



Note

Should straw polls be banned? ☆

S. Nageeb Ali ^{a,*}, J. Aislinn Bohren ^b^a Department of Economics, Pennsylvania State University, United States of America^b Department of Economics, University of Pennsylvania, United States of America

ARTICLE INFO

Article history:

Received 19 March 2019

Available online 17 September 2019

JEL classification:

D7

D8

Keywords:

Information aggregation

Committees

Deliberation

Collusion

ABSTRACT

A Principal appoints a committee of partially informed experts to choose a policy. The experts' preferences are aligned with each other but conflict with hers. We study whether she gains from banning committee members from communicating or "deliberating" before voting. Our main result is that if the committee plays its preferred equilibrium and the Principal uses a threshold voting rule, then she does not gain from banning deliberation. We show using examples how she can gain if she can choose the equilibrium played by the committee, or use a non-anonymous or non-monotone voting rule.

© 2019 Elsevier Inc. All rights reserved.

1. Introduction

A Dean is considering whether to offer a recruiting slot to the economics department. She would like to hire good professors, not hire bad ones, and incurs a particular configuration of costs from making type 1 and type 2 errors. She delegates her choice to a committee of experts who assess the costs of those errors differently. For example, the committee may not find it so costly to hire a bad economist, or conversely, the committee may be less willing to expand the department. A common practice is for committee members to communicate to each other, perhaps through a straw poll, to aggregate their information before voting and then to present a "united front" to the Dean. These straw polls may be difficult to prevent, but one might envision that a Dean might take steps to limit these private communications. For example, she may insist that she (or a representative of hers) attend all committee discussions.

Our paper studies the degree to which the Dean should try to prevent committee deliberations when there is a conflict of interest between her and the committee. We believe that conflicts of this form arise generally, where a Principal delegates a choice to a committee of experts whose preferences differ from her own. The prior literature, particularly Coughlan (2000) and Gerardi and Yariv (2007), has emphasized how the possibility for communication has an important effect on voting outcomes. This paper poses the following question: *even if the Principal could ban communication and straw polls, would she benefit from doing so?*

We pose this question in a standard information aggregation framework, described formally in Section 2. There is an unknown state of the world, and agents appointed to a committee obtain signals of that state. Each committee member

☆ We thank Mehmet Ekmekci, Ben Golub, Ronen Gradwohl, Emir Kamenica, Navin Kartik, Max Mihm, Lucas Siga, Joel Sobel, and various seminar participants for helpful comments. Garima Singhal provided expert research assistance. Ali gratefully acknowledges financial support from the NSF (SES-1530639).

* Corresponding author.

E-mail addresses: nageeb@psu.edu (S.N. Ali), abohren@gmail.com (J.A. Bohren).

makes a binary recommendation that is aggregated by a voting rule that selects a binary action, i.e. whether to hire a job candidate. The Principal and the committee members have completely aligned preferences when the state is known, but *ex interim* there is imperfect information and signal profiles exist at which the Principal and committee members prefer different actions. All committee members are likeminded.

The Principal can choose the committee's size, the voting rule, and whether to permit committee members to *deliberate*—e.g., take a straw poll or secretly communicate—before they vote. Straw polls or secret communications between committee members permit the committee to perfectly coordinate their recommendations based on the joint signal profile and implement their (*ex interim*) preferred outcome. This is a form of collusion that arises in a one-shot setting when committee members can communicate. In contrast, banning deliberation prevents the committee from colluding in this manner. Nevertheless, committee members may still *tacitly collude* by playing their favorite equilibrium in the voting game without communication. Our main result speaks to whether the Principal gains from banning deliberation given this possibility for tacit collusion. We prove the following result in Section 3:

Proposition. *If a committee adopts its most preferred equilibrium, the highest payoff that the Principal achieves from banning deliberation is no more than her payoff from permitting deliberation.*

Here is the logic of our result: generically, in a setting without communication, the committee-optimal equilibrium is in pure strategies. These pure strategies involve some voters fully revealing their signals with their votes and others voting completely uninformatively (i.e. independent of their signals). The voting outcome is then isomorphic to an election in which the uninformative voters are completely absent. Optimizing the committee size and threshold voting rule with this issue in mind leads a Principal to choose a committee design where all voters are voting informatively. Such a committee implements the same outcome as it would if it could have a (truthful) straw poll. Therefore, once a Principal faces the possibility of tacit collusion, she may as well allow committee members to communicate freely.

We view this result as one rationale for the prevalence of practices that allow committee members to secretly communicate or take straw polls prior to voting. In certain instances (e.g., a jury), one may view the Principal as sharing committee members' preferences, and so the use of deliberative practices appears natural. But in other cases where there is a clear conflict of interest, such as that of hiring and budgetary decisions, it may be puzzling as to why we see an absence of measures that impede committee members from secretly communicating with each other. One response is that perhaps it is technologically costly to prohibit communication; another is that perhaps the binary voting environment coarsely filters the potentially rich information of experts, and so communication is needed. To isolate a new tension – namely, that the Principal cannot choose the equilibrium that is played – we preclude both of these rationales. We assume that it is technologically feasible to prohibit deliberation and that voters obtain identically distributed binary signals, so that there is no need for a finer communication language. Our result suggests that this new tension provides an additional reason for allowing communication: if committee members can tacitly collude, the Principal does not gain from banning deliberation.

Section 4 further explores this result by describing conditions under which it fails. First, the result pertains to committees that can coordinate (perhaps through pre-play communication) on their preferred equilibrium. This is an issue that applies broadly and is studied in work on collusion and, more recently, adversarial equilibrium selection in mechanism design (see the Related Literature below for a discussion). In our context, we view the possibility for tacit collusion to emerge more realistically in small committees, say in organizations, where committee members can communicate and coordinate about their preferred way to play the game. In certain cases, committees may be unable to coordinate on their preferred equilibrium. If the Principal is the one who selects the equilibrium, then we show that there are cases where she benefits from banning deliberation. Alternatively, if committee members cannot coordinate on asymmetric strategies and must instead play symmetric equilibria, there are parameter values under which the Principal gains from banning deliberation.

Second, even if the committee coordinates on its favored equilibrium, we show that the result can fail if the Principal can choose voting rules that are non-monotone or non-anonymous. A non-monotone or non-anonymous voting rule gives the Principal a way to “sway the committee” because it allows the Principal to pool multiple events under which a committee member is pivotal. In such cases, the Principal is better off when the committee members cannot reveal their signals to each other.

Finally, our result applies when the Principal is optimally choosing *both* the committee size and its voting rule. However, if she is restricted to a fixed committee size, she may gain from banning deliberation. Thus, the equivalence identified in Proposition 1—of her payoff from banning and permitting deliberation—holds at her optimal committee size and voting rule but not necessarily away from it.

