Monitoring experts*

Yaron Azrieli†

January 16, 2019

Abstract

A decision maker faces a choice problem under uncertainty and may hire experts to collect information regarding the realized state. The experts choose how much (costly) effort to exert, which determines the quality of information they obtain. Efforts and signal realizations are unobservable; moreover, payments can't be contingent on the realized state. The decision maker thus has to design a contract that induces the experts to 'monitor each other' by making payments contingent on the entire vector of reports. We characterize the information structures that the decision maker can implement. In the special case of binary states and signals we characterize the least costly contract that implements a given information structure and study the tradeoff between the value of information and its cost. In particular, we show that discriminating between the experts is a common feature of an optimal contract.

Keywords: Moral hazard, information acquisition, monitoring, value of information.

JEL Classification: D82, D86.

*I thank comments and suggestions from James Best, Aislinn Bohren, Ozan Candogan, PJ Healy, Teddy Kim, Dan Levin, Jim Peck, Edward Schlee, Eran Shmaya, Richard van Weelden, Alex Wolitzky, and from participants at several seminars and conferences where this work has been presented.

†Dept. of Economics, The Ohio State University, 1945 North High Street, Columbus, Ohio 43210, U.S.A. Email: azrieli.2@osu.edu

1 Introduction

In the classic principal-agent problem, the principal relies on the correlation between the agent's unobservable effort and the observed output to motivate the agent to work. Consider the case where the principal is a decision maker (DM, henceforth) facing a choice among several alternatives whose attractiveness is uncertain; and where the agent is an expert who has no stake in the decision but is capable of collecting relevant information at a cost that increases with informativeness. The 'output' in this case can be thought of as the combination of the expert's signal and the ex-post realized state: If the expert's recommendation turned out to be a good one then output is high, while a bad advice corresponds to low output. Incentives can then be provided based on the correlation between effort and outcomes in a way that resembles the classic case (see below for references to this type of models).

But what if the expert privately observes the data he collects, and, moreover, compensations cannot wait until the DM learns the state or her payoff? To illustrate, suppose that a policy maker contemplates between several proposals to reduce global warming. An expert is tasked with predicting the effectiveness of each of the alternatives. How can the policy maker guarantee that the expert conducts a thorough investigation if the latter has exclusive access to the collected data, and given that uncertainty will be resolved only in the far future? As a second example, consider a political candidate who hires a pollster to estimate public sentiment on a certain issue. The true 'state' will likely never be revealed so is not contractible, and the pollster may be able to misrepresent the data when reporting the results.

A potential solution for the DM, and the one we study in this paper, is to hire several experts and have them 'monitor each other'. The basic idea goes as follows: When an expert exerts a high effort he gets an accurate signal of the state, so if all experts exert high efforts and truthfully report their signals then (under reasonable assumptions) these signals are likely to be close to each other. Thus, by paying high compensations in the event of matching signals and low compensations when a mismatch occurs the DM can incentivize the experts to work hard and to reveal what they find. Put differently, the contract creates a coordination game between the experts, and nature's unknown realized state serves as a focal point; if an expert believes that other experts' reports are likely to concentrate around this focal point, then he has an incentive to collect information so that his report will match the state as well; the DM in turn learns about the state through the experts' reports.[1]

In our model, uncertainty is captured by a finite state-space with a common prior belief shared by the DM and all the experts. All parties are risk-neutral. The DM offers a contract which specifies the payment for each expert as a function of the vector of reported signals she receives from the experts. We assume limited liability – negative payments are not allowed. Each expert then chooses what experiment to conduct and, upon obtaining the results, what signal to report to the DM. Throughout we assume that experts' signals are independent conditional on the state of nature. We say that a contract implements a given vector of information structures if in the game it induces it is an equilibrium for the experts to choose their respective structures and to truthfully report their signals.

The contribution of this paper can be divided into two parts. The first is a detailed study of the

[1] Similar ideas of peer-monitoring appeared in papers by Gromb and Martimort [18], Rahman [33], and Bohren and Kravitz [7]; we discuss the relation to these papers in the related literature section below.

case in which the state-space and signal-space for each expert are binary, and the information structures available to the experts are symmetric between the two states. The information structure of each expert in this case is summarized by a single number – the probability of the signal matching the state, which we assume to be increasing in effort. We are interested in properties of the optimal contract in this environment. Our analysis is based on the Grossman-Hart [19] approach: As a first step, for each vector of information structure find the least costly way for the DM to implement it. Second, once the cost function is obtained, maximize the difference between the value of information and its cost over all implementable vectors of information structures.

Even in this simple environment, and with a convex cost function for effort, the first-order condition with respect to effort is not sufficient to guarantee incentive compatibility. The problem is the manipulability of the experts' reports: Even if there is no profitable deviation from the required effort level under truthful reporting, an expert may nevertheless find it profitable to reduce his effort and misreport his signal. In particular, the first-order approach does not apply in our environment. This is reminiscent of the situation in the standard moral hazard setup when the agent can 'burn' output: Incentive compatibility forces the contract to be monotonic in output even when monotonicity is not implied by the first-order condition (See, e.g., Bolton and Dewatripont [8, page 148]).

Despite this difficulty, it turns out that finding the DM's cost function boils down to solving a linear program whose solution can be characterized: The least costly way to implement any vector of information structures involves paying the experts only in the event where they all report the same signal. This gives an expert the maximal incentive to work relative to the expected cost of the contract with that expert. From this we derive an explicit formula for the cost function. An interesting observation that we can already make at this point is that the expected payment to an expert in the least costly contract increases in that expert's own effort, but decreases in other experts' efforts. This implies in particular that inducing higher effort (and thus obtaining better information) is not necessarily more costly for the DM, a somewhat counterintuitive result.

We then move on to study the value of information for the DM. Instead of focusing on a particular decision problem, our approach here is to compare informativeness according to the Blackwell [5] criterion, i.e., to look for properties of the value of information that all decision makers agree on. The main result of this part is the following: If two sets of experts have the same average accuracy[2] of signals, and in one of these sets the spread of accuracies is larger than the other, then the former is more Blackwell informative than the latter. For example, ignoring the cost, every decision maker in every decision problem prefers two experts with respective accuracies $\frac{7}{8}$ and $\frac{5}{8}$ over two experts each with accuracy of $\frac{3}{4}$, and the latter over 3 experts with accuracy of $\frac{2}{3}$ each. This result, which is of independent interest and may be useful in other applications,[3] expresses non-concavity in the value of information of a different type than in the classic result of Radner and Stiglitz [31] (see also Chade and Schlee [11]).

This non-concavity in the value of information makes solving for the optimal contract a particularly

[2]Accuracy is measured as the increase in the probability that the signal matches the state relative to the uninformative structure where this probability is $\frac{1}{2}$.

[3]In voting, for example, this result implies that few well-informed voters outperform many little-informed voters, given that the average accuracy in the two groups is the same.

difficult task, as it implies that first-order conditions are not sufficient for optimality. We emphasize that this is a separate problem from the failure of the first-order approach discussed above; here the issue is that the mapping from efforts' vectors to DM's utility is not well-behaved. Yet, there are several properties of optimal contracts that we can deduce from comparing the cost and value of information. For the case of two experts, we show that arbitrarily close to any decision problem there is a decision problem in which the optimal contract requires uneven compensations to experts. This conclusion holds uniformly across all cost-of-effort functions. The intuition comes directly from the non-concavity result described above: A given total effort generates the least amount of information when divided equally between the experts.[4] Thus, discrimination between experts naturally follows from optimality considerations and need not be the result of prejudice or bias.[5]

Another property of optimal contracts is that they never involve many low-effort experts. More precisely, we show that if the derivative of the cost function is positive at zero effort, then the cost of hiring n experts uniformly diverges to $+\infty$ as n grows. This implies that a given "budget of effort" should never (i.e., for no decision problem) be divided among many experts, as this is both more costly and less informative than dividing it between a small number of experts. In particular, for the case in which full learning of the state is implementable, we derive a uniform upper bound on the optimal number of experts.

In the second part of the paper we leave the binary-binary model in order to obtain a general characterization of the information structures that the DM can induce the experts to produce. This result is achieved with minimal structural assumptions about the environment. Instead of modeling the experts as choosing efforts and reporting strategies, we assume that each expert simply chooses an information structure from a given set of such structures; in subsection 3.1 and Appendix B we show that an effort and reporting model, as in the binary-binary framework, can be translated into this more abstract formulation. The characterization of implementable vectors of information structures follows Rahman ([32], [33]): An obvious necessary condition for $m = (m_1, \ldots, m_n)$ to be implementable is that no expert i can deviate to another structure m'_i which is less costly than m_i and generates the same distribution over signal vectors as m_i (when combined with m_{-i}). We show that if this condition holds then, under our assumptions, for each deviation m'_i there is a contract which renders it unprofitable. A version of the minmax theorem then implies that there is a single contract that simultaneously discourages all deviations. Hence, this condition is also sufficient.

From this basic characterization, and using the particular structure of our environment (conditional independence, especially), we derive easy-to-check sufficient conditions for implementability. Suppose that m_{-i} is such that the conditional distributions of s_{-i} given the various states of nature, viewed as vectors in $\mathbb{R}^{S_{-i}}$, are linearly independent. Then it is easy to see that the mapping from i's choice m_i to the distribution of the vector s of all n signals is one-to-one. Thus, if m is such that this condition holds for all i, then it follows from the above characterization that m is implementable *regardless of the cost function*. Note that linear independence 'typically' holds whenever the number of states is smaller

[4] On the other hand, the cost is typically convex and hence minimized at the equal split point. This is why our result holds only for a dense set of decision problems and not everywhere.

[5] A similar point in a very different context is made in Winter [39].

than the number of possible s_{-i}'s for each i.[6] An interesting corollary of this result is that in a sense implementation becomes easier when more information is requested. Namely, we show that if m satisfies the above linear independence conditions (and hence is implementable with every cost), then so is every m' in which m'_i is more informative than m_i, $i = 1, \dots, n$. Thus, if a certain m is not implementable then it's not because too much is asked from the experts but rather too little.

An obvious shortcoming of the type of contracts we study is that they do not implement the desired equilibrium uniquely; after all, the experts can coordinate their reports without collecting any information. Full implementation is impossible here, since for every contract a zero-effort equilibrium exists. However, the 'good' equilibrium with information acquisition has an appeal that is not captured by the abstract description of the game, but stems from the context of the problem. Collecting information creates a focal course of action even when ex-ante none of the options stands out. If the likelihood of successful coordination increases in the focality of one of the actions, as the theory of Schelling [36] suggests, then we should expect agents to be willing to incur a cost in order to create a focal point. We further discuss this issue in the concluding section.

