Rewarding Disagreement:

Optimal Contracts for Collective Decision Making

Naveen Khanna and Mark Schroder*

February 15, 2015

Abstract

We study the problem of optimal information acquisition and truthful signal reporting by committee members when it is costly to acquire information and members can communicate before reporting. We identify a payoff structure which rewards a member for being correct while also disagreeing with others, and show that it achieves first best. The reward is monotonically increasing in both the quality of information and in the number of members. The payoff structure generalizes easily to committees of different sizes and to states that are discretely or continuously distributed.

*Khanna is the A. J. Pasant endowed chair professor in Finance and Schroder is the Philip J. May endowed professor in Finance at the Broad College of Business, Michigan State University, East Lansing, MI 48824. Khanna can be contacted via email at khanna@broad.msu.edu and Schroder at schroder@broad.msu.edu. We are grateful to the Broad school for summer research funding towards this paper.

Many decisions are made collectively with the expectation that each agent will bring valuable information to the decision-making process and its aggregation across agents will result in superior decisions. Examples include, a board of directors collectively deciding whether to accept merger or divestiture proposals, hire/fire a CEO, or choose between different investment opportunities; medical teams deciding on the best treatment for a patient; juries determining whether to acquit or convict a defendant; or hiring committees seeking the best candidate from a pool of applicants.

However, there are two problems that make it difficult to achieve efficient aggregation. Firstly, if committee members must exert costly effort to contribute to the joint decision, there is a free rider problem because a committee member knows that in many situations his information may not be pivotal, causing him to shirk and save his personal effort costs. Secondly, we know, from Bikhchandani, Hirshleifer, and Welch (1992), Bannerjee (1992), Welch (1992) and Pendergast (1993), that there is a tendency towards conformity when each participant's information is imperfect and each can infer others' information through their observable behavior or actions. Thus, even after incurring the cost of acquiring information, a member may choose to ignore his own information and vote strategically, causing valuable information to be lost.

Not surprisingly, both problems become worse when members can communicate before voting, as this increases the opportunity to free-ride on others' information, further reducing information collection effort. Thus, papers like Persico (2004) and Gershkov and Szentes (2009), which allow for costly information collection, disallow pre-vote communication. However, this is a strong constraint to impose, because most committee decisions are made through a vote which generally follows a period of discussion among members.[1] Thus, an optimal mechanism for efficient collective decision must achieve two objectives: 1) Induce each agent to expend the effort needed to obtain the desired signal quality, and 2) Induce each agent to report that signal truthfully, despite the opportunity for free communication with other members. Ideally, it should apply to committees with many members who may be heterogeneous in their ability to collect information, and generalize to discrete and continuous information-distribution settings.

To achieve these objectives, we depart from existing literature in the following significant manner. While most prior research has focused on designing optimal voting mechanisms to achieve truthful reporting and efficient information aggregation, we instead focus on designing a state and vote contingent payoff function for each agent to induce optimal effort and truthful reporting.[2] To derive such a payoff function, we assume that each agent maximizes his expected payoff by trading off the explicit costs and benefits of information collection and truthful reporting. The resulting payoff function (contract) enables us to separate the agents' signal acquisition and reporting functions from the decision-making problem of the firm, thus allowing us to solve the firm's problem in two distinct steps. First, identify a payoff structure that results in truthful reporting and optimal effort choice in a general decision-making

[1] Papers on Boards by Adams and Ferreira (2007), Harris and Raviv (2008) and Bond and Eraslan (2010) show information sharing to be value enhancing. However, they do not consider the potential impact of discussion on effort choice.

[2] Most existing work also assumes some utility/payoff from correct and incorrect decisions. However, these are generally imposed to capture the common sense notion that agents get a positive utility from a good decision/vote and a negative utility from a bad decision/vote. We explicitly design the agents' payoff functions to achieve truthful reporting.

setting with unrestricted pre-vote communication, and possibly heterogeneous committee members. Second, determine the optimal firm action based on the information revealed by the votes.

In a simple binary-signal setting, the decision can be implemented by a "voting rule," which simply expresses the optimal decision as a mapping from the vote to a decision. In a simple example with homogeneous agents, we show that the optimal decision can be based only on the difference in the number of yes and no votes. Whether a simple majority or supermajority vote is needed depends on the expected profitability of the decision under consideration. For higher expected profitability, a simple majority vote may be optimal as a less precise aggregate signal may suffice, while for lower expected profitability, different extents of supermajorities may be needed. With heterogeneous agents, because the optimal information quality will differ across agents, the votes of better informed agents are given greater weight.[3]

The optimal committee-member payoffs that we derive depend only on the realized state (assumed observed by all agents and the firm after its decision has been taken) and the signal reports (or votes) of all the members. Each member receives a payoff if and only if his reported signal is correct ex-post (which we refer to as the *basic component* of compensation), and an additional payoff if his reported signal differs from the reports of the other agents' (the *disagreement component* of compensation). The disagreement component of the payoff function is increasing in both the number and signal-quality of the agents with whom his report disagrees,

[3]Thus unlike papers like Couglan (2000) and Austen-Smith and Feddersen (2006), here pre-vote communication does not alter collective decisions because all agents vote truthfully with the appropriate payoff structure. For the same reason, unlike in Gerardi and Yariv (2007), voting rules impact decisions.

and incentivizes truthful reporting despite possible pre-report communication with other agents. In equilibrium, no agent has an incentive to be untruthful in his communication with other agents because no agent is influenced by the others' signals. Furthermore, each agent's equilibrium ex-ante expected payoff is the same as that in a single-agent case, depending only on the agent's own signal quality, and a scale factor in the basic component of compensation, which can be set by the firm to elicit first-best agent effort, even though effort is unobserved by the firm.

Our paper builds on a number of important contributions in the field of collective decision. Condorcet (1785) considers a problem in which each juror receives an independent and imperfect signal about a defendant's guilt or innocence and casts a vote consistent with his information (i.e., without strategic considerations). He finds that a decision made collectively by a jury is better than that made by an individual. However, when the possibility of strategic voting is considered by Austen-Smith and Banks (1996), Myerson (1998), Feddersen and Pesendorfer (1996, 1997), among others, Condorcet's conclusion of informed voting can be violated when a jurist considers his vote to be pivotal to the jury's verdict. This can result in inefficient aggregation and potentially even uninformed decisions. Coughlan (2000) stresses the significance of pre-vote communication in jury deliberations, and shows that such pre-vote communication can restore efficient aggregation at least under the assumption of common preferences. Gerardi and Yariv (2007) go further and show that with pre-vote deliberation, voting rules can in fact become irrelevant, thereby emphasizing the importance of deliberation relative to voting rules. Goeree and Yariv (2011) provide strong experimental support for these results. They record that pre-vote deliberation results in similar outcomes under different voting rules, and communication tends to be public and generally truthful. However, Austen-Smith and Feddersen (2006) show that even a small uncertainty about preferences makes it impossible to get informative/non-strategic voting under any voting rule, even with pre-vote deliberation.

An important element of our model is costly endogenous information collection by members. The potential to free ride on others' information reduces the incentive to collect information, thus diminishing the aggregate quality of information on which the firm's decision is made. Recent papers have designed optimal voting rules that overcome free riding by making each voter pivotal, thereby incentivizing each to acquire costly information. For instance, Gershkov and Szentes (2009) develop a sequential voting mechanism with a stochastic stopping time, in which members are selected randomly to acquire information and make a recommendation. This makes every new voter pivotal. However, this mechanism is optimal only as long as the order in the sequence is not known; furthermore, pre-vote communication is not permitted. Persico (2004) allows committee members to simultaneously decide whether to acquire a costly signal and then vote. He identifies an optimal threshold voting rule and optimal committee size. However, he too prohibits pre-vote communication. As in our paper, Gerardi and Yariv (2008) allow for both endogenous information collection and pre-vote communication. However, they focus on voting rules and committee size, and show that conditioning on just these variables may not deliver first best outcomes.

Reward structures similar to that developed here appear to be prevalent in real life, though in more subtle guises. Perhaps the closest would be where a committee is called upon repeatedly to make decisions collectively. Examples include board of directors, investment committees, recruiting committees, promotion and tenure committees, and congressional subcommittees. In all of these instances, repeatedly observing a member's voting decisions would enable an institution/firm to develop a posterior about an agent's information generating/analyzing skills, and establish incentives to induce the optimal amount of effort. The premium for correctly disagreeing could take the form of an increase in reputation, which would likely result in increased demand for the agent's services and/or an increase in the probability of getting promoted. For tractability, this paper uses a static framework in which agents' information-acquisition skills (i.e., their effort cost functions) are public knowledge.

