$\Pr[\theta = G \mid h_{-i} \cap s_i = b] \leq \frac{1}{2}$$

then

$$\Pr[\theta = G \mid h_{-i} \cap s_i = g] \leq \frac{p^2}{p^2 + (1-p)^2}.$$

The intuition behind the lemma is simple. A g signal is equivalent to two g signals and one b signal. So going from the premise of the lemma to the conclusion is analogous to receiving two additional g signals. But two g signals move a prior of $1/2$ to a posterior of $p^2/(p^2 + (1-p)^2)$. And of course, one can ignore the conditioning on h_{-i}.

Proof of Lemma 10: First, observe that

$$\begin{aligned}
&\Pr[s_i = g \cap h_{-i}] \\
&\quad = \Pr[s_i = g \cap h_{-i} \cap \theta = G] + \Pr[s_i = g \cap h_{-i} \cap \theta = B] \\
&\quad = \Pr[s_i = g \mid h_{-i} \cap \theta = G] \cdot \Pr[h_{-i} \cap \theta = G] \\
&\qquad\quad + \Pr[s_i = g \mid h_{-i} \cap \theta = B] \cdot \Pr[h_{-i} \cap \theta = B] \\
&\quad = p_i^g \cdot \Pr[h_{-i} \cap \theta = G] + (1 - p_i^b) \cdot \Pr[h_{-i} \cap \theta = B]
\end{aligned}$$

since the probability of obtaining signal $s_i = g$ or $s_i = b$ depends only on the state θ, and on h_{-i} also only through θ.

An analogous calculation with b replacing g yields

$$\Pr[s_i = b \cap h_{-i}] = (1 - p_i^g) \cdot \Pr[h_{-i} \cap \theta = G] + p_i^b \cdot \Pr[h_{-i} \cap \theta = B].$$

Next,

$$\begin{aligned}
&\Pr[\theta = G \mid h_{-i} \cap s_i = g] \\
&\quad = \frac{\Pr[s_i = g \mid \theta = G \cap h_{-i}] \cdot \Pr[\theta = G \cap h_{-i}]}{\Pr[s_i = g \cap h_{-i}]} \\
&\quad = \frac{p_i^g \cdot \Pr[\theta = G \cap h_{-i}]}{\Pr[s_i = g \cap h_{-i}]} \\
&\quad = \frac{p_i^g \cdot \Pr[\theta = G \cap h_{-i}]}{p_i^g \cdot \Pr[h_{-i} \cap \theta = G] + (1 - p_i^b) \cdot \Pr[h_{-i} \cap \theta = B]}.
\end{aligned} \tag{5}$$

Again, an analogous calculation with b replacing g yields

$$\begin{aligned}
&\Pr[\theta = G \mid h_{-i} \cap s_i = b] \\
&\quad = \frac{(1 - p_i^g) \cdot \Pr[\theta = G \cap h_{-i}]}{(1 - p_i^g) \cdot \Pr[h_{-i} \cap \theta = G] + p_i^b \cdot \Pr[h_{-i} \cap \theta = B]}.
\end{aligned}$$

The premise of the lemma,

$$\Pr\left[\theta = G \mid h_{-i} \cap s_i = b\right] \leq \frac{1}{2},$$

thus implies that

$$\frac{(1-p_i^g) \cdot \Pr\left[\theta = G \cap h_{-i}\right]}{(1-p_i^g) \cdot \Pr\left[h_{-i} \cap \theta = G\right] + p_i^b \cdot \Pr\left[h_{-i} \cap \theta = B\right]} \leq \frac{1}{2},$$

and so

$$\Pr\left[\theta = G \cap h_{-i}\right] \leq \frac{p_i^b}{1-p_i^g} \cdot \Pr\left[\theta = B \cap h_{-i}\right].$$