1.1 Related literature

This paper combines elements from several strands of literature, including moral hazard, monitoring design, value of information, and costly information acquisition. There are two previous papers we are aware of that explicitly consider peer-monitoring of information providers: In Bohren and Kravitz [7] the principal faces an infinite stream of identical decision problems, each with a fixed positive payoff if her action matches the state and a payoff of 0 otherwise. The principal can hire workers to verify the state at a cost, and the main interest is in the optimal rate of monitoring – how often should two workers (and not just one) be assigned to the same problem to make sure that reports about the state are genuine, and how the optimal monitoring structure depends on the commitment power of the principal. The second paper is Gromb and Martimort [18] who study a model of delegated expertise and compare the case of a single expert with two signals to the case of two experts with one signal each. The DM in their model can either undertake a project or not, and the ex-post outcome (the state) is observable and contractible when the project is undertaken. In the two experts case the optimal contract involves payments contingent not only on the outcome of the project, but also on whether the two reports agree. This additional instrument makes hiring two experts better for the DM than hiring just one. Gromb and Martimort's focus is on the implications of the possibility of collusion between the experts for the optimal contract and the principal's payoff.[7]

In the above two papers the workers/experts face a binary choice of either exerting effort or not, so there is no scope to study the tradeoff between the quality of information and its cost. The richness of our environment uncovers properties such as the non-monotonicity of the DM's cost function and the asymmetry of the optimal contract that cannot be discussed in these previous models. Moreover, we do not restrict attention to a particular decision problem as the other papers do, and instead study general

[6] Our linear independence condition resembles that of Crémer and McLean [13], see below for more on the differences between the results.

[7] Collusion here means that the experts play other equilibria than the one intended by the principal.

properties of optimal contracts that are satisfied uniformly across all problems. Finally, the second part of this paper on implementability has no parallel in these previous works.

Rahman [33] emphasizes the ability of a principal to monitor workers by secretly recommend actions and base compensations on reported signals as well as on these recommendations.[8] In the leading example he shows how this can be used to 'monitor the monitor': If a worker was (secretly) asked to shirk and the monitor reported otherwise then the monitor gets punished. The current paper suggests that an alternative solution for the principal is to hire two monitors and pay them only if their reports match. While this still requires the principal to secretly ask the worker to shirk sometimes (to keep the monitors uncertain), the payments need not depend on the recommendation.

The second part of this paper is based on Rahman's characterization of implementability in [32] and [33]. Specifically, our Theorem 3 can be deduced from the results in [32], see in particular Theorem 2 of that paper. For completeness we provide a full proof that uses the additional structure of our environment. This structure also allows us to get the simple sufficient conditions for implementability in Proposition 5 and its corollaries. The linear independence condition of Proposition 5 resembles that of Crémer and McLean [13] in their work on full extraction of surplus in auctions. While the results are related, note that our proposition does not pertain to efficient or revenue-maximizing implementation. Moreover, in our model the experts *choose* the information structure, while in Crémer and McLean the information structure is exogenously given and agents only choose which signal to report to the auctioneer. A recent work of Bikhchandani and Obara [4] extends the results of Crémer and McLean to environments where agents can acquire additional information, but their interest is also in efficiency and full surplus extraction. We refer the reader to [33, Section VI] for a thorough discussion of the connection between the characterization of implementability and previous works in mechanism design and repeated games.

There are two additional papers in which a decision maker hires a group of disinterested experts to collect information. In Dewatripont and Tirole [14] two 'advocates' may bring information in favor of opposite alternatives. Since each of the advocates is restricted to investigate a different (independent) dimension of the state, it is impossible to generate incentives by comparing their messages. Instead, Dewatripont and Tirole assume that collecting information sometimes generates 'hard evidence' that cannot be presented otherwise. Khanna et al. [22] suggest a mechanism that induces a group of experts to acquire costly information and truthfully reveal their signals. They further show that if agents can communicate (using cheap-talk messages) before submitting their reports then there is no equilibrium in which agents lie about their signals. However, their mechanism uses the ex-post realized state, so cannot be applied in our setup.

There are quite a few papers in which a single expert is hired to collect (costly) information. See Osband [29], Zermeño, [40], Rappoport and Somma [34], Chade and Kovrijnykh [10], Carroll [9], Clark [12], and Häfner and Taylor [20], among others. These models differ from each other in various dimensions, but in all of them the set of contractible variables is sufficient to provide incentives, so monitoring the expert is not necessary. Since in our model payments can only depend on the unverifiable reports

[8]See also Strausz [37] on the connection between Rahman's paper and the classic mechanism design framework of Myerson [28].

of the experts, clearly there is no way for the DM to provide incentives with a single expert.

The experts in our model have no stake in the choice of the DM, which separates our framework from most of the extensive literature on 'cheap talk' and 'Bayesian persuasion'. But there are several papers in this literature demonstrating that the receiver can significantly improve the quality of information she gets by comparing messages from multiple senders. Some examples are Krishna and Morgan [23], Battaglini [3], and Gentzkow and Kamenica [15]. Another strand of relevant literature studies the design of committees when committee members may acquire information prior to voting on an issue they all care about, e.g. Persico [30], Martinelli [26], Gerardi and Yariv [16], and Gershkov and Szentes [17]. In these papers incentives to collect information are provided through the voting rule and not through transfers.

Finally, the classic works of Alchian and Demsetz [1] and Holmstrom [21] offer different views regarding the role of principals. The first paper argues that monitoring workers is one of the main roles of the owner of a firm. The DM in our model does not directly monitor the experts e.g. for lack of appropriate knowledge/technology, and instead relies on the experts to monitor each other. However, our results on optimal contracts (Corollary 2, in particular) imply that sometimes experts should be hired *only for the sake of monitoring other experts*, i.e., their input is not taken into account at all in the decision-making process. This can be seen as a form of 'specialization in monitoring' that the theory of Alchian and Demsetz implies. The second paper argues that the principal role is to break the budget-balance constraint and shows that group penalties when output is low can incentivize effort. The optimal contract we derive has the same flavor of group penalties, where no expert is getting paid unless all signals agree, though budget-balance is not straightforward to define in our setup of information provision.

2 A binary-binary model

2.1 The setup

A decision maker (DM) faces a decision problem under uncertainty. There are two possible states − Black (B) or White (W). The prior probability of state B is γ, and we assume that $\frac{1}{2} \leq \gamma < 1$ (the case $0 < \gamma < \frac{1}{2}$ can be obtained by reversing the labels of the states).

The DM hires a set N of $|N| = n$ risk-neutral experts to collect information about the realized state. Each expert $i \in N$ chooses an effort level $e_i \in \left[0, \frac{1}{2}\right]$. The cost of effort is described by the function $c : \left[0, \frac{1}{2}\right] \to \mathbb{R}$. We assume that c is strictly increasing, strictly convex, three times continuously differentiable, and that $c(0) = 0$. Let $\mathcal{C}$ be the set of all cost functions with these properties.

Each expert privately observes a signal from $S_i = \{b, w\}$, where the distribution over signals conditional on each state depends on the effort level that the expert exerts. Specifically, if i chooses e_i then he observes the 'correct' signal with probability $0.5 + e_i$ and the 'wrong' signal with probability $0.5 - e_i$. Thus, i's information structure is described by the stochastic matrix

$$m_i(e_i) := \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 0.5+e_i & 0.5-e_i \\ \hline W & 0.5-e_i & 0.5+e_i \\ \hline\end{array}$$

Note that no effort leads to uninformative signal, and that informativeness increases with effort. We assume that signals for different experts are independent conditional on the state. Given the vector of effort levels $e = (e_1, \ldots, e_n)$, denote by $m(e)$ the information structure obtained by observing the signals of all the experts.

The experts have no stake in the decision, and the DM may offer monetary compensation for their efforts. However, effort is unobservable and realized signals are privately observed by the experts and are unverifiable. Moreover, compensations occur immediately after the experts report their signals, so transfers can't be contingent on the true state. We consider 'direct mechanisms' in which each expert submits a report $s_i \in S_i$ and gets compensated based on the entire vector of reports $s \in S := \times_{i=1}^n S_i$. Thus, a contract is a list $x = (x_1, \ldots, x_n)$ with each $x_i : S \to \mathbb{R}_+$. Note that we assume that payments are non-negative, which captures 'limited liability' on the part of the experts.

A contract x induces a game between the experts. A pure strategy for expert i in this game is a pair (e_i, r_i), where $e_i \in \left[0, \frac{1}{2}\right]$ is i's effort level and $r_i : S_i \to S_i$ is the report that i sends to the DM as a function of the signal he observed. The payoff to expert i given strategy profile $(e, r) = ((e_1, \ldots, e_n), (r_1, \ldots, r_n))$ is

$$U_i(e, r;\ x_i) := \mathbb{E}_{(e,r)}\left[x_i(s)\right] - c(e_i), \tag{1}$$

where the distribution of s used to calculate the expectation is derived from the strategies (e, r) by

$$\mathbb{P}_{(e,r)}(s) = \sum_{s' \in r^{-1}(s)} \left[\gamma \prod_{\{j\ : s'_j = b\}} (0.5 + e_j) \prod_{\{j\ : s'_j = w\}} (0.5 - e_j) + (1-\gamma) \prod_{\{j\ : s'_j = w\}} (0.5 + e_j) \prod_{\{j\ : s'_j = b\}} (0.5 - e_j)\right].$$

It will be convenient to introduce the following notation. For every subset of experts $A \subseteq N$ and every efforts' vector e let $e(A) = \prod_{j \in A}(0.5 + e_j)$ and $\bar{e}(A) = \prod_{j \in A}(0.5 - e_j)$. Thus, $e(A)$ is the probability that all experts in A obtain the 'correct' signal, and $\bar{e}(A)$ is the probability they all obtain the 'wrong' signal. Given a vector of signals $s \in S$ denote $s^b = \{j\ : s_j = b\}$ and $s^w = N \setminus s^b = \{j\ : s_j = w\}$. Finally, let $r^* = (r_1^*, \ldots, r_n^*)$ denote the vector of truthful reporting strategies. Using this notation we have that

$$\mathbb{P}_{(e,r^*)}(s) = \gamma e\left(s^b\right) \bar{e}\left(s^w\right) + (1-\gamma) e\left(s^w\right) \bar{e}\left(s^b\right).$$

Say that a contract x implements the vector of efforts $e = (e_1, \ldots, e_n)$ if (e, r^*) is an equilibrium of the game induced by x with payoff functions as in (1); using Myerson's [28] terminology, x implements e if honesty (truthful reporting) and obedience (choosing the desired effort level) is a best response for each expert given that all other experts are honest and obedient. Efforts' vector e is implementable if there exists a contract x that implements it.