The outline of the paper is as follows. The setting and problem are explained in Section 1. The design of the compensation contracts, is presented in Section 2. Section 3 characterizes all the alternative pure-strategy signal-reporting equilibria. Section 4 shows that it is simple to scale the payoffs to elicit the desired agent effort. Section 5 demonstrates the impact of free ridership in the absence of a disagreement component to the agents' payoffs. Section 6 examines a simple firm decision problem. Extensions to arbitrary discrete states and continuously-distributed states are presented in Section 7. The conclusion in Section 8 is followed by the proofs in Section 9.

1 Setting and Problem

We assume throughout that the firm/principal and all n agents are risk neutral; and because there is a competitive labor market, the firm must satisfy a participation constraint for each agent (a common assumption in the contracting literature). The time sequence of events is as follows:

1. The principal presents the payoff functions to the agents.

2. Each agent a exerts effort q^a at a cost $h^a(q^a)$, where $h^a()$ is increasing and convex, to acquire a private signal (neither the agents' efforts nor signals are observed by the firm) with a precision of q^a. The cost functions $h^1, \ldots, h^n$ are common knowledge.

3. Each agent a receives the signal S^a, which he/she may share with other agents (the firm cannot prevent this).

4. The agents vote, and their votes are observed by the principal.

5. The principal chooses some action (explained below) to maximize expected firm value conditional on the observed votes.

6. The state is revealed to the firm and all agents, and contractual payments are made to the agents.

We assume that there are only two states of nature: $\theta \in \{0, 1\}$, with probabilities $p(0)$ and $p(1) = 1 - p(0)$. Each agent $a \in \{1, \ldots, n\}$ receives a conditionally independent noisy signal $S^a \in \{0,1\}$ of the true state; the signal is correct with probability $q^a \in [1/2, 1)$. Each agent a exerts costly effort to increase the signal precision q^a. We assume that the agents cannot write enforceable contracts to collude in reporting signals.[4] The information sharing is assumed to be unrestricted because the principal cannot monitor the communication between agents during the time interval between the collection of the signals and the vote.

The focus of this paper is the structure of the agents' rewards; specifically, payoff functions that induce truthful reporting and first-best effort exertion by the agents. This can be used for a general class of firm optimization problems, which we outline in the next paragraph, followed by a characterization of the first-best problem. Our contract allows for a first-best solution, thus avoiding the potential difficulties of the second-best problem.

After the signals are reported, but before the state is revealed, the principal chooses some action α from a set $\mathcal{A}$ to maximize the conditional expectation of firm payout function $F(\theta;\alpha)$. Our applications in Sections 5 and 6 will examine binary choice problems, $\mathcal{A} = \{0,1\}$, of whether or not to accept a project; but it would also be straightforward to extend to some continuous choice, $\mathcal{A} = \mathbb{R}$, of how much to invest/lend/borrow. With truthfully reported signals, the firm chooses the action that maximizes the expected payoff conditional on the signals:

$$\max_{\alpha \in \mathcal{A}} E^q \left(F(\theta;\alpha) \middle| S^1, \ldots, S^n \right).$$

[4] Otherwise the agents could agree to cast the votes that maximize the total payout, and then share the payout so that each agent receives (contingent on casting the required vote) at least as much as in the truthful reporting equilibrium. By requiring a secret vote, the firm could make it impossible for the agents to confirm that each voted according to plan.

where E^q denotes the expectation operator with signal precision levels $q = (q^1, \ldots, q^n)$. The principal's first-best problem (if the agents' efforts and signals were observable by the firm) is to choose the time-0 effort levels q to maximize expected firm value net of the total cost of the agents' effort:

$$\max_{q^1,\ldots,q^n \in [1/2,1)^n} E^q \left\{ \max_{\alpha \in \mathcal{A}} E^q \left(F(\theta;\alpha) \middle| \, S^1, \ldots, S^n \right) \right\} - \sum_{a=1}^{n} h^a (q^a) \, .$$

Note that because of the assumed participation constraint,[5] the firm considers only the cost of agents' effort, and not the expected payout to the agents (a fixed component of agent pay would always adjust to ensure that agent utility always equals the reservation utility).

The potential difficulties of the second-best problem (efforts and signals not observed by the principal) are 1) characterizing the Nash equilibria of signal reports (allowing for possible signal sharing) for any feasible set of compensation contracts, 2) characterizing the Nash equilibria of agent efforts given the same contracts, and 3) solving for the optimal set of contracts. The set of contracts could be large (functions of the reported signals and the realized state), and existence and uniqueness of pure-strategy signal-reporting and effort-exertion Nash equilibria is not always guaranteed.

[5] The participation constraint is easily shown to be binding in this context.

2 A Contract that Induces Truthful Revelation

We assume in this section that the precisions $(q^1, \ldots, q^n)$ are common knowledge. In Section 4, we relax this assumption and explicitly introduce the effort-choice problem.

We start by providing the basic intuition in the case of two agents, A and B, and equally likely prior state probabilities. Consider first a contract promising agent A the amount $\delta^A > 0$ if his reported vote is correct (matches the state), and nothing if incorrect. Suppose the two agents share their signals. If both agents receive the same signal, then obviously truthful voting is optimal, but if the signals disagree, there might be an incentive for one agent to report the other's signal. For example, if $S^A = 1$ and $S^B = 0$, then Bayes' rule together with conditional independence of the signals implies

$$E\left(\text{payout to agent } A \mid S^A, S^B\right) = \frac{\delta^A}{2P\left(S^A, S^B\right)} \times \begin{cases} q^A\left(1 - q^B\right) & \text{if } A \text{ reports own signal,} \\ \left(1 - q^A\right) q^B & \text{if } A \text{ reports } B\text{'s signal.} \end{cases}$$

If agent B is better informed (i.e., $q^B > q^A$), then A will optimally report B's signal instead of reporting his own signal to maximize the probability of receiving δ^A. Thus, because of free riding by A some information is lost to the firm, possibly resulting in a worse decision. Free-riding by A can be discouraged by making the reward depend also on whether the reported votes disagree:

$$\text{payout to agent } A = \begin{cases} \delta^A / \left(1 - q^B\right) & \text{if } A\text{'s reported signal is correct and disagrees with } B\text{'s} \\ \delta^A / q^B & \text{if } A\text{'s reported signal is correct and agrees with } B\text{'s.} \end{cases} \tag{1}$$

This adjustment removes the influence of B's signal on A's expected payoff. Assuming B truthfully reports, the expected payoff to A under the modified contract (1) is

$$E\left(\text{payout to agent } A \mid S^A, S^B\right) = \frac{\delta^A}{2P\left(S^A, S^B\right)} \times \begin{cases} q^A & \text{if} \quad A \text{ reports own signal} \\ 1 - q^A & \text{if} \quad A \text{ reports } B\text{'s signal,} \end{cases} \tag{2}$$

regardless of B's signal. Furthermore, $q^A > 1/2$ implies that truthful reporting is strictly preferred by A.[6] If agent B is offered an analogous contract, truthful signal reporting by both agents is a Nash equilibrium. Each agent optimally ignores the other agent's signal because he essentially receives a bonus for disagreeing; and the bonus is strictly increasing in the precision of the other agent's signal. But he will not misreport his signal to generate disagreement because there is no reward if the reported signal is incorrect.

The generalization of this simple example to different state prior probabilities and n agents is given in the following proposition. It is convenient to scale the payout (1) by the probabilities of the other agents' reported signals to simplify the ex-ante expected agent payout. Note that $P\left(S^k = s^k\right)$ denotes the probability that agent k receives a signal of value s^k (where $s^k \in \{0, 1\}$), and $P\left(S^k = s^k \middle| \theta = s^a\right)$ denotes the conditional probability of the same event given that the true state is s^a. Finally, the indicator function $1_{\{\theta = s^a\}}$ is one if $\theta = s^a$ and zero otherwise.

[6] If $q^A = 1/2$ then agent A is indifferent between reporting S^A and $1 - S^A$, but A's signal contains no information in this case.

Proposition 1 *Suppose each agent a is offered the following payout, contingent on the reported signals $s = (s^1, \ldots, s^n)$ and true state θ:*

$$\begin{aligned} g^a(\theta; s) &= \delta^a \frac{1_{\{\theta = s^a\}}}{p(s^a)} \prod_{k \neq a} \frac{P\left(S^k = s^k\right)}{P\left(S^k = s^k \middle| \theta = s^a\right)} \\ &= \hat{\delta}^a \frac{1_{\{\theta = s^a\}}}{p(s^a)} \prod_{\left\{k:\, s^k \neq s^a\right\}} \frac{q^k}{1 - q^k}, \quad \delta^a > 0, \quad a = 1, \ldots, n. \end{aligned} \tag{3}$$

where $\hat{\delta}^a = \delta^a \prod_{k \neq a} P\left(S^k = s^k\right)/q^k$ (note that $\hat{\delta}^a$ does not depend on s^a). Then truthful reporting is a Nash equilibrium regardless of how many other signals each agent observes. The ex-ante expected payout to each agent a in this equilibrium is

$$Eg^a\left(\theta, S^1, \ldots, S^n\right) = 2\delta^a q^a, \quad a = 1, \ldots, n, \tag{4}$$

which is the same as in the case when a is the only agent.