Plugging this into equation (5) yields

$$\begin{aligned}
&\Pr\left[\theta = G \mid h_{-i} \cap s_i = g\right] \\
&\quad = \frac{p_i^g \cdot \Pr\left[\theta = G \cap h_{-i}\right]}{p_i^g \cdot \Pr\left[h_{-i} \cap \theta = G\right] + (1-p_i^b) \cdot \Pr\left[h_{-i} \cap \theta = B\right]} \\
&\quad \leq \frac{p_i^g \cdot \frac{p_i^b}{1-p_i^g} \cdot \Pr\left[\theta = B \cap h_{-i}\right]}{p_i^g \cdot \frac{p_i^b}{1-p_i^g} \cdot \Pr\left[\theta = B \cap h_{-i}\right] + (1-p_i^b) \cdot \Pr\left[h_{-i} \cap \theta = B\right]} \\
&\quad = \frac{p_i^g p_i^b}{p_i^g p_i^b + (1-p_i^g)(1-p_i^b)} \\
&\quad \leq \frac{p^2}{p^2 + (1-p)^2},
\end{aligned}$$

where the last inequality follows from inequality (3). This completes the proof. ■

Finally, we prove the following lemma:

Lemma 11 *If a process is feasible, persuasive, and transparent then $\beta_D \leq \overline{\beta}(p)$.*

Proof of Lemma 11: By Lemma 9, there exists an agent i and an $h_{-i} \in \text{piv}_i$ such that

$$\Pr\left[\theta = G \mid h_{-i} \cap s_i = b\right] \leq \frac{1}{2}.$$

By Lemma 10, it thus holds that

$$\Pr\left[\theta = G \mid h_{-i} \cap s_i = g\right] \leq \frac{p^2}{p^2 + (1-p)^2}.$$

Additionally, observe that

$$\Pr\left[\theta = G \mid h_{-i} \cap s_i = s\right] = \Pr\left[\theta = G \mid h_{-i} \cap s_i = s \cap m_i = m\right]$$

for any $s \in S$ and $m \in M$, since, given the signal s_i, the posterior on θ does not depend on m_i.

Since $h_{-i} \in \mathrm{piv}_i$, there exists some $h_i \in A_i$ for which $r(h_i, h_{-i}) = y$. We will show that conditional on observing the history $h = (h_i, h_{-i})$, the DM's posterior on $\theta = G$ is at most $\overline{\beta}$:

$$\begin{aligned}
&\Pr\left[\theta = G \mid h\right] \\
&= \Pr\left[\theta = G \mid h_{-i} \cap h_i\right] \\
&= \Pr\left[\theta = G \cap s_i = g \mid h_{-i} \cap h_i\right] + \Pr\left[\theta = G \cap s_i = b \mid h_{-i} \cap h_i\right] \\
&= \Pr\left[\theta = G \mid h_{-i} \cap h_i \cap s_i = g\right] \cdot \Pr\left[s_i = g \mid h_{-i} \cap h_i\right] \\
&\quad + \Pr\left[\theta = G \mid h_{-i} \cap h_i \cap s_i = b\right] \cdot \Pr\left[s_i = b \mid h_{-i} \cap h_i\right] \\
&\leq \frac{p^2}{p^2 + (1-p)^2} \cdot \Pr\left[s_i = g \mid h_{-i} \cap h_i\right] + \frac{1}{2} \cdot \Pr\left[s_i = b \mid h_{-i} \cap h_i\right] \\
&\leq \frac{p^2}{p^2 + (1-p)^2} = \overline{\beta}.
\end{aligned}$$

Thus, $\Pr\left[\theta = G \mid h\right] \leq \overline{\beta}$. However, $r(h)$ should be an optimal decision for the DM on history h, which implies that $\beta_D \leq \overline{\beta}$ as claimed. ■

Lemma 11 immediately implies Theorem 8 for normal-form mechanisms.

General mechanisms

We now prove Lemma 8 for processes that involve extensive games. We begin by defining a leaf of a game—a history after which all histories are terminal.

Definition 8 (leaf) *A* leaf $\ell \in H$ *of a game* $\Gamma = (H, P)$ *is a history that satisfies* $\left\{\ell \circ h : h \in \times_{i \in P(\ell)} A_i(\ell)\right\} \subseteq Z$.

The *descendants* of a leaf ℓ, denoted $Z(\ell)$, are all terminal histories that pass through that leaf, namely $Z(\ell) = \{z \in Z : z = \ell \circ h \text{ for some } h \neq \emptyset\}$.