2.2 Cost of information

There would typically be many contracts x that implement a given e. Let $\psi_i(e)$ be the minimal expected payment that the DM would need to make to expert i in a contract that implements e. Formally, $\psi_i(e)$ is the value of the minimization problem (COST) given by

$$\psi_i(e) = \min_{x_i} \quad \mathbb{E}_{(e,r^*)}[x_i(s)] \tag{COST}$$

$$s.t. \qquad (e_i, r_i^*) \in \arg\max_{(e_i', r_i')} \left\{ U_i\left((e_i', r_i'), (e_{-i}, r_{-i}^*);\ x_i\right)\right\} \quad \textit{and} \quad x_i(s) \geq 0 \ \ \forall s \in S.$$

The following theorem gives the solution to program (COST) and the cost function of the DM $\psi_i(e)$. To state the result it will be useful to write $\underline{b}$ ($\underline{w}$) for the vector of reports $s \in S$ in which $s_i = b$ ($s_i = w$) for all i. Also, we will use the shorter notation $N_{-i} = N \setminus \{i\}$ and $N_{-ij} = N \setminus \{i, j\}$.

Theorem 1. *Let $e = (e_1, \ldots, e_n)$ be such that $0 < e_i < 0.5$ for every i. Then e is implementable, and a solution to program (COST) is given by*

$$\begin{aligned} x_i^*(s) \quad = \quad & \frac{1}{\gamma(1-\gamma)\left[e(N_{-i})^2 - \bar{e}(N_{-i})^2\right]} \times \\ & \begin{cases} \gamma\bar{e}(N_{-i})[(0.5-e_i)c'(e_i)+c(e_i)] + (1-\gamma)e(N_{-i})[(0.5+e_i)c'(e_i)-c(e_i)] & \textit{if } s = \underline{b}, \\ \gamma e(N_{-i})[(0.5-e_i)c'(e_i)+c(e_i)] + (1-\gamma)\bar{e}(N_{-i})[(0.5+e_i)c'(e_i)-c(e_i)] & \textit{if } s = \underline{w}, \\ 0 & \textit{otherwise.} \end{cases} \end{aligned}$$

Furthermore, the cost function for the DM is given by

$$\psi_i(e) = \frac{\left[\gamma\bar{e}(N) + (1-\gamma)e(N)\right]\left[\gamma e(N_{-i}) + (1-\gamma)\bar{e}(N_{-i})\right]c'(e_i) + (2\gamma-1)e(N_{-i})\bar{e}(N_{-i})c(e_i)}{\gamma(1-\gamma)\left[e(N_{-i})^2 - \bar{e}(N_{-i})^2\right]}.$$

Proof. We break the proof into four steps. Proofs of auxiliary lemmas appear in Appendix A.

Step 1: Simplifying the constraints

Lemma 1 below shows that one can replace the incentive compatibility constraint in (COST) by three linear constraints: Equation (2) is the first-order condition with respect to effort at e_i; by convexity of c it is necessary and sufficient for deviations to other effort levels to be unprofitable (assuming honest reporting). Inequality (3) guarantees that deviating to zero effort and constant reporting $r_i \equiv b$ is not profitable. Similarly, inequality (4) is the constraint associated with the deviation to zero effort and constant reporting $r_i \equiv w$.

Lemma 1. *A contract $x_i : S \to \mathbb{R}_+$ is feasible for program (COST) if and only if it satisfies the following* constraints:[9]

$$\sum_{s_{-i}\in S_{-i}}\left[\gamma e\left(s_{-i}^{b}\right)\bar{e}\left(s_{-i}^{w}\right)-(1-\gamma)e\left(s_{-i}^{w}\right)\bar{e}\left(s_{-i}^{b}\right)\right]\left[x_i\left(b,s_{-i}\right)-x_i\left(w,s_{-i}\right)\right] \quad = \quad c'(e_i), \tag{2}$$

$$\sum_{s\in S}\mathbb{P}_{(e,r^*)}(s)x_i(s)-c(e_i)\geq\sum_{s_{-i}\in S_{-i}}\mathbb{P}_{\left(e_{-i},r_{-i}^*\right)}\left(s_{-i}\right)x_i\left(b,s_{-i}\right), \tag{3}$$

$$\sum_{s\in S}\mathbb{P}_{(e,r^*)}(s)x_i(s)-c(e_i)\geq\sum_{s_{-i}\in S_{-i}}\mathbb{P}_{\left(e_{-i},r_{-i}^*\right)}\left(s_{-i}\right)x_i\left(w,s_{-i}\right). \tag{4}$$

The constraints (2)-(4), together with the objective $\mathbb{E}_{(e,r^*)}[x_i(s)]$ and the non-negativity constraints define a linear program that, given Lemma 1, is equivalent to (COST). We refer to this auxiliary program as (AUX):

$$\min_{x_i} \quad \mathbb{E}_{(e,r^*)}[x_i(s)] \tag{AUX}$$

$$s.t. \qquad (2)-(4) \quad \textit{and} \quad x_i(s)\geq 0 \quad \forall s\in S.$$

Step 2: x_i^* is feasible for (AUX)

First, since c is convex and satisfies $c(0)=0$ we have that $c(e_i)\leq e_i c'(e_i)$. This implies that x_i^* in the statement of the theorem is non-negative. Second, plugging x_i^* to the constraints (2)-(4) gives

$$\left[\gamma e\left(N_{-i}\right)-(1-\gamma)\bar{e}\left(N_{-i}\right)\right]x_i^*\left(\underline{b}\right)-\left[\gamma\bar{e}\left(N_{-i}\right)-(1-\gamma)e\left(N_{-i}\right)\right]x_i^*\left(\underline{w}\right)=c'(e_i),$$

$$\left[\gamma e(N)+(1-\gamma)\bar{e}(N)\right]x_i^*\left(\underline{b}\right)+\left[\gamma\bar{e}(N)+(1-\gamma)e(N)\right]x_i^*\left(\underline{w}\right)-c(e_i)\geq\left[\gamma e\left(N_{-i}\right)+(1-\gamma)\bar{e}\left(N_{-i}\right)\right]x_i^*\left(\underline{b}\right),$$

and

$$\left[\gamma e(N)+(1-\gamma)\bar{e}(N)\right]x_i^*\left(\underline{b}\right)+\left[\gamma\bar{e}(N)+(1-\gamma)e(N)\right]x_i^*\left(\underline{w}\right)-c(e_i)\geq\left[\gamma\bar{e}\left(N_{-i}\right)+(1-\gamma)e\left(N_{-i}\right)\right]x_i^*\left(\underline{w}\right),$$

respectively. We leave it for the interested reader to verify that these all indeed are satisfied.[10] It follows that x_i^* is feasible for (AUX).

Step 3: The dual of (AUX) and a feasible solution

The dual of program (AUX) is given by the following:

$$\max_{z_1,z_2,z_3} \quad \{c'(e_i)z_1+c(e_i)(z_2+z_3)\} \tag{DUAL}$$

$$s.t. \qquad z_2,z_3\geq 0, \quad \textit{and for every} \quad s_{-i}\in S_{-i}$$

$$\left[\gamma e(s_{-i}^{b})\bar{e}(s_{-i}^{w})-(1-\gamma)e(s_{-i}^{w})\bar{e}(s_{-i}^{b})\right]z_1-\mathbb{P}_{(e,r^*)}(w,s_{-i})z_2+\mathbb{P}_{(e,r^*)}(b,s_{-i})z_3\leq\mathbb{P}_{(e,r^*)}(b,s_{-i}) \tag{5}$$

[9] As is standard, (s_i,s_{-i}) with $s_i=b$ or $s_i=w$ denotes the vector of reports (or signals) in which i reports s_i and all other experts report according to s_{-i}.

[10] We note that the middle constraint holds as equality and that the last inequality boils down to $c(e_i)\leq e_i c'(e_i)$.

and

$$\left[\gamma e(s^b_{-i})\bar{e}(s^w_{-i}) - (1-\gamma)e(s^w_{-i})\bar{e}(s^b_{-i})\right] z_1 - \mathbb{P}_{(e,r^*)}(w,s_{-i})z_2 + \mathbb{P}_{(e,r^*)}(b,s_{-i})z_3 \geq -\mathbb{P}_{(e,r^*)}(w,s_{-i}). \quad (6)$$

Lemma 2 below introduces a particular vector $z^* = (z_1^*, z_2^*, z_3^*)$ and argues that it is feasible for (DUAL). To prove the lemma we first show that the constraint (5) associated with $s_{-i} = \underline{b}_{-i}$ is satisfied at z^* (with equality), and then argue that constraints (5) associated with other s_{-i}'s are less stringent at z^* and hence satisfied as well. A similar argument applies for the set of constraints (6), with $\underline{w}_{-i}$ taking the role of $\underline{b}_{-i}$.

Lemma 2. Let

$$\begin{aligned} z_1^* &= \frac{(\gamma\bar{e}(N) + (1-\gamma)e(N))\,(\gamma e(N_{-i}) + (1-\gamma)\bar{e}(N_{-i}))}{\gamma(1-\gamma)\left[e(N_{-i})^2 - \bar{e}(N_{-i})^2\right]}, \\ z_2^* &= \frac{(2\gamma-1)e(N_{-i})\bar{e}(N_{-i})}{\gamma(1-\gamma)\left[e(N_{-i})^2 - \bar{e}(N_{-i})^2\right]}, \quad \textit{and} \;\; z_3^* = 0. \end{aligned}$$

Then $z^* = (z_1^*, z_2^*, z_3^*)$ is feasible for (DUAL).

Step 4: x_i^* is optimal for (COST)

The value of the objective of (AUX) at x_i^* is

$$\mathbb{E}_{(e,r^*)}[x_i^*(s)] = \mathbb{P}_{(e,r^*)}(\underline{b})x_i^*(\underline{b}) + \mathbb{P}_{(e,r^*)}(\underline{w})x_i^*(\underline{w}),$$

and that of the objective of (DUAL) at z^* is

$$c'(e_i)z_1^* + c(e_i)z_2^*.$$

It is immediate to check that these two values coincide (and also coincide with the formula for $\psi_i(e)$ given in the statement of the theorem). Therefore, by the weak duality theorem of linear programming, x_i^* is optimal for (AUX). From Lemma 1 it now follows that x_i^* is optimal for (COST) as well, which completes the proof. □

Before proceeding it is worth making several remarks regarding Theorem 1:

Remark 1. If the cost function c satisfies $c'(0.5) < +\infty$ then the theorem remains true for any vector e with $0 < e_i \leq 0.5$, i.e., even if some experts' signals fully reveal the state. The only difference in the proof is that the equality in the first-order condition (2) is replaced by a greater-or-equal inequality, but at the optimum this constraint binds so the result is unchanged.