Agent a receives the payment $\hat{\delta}^a/p(s^a)$ if his/her signal is correct, scaled by a "bonus" $q^k/\left(1 - q^k\right)$ for each agent k reporting a different signal. We have divided by the prior state probability $p(s^a)$ because the principal wants the agent to report his signal, not the state that is most likely conditional on that signal. The bonus for disagreeing with agent k is increasing with the precision of k's signal (i.e., as q^k increases). Disagreeing with an uninformed agent (agent k is uninformed if $q^k = 1/2$) results in no additional reward. The Nash equilibrium holds for any set of strictly positive constants δ^a, $a = 1, \ldots, n$, no matter how small. Of course, agent a's time-0 optimal effort choice will be affected by δ^a, which allows the principal to induce the first-best effort level (see Section 4).

The contract's payoff in the truthful-reporting equilibrium can also be expressed in terms of a measure of aggregate signal strength. Denote by $\mathcal{A} = \{1, \ldots, n\}$ the set of all agents, and for any nonempty subset of agents $B \subset \mathcal{A}$ let

$$M^B(\theta) = \prod_{\left\{k \in B:\, S^k = \theta\right\}} \frac{q^k}{1 - q^k}, \quad \theta = 0, 1$$

(for the empty subset we define $M^{\emptyset}(\theta) = 1$). The quantities $M^B(0)$ and $M^B(1)$ represent aggregate measures of the zero and unit signals of agents in the set B. The state-θ payoff (3) in the truthful-revelation equilibrium is then proportional to the aggregate measure of the opposite signal:

$$g^a\left(\theta; S^1, \ldots, S^n\right) = \hat{\delta}^a \frac{1_{\{\theta = S^a\}}}{p(S^a)} M^{\mathcal{A}}(1 - S^a), \quad \delta^a > 0, \quad a = 1, \ldots, n.$$

The ratio $M^{\mathcal{A}}(0) / M^{\mathcal{A}}(1)$ is a sufficient statistic for computing the conditional state distribution, which follows from

$$P\left(\theta | S^1, \ldots, S^n\right) = \frac{p(\theta) M^{\mathcal{A}}(\theta)}{p(0) M^{\mathcal{A}}(0) + p(1) M^{\mathcal{A}}(1)}, \quad \theta = 0, 1. \tag{5}$$

3 Other Signal-Reporting Equilibria

In addition to the truthful reporting equilibrium above, other pure-strategy Nash equilibria with possible misreporting by some/all agents also exist. However, we show that these equilibria generally result in payoffs that make some agents better off and other agents worse off compared to their payoffs in the truthful equilibrium. The better informed agents tend to be the losers. That leads us to believe the misreporting equilibria are unlikely to occur, and we therefore rule out these equilibria in the following sections of the paper.

We also show that among the misreporting agents, the difference in the aggregate strengths of the unit and zero signal cannot be too large (otherwise the misreporting equilibria will not be sustained), which implies that the information lost in these equilibria is unlikely to be large (and is zero if agents are homogeneous).

The following proposition characterizes all pure-strategy voting equilibria given the payoff function (3).

Proposition 2 *Let S be the vector of observed signals and s the vector of reported signals. Also, let $L = \{a \in \mathcal{A} : s^a \neq S^a\}$ denote the set of agents whose reported and true signal disagree (the agents who lie). Then s is a Nash equilibrium if and only if*

$$\sqrt{\frac{q^a}{1-q^a}} M^L(s^a) \geq M^L(1-s^a), \quad a = 1, \ldots, n. \tag{6}$$

Each misreporting agent's conditional expected payoff in the equilibrium s compared to the truthful equilibrium satisfies

$$E\{g^a(\theta, s)|\, S\} > E\{g^a(\theta, S)|\, S\} \quad \Longleftrightarrow \quad M^L(1-S^a) > M^L(S^a), \quad \textit{every } a \in L. \tag{7}$$

The condition (6) is equivalent to

$$\min_{\{a\in\mathcal{A}:\, s^a=0\}}\sqrt{\frac{q^a}{1-q^a}}\geq\frac{M^L(1)}{M^L(0)}\geq\left(\min_{\{a\in\mathcal{A}:\, s^a=1\}}\sqrt{\frac{q^a}{1-q^a}}\right)^{-1}$$

(we interpret the minimum as 1 if the set is empty). That is, s is an equilibrium if the aggregate qualities of the zero and unit signals among the misreporting agents are not too different.

Consider first the case when all the agents are equally informed (identical q's). Condition (6) implies that s is an equilibrium if and only if the number of agents with true signal zero, but reporting one, equals the number of agents with signal one, but reporting zero (e.g., out of ten agents, two receiving a unit signal report zero and two receiving a zero signal report one). From (7), each agent's conditional expected payoff in the equilibrium s is the same as in truthful-reporting equilibrium; furthermore, equilibrium misreporting will not affect the value of the information reported to firm (because $M^L(1)=M^L(0)$).

If the agents' q's are heterogeneous, then the ratio of the aggregate signal qualities among the misreporting agent cannot deviate too far from one. If $M^L(1)/M^L(0)$ were too large (a disproportionate misreporting of the zero signal compared to the unit signal), then an agent $a\in L$ who receives a unit signal will prefer to instead report truthfully to increase his disagreement reward and probability of being correct; the more informed a is, the more likely he is to deviate to truth telling (note that $a\in L$ implies that $M^L(1)$ includes a's signal-quality term). Also, a low-precision agent $a\in\mathcal{A}\backslash L$ (i.e., from the truthful set of agents) who receives a zero signal might prefer to misreport a signal of one to boost his disagreement reward. The case when

$M^L(1)/M^L(0)$ is too small is analogous.

Condition (7) says that the group of agents receiving a zero signal and misreporting one are better off in the misreporting equilibrium only if the aggregate quality of this group's signal is less than that of the other group, which receives a unit signal and misreports zero. That is, the collectively less (more) informed misreporting agents are better (worse) off in the misreporting equilibrium compared to the truthful equilibrium.

In the case of two agents ($n = 2$), which we label A and B, then (6) implies that a misreporting equilibrium exists only when the agents receive different signal ($S^A \neq S^B$) and

$$\left(\frac{q^A}{1-q^A}\right)^2 \geq \frac{q^B}{1-q^B} \geq \sqrt{\frac{q^A}{1-q^A}}.$$

Furthermore, (7) implies that the more informed agent will be worse off and the less informed agent better off in the misreporting equilibrium compared to the truthful equilibrium.

4 Eliciting the Desired Agent Effort

We assume for simplicity that each agent a has an additive effort-cost function $h^a(q^a)$, which is assumed increasing, strictly convex and differentiable, with a derivative satisfying $\lim_{q\to 1} \frac{d}{dq} h^a(q) = \infty$. From Proposition 1, agent a's ex-ante expected payoff in the truthful-reporting equilibrium is $Eg^a(\theta, S) = 2\delta^a q^a$, and therefore agent

a's effort (uniquely) solves the problem

$$\max_{q^a \in [1/2,1)} 2\delta^a q^a - h^a(q^a).$$

Optimal effort is then a continuous and strictly increasing function of δ^a, which we denote $\hat{q}^a(\delta^a)$.

The principal cannot observe agent effort, but knows the agents' effort-cost functions. Therefore the principal can deduce $\hat{q}^a(\delta^a)$ for each agent a and choose the scaling factor δ^a in (3) to obtain the desired effort. Because the contracts are public knowledge, each agent can deduce the optimal efforts of all the other agents, and therefore the conditional signal distributions of all the agents.