Fix a feasible, persuasive, and transparent process, and again assume without loss of generality that $\sigma(\cdot)$ has full support. Fix a leaf ℓ for which there exist $z, z' \in Z(\ell)$ such that $r(z) \neq r(z')$. If there is no such leaf, then this implies that for every leaf ℓ the outcome is always the same $o \in \mathcal{O}$ regardless of the agents' actions. But then we can prune the game tree, replacing each leaf ℓ and its descendants by a terminal node $z(\ell)$ for which $r(z(\ell)) = o$, and modifying the strategy profile accordingly. This leaves a new process, but still one that is feasible, persuasive, and transparent.

Repeating this procedure yields a feasible, persuasive, and transparent process $R = (\Gamma, d, \sigma)$ that has a leaf ℓ for which there exist $z, z' \in Z(\ell)$ such that $r(z) \neq r(z')$.

We will now condition on reaching leaf ℓ. Denote by $\pi_0 \overset{\text{def}}{=} \Pr[\theta = G|\ell]$, and note that $\pi_0 \in (0, 1)$ by our choice of ℓ. For each agent i, define $p_i^g = \Pr[s_i = g|\theta = G \cap \ell]$ and $p_i^b = \Pr[s_i = b|\theta = B \cap \ell]$. The following lemma states that inequality (3) is satisfied:

Lemma 12 *The accuracies* p_i^g *and* p_i^b *defined above satisfy*

$$\frac{p_i^g p_i^b}{p_i^g p_i^b + (1 - p_i^g)(1 - p_i^b)} \leq \frac{p^2}{p^2 + (1 - p)^2}.$$

Proof: Denote by ℓ_i actions of agent i that lead to leaf ℓ and by ℓ_{-i} the actions of agents other than i that lead to ℓ. With a slight abuse of notation, denote by ℓ_i also the event that agent i played the actions in ℓ_i (conditional on reaching preceding histories), and by ℓ_{-i} the event that agents other than i played the actions in ℓ_{-i} (conditional on reaching preceding histories). Note that ℓ_i is dependent on s_i, and ℓ_{-i} is dependent on θ, but that ℓ_i and ℓ_{-i} are conditionally independent (conditioning on s_i, θ, and reaching preceding histories). Then,

$$\begin{aligned}
p_i^g = \Pr\left[s_i = g | \theta = G \cap \ell\right] &= \frac{\Pr\left[s_i = g \cap \theta = G \cap \ell\right]}{\Pr\left[\theta = G \cap \ell\right]} \\
&= \frac{\Pr\left[\theta = G\right] \cdot \Pr\left[\ell_i | \theta = G\right] \cdot \Pr\left[s_i = g | \theta = G\right] \cdot \Pr\left[\ell_i | s_i = g\right]}{\Pr\left[\ell | \theta = G\right] \cdot \Pr\left[\theta = G\right]} \\
&= \frac{p \cdot \Pr\left[\ell_i | s_i = g\right]}{\Pr\left[\ell_i | \theta = G\right]} \\
&= \frac{p \cdot \Pr\left[\ell_i | s_i = g\right]}{p \cdot \Pr\left[\ell_i | s_i = g\right] + (1-p) \cdot \Pr\left[\ell_i | s_i = b\right]}.
\end{aligned}$$

Similarly,

$$p_i^b = \frac{p \cdot \Pr\left[\ell_i | s_i = b\right]}{p \cdot \Pr\left[\ell_i | s_i = b\right] + (1-p) \cdot \Pr\left[\ell_i | s_i = g\right]}.$$

Straightforward algebra confirms that

$$\frac{p_i^g p_i^b}{p_i^g p_i^b + (1-p_i^g)(1-p_i^b)} = \frac{p^2}{p^2 + (1-p)^2},$$

satisfying inequality (3). ■

Thus, we are now in the setting of the previous section – arbitrary prior π_0 and signal accuracies that satisfy inequality (3). Next, define a new process $R|_\ell = (\Gamma|_\ell, d|_\ell, \sigma|_\ell)$ that involves a normal-form mechanism. $R|_\ell$ is the process R, but conditioning on reaching leaf ℓ throughout. $R|_\ell$ is defined as follows:

- $\Gamma|_\ell = (H|_\ell, P|_\ell)$, where $H|_\ell = \{\epsilon\} \cup \{z \in H : z = \ell \circ h \text{ for some } h \neq \emptyset\}$ and $P|_\ell(\epsilon) = P(\ell)$;
- $d|_\ell$ is the transparent disclosure rule.
- $\sigma|_\ell(\epsilon) \equiv \sigma(\ell)$.