Related literature The tension that we study—namely, that even if a designer is involved in designing the rules of the game, she may be unable to force agents to play her preferred equilibrium of that game—has been posed broadly in mechanism and organizational design (Laffont and Martimort, 2000; Mookherjee, 2006). Generally, in a mechanism, agents may collude in reports made to a Principal and in participation decisions. One notion of collusion, typically described as “overt collusion,” is where the agents collude at an *ex interim* stage (after observing their types). Another notion, typically described as “tacit collusion,” is where the agents collude at an *ex ante* stage on how to behave in the mechanism. In our context, “overt collusion” corresponds to deliberation, where committee members share all information with each other before sending reports to the Principal. By contrast, “tacit collusion” corresponds to committee members being unable to share information

with each other, but being able to coordinate their play on their most preferred equilibrium. We identify an equivalence from the Principal's perspective: after designing the committee optimally, she is affected identically by both forms of collusion and incurs no additional cost from overt collusion above and beyond tacit collusion. Thus, in our context, if voters can coordinate on how to play the game at an ex ante stage, there is no harm in letting them communicate with each other.

The concern of tacit collusion relates to the recent interest in “adversarial equilibrium selection” in mechanism design (Hoshino, 2017; Inostroza and Pavan, 2018) where the Principal evaluates mechanisms by the worst equilibrium or rationalizable outcome, as well as the classical study of “full implementation” (Maskin, 1999) where the Principal wishes to guarantee certain outcomes across all equilibria. These analyses and ours share the concern that an institutional designer may not be able to perfectly coordinate the play of agents after designing the institution. Rather than thinking about her worst equilibrium outcome or the entire set of equilibrium outcomes, the Principal here evaluates each design according to the most preferred equilibrium of the agents for that design.

We build on studies of committee decision-making with common interests, particularly McLennan (1998). Our focus is on the best equilibrium (from the committee's standpoint) with and without communication, and with such behavior in mind, how the Principal forms a committee and selects a voting rule. Our message complements that of Coughlan (2000) and Gerardi and Yariv (2007) who establish equivalence results across threshold rules when committee members can communicate with each other before voting. Austen-Smith and Feddersen (2006) study when truthful communication among committee members is incentive-compatible, given the potential for disagreements within the committee. In such cases, deliberation may not be so costly to the Principal. By contrast, we show that even if committee members share perfectly aligned preferences, the Principal does not gain from banning deliberation.

Our paper also connects to work on communication with multiple senders. The study of cheap talk with multiple senders (Gilligan and Krehbiel, 1987; Krishna and Morgan, 2001; Battaglini, 2002) focuses on how the conflict of preferences within a committee of experts can be useful to elicit information. By contrast, we study the conflict between a Principal and a committee of like-minded experts where the Principal has already committed to a voting rule. A series of papers study a setting similar to ours but focus on behavior when the Principal cannot commit to a voting rule, and hence, messages are payoff-irrelevant; see Wolinsky (2002), Morgan and Stocken (2008), Battaglini (2017), and Gradwohl and Feddersen (2018). In these papers, truthful communication is impeded by each agent anticipating that her report matters only on the margin for the Principal, and given the conflict of interest, agents are unwilling to report truthfully on that margin. Battaglini (2017) illustrates how allowing individuals to share communication—i.e., overtly collude—using social media can facilitate information aggregation and make both the agents and the Principal better off. Gradwohl and Feddersen (2018) consider the question of transparency, namely whether a Principal who does not commit to a voting rule should permit herself to observe all communication within the committee. Because transparency distorts communication within the committee, they find that the Principal is better off without transparency.¹ We view this issue as being related and complementary to that which we consider, namely, whether a Principal who can commit to a voting rule should allow the committee members to deliberate beforehand.

2. Model

2.1. Environment

Payoffs A Principal is choosing to accept or reject an alternative and delegates that choice to a committee of experts. The quality of the alternative is either low, $\omega = L$, or high, $\omega = H$. All players share a common prior that attributes probability $\pi \in (0, 1)$ to the alternative being of high quality.

The Principal and committee share the same ordinal preferences in that they both prefer to accept high quality alternatives and reject low quality alternatives, but they differ in the intensity of their preferences. All players receive a payoff of zero from accepting a high quality alternative or rejecting a low quality alternative. The Principal's payoff from accepting a low quality alternative is $-q_P$ and her payoff from rejecting a high quality alternative is $-(1 - q_P)$. The analogous payoffs for a committee member are $-q_C$ and $-(1 - q_C)$, respectively (see Fig. 1).² We assume that $1 > q_P > q_C > 0$. Because the Principal is worse off from accepting the low quality alternative, for every interior belief about the quality of the alternative, the committee is more willing than the Principal to accept the alternative.

	Low	High
Reject	0, 0	$-(1 - q_P), -(1 - q_C)$
Accept	$-q_P, -q_C$	0, 0

Fig. 1. The Principal's and committee members' payoffs.

¹ Feddersen and Gradwohl (2018) show that this conclusion may be reversed if there are multiple receivers with different preferences.

² This parameterization of the payoff matrix is without loss of generality when the Principal and the committee have the same state-dependent ordinal preferences.

Committee design If chosen to be on the committee, expert i obtains private information $s_i \in \{l, h\}$ about the quality of the alternative, where $\gamma^l = \Pr(s^i = l|L)$ and $\gamma^h = \Pr(s^i = h|H)$ denote the precision of his signal in each state. Without loss of generality, we assume that $\gamma^l \geq 1/2$ and $\gamma^h \geq 1/2$. No signal is perfectly informative ($\gamma^l < 1$ and $\gamma^h < 1$) and conditional on ω , signals are independent across experts. Let N be the number of available experts, which we assume to be finite. The Principal chooses a committee size $C \leq N$. Let \mathcal{C} denote the set of committee members and $s \equiv (s_i)_{i \in \mathcal{C}}$ denote the committee signal profile.