Example 1 *Let the effort-cost function be*

$$h^a(q^a) = \alpha^a \frac{q^a}{1-q^a}, \quad a = 1, \ldots, n,$$

which is increasing from 1 *to* ∞ *as* q^a *increases from* $1/2$ *to* 1. *Agent* a*'s effort-choice problem is therefore*

$$\max_{q^a \geq 1/2} 2\delta^a q^a - \alpha^a \frac{q^a}{1-q^a},$$

and optimal effort is

$$\hat{q}^a = \max\left(1 - \sqrt{\frac{\alpha^a}{2\delta^a}}, \frac{1}{2}\right).$$

Any $\hat{q}^a \in [0.5, 1)$ *can be induced by the principal by letting* $\delta^a = (\alpha^a/2)(1-\hat{q}^a)^{-2}$. *Suppose the first-best effort levels are* q^a_{FB}, $a = 1, \ldots, n$. *By offering the agents the payoff functions* (3) *with the scaling parameters* $\left\{\delta^a = (\alpha^a/2)(1-q^a_{FB})^{-2}; a = 1, \ldots, n\right\}$, *the principal is able to elicit first-best efforts and sustain a Nash equilibrium with truthful reporting of the signals.*

5 The Impact of Free Riding

To illustrate the impact of free riding when there is no disagreement premium, suppose each agent a's payoff depends only on being correct:

$$g^a\left(\theta; s^1, \ldots, s^n\right) = \delta^a 1_{\{\theta = s^a\}}, \quad \delta^a > 0, \quad a = 1, \ldots, n. \tag{8}$$

That is, agent a receives δ^a if the signal is correct and nothing otherwise. For the following, we will assume that each agent discloses his signal to the other agents.[7] Then it is optimal for each agent to report the mode (or one of the modes) of the conditional distribution (5):

$$s^a \in \arg \max_{s \in \{0,1\}} p(s) M^{\mathcal{A}}(s), \quad a = 1, \ldots, n.$$

But there are two problems: 1) The mode of the conditional distribution is not as informative as the signals themselves, and is generally insufficient for the firm to compute $M^{\mathcal{A}}(0)/M^{\mathcal{A}}(1)$; and 2) each agent will tend to exert less effort at time zero.

The following example shows that contract (8) can be strictly worse for the firm than contract (3). Free riding under (8) results in only one of the two agents

[7] Each agent is indifferent to disclosing his signal to the other agents, and each agent to whom he discloses is better off (sometimes strictly so).

acquiring an informative signal, but that signal is not precise enough at any effort level to result in positive expected profit. The firm will therefore not offer any payment to either agent ($\delta^A = \delta^B = 0$), and no effort will be exerted. Under contract (3), on the other hand, the firm will incentivize both agents to invest in gathering information because the optimal combined signal will be sufficiently precise that the net present value of the project (net of effort cost) will be positive.

Example 2 *Suppose* $p(0) = p(1) = 1/2$ *and a common effort-cost function* $h^a(q) = \frac{1}{2}\left(q - \frac{1}{2}\right)^2$ *for the two agents* $a \in \{A, B\}$*. Also assume that each agent's effort cannot exceed* $3/4$ *(that is,* $q^A, q^B \in [1/2, 3/4]$*).*[8] *After the signals are reported, the principal can choose to undertake an investment with (ex-post) payout* $\theta - c$*, where the fixed cost* c *is assumed to satisfy* $c > 11/16$*.*

With the payoff (8)*, there are two possible time-*0 *Nash effort equilibria: only agent* A *exerting effort or only agent* B *exerting effort. The Nash signal-reporting equilibrium is for both agents to report the signal of the only informed agent. But because only one informative signal is received, and the precision is constrained (by* $3/4$*), the firm will always forgo the investment because the conditional expected payout is negative.*[9] *The firm optimally sets* $\delta^A = \delta^B = 0$ *to induce the minimum signal precision (yielding* $q^A = q^B = 1/2$*).*

Now consider the promised payoff (3)*, which induces both agents to report truthfully. It can be shown that under the optimal policy, the principal accepts the project only if* $S^A = S^B = 1$*, and the firm chooses the same effort level* q *for each agent. Substituting* $P\left(\theta = 1 \,\middle|\, S^A = S^B = 1\right) = q^2/\left\{q^2 + (1-q)^2\right\}$*, the principal's problem*

[8] Or, equivalently, infinite cost for effort in excess of 3/4.

[9] See the appendix for the derivations.

is to choose effort q to maximize ex-ante firm value Π:

$$\Pi = \max_{q \in [1/2, 3/4]} \left\{ P\left(S^A = S^B = 1\right)\left(\frac{q^2}{q^2 + (1-q)^2} - c\right) - \left(q - \frac{1}{2}\right)^2 \right\} \qquad (9)$$

Substituting $P\left(S^A = S^B = 1\right) = \frac{1}{2}\left(q^2 + (1-q)^2\right)$ and solving we get optimal effort is $\hat{q} = \frac{1}{2}(c+1) / \left(c + \frac{1}{2}\right)$, which results in positive firm value, $\Pi > 0$. Optimal effort is elicited by letting $\delta^A = \delta^B = (1/2)(\hat{q} - 1/2) = 1/\left\{8(c + 1/2)\right\}$.

6 The Firm's Decision

We suppose, for simplicity, that the firm must choose whether to accept or reject a project. The firm's decision is determined solely by the value of the sufficient statistic $M^{\mathcal{A}}(1) / M^{\mathcal{A}}(0)$, which measures the aggregate strength of the unit signal relative to the zero signal. We interpret the report of a unit signal as an accept vote, and zero signal as a reject vote. In our two-state setting ($\theta \in \{0, 1\}$) we can assume without loss of generality[10] an ex-post payout $\theta - c$ if the project is accepted. The payout is zero if rejected. We interpret c as the investment required to undertake the project; to make the decision nontrivial we assume $c \in (0, 1)$.

The firm should accept the project, after receiving the agent votes, only if $P\left(1 | S^1, \ldots, S^n\right) > c$; that is,

$$\frac{M^{\mathcal{A}}(1)}{M^{\mathcal{A}}(0)} > \frac{cp(0)}{(1-c)\,p(1)},$$

[10] Because the firm accepts the project only if the expected payout is positive, we can always scale so that the payout to $1 - c$ in the unit state.

or equivalently

$$\sum_{\{a\in\mathcal{A}:\, S^a=1\}} \ln\left(\frac{q^a}{1-q^a}\right) - \sum_{\{a\in\mathcal{A}:\, S^a=0\}} \ln\left(\frac{q^a}{1-q^a}\right) > \ln\left(\frac{cp\,(0)}{(1-c)\,p\,(1)}\right).$$

If the agents are homogeneous in information acquisition costs, and therefore equally informed (identical q's) in equilibrium, then the project-acceptance rule is of the form

$$\#\text{ unit signals } - \#\text{ zero signals } > K$$

for some constant K (which is increasing in c). This can be implemented using a simple voting rule: majority rule if $K = 0$ (that is $cp\,(0) = (1-c)\,p\,(1)$), and supermajority or unanimity if $K > 0$. In the heterogeneous case, each agent a's vote is weighted by $\ln\left(\frac{q^a}{1-q^a}\right)$, and the difference in the weighted votes must exceed the threshold $\ln\left(\frac{cp(0)}{(1-c)p(1)}\right)$, which is increasing in c as well as in $p\,(0)\,/p\,(1)$, the relative prior probability of the zero-payoff to unit-payoff state. The votes of a more informed agent therefore receive more weight.

7 Extensions to Discrete and Continuous States

Suppose there are m possible states $\{\theta_1, \ldots, \theta_m\}$, with probabilities $p\,(\theta_j)$, $j = 1, \ldots, m$. Each agent $a \in \{1, \ldots, n\}$ observes a costly signal S^a of the state. We denote the conditional distribution of a's signal $P\left(S^a = s^a \middle|\, \theta\right) = f^a\left(s^a \middle|\, \theta\right)$, and the marginal distribution $P\,(S^a = s^a) = f^a\,(s^a)$, where the dummy variable s^a is also valued in $\{\theta_1, \ldots, \theta_m\}$. We assume for now that the efforts q^a, $a \in \{1, \ldots, n\}$, exerted before receiving the signals, are given and known by all, as are the conditional distribution functions $f^a(\cdot|\,\theta)$, $a \in \{1,\ldots,n\}$. The signal distributions depend on these efforts, but we omit this dependence in the notation until we consider the optimal effort problem in Section 7.1 below. There we explicitly introduce the effort parameters, and assume that the function form $f^a(\cdot|\,\theta; q^a)$, for each a are common knowledge, but the q^as are private information. Assumption 1 below is the same as in the binary case, and Assumption 2 is for simplicity (Remark 1 shows that it can be relaxed).

Assumptions:

1. Conditional on the true state θ, the signals are independent:

$$P\left(S^1,\ldots,S^n\right|\theta\right) = \prod_{a=1}^{n} f^a\left(S^a\right|\theta\right).$$

2. The probability of any agent a receiving a signal of θ_j is conditionally most likely when the state is θ_j:

$$f^a\left(\theta_j\right|\theta_j\right) \geq f^a\left(\theta_j\right|\theta_k\right), \quad j,k \in \{1,\ldots,m\}, \quad \text{for every agent } a. \tag{10}$$

The extension to Proposition 1 follows.