Remark 2. If $e_i = 0$ for some expert i then $x_i \equiv 0$ solves (COST) and $\psi_i(e) = 0$. In addition, the signals obtained from zero-effort experts are uninformative. We can therefore restrict attention only to experts that exert strictly positive efforts. However, for e to be implementable it is necessary (and sufficient) that at least two experts exert effort. See Section 3 for a more general characterization of implementable information structures.

Remark 3. Consider the case of a uniform prior $\gamma = 0.5$. The cost function $\psi_i(e)$ in this case boils down to

$$\psi_i(e) = \frac{e(N) + \bar{e}(N)}{e(N_{-i}) - \bar{e}(N_{-i})} c'(e_i).$$

Moreover, the cost-minimizing contract is not unique, and one of the solutions to program (COST) is given by

$$x_i(\underline{b}) = x_i(\underline{w}) = \frac{c'(e_i)}{e(N_{-i}) - \bar{e}(N_{-i})},$$

and $x_i(s) = 0$ otherwise. Note that this solution is different than x_i^* (with $\gamma = 0.5$) given in the statement of the theorem. With the above solution both constraints (3) and (4) do not bind, i.e., there is an optimal solution to the relaxed program containing only the local effort constraint (2) which is feasible for the original program (COST). This however is not true whenever $\gamma \neq 0.5$: The optimal solution in the relaxed program with constraint (2) alone violates one of the 'global' constraints (3) or (4).

Remark 4. Consider a variant of our model in which the DM directly observes the experts' signals, or, alternatively, that experts' reports are freely verifiable. Experts here only choose efforts (and not reporting strategies), so incentive compatibility is characterized by the first-order condition (2) only. We distinguish between two cases: If $\gamma = 0.5$ then it should be clear from the previous Remark 3 that the cost function $\psi_i(e)$ remains unchanged, and so nothing changes in our subsequent analysis. But when $\gamma > 0.5$ it is not hard to show that the cost-minimizing contract is given by $x_i(\underline{b}) = \frac{c'(e_i)}{\gamma e(N_{-i}) - (1-\gamma)\bar{e}(N_{-i})}$ and $x_i(s) = 0$ for all other $s \in S$. That is, experts are paid only if they all obtain the ex-ante more likely signal. This in turn gives an expected payment to expert i of

$$\frac{\gamma e(N) + (1-\gamma)\bar{e}(N)}{\gamma e(N_{-i}) - (1-\gamma)\bar{e}(N_{-i})} c'(e_i).$$

The (strictly positive) difference between this last expression and $\psi_i(e)$ of Theorem 1 is the additional rent that expert i derives due to the exclusive access he has to the signal.

Remark 5. As the optimal contract x_i^* pays positive amounts only in the events $s = \underline{b}$ and $s = \underline{w}$, it is interesting to think in which of these two the payment is higher. A direct comparison shows that $x_i^*(\underline{b}) > x_i^*(\underline{w})$ if and only if $c(e_i) < c'(e_i)[e_i + 0.5 - \gamma]$. Thus, for a given e_i, if state B is sufficiently likely (γ large enough) then the payment at the less likely event $\underline{w}$ is higher, while if the prior is close to uniform then the payment at the more likely event $\underline{b}$ is higher (recall that $c(e_i) < e_i c'(e_i)$). To understand why, note that the binding constraint which determines how the total payment is divided between these two events is (3, which requires that a deviation to no effort and constant reporting of b is not profitable. As γ increases this deviation becomes more attractive, which forces the relative payment in the less likely event $s = \underline{w}$ to increase.

A corollary of Theorem 1 is that the total expected payment to an expert in the least costly contract increases in that expert's own effort and decreases in other experts' efforts (the proof is in Appendix A).

Corollary 1. For any interior e and two experts $i \neq j$ it holds that $\frac{\partial \psi_i(e)}{\partial e_i} > 0$ and $\frac{\partial \psi_j(e)}{\partial e_i} < 0$.

Let $\psi(e) = \sum_{i=1}^{n} \psi_i(e)$ be the total expected cost for the DM in the least costly contract that implements e. Corollary 1 implies that in $\frac{\partial \psi(e)}{\partial e_i} = \frac{\partial \psi_i(e)}{\partial e_i} + \sum_{j \neq i} \frac{\partial \psi_j(e)}{\partial e_i}$ the first term is positive and the second is negative. As we shall see below (e.g. in Proposition 2), it may happen that the second effect is stronger, so that the expected cost to the DM is decreasing in the effort required from i.

2.3 Value of information

Finding the cost of obtaining information, we now consider the value of information for the DM. A decision problem can be described by a set of possible alternatives and a utility function that maps each alternative-state pair to the reals. Any decision problem induces a function $v : \Delta(\{B, W\}) \to \mathbb{R}$ that assigns to each belief the maximal achievable expected utility for the DM given that belief.[11] For any decision problem the induced function v is the point-wise maximum of a family of linear functions and is therefore convex and continuous. Conversely, any convex and continuous v can be obtained from some decision problem (see, e.g., Azrieli and Lehrer [2]). It will be convenient to work directly with the 'value function' v rather than explicitly modeling decision problems. We denote by $q \in [0, 1]$ the DM's belief that the state is B and identify $\Delta(\{B, W\})$ with [0,1]. Let $\mathcal{V}$ be the set of all convex and continuous functions $v : [0, 1] \to \mathbb{R}$.

Example 1. Suppose that the set of available alternatives is $\{B, W\}$ (same as the set of states) and that the DM gets a utility of 1 if her choice matches the state and a utility of 0 otherwise. Then the induced v is given by $v(q) = \max\{q, 1 - q\}$.

Example 2. Suppose that the DM needs to choose between a safe alternative S and a risky alternative R. Choosing S yields a sure utility of 0, while choosing R yields a utility of 1 in state B and a utility of -1 in state W. The corresponding v is then $v(q) = 0$ for $0 \leq q \leq 0.5$ and $v(q) = 2q - 1$ for $0.5 < q \leq 1$.

Example 3. Let the set of alternatives be the unit interval $[0, 1]$, and the utility function be $u(a, B) = -(1 - a)^2$ and $u(a, W) = -a^2$ for every alternative $a \in [0, 1]$. Then it is well-known and easy to check that when the DM's belief is q her optimal choice is $a = q$. This gives $v(q) = -q(1 - q)$.

After receiving the vector of signals s from the information structure $m(e)$, the DM updates her belief using Bayes rule and chooses the alternative that maximizes her expected utility. If we let M_e denote the distribution over posterior beliefs induced by $m(e)$ (and the prior γ), then the value of information structure $m(e)$ in decision problem $v \in \mathcal{V}$ is

$$V_v(e) := \int_0^1 v(q) dM_e(q).$$

Our goal in this section is to formulate a condition on a pair of efforts' vectors e, e' which guarantees that $V_v(e) \geq V_v(e')$ for every $v \in \mathcal{V}$ (i.e., for every decision problem); this will later allow us to

[11] Existence of a maximum is guaranteed under appropriate compactness and continuity assumptions.

draw general conclusions about which effort vectors may be optimal for the DM. As is well-known since Blackwell [5], this relation between information structures can also be described through their stochastic matrices (the 'garbling' condition), or by the distributions over posteriors $M_e, M_{e'}$ they induce (the 'mean-preserving spread' condition). We will describe this relation by saying that $m(e')$ is a garbling of $m(e)$, or that $m(e)$ is more informative than $m(e')$.

For the rest of this section we assume (without loss) that efforts are ordered in decreasing order from highest to lowest. Consider two efforts' vectors $e = (e_1 \geq e_2 \geq \ldots \geq e_n)$ and $e' = (e'_1 \geq e'_2 \geq \ldots \geq e'_m)$. Say that e *dominates* e' if $e_i \geq e'_i$ for every $i = 1, \ldots, \max\{m, n\}$, and that e *weakly majorizes* e' if $\sum_{i=1}^k e_i \geq \sum_{i=1}^k e'_i$ for every $k = 1, \ldots, \max\{m, n\}$, where, in case $m \neq n$, the shorter of the two vectors is appended with zeroes.[12] Domination clearly implies weak majorization, but the converse is not true: $e = (3/8, 1/8)$ majorizes $e' = (1/4, 1/4)$ but does not dominate it.

A classic result of Blackwell and Girschick [6, Theorem 12.3.1 on page 332] says that if information structure P is more informative than P' and Q is more informative than Q', and if each of the pairs (P, Q) and (P', Q') are independent conditional on the state, then the combined information (P, Q) is more informative than the combined information (P', Q'). This implies that if e dominates e' then $m(e)$ is more informative than $m(e')$. Thus, for every decision problem v, V_v is non-decreasing in each expert's effort. The next Theorem 2, which may be of independent interest, strengthen this conclusion by showing that if e weakly majorizes e' then $m(e)$ is more informative than $m(e')$. In other words, for every $v \in \mathcal{V}$ the value of information V_v is a Schur-convex function of effort vectors. Roughly speaking, a function is Schur-convex if it is (1) symmetric and (2) convex in a restricted set of directions.[13] In particular, except for pathological examples, the value of information is not concave in efforts. We note that a different kind of non-concavity in the value of information has been shown by Radner and Stiglitz [31] (see also Chade and Schlee [11]).

Theorem 2. If e weakly majorizes e' then $m(e)$ is more informative than $m(e')$.

Proof. Suppose that e weakly majorizes e'. First, we may assume without loss of generality that both have the same number n of experts; otherwise, add zero-effort experts to the shorter of the two. Second, it is without loss to assume that $\sum_i e_i = \sum_i e'_i$: If e weakly majorizes e' and $\sum_i e_i > \sum_i e'_i$ then there exists e'' such that (i) e'' (exactly) majorizes e', and (ii) e dominates e'' (Marshal et al. [25, Proposition A.9 on page 177]). By Blackwell and Girschick's result $m(e)$ is more informative than $m(e'')$, so the case of unequal total effort follows from the case of equal total effort.

Now, for two vectors $z, z' \in \mathbb{R}^n$ say that z' is obtained from z by a Pigou-Dalton (PD) transfer if there are coordinates i, j with $z_i \geq z_j$ and $0 \leq \delta \leq z_i - z_j$ such that $z'_j = z_j + \delta$, $z'_i = z_i - \delta$, and $z'_k = z_k$ for every $k \neq i, j$. Also, say that z' can be obtained from z by a sequence of PD transfers if there are L and vectors $z_1, \ldots, z_L$ such that $z_1 = z$, $z_L = z'$, and z_l is obtained from z_{l-1} by a PD transfer for every $l = 2, \ldots, L$. It is well known (see, e.g., Marshal et al. [25, Proposition A.1 on page 155]) that if z (exactly) majorizes z' then z' can be obtained from z by a sequence of PD transfers.[14]

[12] The definition of weak majorization allows the total sum $\sum_{i=1}^n e_i$ to be strictly larger than $\sum_{i=1}^n e'_i$, while majorization requires that the two sums are equal. For our results it is sufficient to assume weak majorization.