Each committee member i simultaneously votes to accept or reject, $v_i \in \{A, R\}$. Denote the committee voting profile by $v \equiv (v_i)_{i \in \mathcal{C}}$. The Principal commits to a voting rule $f : \{A, R\}^C \rightarrow \{0, 1\}$ to aggregate votes, where an alternative is accepted if $f(\cdot) = 1$ and otherwise is rejected. We refer to (f, C) as the *committee design*. A voting rule f is *monotone* if $f(A, v_{-i}) = 1$ whenever $f(R, v_{-i}) = 1$; *anonymous* if $f(v) = f(v')$ for all v' a permutation of v ; and *non-constant* if there exist v and v' in $\{A, R\}^C$ such that $f(v) \neq f(v')$. We refer to monotone and anonymous voting rules as *threshold rules*, given that each such rule is equivalent to accepting the alternative whenever the number of accept votes exceeds some threshold k in $\{0, \dots, C + 1\}$. Let \mathcal{F}_C be the set of *allowable* voting rules for a committee of size C , and let $\mathcal{F} \equiv \{\mathcal{F}_C\}_{C \leq N}$ denote this set for each feasible committee size. We take \mathcal{F} to be a primitive of our model and assume that \mathcal{F}_C is a subset of the set of non-constant voting rules for each $C \leq N$. The allowable set of voting rules may contain additional restrictions such as, for example, a restriction to threshold rules.

We compare behavior under two communication protocols for the committee: (i) no deliberation, in which the committee cannot share private information before voting, and (ii) deliberation, in which the committee can freely communicate before voting.

No deliberation Suppose that the committee cannot share private information in any way before voting. In other words, there is *no deliberation*. A strategy for committee member i is a mapping $\sigma_i : \{l, h\} \rightarrow \Delta\{A, R\}$ from his private signal to (a distribution over) his vote. To simplify notation, let $\sigma_i(s_i)$ denote the probability with which he votes to accept following signal s_i . Following the literature, we say that committee member i votes *fully informatively* if $\sigma_i(h) = 1$ and $\sigma_i(l) = 0$.

Fixing a committee \mathcal{C} of size C , let $\sigma \equiv (\sigma_i)_{i \in \mathcal{C}}$ denote a strategy profile and Σ_C denote the set of strategy profiles. Since committee members share common interests, they each earn an identical *ex ante* expected payoff $W(\sigma; f, C)$ from strategy profile σ when the Principal uses design (f, C) . Let $\Sigma^*(f, C)$ denote the set of equilibrium strategy profiles for design (f, C) in the game without deliberation.

Definition 1. A committee **tacitly colludes** if it plays an equilibrium in the game without deliberation that maximizes the committee's payoff, $\sigma^* \in \arg \max_{\sigma \in \Sigma^*(f, C)} W(\sigma; f, C)$.

Tacit collusion corresponds to the committee behaving in a way that is committee-optimal (among equilibria). If the committee has multiple optimal equilibria, we assume that it resolves indifference in favor of an equilibrium that is optimal for the Principal. Let $\Sigma_T^*(f, C)$ denote the set of tacit collusive equilibria.

The Principal earns *ex ante* expected payoff $V(\sigma; f, C)$ from strategy profile σ when she chooses design (f, C) . Therefore, when the committee engages in tacit collusion, the Principal's expected payoff is $V_T(f, C) \equiv V(\sigma_T^*; f, C)$ for any $\sigma_T^* \in \Sigma_T^*(f, C)$. Given a set of allowable voting rules \mathcal{F} , the best payoff that the Principal can achieve in any committee design under tacit collusion is

$$V_T^*(\mathcal{F}) = \max_{C \leq N} \max_{f \in \mathcal{F}_C} V_T(f, C). \tag{1}$$

Deliberation Suppose the committee can freely communicate their private information with each other before voting, in other words, *deliberate*. One may envision a range of communication protocols, but given the simplicity of the environment that we consider, it suffices to study the following simple protocol. Suppose that, as in a *straw poll*, each committee member simultaneously sends the message h or l . Messages are publicly observed by all members of the committee, but not the Principal. Based on these messages, committee members vote on whether to accept the alternative. In this game, we define overt collusion as follows.

Definition 2. A committee engages in **overt collusion** if it selects its most preferred equilibrium in the game with deliberation.

Let $V_O(f, C)$ be the Principal's expected payoff from using design (f, C) when the committee engages in overt collusion. Given a set of allowable voting rules \mathcal{F} , the best payoff that the Principal can achieve in any committee design is

$$V_O^*(\mathcal{F}) = \max_{C \leq N} \max_{f \in \mathcal{F}_C} V_O(f, C). \tag{2}$$

2.2. Discussion of model

We make a number of assumptions to simplify and sharpen the analysis in this stylized model of deliberation and voting. Importantly, committee members share pure common values, and information is binary. We view the pure common values environment as being appropriate to elucidate whether a Principal may wish to allow deliberation, even if all committee members truthfully reveal information to each other and unanimously agree on a decision after doing so. If committee members had different preferences from each other, they might not truthfully reveal information to each other (Austen-Smith and Feddersen, 2006; Gerardi and Yariv, 2007), and thus be less able to collude even with communication.³

We restrict attention to identically distributed binary signals to isolate the effect of tacit collusion. An important complementary motive for permitting deliberation is that the binary action of voting may be too coarse to appropriately reflect the richness of voters' information. In a model with a richer or heterogeneous information structure, the power of tacit collusion would then be confounded with the gains from allowing a more expressive language through deliberation. To isolate the particular force that we study and obviate this orthogonal motive for deliberation, we assume that the action space of voting is as rich as the information space.

3. Banning deliberation: an irrelevance result

3.1. Some preliminary results

Before proceeding to our main result, we prove some preliminary results about overt and tacit collusion. These results are used both in our main result and the examples in Section 4.

Overt collusion Let us turn first to the case of overt collusion. The committee faces a common interest game and so the committee-optimal strategy is to truthfully reveal private information and then vote unanimously for the committee-preferred action given this information. In other words, for any committee design (f, C) , the committee plays a voting profile v such that $f(v) = 1$ if and only if, given signal profile s , $P(\omega = H|s) \geq q_C$. The outcome, and therefore, the Principal's value from overt collusion, is the same for all (non-constant) voting rules.

Lemma 0. For every committee size C , $V_O(f, C) = V_O(f', C)$ for all $f, f' \in \mathcal{F}_C$.

We omit a formal proof, given the logic explicated in the preceding discussion.

Tacit collusion Our novel results are in the characterization of tacit collusion. Lemma 1 establishes that there exists a committee-optimal pure strategy equilibrium for any committee design. Moreover, for generic parameters—that is, for all but a Lebesgue measure-zero set of parameters $(\pi, q_C, \gamma^l, \gamma^h)$ —such an equilibrium is unique within the class of all strategy profiles (up to a re-ordering of player labels). Lemma 2 shows that once the committee is selecting its preferred equilibrium, it is without loss for the Principal to restrict attention to committee designs in which fully informative voting is the committee-optimal equilibrium.⁴ Both of these results apply to any set of allowable voting rules \mathcal{F}_C .