Proposition 3 *Suppose each agent a is offered the following payout, contingent on the reported signals and true state:*

$$g^a\left(\theta; s^1,\ldots,s^n\right) = \delta^a \frac{1_{\{\theta=s^a\}}}{p\left(s^a\right)} \prod_{k\neq a} \frac{f^k\left(s^k\right)}{f^k\left(s^k\right|s^a\right)}, \quad \delta^a > 0, \quad a = 1,\ldots,n. \tag{11}$$

Then truthful reporting is a Nash equilibrium regardless of how many other signals each agent observes. The ex-ante expected payout to each agent a in this equilibrium is

$$Eg^a\left(\theta, S^1, \ldots, S^n\right) = \delta^a \sum_{j=1}^{m} f^a\left(\theta_j \middle| \theta_j\right), \quad a = 1, \ldots, n, \tag{12}$$

and is the same as the single-agent case, regardless of the number of agents.

As in the binary case, the contract removes any incentive for agent a to be influenced by agent k's signal by dividing the payment to agent a by (applying Bayes rule)

$$\frac{f^k\left(s^k \middle| s^a\right)}{f^k\left(s^k\right)} = \frac{P\left(\theta = s^a \middle| S^k = s^k\right)}{p\left(s^a\right)},$$

which represents the relative impact of agent k's signal on the likelihood that agent a's signal is correct. Therefore if agent k's signal significantly reduces the likelihood that a's signal is correct, relative to the prior probability, then a large payment to a is required to induce a to ignore k's signal. On the other hand, if k's signal is uninformative, then the ratio is one, and no compensation for ignoring k's signal is required. The Nash equilibrium holds for any set of strictly positive constants δ^a, $a = 1, \ldots, n$, no matter how small.

The following remark shows that Assumption 2 above can be omitted by replacing the "basic component" of the contract, which we denote $\hat{g}^a\left(\theta, s^a\right)$. In Propositions 1 and 3, the basic part of the contract is

$$\hat{g}^a\left(\theta, s^a\right) = \delta^a \frac{1_{\{\theta = s^a\}}}{p\left(s^a\right)}. \tag{13}$$

Remark 1 (generalizations) *Suppose the basic payoff* $\hat{g}^a(\theta, s^a)$ *induces truthful reporting of each agent a's signal when there is no sharing of information (e.g., as in the single-agent case). Then truthful reporting is a Nash equilibrium under the set of contracts*

$$g^a\left(\theta; s^1, \ldots, s^n\right) = \hat{g}^a(\theta, s^a) \prod_{k \neq a} \frac{f^k\left(s^k\right)}{f^k\left(s^k \middle| \theta\right)}, \quad a = 1, \ldots, n, \tag{14}$$

regardless of how many other signals each agent observes. For example, we can eliminate Assumption 2 and obtain truthful revelation with the basic payoff[11]

$$\hat{g}^a(\theta, s^a) = \delta^a \frac{P(\theta | S^a = s^a)}{\sqrt{\sum_{k=1}^m P(\theta_k | S^a = s^a)^2}}, \quad \delta^a > 0. \tag{15}$$

The basic payoff (15) *can be easily generalized to absolutely continuous distributions (the basic payoff* (13) *used in the contract* (3) *applies only to discrete distributions). The ex-ante expected payoff to agent a with payoff* (14) *with the basic component* (15) *is*

$$Eg^a(\theta, S) = \delta^a \sum_{j=1}^m f^a(\theta_j) \sqrt{\sum_{k=1}^m P(\theta = \theta_k | S^a = \theta_j)^2}. \tag{16}$$

Notice that the expected payoff is maximized when a's signal is perfect (i.e., $P(\theta = \theta_j | S^a = \theta_j) = 1$ *for* $j = 1, \ldots, m$*) and minimized when the signal is uninformative (i.e.,* $P(\theta = \theta_k | S^a = \theta_j) = 1/m$ *for all* i, j*).*

[11] The proof is essentially the same as the continuous-state case in Proposition 5.

7.1 Eliciting Agent Effort

In this section we examine the agents' effort choice problem in the discrete distribution setting (the modifications for the continuous distribution setting are minor). The effort exerted determines the quality of the signals they receive. The conditional probability that agent a's signal is s^a conditional on state θ given effort q^a will be denoted $f^a\left(s^a \middle| \theta; q^a\right)$. We assume for simplicity that effort satisfies $q^a \in [0, \bar{q}]$ for some $\bar{q} > 0$. Proposition 3 above shows that with the payout form (11), the ex-ante expected payoff (12) depends only on the sum, over all states, of getting the right signal conditional on the state. Therefore we need only model the dependence of this sum on q^a, which we denote

$$F^a\left(q^a\right) = \sum_{k=1}^{m} f^a\left(\theta_k \middle| \theta_k; q^a\right), \quad a = 1, \ldots, n;$$

and assume $F^a\left[0, \bar{q}\right] \rightarrow \mathbb{R}_+$ is increasing, concave, and differentiable. The functions F^a and $h^a, a = 1, \ldots, n$, are known by all the agents and the principal.

As in Section 4, we also assume an additive effort cost $h^a\left(q^a\right)$, which is assumed increasing, strictly convex and differentiable, with a derivative satisfying $\lim_{q \rightarrow \bar{q}} \frac{d}{dq} h^a\left(q\right) > \delta^a F^a\left(q^a\right)/q^a$ (in the binary setting, $F^a\left(q^a\right) = 2q^a$).

Each agent a solves the problem

$$\max_{q^a \in [0, \bar{q}]} \delta^a F^a\left(q^a\right) - h^a\left(q^a\right).$$

Under the above assumptions, optimal effort, which we denote $\hat{q}^a\left(\delta^a\right)$, is a continuous and strictly increasing function of δ^a, and therefore any agent effort levels (in the range $[0, \bar{q}]$) can be elicited by the principal through the appropriate choice of δ^a's.

Proposition 4 *Suppose the first-best effort levels are q^a_{FB}, $a = 1, \ldots, n$. Under the truthful revelation equilibrium, the first-best effort levels can be implemented by the principal by offering the contract* (3) *with the scaling parameters*

$$\delta^a = (\hat{q}^a)^{-1}(q^a_{FB}), \quad a = 1, \ldots, n;$$

(i.e., choosing δ^a so that $\hat{q}^a(\delta^a) = q^a_{FB}$).

7.2 The Case of Continuous States

Let $f^a(S^a|\theta)$ denote the density of agent a's signal conditional on the true state θ, let $f^a(S^a)$ denote the marginal density of S^a, and let $p(\theta)$ denote the prior density of state θ. To minimize notation, we will also let $f^a(\theta|S^a)$ denote the true state density conditional on agent a's signal. It is always clear from the context always which function $f^a(\cdot|\cdot)$ we mean.

Proposition 5 *Suppose each agent a is offered the following payout, contingent on the true state and reported signals:*

$$g^a\left(\theta; s^1, \ldots, s^n\right) = \delta^a \frac{f^a(\theta|s^a)}{\sqrt{\int f^a(\theta|s^a)^2 d\theta}} \prod_{k \neq a} \frac{f^k\left(s^k\right)}{f^k(s^k|\theta)}, \quad a = 1, \cdots, n. \tag{17}$$

Then truthful reporting is a Nash equilibrium regardless of how many other signals each agent observes. The ex-ante expected payout to each agent a in this equilibrium *is*

$$E\left(g^a\left(\theta, S^1, \ldots, S^n\right)\right) = \delta^a \int_{s^a} f^a\left(s^a\right) \sqrt{\int f^a\left(\theta | s^a\right)^2 d\theta}\, ds^a, \tag{18}$$

and is the same as in the single-agent case, regardless of the number of agents.

The key difference in the payout (17) compared to (11) is the change in the basic part of the payout, which the continuous-state analog of the discrete-time case in Remark 1. The following example considers multivariate normal state and signals.