[13] Thus, symmetry and convexity imply Schur-convexity. Schur-convexity implies symmetry but not convexity.

[14] The converse of this statement is true as well, but is not needed for our purposes.

Therefore, to complete the proof we only need to show that if e' is obtained from e by a PD transfer then $m(e)$ is more informative than $m(e')$. But since a PD transfer changes the efforts of only two experts, it follows from Blackwell and Girschick's result that we may ignore all other experts and consider only the case $n = 2$. This is established in the following lemma, whose proof appears in the appendix.

Lemma 3. Suppose that $e_1 \geq e_2$ and $e'_1 \geq e'_2$ are such that $e_1 + e_2 = e'_1 + e'_2$ and $e_1 \geq e'_1$ (i.e., e' is obtained from e by a PD transfer). Then $m(e_1, e_2)$ is more informative than $m(e'_1, e'_2)$.

□

2.4 Optimal contracts

We now consider the DM's problem of maximizing the difference between the value of information and its cost. Recall that the primitives of the model are the cost of effort function $c \in \mathcal{C}$ from which ψ is derived, and the value function $v \in \mathcal{V}$ from which V is derived. For any e let $\pi_{c,v}(e) = V_v(e) - \psi_c(e)$ be the net expected utility of the DM given efforts' vector e, where, by convention, $\pi_{c,v}(e) = -\infty$ when e is not implementable. We sometimes omit the subscripts c and v when no confusion may arise. For expositional reasons we focus here on the case of a uniform prior $\gamma = 0.5$; similar results continue to hold when $\gamma \neq 0.5$.

Even though the arguments of π are efforts' vectors, we refer to a maximizer of this function as an *optimal contract*. One can think of a contract as specifying both the required efforts e and the payments x that implement e in the least costly way.

2.4.1 Two experts

We start with the case in which the DM is able to hire only two experts ($n = 2$). From Theorem 1 we know that every $0 < e = (e_1, e_2) < 0.5$ is implementable, and that the cost of implementing any such e is

$$\psi(e_1, e_2) = \left(\frac{1}{2} + 2e_1e_2\right)\left(\frac{c'(e_1)}{2e_2} + \frac{c'(e_2)}{2e_1}\right).$$

If $c'(0.5) < +\infty$ then this formula remains true for $0 < e_1, e_2 \leq 0.5$. If $e_i = 0$ and $e_j > 0$ then e is not implementable, while the no-effort vector can be implemented at zero cost: $\psi(0, 0) = 0$.

Our first result here is that 'typically' the optimal contract involves discriminating between the experts. That is, in the optimal contract the DM will offer different compensations to the two experts, and this will result in the experts exerting different effort levels. The intuition for this result comes directly from Theorem 2: Getting two signals of the same accuracy from the experts is less valuable than getting one more accurate and one less accurate signals, subject to the two combinations having the same average accuracy. However, the cost of the former option is lower than the cost of the latter, so we cannot immediately conclude that equal efforts are not optimal. Nevertheless we show that near any decision problem v there is another decision problem $\tilde{v}$ such that, for every cost function c, the maximizer of $\pi_{c,\tilde{v}}$ is not on the diagonal $e_1 = e_2$.

Proposition 1. Fix $v \in \mathcal{V}$ and $\epsilon > 0$. Then there is $\tilde{v} \in \mathcal{V}$ such that
(i) $|\tilde{v}(q) - v(q)| \leq \epsilon$ for all $q \in [0,1]$; and
(ii) For every $c \in \mathcal{C}$, if $e = (e_1, e_2) > 0$ is a maximizer of $\pi_{c,\tilde{v}}$ then $e_1 \neq e_2$.

Proof. Given $v \in \mathcal{V}$ and $\epsilon > 0$, let $v' \in \mathcal{V}$ be given by $v'(q) = \epsilon \max\{q, 1-q\}$, and let $\tilde{v} = v + v'$.[15] Then $\tilde{v} \in \mathcal{V}$ as the sum of two convex and continuous functions, and $|\tilde{v}(q) - v(q)| = |v'(q)| \leq \epsilon$ for all $q \in [0,1]$.

Fix some $c \in \mathcal{C}$ and $0 < e_1 = e_2 < 0.5$. Let $\bar{\delta} > 0$ be a small number and consider $\pi_{c,\tilde{v}}(e_1 + \delta, e_2 - \delta)$ for $\delta \in [-\bar{\delta}, \bar{\delta}]$. We show that this function has a strict local minimum at $\delta = 0$, which implies that (e_1, e_2) is not a maximizer of $\pi_{c,\tilde{v}}$.

First, from Theorem 2 we know that for all $\delta \in [-\bar{\delta}, \bar{\delta}]$

$$V_v(e_1 + \delta, e_2 - \delta) \geq V_v(e_1, e_2). \tag{7}$$

Second, a direct calculation gives

$$V_{v'}(e_1 + \delta, e_2 - \delta) = \begin{cases} \epsilon(0.5 + e_1 + \delta) & \textit{if } 0 \leq \delta \leq \bar{\delta} \\ \epsilon(0.5 + e_2 - \delta) & \textit{if } -\bar{\delta} \leq \delta \leq 0. \end{cases}$$

Also, since ψ_c is symmetric and differentiable, the derivative $\frac{d\psi_c(e_1+\delta, e_2-\delta)}{d\delta}$ is zero at $\delta = 0$. It follows that the right-derivative at $\delta = 0$ of the difference

$$V_{v'}(e_1 + \delta, e_2 - \delta) - \psi_c(e_1 + \delta, e_2 - \delta)$$

is $+\epsilon$ and the left-derivative of this difference at $\delta = 0$ is $-\epsilon$. Thus, for all $\delta \neq 0$ sufficiently close to zero,

$$V_{v'}(e_1 + \delta, e_2 - \delta) - \psi_c(e_1 + \delta, e_2 - \delta) > V_{v'}(e_1, e_2) - \psi_c(e_1, e_2). \tag{8}$$

Summing up, for all $\delta \neq 0$ sufficiently close to zero,

$$\begin{aligned} \pi_{c,\tilde{v}}(e_1 + \delta, e_2 - \delta) \quad &= \quad V_v(e_1 + \delta, e_2 - \delta) + V_{v'}(e_1 + \delta, e_2 - \delta) - \psi_c(e_1 + \delta, e_2 - \delta) \geq \\ & \quad\quad V_v(e_1, e_2) + V_{v'}(e_1 + \delta, e_2 - \delta) - \psi_c(e_1 + \delta, e_2 - \delta) > \\ & \quad\quad V_v(e_1, e_2) + V_{v'}(e_1, e_2) - \psi_c(e_1, e_2) = \pi_{c,\tilde{v}}(e_1, e_2), \end{aligned}$$

where the first equality is by linearity of expectation, the weak inequality is by (7), the strict inequality is by (8), and the last equality is again by linearity of expectation.

Finally, $e_1 = e_2 = 0.5$ is not optimal (when implementable) since the DM can learn the state at a lower cost by choosing $e_1 = 0.5$ and $e_2 < 0.5$ (this is easy to verify directly). $\square$

In the next proposition we derive an additional property of optimal contracts, namely that there is a lower bound on the effort that an expert should be asked to exert. More precisely, given the cost

[15] Recall that v' is obtained from a decision problem with two alternatives as in Example 1.

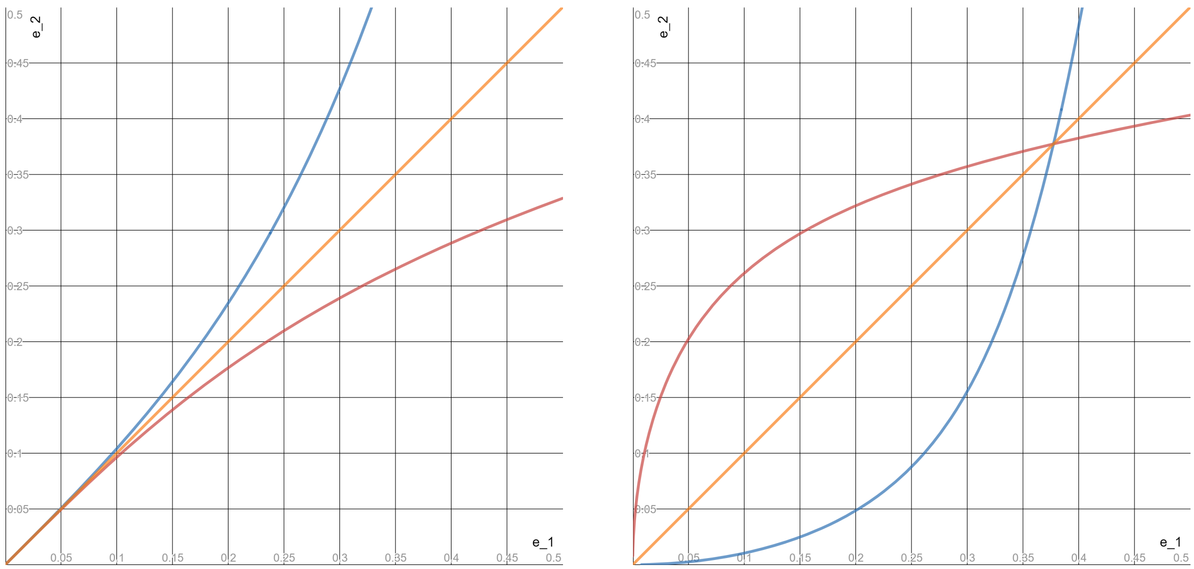

Figure 1: The left and right panels show the bound f for the cost functions $c(e_i) = \frac{e_i^2}{2}$ and $c(e_i) = \frac{e_i^2}{2} + e_i$, respectively. The optimal contract is never (for no decision problem) below the red curves or to the left of the blue curves. If v is the maximum of two linear functions that cross at the prior $\gamma = 0.5$ then the optimal contract is on these curves.

function c, if expert 1 is asked to exert effort $e_1 > 0$ then there is a positive number $f(e_1)$ such that it is *never* (i.e., for no decision problem) optimal to ask expert 2 to exert effort less than $f(e_1)$. The same bound applies for e_1 when e_2 is held fixed. These bounds are illustrated in Figure 1 for the cost functions $c(e_i) = \frac{e_i^2}{2}$ (left panel) and $c(e_i) = \frac{e_i^2}{2} + e_i$ (right panel).