Lemma 1. For any committee design (f, C) , there exists a committee-optimal equilibrium that is in pure strategies. Generically, in every committee-optimal equilibrium, all members who are pivotal with positive probability play a pure strategy.

Proof. Given a committee (f, C) , a strategy profile σ is committee-optimal if $W(\sigma; f, C) \geq W(\sigma'; f, C)$ for all $\sigma' \in \Sigma_C$. Such a strategy profile exists because W is continuous in $(\sigma_i(l), \sigma_i(h))$ for each i . McLennan (1998) establishes that every committee-optimal strategy profile is a Nash equilibrium.

First we show that there exists a committee-optimal strategy profile in pure strategies. Suppose that there is a committee-optimal strategy profile σ in which member i follows a strictly mixed strategy. We construct a pure strategy profile that achieves the same payoff. Let us denote the four different pure strategies for member i : $\tilde{\sigma}_{AA}$ is the strategy that involves voting to accept regardless of signal; $\tilde{\sigma}_{AR}$ is the strategy of informative voting; $\tilde{\sigma}_{RA}$ is the strategy of inverted voting, $\sigma_i(h) = 0$ and $\sigma_i(l) = 1$; and $\tilde{\sigma}_{RR}$ is the strategy that involves voting to reject regardless of signal. If member i is playing a mixed strategy, then he is randomizing between at least two of these pure strategies, $\tilde{\sigma}$ and $\tilde{\sigma}'$. For member i to be willing to randomize, it must be that $W(\tilde{\sigma}, \sigma_{-i}; f, C) = W(\tilde{\sigma}', \sigma_{-i}; f, C)$. But then the strategy $(\tilde{\sigma}, \sigma_{-i})$ is also committee-optimal. Iterating through each committee member who randomizes results in a pure strategy profile $\tilde{\sigma}$ that is payoff-equivalent to σ , and therefore also a committee-optimal strategy profile. But then $\tilde{\sigma}$ is also an equilibrium strategy profile, and therefore there also exists a committee-optimal strategy profile in which all members follow a pure strategy.

We prove the second statement of Lemma 1 in Appendix A. \square

³ In a different context where players have misaligned incentives—namely that of auctions—Agranov and Yariv (2018) show that communication alone results in little collusion.

⁴ Persico (2003) establishes an analogue for Lemma 2 in the context of threshold rules (i.e. monotone and anonymous voting rules), restricting attention to “monotone” pure-strategy equilibria. We show that a more general conclusion holds across voting rules and equilibria.

Lemma 2. *Generically, for any committee design (f, C) , there exists a committee design (f', C') with $C' \leq C$ in which the committee-optimal equilibrium is for all committee members to vote informatively and all committee members are pivotal with positive probability, such that $V_T(f, C) = V_T(f', C')$. Further, if f is a threshold rule, then f' is a threshold rule.*

Proof. Fixing committee design (f, C) , by Lemma 1, there exists a committee-optimal strategy profile in pure strategies. Let σ^* denote this strategy profile. Recall that this implies that σ^* is a committee-optimal equilibrium for (f, C) . This strategy yields ex ante expected payoff $W(\sigma^*; f, C)$ to the committee and $V_T(f, C)$ to the Principal. In a pure strategy profile, a committee member either (i) always votes to accept; (ii) always votes to reject or (iii) votes informatively or inverted. When the committee plays σ^* , let $\mathcal{I} \subset C$ denote the subset of committee members who vote informatively or inverted and are pivotal with positive probability, and let $\mathcal{U} \equiv C \setminus \mathcal{I}$ denote the set of voters who always accept, always reject or are not pivotal with positive probability. For any subset of committee members $C' \subset C$ and strategy profile σ , let $\sigma_{C'} \equiv (\sigma_i)_{i \in C'}$ denote the strategy profile and $v_{C'} \equiv (v_i)_{i \in C'}$ denote the voting profile for members in C' .

First, we show that when some committee members who play σ^* vote uninformatively or are pivotal with probability zero, there exists a design and a strategy profile in which no members vote uninformatively and all members are pivotal with positive probability that yields an equivalent expected payoff to the Principal and the committee as design (f, C) and strategy σ^* . Define a voting rule $f'' : \{A, R\}^{|\mathcal{I}|} \rightarrow \{0, 1\}$ such that f'' implements the same outcome as voting rule f for any vote profile $v_{\mathcal{I}}$ played by informative and inverted voters who are pivotal with positive probability, given uninformative and non-pivotal voters select vote profile $v_{\mathcal{U}}$ corresponding to uninformative voters' constant strategy in σ^* and an arbitrary vote profile for non-pivotal voters. In other words, $f''(v_{\mathcal{I}}) = f(v_{\mathcal{I}}, v_{\mathcal{U}})$ for all $v_{\mathcal{I}} \in \{A, R\}^{|\mathcal{I}|}$. Therefore, design $(f'', |\mathcal{I}|)$ and strategy profile $\sigma_{\mathcal{I}}^*$ yield an equivalent payoff to the Principal and the committee as design (f, C) and strategy profile σ^* , $V(\sigma_{\mathcal{I}}^*; f'', |\mathcal{I}|) = V(\sigma^*; f, C)$ and $W(\sigma_{\mathcal{I}}^*; f'', |\mathcal{I}|) = W(\sigma^*; f, C)$.

Next, we show that $\sigma_{\mathcal{I}}^*$ is a committee-optimal strategy profile for design $(f'', |\mathcal{I}|)$. Suppose there exists a $\sigma'_{\mathcal{I}} \in \Sigma_{|\mathcal{I}|}$ such that $W(\sigma'_{\mathcal{I}}; f'', |\mathcal{I}|) > W(\sigma_{\mathcal{I}}^*; f'', |\mathcal{I}|)$. The set of vote profiles that occur with positive probability under $\sigma'_{\mathcal{I}}$ is a subset of the set of vote profiles that occur with positive probability under $\sigma_{\mathcal{I}}^*$, since all feasible vote profiles $v_{\mathcal{I}} \in \{A, R\}^{|\mathcal{I}|}$ occur with positive probability under $\sigma_{\mathcal{I}}^*$. By extension, this same relationship holds for strategy profiles $\sigma' = (\sigma'_{\mathcal{I}}, \sigma_{\mathcal{U}}^*)$ and σ^* . Therefore, strategy profile σ' yields payoff $W(\sigma'; f, C) = W(\sigma'_{\mathcal{I}}; f'', |\mathcal{I}|)$, since by definition f'' implements the same outcome as f on all vote profiles that occur with positive probability under σ^* . This implies $W(\sigma'; f, C) > W(\sigma_{\mathcal{I}}^*; f'', |\mathcal{I}|) = W(\sigma^*; f, C)$, which is a contradiction since σ^* is a committee-optimal strategy profile for design (f, C) . Therefore, $\sigma_{\mathcal{I}}^*$ is a committee-optimal strategy profile for design $(f'', |\mathcal{I}|)$. This implies that it is also a committee-optimal equilibrium for $(f'', |\mathcal{I}|)$.