Furthermore, define the DM's decision rule $r|_\ell$ as $r|_\ell(z) = r(\ell \circ z)$ for all $z \in H \backslash \{\epsilon\}$. We now claim that $R|_\ell$ is feasible, persuasive, and transparent. It is feasible since R is feasible, and so σ is an equilibrium at every history h that is reached with positive probability, and in particular at ℓ. It is persuasive by our choice of ℓ: there exist $z, z' \in Z(\ell)$ such that $r(z) \neq r(z')$. Finally, it is transparent since $d|_\ell$ is the transparent disclosure rule. Applying Lemma 11 to the process $R|_\ell$ proves Theorem 8.

General mechanisms: small β_D and verifiable signals

In this section we show that transparency may be beneficial to the DM when $\beta_D \leq \overline{\beta}$ or when signals are verifiable, but when the committee may choose any process. It is straightforward to see that the results of Theorems 2, 3, and 4 generalize to arbitrary normal-form mechanisms. However, when processes can be arbitrary, possibly extensive form, it is much more difficult to determine which process is committee-optimal under transparency, and so we do not have more general versions of these theorems. Instead, we provide examples with fixed N and β_D for which transparency is strictly beneficial to the DM.

Small β_D

The following claim provides an example in which the DM's expected utility is strictly higher when transparency is required than when it is not, for small β_D.

Claim 4 *Fix $N = 2$, any $p > 1/2$, and $\beta_D \leq \overline{\beta}(p)$. Then the DM* strictly *prefers transparency.*

Proof: Observe first that $\beta_D < \overline{\beta}$. One optimal non-transparent process for the committee is a voting process that discloses the majority: the outcome is m_y whenever at least one of the committee members obtains a g signal. This process is persuasive by the assumption on β_D, and under this process the DM's expected utility is some

$\alpha \geq 0$. However, we claim that under transparency every optimal process of the committee gives the DM expected utility *strictly greater than* α. To see this, fix a transparent process that is optimal for the committee. When the signal profile $s = (g, g)$ the outcome must be y with probability 1, and when $s = (b, b)$ the outcome must be n with probability 1, and so the process yields utility at least α to the DM. If the process yields expected utility exactly α to the DM, then it implements opaque majority—namely, the outcome is y whenever $s = (g, b)$ or $s = (b, g)$, but never when $s = (b, b)$. Lemma 13 below shows that this is impossible, completing the proof of the claim. ■

Lemma 13 *When $N = 2$, no transparent process can implement opaque majority.*

Proof: Suppose towards a contradiction that some transparent process could implement opaque majority. Without loss of generality, assume that the process has no simultaneous moves—if there are, simply separate them into sequential moves. On a terminal node $z \in Z$ for which $r(z) = n$, the DM infers that $\Pr[s = (b, b)|z] = 1$. Thus, there is a history h on the path to z that is the first time he assigns probability 1 to that event. Suppose committee member 1 makes the last move leading to h. Before this move, the DM can infer that $\Pr[s_2 = b] = 1$. But this means that regardless of 1's action, the DM will choose outcome n. Thus, with positive probability the outcome is n when $s = (g, b)$. This is a contradiction. ■

Verifiable signals

In the following example, the DM's expected utility is strictly higher when transparency is required than when it is not, when signals are verifiable.

Example 3 Fix $N = 4$ and $\beta_D = \Pr[\theta = G \mid \#\{i : s_i = g\} \geq 3]$: that is, the DM prefers y only when at least 3 agents have a good signal. Note that $\beta_D > \overline{\beta}$.