Example 3 (normally distributed signals) *Suppose θ is distributed standard normal, let each agent a's signal has mean θ and standard deviation σ^a. Then the contract* (17) *is*[12]

$$g^a\left(\theta; s^1, \ldots, s^n\right) = \frac{\delta^a}{\sqrt{v^a}\pi^{1/4}} \exp\left(-\frac{1}{2}\left(\frac{\theta - \mu^a\left(s^a\right)}{v^a}\right)^2\right) \prod_{k \neq a} v^k \exp\left(\frac{1}{2}\left[\left(\frac{\theta - \mu^k\left(s^k\right)}{v^k}\right)^2 - \theta^2\right]\right).$$

where the conditional moments $\mu^a\left(s^a\right) = E\left(\theta | S^a = s^a\right)$ and $v^a = \sqrt{Var\left(\theta | S^k = s^k\right)}$ are given by

$$\mu^a\left(s^a\right) = \frac{s^a}{1 + \left(\sigma^a\right)^2} \quad v^a = \sqrt{\frac{\left(\sigma^a\right)^2}{1 + \left(\sigma^a\right)^2}}, \quad a = 1, \cdots, n.$$

The conditional and unconditional expected payouts (from (22)*) are*

$$E\left(g^a\left(\theta, S^a\right)\right) = E\left(g^a\left(\theta, S^a\right) | S^a\right) = \frac{\delta^a}{\sqrt{2v^a}\pi^{1/4}},$$

[12] See Section 9 for the derivations.

which is monotonically decreasing in σ^a. *In the single-agent case, agent a's payout is maximized when* $\theta = \mu^a(s^a)$; *that is, when the realized state equals the expected value of the state conditional on the agent's reported signal. The payout is more concentrated for states close to* $\mu^a(s^a)$ *the more precise agent a's signal. The "disagreement" component of pay is increasing in distance between the realized state and* $\mu^k(s^k)$, *for each* $k \neq a$.

Defining agent a's signal precision $q^a = (\sigma^a)^{-2}$, *then agent a's time-*0 *effort-choice problem is* $\max_{q^a \in [0,\bar{q}]} E(g^a(\theta, S^a)) - h^a(q^a)$ *where* $E(g^a(\theta, S^a)) = \frac{\delta^a}{\sqrt{2}\pi^{1/4}}(1+q^a)^{1/4}$. *The principal can choose* δ^a *to elicit any effort/precision level* $\hat{q}^a \in [0, \bar{q}]$.

8 Conclusion

We show that a simple set of payoff functions generates a Nash equilibrium with truthful voting/reporting and first-best effort in committees with heterogeneous agents and pre-vote communication. The payoff functions are characterized by a positive payoff for a correct vote together with a bonus for correctly disagreeing with other agents (or, equivalently, a bonus for the incorrect votes of others), and can be scaled by the firm to elicit whatever information-aquisition effort is desired by the firm. In the case of a firm choosing whether to accept a project, the optimal decision can be implemented with the following voting rule: accept only if the difference in the weighted votes for and against exceeds some threshold, with weights increasing in the agent's signal precision.

Pre-communication does not impact outcomes because the disagreement component of our payoff function nullifies the advantage from free-riding on others' information. In the truthful equilibrium, this component is increasing in the aggregate signal strength of all the agents receiving the opposite signal. There are also Nash equilibria that support non-truthful voting; however in the binary-signal case, we show that the better informed of the misreporting agents tend to be worse off compared to the truthful-reporting equilibrium. These agents might therefore prefer to keep their signals secret (thus eliminating the possibility of a misreporting equilibrium). While these results make us believe that misreporting equilibria are unlikely, they may also explain empirical findings that agents communicate less than theory predicts.

An interesting extension of our model is to a dynamic setting in which each agent must report an incremental signal on the evolution of some unobserved process in exchange for an incremental payment based only on the reports and the terminal value of the process. A dynamic model could also include uncertainty about each agent's cost-function parameter. Repeated votes on different decisions would allow the principal to update his priors on the parameters and fine-tune the payoff functions to achieve optimality. We leave this to future research.

9 Proofs and Derivations

Proof of Proposition 1

We consider agent n's decision on whether to vote his signal. The arguments for the other agents' choices are the same. Assuming agents $1, \ldots, n-1$ report truthfully, the conditional expected payoff from agent n reporting signal s^n is

$$E\left\{ g^n\left(\theta, S^1, \ldots, S^{n-1}, s^n\right) \middle| S\right\} = \delta^n \frac{P\left(\theta = s^n \middle| S\right)}{p\left(s^n\right)} \prod_{a<n} \frac{P\left(S^a\right)}{P\left(S^a \middle| \theta = s^n\right)}.$$

Now substitute $P\left(\theta = s^n \middle| S\right) = \frac{p(s^n)}{P(S)} P\left(S^n \middle| \theta = s^n\right) \prod_{a<n} P\left(S^a \middle| \theta = s^n\right)$ to get

$$E\left\{ g^n\left(\theta, S^1, \ldots, S^{n-1}, s^n\right) \middle| S\right\} = \left\{\delta^n \frac{P\left(S^n \middle| \theta = s^n\right)}{P\left(S^n\right)}\right\} \frac{P\left(S^n\right) \prod_{a<n} P\left(S^a\right)}{P\left(S\right)}. \quad (19)$$

Because $P\left(S^n \middle| \theta = S^n\right) = q^n > P\left(S^n \middle| \theta = 1 - S^n\right) = 1 - q^n$, then reporting $s^n = S^n$ is optimal given the contract (3) with all signals observed. Letting

$$f^n\left(\theta; s^n\right) = \delta^n \frac{1_{\{\theta = s^n\}}}{p\left(s^n\right)}$$

denote the basic component of agent n's promised payoff, note that

$$E\left\{ f^n\left(\theta, s^n\right) \middle| S^n\right\} = \delta^n \frac{P\left(S^n \middle| \theta = s^n\right)}{P\left(S^n\right)},$$

which matches the braced term on the right side of (19).

If only a subset of the other agents' signals is observed, then iterated expectation implies that truthful signal revelation again results in a higher expected payout (conditional on the observed signals).

To obtain the expected payout conditional only on S^n we use

$$\begin{aligned} E\left\{ \frac{P\left(S^n\right)\prod_{a<n}P\left(S^a\right)}{P\left(S\right)} \right| \left. S^n \right\} &= \sum_{S^1,\ldots S^{n-1}\in\{0,1\}^{n-1}} \frac{P\left(S\right)}{P\left(S^n\right)} \frac{P\left(S^n\right)\prod_{a<n}P\left(S^a\right)}{P\left(S\right)} \\ &= \prod_{a<n}\sum_{S^a\in\{0,1\}} P\left(S^a\right) = 1 \end{aligned}$$

to get

$$E\left\{ g^n\left(\theta, S^1,\ldots,S^{n-1},s^n\right)\right| \left. S^n\right\} = E\left\{ g^n\left(\theta,s^n\right)\right| \left. S^n\right\} = E\left\{ f^n\left(\theta,s^n\right)\right| \left. S^n\right\},$$

which is the same as in the case when agent n is the only agent. The ex-ante expected payout is.

$$\delta^n E\left(\frac{P\left(S^n\right| \left.\theta = S^n\right)}{P\left(S^n\right)}\right) = \delta^n \sum_{S^n\in\{0,1\}} P\left(S^n\right)\frac{P\left(S^n\right| \left.\theta = S^n\right)}{P\left(S^n\right)} = 2\delta^n q^n.$$

Proof of Proposition 2

For any signal-vector S and report s, we get, from (3),

$$\begin{aligned} E\left\{ g^n\left(\theta,s\right)\right| \left. S\right\} &= \hat{\delta}^n\frac{P\left(\theta = s^n\right| \left. S\right)}{p\left(s^n\right)} \prod_{\{a\in\mathcal{A}:\, s^a\neq s^n\}} \frac{q^a}{1-q^a} \\ &= \frac{\hat{\delta}^n}{P\left(S\right)}\left(\prod_{a\in\mathcal{A}} q^a\right)\prod_{a\in\mathcal{A}}\frac{P\left(S^a\right| \left.\theta = s^n\right)}{q^a} \prod_{\{a\in\mathcal{A}:\, s^a\neq s^n\}} \frac{q^a}{1-q^a}. \end{aligned}$$

Defining ψ

$$\psi = \hat{\delta}^n\left(\prod_{a\in\mathcal{A}} q^a\right)/P\left(S\right),$$

(note that ψ does not depend on s^n) this simplifies to

$$E\left\{ g^n(\theta,s) \middle| S\right\} = \psi \left(\prod_{\{a\in\mathcal{A}:\, s^a\neq s^n\}} \frac{q^a}{1-q^a}\right)\left(\prod_{\{a\in\mathcal{A}:\, S^a\neq s^n\}} \frac{q^a}{1-q^a}\right)^{-1} \tag{20}$$
$$= \psi \left(\prod_{\{a\in L:\, s^a\neq s^n\}} \frac{q^a}{1-q^a}\right)\left(\prod_{\{a\in L:\, S^a\neq s^n\}} \frac{q^a}{1-q^a}\right)^{-1},$$