By definition, design (f, C) and strategy σ^* result in the same outcome as design $(f'', |\mathcal{I}|)$ and strategy $\sigma_{\mathcal{I}}^*$ for any realized profile of private signals. Therefore, when the committee plays a committee-optimal equilibrium, the Principal's value of committee design (f, C) is equal to the value of design $(f'', |\mathcal{I}|)$.

Finally, we show that if some committee members vote inverted in strategy profile $\sigma_{\mathcal{I}}^*$ under design $(f'', |\mathcal{I}|)$, there exists a design of equivalent value to the Principal in which all members vote informatively in the committee-optimal strategy profile. Given design $(f'', |\mathcal{I}|)$ and strategy profile $\sigma_{\mathcal{I}}^*$, let $\mathcal{J} \subset \mathcal{I}$ be the set of members who play the inverted voting strategy and let $\alpha_i = \{v_{-i} | f(A, v_{-i}) \neq f(R, v_{-i})\}$ be the set of voting profiles at which member i is pivotal. Member i is willing to play the inverted voting strategy if he prefers to reject when he observes a high signal,

$$\sum_{a_{-i} \in \alpha_i} P(v_{-i} | h) (2f''(A, v_{-i}) - 1) (\mu(h, v_{-i}) - q_C) \leq 0$$

and he prefers to accept when he observes a low signal,

$$\sum_{a_{-i} \in \alpha_i} P(v_{-i} | l) (2f''(A, v_{-i}) - 1) (\mu(l, v_{-i}) - q_C) \geq 0$$

where $\mu(s, v_{-i})$ the posterior belief that the state is H at action profile v_{-i} when other players play strategy $\sigma_{\mathcal{I}}^*$ and member i observes signal s , and $P(v_{-i} | s)$ is the probability that member i believes that other players play v_{-i} when he observes signal s . Define a voting rule $f' : \{A, R\}^{|\mathcal{I}|} \rightarrow \{0, 1\}$ such that for $i \in \mathcal{J}$ and $v_{-i} \in \alpha_i$, $f'(A, v_{-i}) = f''(R, v_{-i})$ and $f'(R, v_{-i}) = f''(A, v_{-i})$. In other words, at the action profiles at which inverted voters are pivotal, invert the voting rule. Then under f' , member i is willing to play an informative voting strategy, as this reverses the inequalities in the above incentive constraints. For any signal profile s , the same decision is taken under f' and f'' . Therefore, changing f'' does not interfere with the incentives of other members and the Principal's value of committee design $(f', |\mathcal{I}|)$ is equal to the value of design $(f'', |\mathcal{I}|)$. In design $(f', |\mathcal{I}|)$, all committee members vote informatively in the committee-optimal equilibrium.

To prove the final statement, suppose that f is a threshold rule in which an alternative is accepted if at least k committee members vote to accept. Under a threshold rule, conditional on being pivotal, the payoff from accepting with a positive signal is higher than that from accepting with a negative signal and the payoff from rejecting with a negative signal is higher than the payoff from rejecting with a positive signal, regardless of the behavior of others. Hence, there is no equilibrium strategy profile in which a committee member is pivotal with positive probability and adopts an inverted voting strategy. Therefore, given committee-optimal strategy σ^* , $\sigma_{\mathcal{I}}^*$ involves all committee members voting informatively.

It remains to show that when f is a threshold rule, then voting rule f'' constructed above is also a threshold rule. Let $\mathcal{U}_A \subset \mathcal{U}$ denote the set of committee members who vote to always accept when playing σ^* . Then the threshold rule in which an alternative is accepted if at least $k'' \equiv \max\{0, k - |\mathcal{U}_A|\}$ out of $|\mathcal{I}|$ members vote to accept will implement the same outcome as threshold rule f in a committee of size C . Therefore, f'' is a threshold rule with cut-off k'' . From here, the outlined above established that $(f'', |\mathcal{I}|)$ and strategy profile $\sigma_{\mathcal{I}}^*$ yield an equivalent payoff to the Principal as (f, C) and σ^* , and that $\sigma_{\mathcal{I}}^*$ is a committee-optimal equilibrium for $(f'', |\mathcal{I}|)$. Given that we do not need to consider inverted voting for threshold rules, $f' = f''$. \square

3.2. The main result

The Principal's design problem involves choosing whether to delegate to a committee, a design (f, C) , and whether to allow the committee to deliberate, for example, in the form of a straw poll. We study how the Principal makes these choices when she anticipates that if deliberation is banned, then the committee will tacitly collude and play its preferred equilibrium.

Proposition 1. *If the Principal must use a threshold voting rule and the committee tacitly colludes, then the Principal does not gain from banning deliberation: $V_T^*(\mathcal{F}) = V_O^*(\mathcal{F})$ when \mathcal{F} is the set of threshold voting rules.*

Our result offers a strategic rationale for the ubiquity of deliberation. Whenever the Principal uses a threshold voting rule—i.e. an anonymous and monotone rule—and committee members can coordinate on their preferred equilibrium, the Principal may as well permit committee members to take a straw poll.

Proof of Proposition 1. Suppose \mathcal{F} is the set of threshold voting rules. By Lemma 1, for any threshold rule, there exists a committee-optimal pure strategy equilibrium. If the committee selects its preferred equilibrium, by Lemma 2, it is without loss for the Principal to restrict attention to threshold rules in which fully informative voting is a committee-optimal equilibrium. Therefore, under tacit collusion, there exists a committee-design (f^*, C^*) that is optimal within the set of threshold voting rules and induces fully informative voting.

Let (f, C) be a committee design in which f is a threshold rule and there exists a committee-optimal equilibrium in which all committee members vote fully informatively under tacit collusion. Lemma 2 of Austen-Smith and Banks (1996) establishes that informative voting is an equilibrium if and only if the threshold rule is statistically optimal for the committee. Therefore, the decisions reached in design (f, C) are identical under tacit and overt collusion, $V_T(f, C) = V_O(f, C)$. Thus, if \mathcal{F} is the set of threshold voting rules, $V_T^*(\mathcal{F}) \leq V_O^*(\mathcal{F})$. Moreover, it also follows from Lemma 2 of Austen-Smith and Banks (1996) that for every committee C and voting rule f , there exists a threshold voting rule f' such that $V_O(f, C) = V_T(f', C)$, and so $V_T^*(\mathcal{F}) \geq V_O^*(\mathcal{F})$. Therefore, the two are equal. \square

The conditions for Proposition 1 are tight, which we illustrate below in Section 4.