The intuition for this result is that if e_2 is close to zero then the report of expert 2 is almost independent of the state, and hence the distribution over pairs of signals (s_1, s_2) (that determines the expected payment) does not change much when expert 1 increases his effort; in the extreme case where $e_2 = 0$ the distribution over (s_1, s_2) is uniform for every e_1. Thus, in order to induce expert 1 to exert e_1 the payment needs to be very large when the reports match. It follows that for small e_2 the increase in the expected payment needed to incentivize expert 2 to increase his effort is overwhelmed by the resulting decrease in the expected payment to expert 1 (recall Corollary 1). Therefore, increasing e_2 both reduces the cost and provides more information for the DM.

Proposition 2. Let $c \in \mathcal{C}$ be such that the derivative c' is convex on $[0, 0.5)$. Then for every $e_1 \in (0, 0.5)$ there is a unique $f(e_1) \in (0, 0.5)$ such that $\frac{\partial \psi(e_1, f(e_1))}{\partial e_2} = 0$. The function $f : (0, 0.5) \to (0, 0.5)$ is continuous, strictly increasing, and defines a lower bound for the effort in an optimal contract: For any $v \in \mathcal{V}$, if (e_1, e_2) is a maximizer of $\pi_{c,v}$ then $e_2 \geq f(e_1)$ and $e_1 \geq f(e_2)$.

Proof. The partial derivative with respect to e_2 of the cost function is given by

$$\frac{\partial \psi(e_1, e_2)}{\partial e_2} = \frac{\partial \psi_1(e_1, e_2)}{\partial e_2} + \frac{\partial \psi_2(e_1, e_2)}{\partial e_2} = \frac{-c'(e_1)}{4e_2^2} + c'(e_2) + \frac{0.5 + 2e_1 e_2}{2e_1} c''(e_2).$$

Fixing $0 < e_1 < 0.5$, $\frac{\partial \psi}{\partial e_2}$ is clearly negative when e_2 is sufficiently close to zero and positive when e_2 is sufficiently close to 0.5. Furthermore,

$$\frac{\partial^2 \psi(e_1, e_2)}{\partial e_2^2} = 2c''(e_2) + \left(e_2 + \frac{1}{4e_1}\right) c'''(e_2) + \frac{c'(e_1)}{2e_2^3}$$

is strictly positive for any $e_1, e_2 > 0$ (recall that c' is assumed to be convex), so ψ is coordinate-wise convex. It follows that there is a unique $f(e_1)$ satisfying $\frac{\partial \psi(e_1, f(e_1))}{\partial e_2} = 0$. We get that $\psi(e_1, \cdot)$ is strictly decreasing for $e_2 \in [0, f(e_1)]$ and strictly increasing for $e_2 \in [f(e_1), 0.5]$. It is immediate to see from the above derivatives that f is continuous and strictly increasing. By symmetry, the same f applies when e_2 is held fixed and e_1 varies.

Now, suppose that $e = (e_1, e_2)$ satisfies $e_2 < f(e_1)$. Define $e' = (e_1, f(e_1))$. Then $\psi(e) > \psi(e')$ by the definition of $f(e_1)$ and $V_v(e) \leq V_v(e')$ for any $v \in \mathcal{V}$ since e' dominates e. It follows that $\pi_{c,v}(e) < \pi_{c,v}(e')$. A similar argument applies when $e_1 < f(e_2)$. □

In the special case where v is the maximum of two linear functions that cross at the prior $\gamma = 0.5$ (as in Examples 1 and 2 above) we can say more about the location of the optimal contract, since the value of information in this case depends only on the maximum of the two efforts. This implies that one of the experts will be hired only for monitoring purposes, and the effort required from that expert will be chosen to minimize the total cost of the contract. In Figure 1 this means that the optimal contract is located on the red or blue curves.

Corollary 2. Suppose $c \in \mathcal{C}$ satisfies the assumption in Proposition 2, and let f be the function defined in that proposition. If $v \in \mathcal{V}$ is the maximum of two linear functions that cross at the prior $\gamma = 0.5$, and if $(e_1, e_2) > 0$ is a maximizer of $\pi_{c,v}$, then either $e_1 > e_2$ and $e_2 = f(e_1)$ or $e_2 > e_1$ and $e_1 = f(e_2)$.

Proof. We start with the following simple lemma.

Lemma 4. If $v \in \mathcal{V}$ is the maximum of two linear functions that cross at the prior $\gamma = 0.5$ then there are $\alpha > 0$, $\beta \in \mathbb{R}$ such that $V(e_1, e_2) = \alpha \max\{e_1, e_2\} + \beta$, i.e., V is increasing and linear in the accuracy of the more accurate signal.

Suppose now that $(e_1, e_2) > 0$ is a maximizer of $\pi_{c,v}$, and consider first the case $e_1 > e_2$. From Proposition 2 we have $e_2 \geq f(e_1)$. We claim that this inequality must hold as equality. Indeed, suppose by contradiction that $e_2 > f(e_1)$. Then $\max\{e_1, e_2\} = \max\{e_1, f(e_1)\}$, so by Lemma 4 $V(e_1, e_2) = V(e_1, f(e_1))$. Also, $\psi(e_1, e_2) > \psi(e_1, f(e_1))$ by the definition of $f(e_1)$. Hence $\pi_{c,v}(e_1, e_2) < \pi_{c,v}(e_1, f(e_1))$, contradicting the optimality of (e_1, e_2). By a similar argument, if $e_2 > e_1$ is a maximizer then $e_1 = f(e_2)$ must hold. Finally, $e_1 = e_2$ can't be optimal when v is the maximum of two linear functions that cross at $\gamma = 0.5$ by the same argument as in the proof of Proposition 1 (for the function v'). □

2.4.2 Many experts

We now consider the case where the DM may hire any number of experts. Our main insight here is that hiring many low-effort experts is dominated (less informative and more costly) by hiring few high-effort experts. To formalize this, for each $\bar{t} \in (0, 0.5)$ and $\bar{n} \geq 2$ denote by $e\left(\bar{t}, \bar{n}\right)$ the constant efforts' vector with $\bar{n}$ experts each of which exerts effort $\bar{t}$.

Proposition 3. Suppose $c'(0) > 0$ and fix some $\bar{t}, \bar{n}$. Every $e = (e_1, \ldots, e_n) > 0$ with $n \geq \frac{2\psi_c(e(\bar{t},\bar{n}))}{c'(0)}$ and $e_i \leq \frac{\bar{n}\bar{t}}{n}$ for all i is more costly and less informative than $e\left(\bar{t}, \bar{n}\right)$. In particular, any such e is not optimal for any decision problem $v \in \mathcal{V}$.

Proof. First, for every $e = (e_1, \ldots, e_n) > 0$ we have

$$\frac{e(N) + \bar{e}(N)}{e(N \setminus \{i\}) - \bar{e}(N \setminus \{i\})} = \left(\frac{1}{2} + e_i\right) \frac{e(N \setminus \{i\}) + \frac{\bar{e}(N)}{\frac{1}{2}+e_i}}{e(N \setminus \{i\}) - \bar{e}(N \setminus \{i\})} \geq \frac{1}{2} + e_i \geq \frac{1}{2},$$

which implies

$$\psi_c(e) = [e(N) + \bar{e}(N)] \sum_{i=1}^{n} \frac{c'(e_i)}{e(N \setminus \{i\}) - \bar{e}(N \setminus \{i\})} \geq \frac{1}{2} \sum_{i=1}^{n} c'(e_i).$$

Convexity of c implies that $c'(e_i) > c'(0)$ for all i, so $\psi_c(e) > \frac{n}{2} c'(0)$. It follows that if $n \geq \frac{2\psi_c(e(\bar{t},\bar{n}))}{c'(0)}$ then $\psi_c(e) > \psi_c\left(e\left(\bar{t}, \bar{n}\right)\right)$.

Second, if e satisfies the assumptions of the proposition then $n \geq \frac{2\psi_c(e(\bar{t},\bar{n}))}{c'(0)} \geq \frac{2}{c'(0)} \frac{\bar{n}}{2} c'\left(\bar{t}\right) > \bar{n}$. This also implies that $e_i \leq \frac{\bar{t}\bar{n}}{n} < \bar{t}$ for all i. Thus, any such e is weakly majorized by $e\left(\bar{t}, \bar{n}\right)$, so by Theorem 2 $e\left(\bar{t}, \bar{n}\right)$ is more informative than e. $\square$

Corollary 3. Suppose $c'(0) > 0$ and $c'(0.5) < +\infty$. Then for every decision problem $v \in \mathcal{V}$ the optimal contract has at most $\frac{4c'(0.5)}{c'(0)}$ experts.

Proof. Apply Proposition 3 for $\bar{t} = 0.5$ and $\bar{n} = 2$. $\square$

We conclude this section with the following result which generalizes Proposition 1 from two experts to any even number of experts.

Proposition 4. Let $n \geq 2$ be even and fix $v \in \mathcal{V}$ and $\epsilon > 0$. Then there is $\tilde{v} \in \mathcal{V}$ such that
(i) $|\tilde{v}(q) - v(q)| \leq \epsilon$ for all $q \in [0, 1]$; and
(ii) For every $c \in \mathcal{C}$, if $e = (e_1, \ldots, e_n) > 0$ is a maximizer of $\pi_{c,\tilde{v}}$ among all vectors of n experts then there are $1 \leq i, j \leq n$ such that $e_i \neq e_j$.

Sketch of proof:
Define the function v' in the same way as in the proof of Proposition 1. Consider any $e_1 = \ldots = e_n > 0$. Then it is not hard to check that the function $V_{v'}(e_1 + \delta, e_2, \ldots, e_{n-1}, e_n - \delta)$ of δ has a strict local minimum at $\delta = 0$, and, moreover, the right derivative is strictly positive and left derivative is strictly negative at that point. Defining $\tilde{v} = v + v'$ and repeating the argument in the proof of Proposition 1 completes the proof. Note that if n is odd then the optimal action for the DM in the decision problem

v' is independent of δ (for δ close to 0), which implies that $V_{v'}(e_1+\delta, e_2, \ldots, e_{n-1}, e_n-\delta)$ is smooth at $\delta=0$, hence the failure of the argument. $\square$

3 Characterization of implementability

We now leave the binary-binary model of the previous section and consider a more general framework. Our goal in this section is to characterize the information structures that the DM can implement. An important observation is that in our setup what matters for the DM is not the efforts that the experts exert per-se, but rather the mappings from states to distributions over reported signals that the experts' strategies induce; indeed, these mappings determine both the transfers from the DM to the experts and the expected utility of the DM in the decision problem she faces. Therefore, in the current section we will assume that experts directly choose information structures (i.e., mappings from states to reported signals) rather than efforts and reporting strategies. In subsection 3.1 below we illustrate using the model of the Section 2 how the question of implementability can be translated into the framework of the current section, and in Appendix B we provide a more general treatment of the connection between the two frameworks.