The second expression shows that when determining agent n's optimal report (vote) s^n, only the signals, precisions, and reports of the set L of misreporting agents matters. To simplify notation we consider the case when all agents misreport (or redefine n as the number of misreporting agents, and reorder the agents so that the first n misreport). Let $\mathbf{1}$ denote a vector of ones, and consider the candidate misreporting equilibrium $s = \mathbf{1} - S$ (all agents with zero signals report one, and vice versa). The expected payoff for agent n is

$$E\left\{ g^n(\theta,\mathbf{1}-S) \middle| S\right\} = \psi \left(\prod_{\{a\in L:\, S^a\neq S^n\}} \frac{q^a}{1-q^a}\right)\left(\prod_{\{a\in L:\, S^a= S^n\}} \frac{q^a}{1-q^a}\right)^{-1} = \psi\frac{M^L(1-S^n)}{M^L(S^n)}$$

If agent n deviates from the candidate equilibrium by reporting truthfully, so that $\tilde{s} = (1-S^1,\ldots,1-S^{n-1},S^n)$, the expected payoff is

$$E\left\{ g^n(\theta,\tilde{s}) \middle| S\right\} = \psi \left(\prod_{\{a\in L:\, S^a= S^n\}} \frac{q^a}{1-q^a}\right)\left(\prod_{\{a\in L:\, S^a\neq S^n\}} \frac{q^a}{1-q^a}\right)^{-1} = \psi\frac{M^L(S^n)}{M^L(1-S^n)}\frac{1-q^n}{q^n}.$$

Therefore agent n does *not* deviate if condition (6) holds (for $a = n$ and $s^n = 1-S^n$).

Now consider the strategy of the typical truthful-reporting agent when other agents misreport. Let $L = \{1, \ldots, n\}$ be the set of misreporting agents, $m \in T$ be some agent in the truthful-reporting set of agents, and $S = (S^1, \ldots, S^n, S^m)$ be the relevant vector of signals. Agent m's conditional expected payout for reporting truthfully is

$$E\left\{ g^m\left(\theta, 1 - S^1, \ldots, 1 - S^n, S^m\right) \middle| S \right\} = \frac{M^L(S^m)}{M^L(1 - S^m)} \frac{q^m}{1 - q^m}.$$

If agent m deviates from the candidate equilibrium by misreporting, his expected payout is

$$E\left\{ g^m\left(\theta, 1 - S^1, \ldots, 1 - S^n, 1 - S^m\right) \middle| S \right\} = M^L(1 - S^m) / M^L(S^m).$$

Therefore agent m will not deviate from the candidate equilibrium if condition (6) holds (for $a = m$ and $S^m = s^m$).

Finally, we consider the conditional expected payoff of any agent $a \in L$ in some misreporting equilibrium compared to the truthful equilibrium. For notational simplicity we again let $L = \mathcal{A}$. Suppose $s = \mathbf{1} - S$ satisfies the condition (6) for an equilibrium. Using (20), the ratio of the truthful reporting to misreporting equilibria are (the disagreement reward is the same in both equilibria)

$$\frac{E\{g^n(\theta, S)|\, S\}}{E\{g^n(\theta, \mathbf{1} - S)|\, S\}} = \left(\prod_{\{a \in \mathcal{A} :\, S^a \neq S^n\}} \frac{q^a}{1 - q^a} \right)^{-1} \left(\prod_{\{a \in \mathcal{A} :\, S^a = S^n\}} \frac{q^a}{1 - q^a} \right) = \frac{M(S^n)}{M(1 - S^n)}.$$

Derivation of Example 2

Consider first the payoff (8), and suppose that only agent a exerts effort. Assuming the optimal effort level $\hat{q}^a$ has already been expended (we will solve for $\hat{q}^a$ below), the project will only be accepted if agent a receives (and therefore reports) the signal $S^a = 1$, (which has probability q^a), and $\hat{q}^a > c$. The principal's time-0 problem, for now ignoring the "individual rationality" constraint $\Pi > 0$ (if $\Pi \leq 0$ the principal will let $\hat{q}^a = 1/2$ and forgo investing altogether), is

$$\Pi = \max_{q^a \in [1/2, 3/4]} \left\{ P\left(S^a = 1\right)\left(q^a - c\right) - \frac{1}{2}\left(q^a - \frac{1}{2}\right)^2 \right\}.$$

Substituting $P\left(S^a = 1\right) = 1/2$ we get $\hat{q}^a = 3/4$ and $\Pi = \frac{1}{2}\left(\frac{3}{4} - c\right) - \frac{1}{2}\left(\frac{3}{4} - \frac{1}{2}\right)^2$. It is easy to confirm that $c \geq 11/16$ implies $\Pi \leq 0$, and therefore the firm will choose $\hat{q}^a = \hat{q}^b = 1/2$, which is implemented by letting $\delta^a = \delta^b = 0$.

Now consider the promised payoff (3). Under the optimal policy, the principal accepts the project only if $S^A = S^B = 1$; otherwise, if, for example, the project is also accepted if $S^A = 1, S^B = 0$, then B's signal is irrelevant to the decision, and it would be cheaper to pay only agent A (which would match the solution above). Furthermore, the concavity of the problem implies that equal effort (recall the agents have identical cost functions) must be optimal. Substituting $P\left(S^a = S^b = 1\right) = \frac{1}{2}\left(q^2 + (1-q)^2\right)$ into the principal's problem (9)

$$\Pi = \max_{q \in [1/2, 3/4]} \left\{ \frac{1}{2}\left[(1-c)\,q^2 - c\,(1-q)^2\right] - \left(q - \frac{1}{2}\right)^2 \right\},$$

we get the FOC

$$(1-c)\,\hat{q}+c\,(1-\hat{q})=2\left(\hat{q}-\frac{1}{2}\right),$$

which yields the expression for the common optimal effort level $\hat{q}$. Substituting $\hat{q}$ into the expression for Π we get

$$\begin{aligned}\Pi &= \frac{1}{2}\left[(1-c)\left(\frac{c+1}{2\left(c+\frac{1}{2}\right)}\right)^2 - c\left(\frac{c}{2\left(c+\frac{1}{2}\right)}\right)^2\right] - \frac{1}{4}\left(\frac{1}{2\left(c+\frac{1}{2}\right)}\right)^2 \\ &= \frac{1}{2}\left(\frac{1}{2\left(c+\frac{1}{2}\right)}\right)^2\left\{(1-c)\,(c+1)^2-c^3-\frac{1}{2}\right\}.\end{aligned}$$

It is easy to confirm that $\Pi > 0$ for c in a neighborhood of $11/16$.

Proof of Proposition 3

We present the proof for the payoff function (14) (with a general basic part $\hat{g}$). We consider agent n's decision on whether to vote his signal. The arguments for the other agents' choices are the same. Suppose the contract $g^n(\theta, s^n)$ induces truthful revelation when agent n observes only S^n (i.e., the case when n is the only agent). For the basic payoff (13), this was shown in Section 2. Now suppose agent n observes all n signals, which we denote $S=(S_1,\ldots,S_n)$. Let $f(S)$ denote the marginal probability of S. The conditional expected payoff from reporting signal s^n is

$$\begin{aligned}E\left\{g^n\left(\theta,S^1,\ldots,S^{n-1},s^n\right)\middle|\,S\right\} &= \sum_{j=1}^{m} P\left(\theta_j\middle|\,S\right) g^n\left(\theta_j,S^1,\ldots,S^{n-1},s^n\right) \\ &= \frac{1}{f(S)}\sum_{j=1}^{m} p(\theta_j)\,P\left(S\middle|\,\theta=\theta_j\right)\hat{g}^n\left(\theta_j,s^n\right)\prod_{k<n}\frac{f^k\left(S^k\right)}{f^k\left(S^k\middle|\,\theta_j\right)}.\end{aligned}$$

By conditional independence (Assumption 1), $P\left(S|\,\theta=\theta_j\right)=f^n\left(S^n|\,\theta_j\right)\prod_{k<n}f^k\left(S^k\middle|\,\theta_j\right)$, and therefore the conditional expected payoff is proportional to that in the single-agent case:

$$\begin{aligned} E\left\{g^n\left(\theta,S^1,\ldots,S^{n-1},s^n\right)\middle|\,S\right\} &= \frac{\prod_{k<n}f^k\left(S^k\right)}{f\left(S\right)}\sum_{j=1}^{m}p\left(\theta_j\right)f^n\left(S^n|\,\theta_j\right)\hat{g}^n\left(\theta_j,s^n\right) \\ &= \frac{f^n\left(S^n\right)\prod_{k<n}f^k\left(S^k\right)}{f\left(S\right)}E\left\{\hat{g}^n\left(\theta,s^n\right)|\,S^n\right\}. \end{aligned}$$

Therefore $s^n=S^n$ is optimal given the contract (14) with all signals observed.