First observe that without transparency, the unique committee-optimal feasible process is the one in which agents vote sincerely, then disclose m_y if and only if at least 3 agents voted y. This yields the committee their preferred outcome for every signal profile s (namely, outcome y whenever $\text{diff}(s) > 0$ and outcome n otherwise). However, it yields expected utility 0 to the DM, since his posterior on recommendation m_y is precisely β_D.

We claim that with transparency, however, the DM does get positive expected utility. This is straightforward if we consider only normal-form processes—in that case, in the committee-optimal feasible process one of the agents is non-pivotal, and the DM chooses y only if all other agents reveal v_i^g (this leads to a posterior of $p^3/(p^3 + (1-p)^3) > \beta_D$ and so positive expected utility). The difficulty arises from the possibility of extensive game-forms.

To see why it still holds, fix a committee-optimal feasible, transparent process, and suppose towards a contradiction that the DM gets expected utility 0. First observe that there can be at most 1 agent who is never pivotal, since if there are more then the maximum posterior on any terminal history would be $\overline{\beta}$. Next, for every $z \in Z$ with $r(z) = y$, there must exist some history h on the path to z and some agent i with $A_i(h) \neq \emptyset$, for whom $\Pr[s_i = g|z] \notin \{0, 1\}$. Otherwise, at history z the DM's posterior would either be $p^4/(p^4 + (1-p)^4)$—if he is certain about the signals of all agents—or $p^3/(p^3 + (1-p)^3)$, if he is certain about the signals of three agents and with no information about the signal of the fourth. Fix such an agent i and histories h and z, and note that because $\beta_D > \overline{\beta}$, at every terminal history $z \in Z$ at which the DM chooses y member i of *both* types wants outcome y—formally, the posterior $\Pr[\theta = G|z, s_i = b] > 1/2$.

Denote by $U_h^g \stackrel{\text{def}}{=} \text{supp}(\sigma_i(g, h))$, $U_h^b \stackrel{\text{def}}{=} \text{supp}(\sigma_i(b, h))$, and $U_h \stackrel{\text{def}}{=} U_h^g \cap U_h^b$. Note that $U_h \neq \emptyset$, since then actions of agent i at h completely disclose s_i, contradicting the assumption that $\Pr[s_i = g|z] \notin \{0, 1\}$. Consider now the following cases, each leading either to contradiction or to a trimming of the game tree. In the latter case, one can repeat the argument with the process that utilizies the trimmed tree. At some point it is no longer possible to trim the tree, which will lead to a contradiction in one of the cases below.

1. If $U_h \subsetneq U_h^b$, then define σ_i' to be the following deviation of agent i from σ: σ' is identical to σ except on history h and when $s_i = b$, with $\sigma_i'(b,h) = \sigma_i(b,h)|_{U_h}$, the distribution of $\sigma_i(b,g)$ restricted to support U_h. This is a profitable deviation: under σ_i there is positive probability that agent i chooses an action in $U_h^b \setminus U_h$, and these actions all lead to outcome n (since they lead to terminal histories z' with $\Pr\left[s_i = g | z'\right] = 0$, and so to a maximal posterior of $\overline{\beta}$ on the state being G). However, actions in U_h lead to outcome y with positive probability—in particular, they lead to history z with positive probability—and so to positive utility for the committee.

2. If $U_h \subsetneq U_h^g$ and the expected utility of the committee when agent i plays an action in $U_h^g \setminus U_h$ is strictly greater than when he plays an action in U_h, then let σ_i' be the profitable deviation that is identical to σ except that $\sigma_i'(g,h) = \sigma_i(g,h)|_{U_h^g \setminus U_h}$.

3. If $U_h \subsetneq U_h^g$ and the expected utility of the committee when agent i plays an action in U_h is strictly greater than when he plays an action in $U_h^g \setminus U_h$, then let σ_i' be the profitable deviation that is identical to σ except that $\sigma_i'(g,h) = \sigma_i(g,h)|_{U_h}$.

4. If $U_h \subsetneq U_h^g$ and the expected utility of the committee when agent i plays an action in U_h is equal to the committee's utility when he plays an action in $U_h^g \setminus U_h$, then one can trim the tree by removing from $A_i(h)$ the actions in $U_h^g \setminus U_h$.