The set of possible states of nature is $\Omega=\{\omega_1,\ldots,\omega_K\}$, and the prior belief over Ω is $\gamma=(\gamma_1,\ldots,\gamma_K)$ with $\gamma_k>0$ for all $k=1,\ldots,K$. The DM may obtain information regarding the state of nature from a group of $n\geq 2$ experts, where $N=\{1,\ldots,n\}$ denotes the set of experts. The finite set of signals that expert i may observe is S_i, with a typical element denoted by s_i. Each expert i chooses a mapping $m_i:\Omega\to\Delta(S_i)$ from a set M_i of such mappings. As is standard, we view each $m_i\in M_i$ as a stochastic matrix with K rows and $|S_i|$ columns, where $m_i(k,s_i)$ is the probability that signal $s_i\in S_i$ realizes conditional on the state being $\omega_k\in\Omega$. The set M_i is viewed as a subset of $\mathbb{R}^{K|S_i|}$ endowed with the standard Euclidean norm $\|\cdot\|$. The cost of choosing $m_i\in M_i$ is described by the function $C_i:M_i\to\mathbb{R}_+$.

We make the following assumptions:[16]

(A1) M_i is a non-empty, closed, and convex polyhedral set.
(A2) C_i is convex on M_i.
(A3) C_i is Lipschitz continuous on M_i: There is $\beta>0$ such that $|C_i(m_i)-C_i(m_i')|\leq\beta\|m_i-m_i'\|$ for every $m_i,m_i'\in M_i$.[17]

Let $S=S_1\times\ldots\times S_n$ with $s=(s_1,\ldots,s_n)\in S$ denoting a vector of signal realizations, and let $M=M_1\times\ldots\times M_n$ with a typical element $m=(m_1,\ldots,m_n)$. Any $m\in M$ induces a distribution over S that we denote by $\mathbb{P}_m$:

$$\mathbb{P}_m(s)=\sum_{k=1}^{K}\gamma_k\prod_{i=1}^{n}m_i(k,s_i). \tag{9}$$

[16] In subsection 3.1 and in Appendix B we show that properties (A1)-(A3) naturally follow from standard assumptions on the technology and cost of collecting information. If one starts from a model of efforts and reports then there are additional properties that the resulting M_i and C_i satisfy. For example, the induced M_i is closed under garblings and the induced C_i is monotonically increasing in the Blackwell ordering. We do not include these assumptions here as they are not needed to prove the results of this section.

[17] It is worth pointing out that convexity of C_i implies that it is Lipschitz on any closed subset of the relative interior of M_i (Rockafellar [35, Theorem 10.4]). Therefore, (A3) only has a bite on the relative boundary of M_i.

The expectation operator relative to this distribution is denoted by $\mathbb{E}_m$. Note that (9) assumes that signals of different experts are independent conditional on the state of nature.

A contract is a list $x = (x_1, \ldots, x_n)$, where $x_i : S \to \mathbb{R}_+$ is the payment to i given $s \in S$. Each contract induces a game between the experts. The set of strategies for expert i is M_i and his payoff given strategy profile m is

$$U_i(m;\ x_i) := \mathbb{E}_m[x_i(s)] - C_i(m_i). \tag{10}$$

Definition 1. Say that $m^* \in M$ is *implementable* if there exists a contract x such that m^* is an equilibrium of the game induced by x.

Example 4. Suppose $\Omega = \{\omega_1, \omega_2, \omega_3\}$, $\gamma = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$, $n = 2$, and $S_1 = S_2 = \{0, 1\}$. Let m_1^*, m_2^*, m_2 be given by

m_1^*:

	0	1
ω_1	3/4	1/4
ω_2	1/2	1/2
ω_3	1/4	3/4

m_2^*:

	0	1
ω_1	3/4	1/4
ω_2	1/2	1/2
ω_3	1/4	3/4

m_2:

	0	1
ω_1	1/2	1/2
ω_2	1	0
ω_3	0	1

It is straightforward to check that $\mathbb{P}_{(m_1^*, m_2^*)} = \mathbb{P}_{(m_1^*, m_2)}$.[18] In other words, given that expert 1 chooses m_1^*, the distribution over pairs of signals is the same when expert 2 chooses m_2 as when he chooses m_2^*. Hence, if $C_2(m_2^*) > C_2(m_2)$, then (m_1^*, m_2^*) can't be implemented since for every contract x_2 it is profitable for expert 2 to deviate from m_2^* to m_2.

Example 4 suggests a necessary condition for implementation of m^*: There can't be an expert i and $m_i \in M_i$ such that both $\mathbb{P}_{m^*} = \mathbb{P}_{(m_i, m_{-i}^*)}$ and $C_i(m_i) < C_i(m_i^*)$. In the following theorem we apply an idea of Rahman [32, 33] to show that this condition is also sufficient. The proof shows that, if the condition is satisfied, then for every possible deviation m_i of i there exists a contract x_i that will make this deviation unprofitable. Of course, the difficulty is to find one x_i that simultaneously discourages all possible deviations. The existence of such contract is a consequence of the minmax theorem.

Theorem 3. Under assumptions (A1)-(A3), m^* is implementable if and only if for every $i \in N$ and every $m_i \in M_i$ either $\mathbb{P}_{m^*} \neq \mathbb{P}_{(m_i, m_{-i}^*)}$ or $C_i(m_i) \geq C_i(m_i^*)$ (or both).

Proof. The 'only if' part is obvious, so we only prove the 'if' part. Fix m^* and i. Let $D > 0$ be a large constant to be determined later. Define the function $f : M_i \times [0, D]^S \to \mathbb{R}$ by

$$f(m_i, x_i) = \left(\mathbb{E}_{(m_i, m_{-i}^*)}[x_i(s)] - \mathbb{E}_{m^*}[x_i(s)]\right) + \left(C_i(m_i^*) - C_i(m_i)\right). \tag{11}$$

Thus, $f(m_i, x_i)$ is the payoff gain (or loss) for expert i when he chooses m_i rather than the prescribed m_i^*, all other experts choose according to m_{-i}^*, and given contract x_i. The following lemma gives some basic properties of f.

[18] For example, if $(s_1, s_2) = (0, 0)$ then $\mathbb{P}_{(m_1^*, m_2^*)}(0, 0) = \frac{1}{3}\left(\frac{3}{4} * \frac{3}{4} + \frac{1}{2} * \frac{1}{2} + \frac{1}{4} * \frac{1}{4}\right) = \frac{7}{24}$ and $\mathbb{P}_{(m_1^*, m_2)}(0, 0) = \frac{1}{3}\left(\frac{3}{4} * \frac{1}{2} + \frac{1}{2} * 1 + \frac{1}{4} * 0\right) = \frac{7}{24}$.

Lemma 5. The function f defined in (11) is concave and continuous in m_i for each fixed x_i, and affine (and continuous) in x_i for each fixed m_i.

By assumption (A1) and Lemma 5 the conditions of the minmax theorem (see Rockafellar [35, Corollary 37.6.2]) are satisfied and therefore

$$\max_{m_i \in M_i} \min_{x_i \in [0,D]^S} f(m_i, x_i) = \min_{x_i \in [0,D]^S} \max_{m_i \in M_i} f(m_i, x_i). \tag{12}$$

We now show that for D large enough the left-hand side of (12) equals zero. Note first that f vanishes whenever $m_i = m_i^*$, and hence that

$$\max_{m_i \in M_i} \min_{x_i \in [0,D]^S} f(m_i, x_i) \geq 0. \tag{13}$$

To show the other inequality, define the mapping $T : M_i \to \mathbb{R}^S$ by $T(m_i)(s) = \mathbb{P}_{(m_i, m_{-i}^*)}(s) - \mathbb{P}_{m^*}(s)$ for each $s \in S$. In words, T sends each m_i to the difference between the distribution over S that it induces (together with m_{-i}^*) and the distribution over S that the desired m_i^* induces (together with m_{-i}^*). Note that T is affine. For each m_i let $\bar{m}_i$ be its projection to the (non-empty, compact, convex, polyhedral) set $M_i \cap Ker(T)$, where $Ker(T)$ is the kernel of T. We will need the following lemma.

Lemma 6. There is $\beta' > 0$ such that $\|T(m_i)\|_S \geq \beta' \|m_i - \bar{m}_i\|$ for every $m_i \in M_i$, where $\| \cdot \|_S$ is the standard Euclidean norm on $\mathbb{R}^S$.

Now, given some $m_i \in M_i$ define the contract x_i by

$$x_i(s) = \begin{cases} D & \textit{if } T(m_i)(s) < 0 \\ 0 & \textit{if } T(m_i)(s) \geq 0. \end{cases}$$

Thus, according to x_i the expert gets paid D at signal realizations that are more likely to occur under m^* than under $(m_i, m^* - i)$, and gets nothing at the other realizations. With these m_i, x_i we have

$$\begin{aligned}
f(m_i, x_i) &= D \left[\sum_{\{s \,:\, T(m_i)(s) < 0\}} \Big(\mathbb{P}_{(m_i, m_{-i}^*)}(s) - \mathbb{P}_{m^*}(s) \Big) \right] + (C_i(m_i^*) - C_i(m_i)) \\
&= -\frac{D}{2} \sum_{s \in S} \Big| \mathbb{P}_{(m_i, m_{-i}^*)}(s) - \mathbb{P}_{m^*}(s) \Big| + (C_i(m_i^*) - C_i(m_i)) \\
&\leq -\frac{D\sqrt{|S|}}{2} \|T(m_i)\|_S + (C_i(m_i^*) - C_i(m_i)) \\
&\leq -\frac{D\sqrt{|S|}}{2} \|T(m_i)\|_S + (C_i(\bar{m}_i) - C_i(m_i)) \\
&\leq -\frac{D\sqrt{|S|}}{2} \beta' \|m_i - \bar{m}_i\| + \beta \|m_i - \bar{m}_i\| = \left(-\frac{D\sqrt{|S|}}{2} \beta' + \beta \right) \|m_i - \bar{m}_i\|,
\end{aligned}$$

where the first equality is by the definition of x_i, the second equality follows from the fact that both $\mathbb{P}_{(m_i, m_{-i}^*)}$ and $\mathbb{P}_{m^*}$ are distributions (sum-up to 1), the first inequality by a standard relation between the l_1 and l_2 norms, the next inequality from the fact that $\bar{m}_i \in Ker(T)$ and the assumption of the theorem, and the last inequality from Lemma 6 and (A3). Therefore, if $D = \frac{2\beta}{\beta'\sqrt{|S|}}$ then for each $m_i \in M_i$ there is $x_i \in [0, D]^S$ such that $f(m_i, x_i) \leq 0$, i.e.,

$$\max_{m_i \in M_i} \min_{x_i \in [0,D]^S} f(m_i, x_i) \leq 0.$$

Combined with (13) we proved that $\max_{m_i \in M_i} \min_{x_i \in [0,D]^S} f(m_i, x_i) = 0$.