4. When there are gains from banning deliberation

Proposition 1 shows that if the set of allowable voting rules is the set of threshold voting rules and committee members can coordinate on their preferred equilibrium, then the Principal does not gain from banning deliberation. Relaxing the conditions of Proposition 1, this section describes scenarios where banning deliberation benefits the Principal. We first show that there are gains from banning deliberation when the Principal can select the equilibrium. We then show that even when the committee selects the equilibrium, there are gains from banning deliberation if the committee size is fixed or the Principal can use a non-threshold voting rule.

We illustrate these possibilities in a setting with three available committee members ($N = 3$), a uniform prior ($\pi = 1/2$), and a symmetric signal precision across states ($\gamma^l = \gamma^h \equiv \gamma$). Assume that the payoff parameters are such that the committee strictly prefers to accept any alternative with at least two favorable signals and strictly prefers to reject any alternative with zero or one favorable signals, while the Principal strictly prefers to accept alternatives with three favorable signals and strictly prefers to reject all other alternatives. These parametric restrictions correspond to $q_C \in (1 - \gamma, \gamma)$ and $q_P \in (\gamma, \gamma^3/(\gamma^3 + (1 - \gamma)^3))$.

4.1. Principal-optimal equilibrium

We show that for certain parameter values, the Principal prefers a mixed strategy equilibrium in the game without deliberation to overt collusion in the game with deliberation. In this case, banning deliberation is optimal for the Principal if she can select the equilibrium played by committee members.

Observation 1. Suppose the Principal can select the equilibrium and must use a threshold voting rule. Then there exist an open set of parameters (q_C, q_P, γ) for which the optimal design bans deliberation and selects a mixed strategy equilibrium.

Note that each committee member receives a strictly lower payoff from this mixed strategy equilibrium relative to both tacit and overt collusion.

We prove Observation 1 as follows. Consider a committee design that (i) bans deliberation, (ii) selects all three available committee members, and (iii) uses a threshold rule that accepts an alternative if and only if all three committee members vote to accept. We show that the Principal-optimal equilibrium for this design is in mixed strategies and it achieves a higher payoff for the Principal than her payoff from overt collusion in the game with deliberation or from any pure strategy equilibrium in any other feasible committee design without deliberation.

Let us first construct a symmetric mixed strategy equilibrium in this design. Each committee member votes to accept following a favorable signal $s_i = h$ and randomizes between accepting and rejecting following an unfavorable signal $s_i = l$. Committee member i is pivotal when the other two committee members vote to accept. He is indifferent between accepting and rejecting, conditioning on being pivotal and observing $s_i = l$, if

$$\left(\frac{1-\gamma}{\gamma}\right)\left(\frac{\gamma+p(1-\gamma)}{1-\gamma+p\gamma}\right)^2 = \frac{q_C}{1-q_C}, \tag{3}$$

where p denotes the probability that a committee member votes to accept following $s_i = l$. The solution $p^*(q_C, \gamma)$ to (3) characterizes the unique symmetric mixed strategy equilibrium of this form.⁵

For an open set of parameters (q_C, q_P, γ) , this design and mixed strategy equilibrium yield a strictly higher payoff for the Principal than her payoff in any design with overt collusion. In the mixed strategy equilibrium, the committee rejects the alternative with positive probability when only two members receive favorable signals, as the member who receives an unfavorable signal rejects with positive probability. The Principal indeed prefers to reject such an alternative, while the committee prefers to accept it (and would do so if it could deliberate). The committee also accepts the alternative with positive probability following zero or one favorable signals, as any member who receives an unfavorable signal votes to accept with positive probability. Both the Principal and the committee prefer to reject such alternatives. There exist parameter values such that: (i) the Principal's gain from the higher probability of rejecting an alternative with two favorable signals is greater than her loss from the higher probability of accepting an alternative with zero or one favorable signals, relative to a three-member committee that engages in overt collusion; and (ii) overt collusion in designs with fewer committee members yield strictly lower payoffs to the Principal due to the informational loss from the smaller committee size.⁶

Fig. 2 graphs the Principal's expected payoff in this mixed strategy equilibrium and under overt collusion in two- and three-person committees. As can be seen in the figure, the Principal gains from banning deliberation and selecting the mixed strategy equilibrium for intermediate values of q_C .

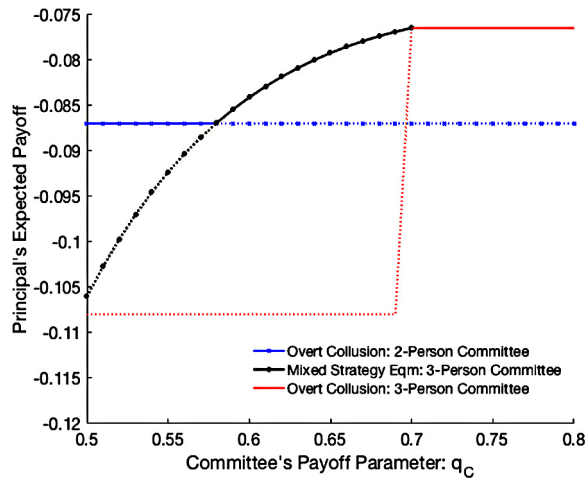


Fig. 2. If the Principal can design the committee and select the equilibrium, for intermediate q_C she chooses a mixed strategy equilibrium in a 3-person committee with no deliberation ($\gamma = 0.7$ and $q_P = 0.8$).

Because there exists a mixed strategy equilibrium of this game without deliberation that improves the Principal's payoff beyond that of overt collusion in the game with deliberation, clearly she prefers to ban deliberation if she selects her favorite

⁵ A solution exists, since the left hand side of (3) is continuous and decreasing in p , is equal to $\gamma/(1-\gamma) > q_C/(1-q_C)$ at $p = 0$, and is equal to $(1-\gamma)/\gamma < q_C/(1-q_C)$ at $p = 1$.

⁶ We show that this holds for an open set of parameters around $\gamma = .7$, $q_C = .65$ and $q_P = .8$. At these exact parameters, the Principal earns $V_N = -.0793$ in the mixed strategy equilibrium, while when allowing deliberation, she would earn $V_{O3} = -.1080$ with a three-member committee, $V_{O2} = -.0870$ with a two-member committee, $V_{O1} = -.1500$ with a one-member committee, and would earn $-.1000$ from not delegating to a committee and always rejecting. The probability of accepting following a low signal is quite small: for these exact parameters, it is $p^*(.65, .7) = .0651$.

equilibrium. It so happens that a mixed strategy equilibrium is in fact the best equilibrium for the Principal. By Proposition 1, when the Principal prefers this design and mixed strategy equilibrium to her best payoff under overt collusion, she also prefers it to any committee size, threshold voting rule and pure strategy equilibrium in the game without deliberation. Therefore, the Principal finds it optimal to ban deliberation and select a mixed strategy equilibrium.