If only a subset of the other agents' signals is observed, then iterated expectation implies that truthful signal revelation again results in a higher expected payout (conditional on the observed signal).

Given truthful reporting, the expected payout conditional on S^n is

$$\begin{aligned} E\left\{g^n\left(\theta,S\right)|\,S^n\right\} &= \sum_{S^1,\ldots S^{n-1}\in\{\theta_1,\ldots,\theta_m\}^{n-1}}\frac{f\left(S\right)}{f^n\left(S^n\right)}\frac{f^n\left(S^n\right)\prod_{k<n}f^k\left(S^k\right)}{f\left(S\right)}E\left\{\hat{g}^n\left(\theta,S^n\right)|\,S^n\right\} \\ &= \prod_{k<n}\left(\sum_{S^k\in\{\theta_1,\ldots,\theta_m\}}f^k\left(S^k\right)\right)E\left\{\hat{g}^n\left(\theta,S^n\right)|\,S^n\right\} \\ &= E\left\{\hat{g}^n\left(\theta,S^n\right)|\,S^n\right\}. \end{aligned}$$

Using

$$E\left(\hat{g}^n\left(\theta,s^n\right)|\,S^n\right)=\delta^n\frac{P\left(\theta=s^n|\,S^n\right)}{p\left(s^n\right)}=\delta^n\frac{f^n\left(S^n|\,s^n\right)}{f^n\left(S^n\right)}, \tag{21}$$

we get

$$E\left\{g^n\left(\theta,S^n\right)\right\}=\sum_{j=1}^{m}f^n\left(\theta_j\right)E\left\{\left.\hat{g}^n\left(\theta,\theta_j\right)\right|S^n=\theta_j\right\}=\delta^a\sum_{j=1}^{m}f^n\left(\left.\theta_j\right|\theta_j\right).$$

Proof of Proposition 5

First consider the single-agent case with agent-a payoff function

$$\hat{g}^a\left(\theta;s^a\right)=\delta^a\frac{f^a\left(\left.\theta\right|s^a\right)}{\sqrt{\int f^a\left(\left.\theta\right|s^a\right)^2d\theta}}.$$

Conditional on observing the signal S^a, agent a's expected payout is

$$E\left(\left.\hat{g}^a\left(\theta,s^a\right)\right|S^a\right)=\delta^a\int\frac{f^a\left(\left.\theta\right|s^a\right)}{\sqrt{\int f^a\left(\left.\theta\right|s^a\right)^2d\theta}}f^a\left(\left.\theta\right|S^a\right)d\theta.$$

The Cauchy-Schwartz inequality implies

$$\int\frac{f^a\left(\left.\theta\right|s^a\right)f^a\left(\left.\theta\right|S^a\right)}{\sqrt{\int f^a\left(\left.\theta\right|s^a\right)^2d\theta}}d\theta\leq\sqrt{\int f^a\left(\left.\theta\right|S^a\right)^2d\theta}$$

with equality at $s^a=S^a$ (inequality is strict unless $f^a\left(\left.s^a\right|\cdot\right)\propto f^a\left(\left.S^a\right|\cdot\right)$ some $s^a\neq S^a$). Therefore truthful reporting is optimal in the single-agent case, and yields the conditional expected payoff

$$E\left(\left.\hat{g}^a\left(\theta,S^a\right)\right|S^a\right)=\delta^a\sqrt{\int f^a\left(\left.\theta\right|S^a\right)^2d\theta}. \tag{22}$$

Now consider the case when agent n observes all n signals. Let $f(S|\theta)$ denote the conditional joint density of the vector $S = (S^1, \ldots, S^n)$ given θ, and $f(S)$ the joint marginal density of S. The conditional expected payoff from reporting signal s^n is

$$\begin{aligned} E\left\{ g^n\left(\theta, S^1, \ldots, S^{n-1}, s^n\right) \middle| S \right\} &= \int \frac{f(S|\theta)\, p(\theta)}{f(S)} g^n\left(\theta, S^1, \ldots, S^{n-1}, s^n\right) d\theta \\ &= \frac{f^n(S^n)}{f(S)} \prod_{k<n} f^k\left(S^k\right) \int \frac{f^n(S^n|\theta)\, p(\theta)}{f^n(S^n)} \hat{g}^n(\theta; s^n)\, d\theta \\ &= \left(\frac{f^n(S^n)}{f(S)} \prod_{k<n} f^k\left(S^k\right) \right) E\left\{ \hat{g}^n(\theta, s^n) \middle| S^n \right\}. \end{aligned}$$

where we have used (from Assumption 1) $f(S|\theta) = f^n(S^n|\theta) \prod_{k<n} f^k\left(S^k \middle| \theta\right)$. Therefore the conditional payoff is proportional to the payoff in the single-agent case, and $s^n = S^n$ is optimal.

If only a subset of the other agents' signals is observed, then iterated expectation implies that truthful signal revelation again results in a higher expected payout (conditional on the observed signal).

Given truthful reporting, the expected payout conditional on S^n is

$$\begin{aligned} E\left\{ g^n(\theta, S) \middle| S \right\} &= \int_{S^1} \cdots \int_{S^{n-1}} \frac{f(S)}{f^n(S^n)} \left(\frac{f^n(S^n)}{f(S)} \prod_{k<n} f^k\left(S^k\right) \right) E\left\{ \hat{g}^n(\theta, S^n) \middle| S^n \right\} dS^1 \cdots dS^{n-1} \\ &= \int_{S^1} f^k\left(S^1\right) dS^1 \cdots \int_{S^1} f^k\left(S^{n-1}\right) dS^{n-1} E\left\{ \hat{g}^n(\theta, s^n) \middle| S^n \right\} \\ &= E\left\{ \hat{g}^n(\theta, S^n) \middle| S^n \right\}. \end{aligned}$$

REFERENCES:

ADAM R. AND D. FERREIRA (2007): "A Theory of Friendly Boards," Journal of Finance, 66, 217-250.

AUSTEN-SMITH, D., AND J. BANKS (1996): "Information Aggregation, Rationality, and the Condorcet Jury Theorem," American Political Science Review, 90, 34-45. [896,901]

AUSTEN-SMITH D., AND T. FEDDERSEN (2006): "Deliberation and Voting Rules," American Political Science Review, 100, 209-217. [893,896]

BOND P. AND H. ERASLAN (2010): "Strategic Voting over Strategic Proposals," Review of Economic Studies, 77, 459-490.

BANNERJEE, A. V. (1992): "A Simple Model of Herd Behavior," Quarterly Journal of Economics, 107, 797-817.

BIKHCHANDANI, S., D. HIRSHLEIFER AND I. WELCH (1992): "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades," Journal of Political Economy, 100, 992-1026.

COUGHLAN, P. (2000): "In Defence of Unanimous Jury Verdicts: Mistrials, Communication, and Strategic Voting," American Political Science Review, 94, 375-393. [896,897]

CONDORCET, MARQUIS DE (1785): "Essai sur l'application de l'analyse a la probabilite des decisions rendues a la probabilite des voix,".Paris: De l'imprimerie royal.

FEDDERSEN, T. J., AND W. PESENDORFER (1996): "The Swing Voter's Curse," American Economic Review, 86, 408-424. [896, 897]

FEDDERSEN, T. J., AND W. PESENDORFER (1997): "Voting Behavior and Information Aggregation in Elections With Private Information," Econometrica, 65, 1029-1058. [896]

GERARDI, D., AND L. YARIV (2007): "Deliberative Voting," Journal of Economic Theory, 134, 317-338. [893,896,919]

GERARDI, D., AND L. YARIV (2008): "Information Acquisition in Committees," Games and Economic Behavior, 62, 436-459. [896]

GERSHKOV A., AND B. SZENTES (2009): "Optimal Voting Schemes with Costly Information Acquisition," Journal of Economic Theory, 144, 36-68. [896]

GOEREE, J. K., AND L. YARIV (2011): "An Experimental Study Of Collective Deliberation," Econometrica, 79, 893-921. [893, 896]

HARRIS M. AND A. RAVIV (2008): "A Theory of Board Control and Size," Review of Financial Studies, 21, 1797-1832.

MYERSON R. (1998): "Extended Poisson Games and the Condorcet Jury Theorem," Games and Economic Behavior, 25, 111-131. [896]

PENDERGAST C. (1993): "A Theory of Yes Men," American Economic Review, 83, 757-770.

PERSICO N. (2004): "Committee Design with Endogenous Information," Review of Economic Studies, 71, 165-191. [896]

WELCH, I. (1992) "Sequential Sales, Learning and Cascades," Journal of Finance, 47, 695-733.