5. Finally, if $U_h^b = U_h^g$, then agent i must be indifferent between all actions in U_h, for each one of his signals. But in that case, one can trim the tree so that only one such action remains. The action of agent i at h is then fixed, and so $A_i(h) = \emptyset$.

Generalization of Theorem 2

The following is the more general statement of Theorem 2, which allows the committee to choose any committee-optimal feasible process under opacity, and in particular to break ties in favor of n instead of y.

Theorem 6 *Fix any $p > 1/2$, $\beta_D \in (p, \overline{\beta}(p)]$,[16] and even N. Then for every committee-optimal, feasible process P_T under transparency there is a committee-optimal, feasible process P_O without transparency, such that the DM is indifferent between P_T and P_O. However, there exists a committee-optimal, feasible process P'_O without transparency such that the DM strictly prefers every committee-optimal, feasible process under transparency to P'_O.*

Two natural ways to compare processes are the worst case and best case: in the first, the DM compares the worst committee-optimal feasible processes under transparency and under non-transparency, and in the second he compares the best such processes. Two implications of Theorem 6 are that from a worst-case perspective the DM strictly prefers transparency, whereas from a best-case perspective he is indifferent.

Proof of Theorem 6: A committee-optimal process will lead to one in which the outcome is y whenever s satisfies $\text{diff}(s) > 0$ and to outcome n when s satisfies $\text{diff}(s) < 0$. Since N is even, there are many such processes, and they differ in whether the outcome on signal profile s with $\text{diff}(s) = 0$ yields outcome y or n.

Now, although the committee is indifferent amongst all these processes, the DM is not. He obtains negative expected utility when the profile satisfies $\text{diff}(s) = 0$, and would thus like to minimize the number of such profiles that lead to outcome y. From his perspective, the worst committee-optimal process is the one that has $d(s) = m_y$ if and only if $\text{diff}(s) \geq 0$ (i.e., opaque, sincere majority). The best process is the one in which $d(s) = m_y$ if and only if $\text{diff}(s) > 0$.

[16] If $\beta_D \in (1/2, p]$ then there are more committee-optimal feasible processes under transparency. A similar result can be shown in that case as well.

Consider now the case of transparency. First, note that the sincere strategy profile is feasible and persuasive here. Furthermore, this strategy profile is committee-optimal, as the committee obtains their desired outcome on every signal profile (outcome y if $\text{diff}(s) > 0$ and outcome n if $\text{diff}(s) < 0$). This implies that the symmetric, mixed strategy profile cannot be committee-optimal, since that profile inevitably leads to some incorrect outcome with positive probability.

Furthermore, the asymmetric strategy profiles à la McLennan (1998) are also not committee-optimal here: if some agent i is non-pivotal, then the signal profile in which $N/2$ agents other than i have good signals leads to outcome n, but the committee desires outcome y if $s_i = g$. The only other committee-optimal, feasible processes under transparency are ones where $N - 1$ agents vote sincerely, and one agent mixes (voting y also on signal $s_i = b$ with some probability). This is feasible and committee-optimal as long as he does not mix too much, so that the posterior when i as well as $N/2$ agents other than i vote y will be at least β_D.

Under opacity, the asymmetric equilibria are feasible (with the one agent either non-pivotal or mixing), as is the sincere profile. The DM is indifferent between the sincere profile and the committee-optimal process in which $d(s) = m_y$ if and only if $\text{diff}(s) > 0$. He is also indifferent between the asymmetric profiles with one agent mixing and the committee-optimal process in which $d(s) = m_y$ if and only if at least $N/2 + 1$ agents vote y. However, the committee-optimal processes with the asymmetric equilibrium à la McLennan (1998), in which $d(s) = m_y$ if and only if at least $N/2$ of the pivotal agents vote y are strictly worse for the DM than any of the profiles under transparency: they lead to outcome y whenever $N/2$ pivotal agents vote y, but the non-pivotal agent has a bad signal, which arises from signal profile s with $\text{diff}(s) = 0$. ■