Symmetric strategies This observation also illustrates that if committee members cannot coordinate on asymmetric strategies and instead are restricted to equilibria in symmetric strategies, then there are parameter values under which the Principal gains from banning deliberation even if the committee selects its most preferred symmetric equilibrium in the game without deliberation. In contrast, the Principal never strictly prefers to allow deliberation when restricted to symmetric equilibria. This is because any outcome in the game with overt collusion can also be implemented in a symmetric committee-optimal equilibrium of the game without deliberation by choosing a committee design in which informative voting is an equilibrium.

4.2. Tacit collusion

This section returns to the case where the committee selects its preferred equilibrium and explores when the Principal gains from banning deliberation.

Fixed committee size This example illustrates how the Principal may prefer to ban deliberation if institutional precedent or legal agreements specify that the committee size cannot be altered. Suppose that the committee size is exogenously fixed at three members. We make the following observation.

Observation 2. Suppose the committee must have three members, the Principal must use a threshold voting rule, and the committee engages in tacit collusion in the game without deliberation. Then the optimal design bans deliberation and uses a threshold of three.

In this design, the committee selects an asymmetric equilibrium in which one committee member votes to accept regardless of his signal, and the remaining two members vote informatively. The Principal's expected payoff in this asymmetric equilibrium equals her expected payoff from overt collusion in a two-member committee. The contrast with our main result is that when the Principal can also vary the committee size, then she no longer gains from banning deliberation as she can simply select a smaller committee.

The logic for establishing Observation 2 is as follows. When the Principal is restricted to a committee size of three, this asymmetric equilibrium sways the committee towards rejecting with positive probability when one member receives an unfavorable signal. If one of the members who votes informatively receives the unfavorable signal, the alternative is rejected, while if the member who votes uninformatively receives the unfavorable signal, the alternative is accepted. Therefore, the alternative is rejected for two out of the three possible signal profiles that have one unfavorable signal. This contrasts with the decision under overt collusion in a three person committee, in which the alternative is accepted following all signal profiles that have one unfavorable signal. The decisions on all other signal profiles are the same in this asymmetric equilibrium and under overt collusion. Therefore, the Principal receives a higher payoff in the asymmetric equilibrium of the game without deliberation compared to overt collusion.

This observation also illustrates that the Principal's payoff can decrease with respect to committee size: if the Principal can vary the committee size and the committee overtly colludes, then she prefers a two-member committee to a three-member committee.

Non-anonymous monotone rule Finally, we illustrate how the Principal gains from banning deliberation if the voting rule can be non-anonymous (but still must be monotone); an analogous example can be used to illustrate the same effect if the voting rule can be non-monotone (even if it must be anonymous). We make the following observation.

Observation 3. Suppose the Principal must use a monotone voting rule and the committee engages in tacit collusion in the game without deliberation. Then there exist an open set of parameters (q_C, q_P, γ) such that the optimal design bans deliberation and uses a non-anonymous voting rule.

Vote Profile: (v_1, v_2, v_3)	Chosen Action: $f(v_1, v_2, v_3)$
$(A, A, A), (A, A, R)^*, (A, R, A)^*$	1 (accept)
$(R, R, R), (R, A, R), (A, R, R), (R, R, A), (R, A, A)^*$	0 (reject)

Fig. 3. A non-anonymous voting rule, where the starred vote profiles indicate the lack of anonymity.

Consider a three-member committee and the monotone and non-anonymous voting rule f^* described in Fig. 3. This voting rule involves player 1 having the ability to veto acceptance: if she votes to reject, then the alternative is rejected. If she votes to accept, then at least one of players 2 and 3 has to vote in favor for the alternative to be accepted. Thus,

$1 = f(A, A, R) = f(A, R, A) \neq f(R, A, A) = 0$, illustrating the lack of anonymity. This voting rule is monotone: switching any player's vote from rejection to acceptance, holding all other votes fixed, cannot switch the chosen action from acceptance to rejection.

Suppose that under f^* , informative voting is a committee-optimal equilibrium of the game without deliberation. Then banning deliberation sways the committee towards rejecting the alternative when player 1 obtains an unfavorable signal and players 2 and 3 each receives favorable signals, i.e. signal profile (l, h, h) . With deliberation, the committee would accept an alternative with this signal profile. But without deliberation, f^* makes player 1 pivotal in both the cases where only one of players 2 and 3 is voting to accept and the case where *both* of them are doing so, i.e. vote profiles $(v_2, v_3) \in \{(A, R), (R, A), (A, A)\}$. We show below that for an open set of parameters, pooling these three vote profiles maintains player 1's incentive to vote informatively and the committee rejects the alternative following signal profile (l, h, h) . For all other signal profiles besides (l, h, h) , the outcome induced by tacit collusion with rule f^* is identical to that induced by overt collusion. Therefore, since the Principal prefers to reject an alternative with signal profile (l, h, h) , banning deliberation and using rule f^* yields a higher payoff than overt collusion in a three-person committee.⁷

Next, we show that there exists an open set of parameters such that under f^* , informative voting is a committee-optimal equilibrium of the game without deliberation. We first show that there exist parameter values such that informative voting is an equilibrium strategy profile. When members 1 and 2 vote informatively, it is straightforward to see that member 3 has a strict incentive to vote informatively conditioning on being pivotal, since member 1 is voting to accept and member 2 is voting to reject. The same holds for member 2. Consider member 1. He is pivotal when at least one of members 2 and 3 votes in favor of the alternative. While he does not have an incentive to vote informatively when both members 2 and 3 vote in favor of the alternative, for an open set of parameters (γ, q_C) , it is strictly optimal for him to vote informatively when averaging across all voting profiles at which he is pivotal.⁸ In this case, informative voting is an equilibrium. We next determine whether informative voting is also the committee-optimal equilibrium. Given Lemma 1, it suffices to compare informative voting to each other pure strategy equilibrium. Under our parametric restrictions, the informative voting profile does indeed yield a higher expected payoff for the committee than any other pure strategy equilibrium. We omit these straightforward calculations.

Finally, we verify that for an open set of parameters, the Principal's payoff from tacit collusion with rule f^* is higher than her payoff from overt collusion in a one or two-member committee and her payoff from rejecting all alternatives.⁹ Therefore, if the Principal can use a non-anonymous voting rule, she benefits from banning deliberation.

5. Conclusion

Many organizations, firms, and legislatures rely on committees to evaluate proposals. It is often the case that the preferences of those who serve on these committees conflict with those of the Principal who appoints the committee. A ubiquitous feature of committee-design is that committees are given opportunities to secretly deliberate, communicate, and coordinate prior to voting. Many committees draw straw polls (which aren't reported to the Principal) before coordinating on a voting decision. This feature may appear puzzling insofar as the Principal may be hurt by allowing the committee to have this informational superiority.

One rationale for this procedure is technological: it may simply be too costly to ban deliberation, and so this is an unavoidable feature of institutional design. Our analysis suggests a different, *strategic* rationale for the prevalence of deliberative practices: if the Principal anticipates that the committee tacitly colludes on a committee-optimal equilibrium if she bans deliberation, then she finds no reason to do so. In other words, she shouldn't ban straw polls, even if she can.

Appendix A

Proof of second statement in Lemma 1 We first show that in any committee-optimal pure strategy profile σ , generically, any player who is pivotal with positive probability strictly prefers σ_i to any alternative strategy σ'_i , i.e. $W(\sigma) > W(\sigma'_i, \sigma_{-i})$. Let α_i be the set of action profiles at which member i is pivotal and let $\mu_\sigma(s_i, v_{-i})$ be the posterior belief that the state is H when other members play action profile v_{-i} and member i observes signal s_i . Let M be the set of posterior beliefs that can be induced by the signal profile for a subset of committee members. Given that all other members are playing a pure strategy, the support of μ_σ is a subset of M . Since M is a measure zero set for any finite committee size, $\text{supp } \mu_\sigma$ is also a measure zero set. Similarly, let P_s be the set of probabilities of subsets of other members' signal profiles. Then the probability of any pure strategy action profile v_{-i} under strategy σ , denoted $P_\sigma(v_{-i})$, is in $P_s \cup \{0, 1\}$. The set P_s is also measure zero for any finite committee size. Member i is willing to mix between choosing to accept and reject at signal s_i if

⁷ This holds for all q_P that satisfy our initial parametric restrictions i.e. $q_P \in (\gamma, \gamma^3/(\gamma^3 + (1 - \gamma)^3))$.

⁸ Voting informatively is an equilibrium for an open set of parameters around $\gamma = .7$ and $q_C = .49$.

⁹ Overt collusion in each smaller committee leads the Principal to accept the alternative on a strictly larger set of signal profiles than in a three-person committee with no deliberation, which leads to a strictly lower payoff. A one-member committee will accept if the member receives a high signal, while a for $q_C < 1/2$, a two-member committee will accept if there is at least one favorable signal. The Principal prefers the three-member committee with no deliberation to always rejecting for an open set of parameters around $\gamma = .7$, $q_C = .49$ and $q_P = .71$. At these exact parameters, she earns $-.1070$ when using this non-anonymous rule and $-.1450$ from always rejecting.

$$\sum_{v_{-i} \in \alpha_i} P_{\sigma}(v_{-i})(2f(A, v_{-i}) - 1)(\mu_{\sigma}(s_i, v_{-i}) - q_C) = 0, \quad (4)$$

which holds for a measure zero set of q_C . If (4) does not hold for either signal, then member i strictly prefers a pure strategy. But then generically, all players strictly prefer the committee-optimal pure strategy σ . Therefore, a mixed strategy profile cannot achieve the same value as a committee-optimal pure strategy profile, as with positive probability, this results in committee members playing pure strategy profiles have a strictly lower payoff than $W(\sigma)$.

References

- Agranov, Marina, Yariv, Leeat, 2018. Collusion through communication in auctions. *Games Econ. Behav.* 107, 93–108.
- Austen-Smith, David, Banks, Jeffrey S., 1996. Information aggregation, rationality, and the Condorcet jury theorem. *Am. Polit. Sci. Rev.* 90 (1), 34–45.
- Austen-Smith, David, Feddersen, Timothy J., 2006. Deliberation, preference uncertainty, and voting rules. *Am. Polit. Sci. Rev.* 100 (2), 209–217.
- Battaglini, Marco, 2002. Multiple referrals and multidimensional cheap talk. *Econometrica* 70 (4), 1379–1401.
- Battaglini, Marco, 2017. Public protests and policy making. *Q. J. Econ.* 132 (1), 485–549.
- Coughlan, Peter J., 2000. In defense of unanimous jury verdicts: mistrials, communication, and strategic voting. *Am. Polit. Sci. Rev.*, 375–393.
- Feddersen, Timothy, Gradwohl, Ronen, 2018. Decentralized Advice. Working Paper.
- Gerardi, Dino, Yariv, Leeat, 2007. Deliberative voting. *J. Econ. Theory* 134 (1), 317–338.
- Gilligan, Thomas W., Krehbiel, Keith, 1987. Collective decisionmaking and standing committees: an informational rationale for restrictive amendment procedures. *J. Law Econ. Organ.* 3 (2), 287–335.
- Gradwohl, Ronen, Feddersen, Timothy, 2018. Persuasion and transparency. *J. Polit.* 80 (3), 903–915.
- Hoshino, Tetsuya, 2017. Using Strategic Uncertainty to Persuade Multiple Agents. Working Paper.
- Inostroza, Nicolas, Pavan, Alessandro, 2018. Persuasion in Global Games With Application to Stress Testing. Working Paper.
- Krishna, Vijay, Morgan, John, 2001. A model of expertise. *Q. J. Econ.* 116 (2), 747–775.
- Laffont, Jean-Jacques, Martimort, David, 2000. Mechanism design with collusion and correlation. *Econometrica* 68 (2), 309–342.
- Maskin, Eric, 1999. Nash equilibrium and welfare optimality. *Rev. Econ. Stud.* 66 (1), 23–38.
- McLennan, Andrew, 1998. Consequence of the Condorcet jury theorem for beneficial information aggregation by rational agents. *Am. Polit. Sci. Rev.* 92 (2), 413–418.
- Mookherjee, Dilip, 2006. Decentralization, hierarchies, and incentives: a mechanism design perspective. *J. Econ. Lit.* 44 (2), 367–390.
- Morgan, John, Stocken, Phillip C., 2008. Information aggregation in polls. *Am. Econ. Rev.* 98 (3), 864–896.
- Persico, Nicola, 2003. Committee design with endogenous information. *Rev. Econ. Stud.* 70, 1–27.
- Wolinsky, Asher, 2002. Eliciting information from multiple experts. *Games Econ. Behav.* 41 (1), 141–160.