It follows from the minmax equality (12) that for D large enough $\min_{x_i \in [0,D]^S} \max_{m_i \in M_i} f(m_i, x_i) = 0$. That is, there exists $x_i \in [0, D]^S$ such that $f(m_i, x_i) \leq 0$ for every $m_i \in M_i$. Repeating the same argument for each expert i we get a contract x that implements m^*. □

When is it the case that a given profile m^* can always be implemented, regardless of the cost functions? Theorem 3 implies that this is the case whenever for every i the mapping $m_i \longmapsto \mathbb{P}_{(m_i, m^*_{-i})}$ is one-to-one on M_i. Denote by $\mathbb{P}_{m^*_{-i}|\omega_k}$ the distribution over S_{-i} induced by m^*_{-i} conditional on state ω_k, that is, $\mathbb{P}_{m^*_{-i}|\omega_k}(s_{-i}) = \prod_{j \neq i} m^*_j(k, s_j)$. In the next proposition we view each $\mathbb{P}_{m^*_{-i}|\omega_k}$ as a vector in $\mathbb{R}^{|S_{-i}|}$.

Proposition 5. Assume (A1)-(A3). Given $m^* \in M$, if for every $i \in N$ the K vectors $\left\{\mathbb{P}_{m^*_{-i}|\omega_k}\right\}_{k=1}^{K}$ are linearly independent then m^* is implementable.

Proof. We show that if the vectors $\left\{\mathbb{P}_{m^*_{-i}|\omega_k}\right\}_{k=1}^{K}$ are linearly independent then $\mathbb{P}_{m^*} \neq \mathbb{P}_{(m_i, m^*_{-i})}$ for every $m_i \neq m^*_i$. By Theorem 3 this is sufficient to complete the proof.

Suppose that for some m_i we have $\mathbb{P}_{m^*} = \mathbb{P}_{(m_i, m^*_{-i})}$. Fix some $s_i \in S_i$. Then for every $s_{-i} \in S_{-i}$ the probability of (s_i, s_{-i}) is the same under m^*_i and m_i, that is

$$\sum_{k=1}^{K} \gamma_k \left(\prod_{j \neq i} m^*_j(k, s_j) \right) m^*_i(k, s_i) = \sum_{k=1}^{K} \gamma_k \left(\prod_{j \neq i} m^*_j(k, s_j) \right) m_i(k, s_i).$$

Denoting $y_k = m^*_i(k, s_i) - m_i(k, s_i)$ $(1 \leq k \leq K)$ we get that for every s_{-i}

$$\sum_{k=1}^{K} \gamma_k y_k \mathbb{P}_{m^*_{-i}|\omega_k}(s_{-i}) = 0,$$

but by linear independence this implies that $y_k = 0$ for all k. It follows that $m_i(k, s_i) = m^*_i(k, s_i)$ for every k, and since s_i was arbitrary we get $m^*_i = m_i$. □

We conclude with the following two corollaries of Proposition 5. The first one shows that increasing informativeness cannot hinder implementability. The second shows that in the special case of $K = 2$ states implementation of any (informative) structure is possible.

Corollary 4. Suppose that m^* satisfies the assumption of Proposition 5, so it is implementable regardless of the cost functions. Let m^{**} be such that, for each $i \in N$, m^{**}_{-i} is (weakly) more informative than m^*_{-i} in the sense of Blackwell. Then m^{**} also satisfies the assumption of Proposition 5 and is therefore implementable regardless of the cost functions. In particular, if m_i^{**} is (weakly) more informative than m_i^* for every $i \in N$ then m^{**} is implementable regardless of the cost functions.

Proof. Let m^* and m^{**} be as in the corollary and fix i. Then there is a stochastic matrix p with dimensions $|S_{-i}| \times |S_{-i}|$ such that $m^{**}_{-i}p = m^*_{-i}$. Suppose that $\{y_k\}_{k=1}^K$ are real numbers such that for all s_{-i}

$$\sum_{k=1}^{K} y_k \mathbb{P}_{m^{**}_{-i}|\omega_k}(s_{-i}) = 0.$$

Then for every s_{-i}

$$\sum_{k=1}^{K} y_k \mathbb{P}_{m^{*}_{-i}|\omega_k}(s_{-i}) = \sum_{k=1}^{K} y_k \sum_{s'_{-i}} \mathbb{P}_{m^{**}_{-i}|\omega_k}(s'_{-i}) p(s'_{-i}, s_{-i}) = \sum_{s'_{-i}} p(s'_{-i}, s_{-i}) \sum_{k=1}^{K} y_k \mathbb{P}_{m^{**}_{-i}|\omega_k}(s'_{-i}) = 0,$$

which by assumption implies that $\{y_k\}_k$ are all zero. Thus, $\{\mathbb{P}_{m^{**}_{-i}|\omega_k}\}_{k=1}^K$ are linearly independent. □

Corollary 5. Suppose $K = 2$ and let m^* be such that, for all $i \in N$, m_i^* is not completely uninformative (constant). Then m^* is implementable.

Proof. Fix i. Since for each $j \neq i$ the distribution over S_j conditional on ω_1 is different than the distribution over S_j conditional on ω_2 under m_j^*, it follows that $\mathbb{P}_{m^*_{-i}|\omega_1} \neq \mathbb{P}_{m^*_{-i}|\omega_2}$. Since both are distributions they must be linearly independent. □

Note that implementability of any positive efforts' vector in the binary-binary model of Section 2 can be viewed as a special case of Corollary 5 (assuming c is Lipschitz).

3.1 Connection between the models

Consider a version of the binary-binary model of Section 2 in which expert i's choice of effort is restricted to [0,0.25] rather than [0,0.5].[19] What is the set of mappings from states to distributions over reports that i can induce by playing some (possibly mixed) strategy? Denote by α the probability of a b report given state B, and by β the probability of a w report in state W, so that mappings from states to distributions over reports can be identified with matrices of the form

	b	w
B	α	$1-\alpha$
W	$1-\beta$	β

If i truthfully reports his signal then we have $\alpha = \beta$ and these increase from 0.5 when $e_i = 0$ to 0.75 with full effort of $e_i = 0.25$. If the reporting strategy is $r_i(b) = w$ and $r_i(w) = b$ then again $\alpha = \beta$ and they range from 0.5 with no effort to 0.25 with full effort. When $r_i \equiv b$ we get $\alpha = 1$ and $\beta = 0$, and when $r_i \equiv w$ we have $\alpha = 0$ and $\beta = 1$ (these are true for any effort level). These are all the mappings from states to distributions over reports that i can generate with pure strategies; they are depicted by

[19] This modification helps illustrate the idea.

the blue interval and two dots in Figure 2. It follows that by playing mixed strategies i can induce any mapping in the convex hull of the blue set, that is, any point inside the black dashed line in the figure. Note that this set is equal to the set of all garblings of the most informative structure $b_i(0.25)$, and that it satisfies all the requirements of assumption (A1). We can think of i as choosing from this set of information structures instead of choosing (mixtures of) efforts and reports.

However, there are typically many different mixed strategies that induce the same information structure. For example, mixing between the pure strategies $(e_i = 1/14, r_i^*)$ and $(e_i = 0, r_i \equiv b)$ with probabilities 7/8 and 1/8, respectively, yields the mapping

$$\frac{7}{8}\times \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 4/7 & 3/7 \\ \hline W & 3/7 & 4/7 \\ \hline \end{array} \;+\; \frac{1}{8}\times \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 1 & 0 \\ \hline W & 1 & 0 \\ \hline \end{array} \;=\; \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 5/8 & 3/8 \\ \hline W & 1/2 & 1/2 \\ \hline \end{array}$$

This mapping is the red dot in Figure 2. The same mapping can be induced by mixing between $(e_i = 1/8, r_i^*)$, $(e_i = 0, r_i \equiv b)$, and $(e_i = 0, r_i \equiv w)$ with probabilities 1/2, 5/16 and 3/16, respectively:

$$\frac{1}{2}\times \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 5/8 & 3/8 \\ \hline W & 3/8 & 5/8 \\ \hline \end{array} \;+\; \frac{5}{16}\times \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 1 & 0 \\ \hline W & 1 & 0 \\ \hline \end{array} \;+\; \frac{3}{16}\times \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 0 & 1 \\ \hline W & 0 & 1 \\ \hline \end{array} \;=\; \begin{array}{|c|c|c|}\hline & b & w \\ \hline B & 5/8 & 3/8 \\ \hline W & 1/2 & 1/2 \\ \hline \end{array}$$

If i chooses the former mixed strategy then his expected cost of effort is $\frac{7}{8}c\left(\frac{1}{14}\right)$, while choosing the latter mixture costs $\frac{1}{2}c\left(\frac{1}{8}\right)$. Clearly, i has no reason to choose the more costly of these two mixtures, as both lead to the same distribution over vectors of signals and hence to the same expected payment. More generally, i would always choose the least costly mixture to induce a given information structure. Therefore, the relevant cost function is the *convexification* of the cost of mappings induced by pure strategies.[20] In particular, the cost function is convex, so (A2) is satisfied. Lipschitz continuity (A3) holds as well whenever c is Lipschitz continuous. See Appendix B for more details on the connection between the two frameworks.

4 Concluding remarks

As already mentioned in the introduction, our notion of implementation only requires that honesty and obedience is an equilibrium, and does not rule out the existence of other equilibria in the game induced by the contract. In particular, if all experts other than i exert zero effort, so that their reports are independent of the state, then regardless of the contract i's best response is to exert zero effort as well. Thus, for any contract there exists a no-effort equilibrium.[21] There may also be other equilibria with more or less effort than the one intended by the planner. However, as the following proposition shows, at least in some environments the problem is less severe than it may seem.

[20]There may be several pure strategies that induce the same mapping, for example all levels of effort induce the same information when combined with constant reporting of b; the relevant cost is obviously the minimal one.

[21]It is possible that no-effort equilibria exist only in mixed strategies. The same argument applies also for indirect (finite) mechanisms.

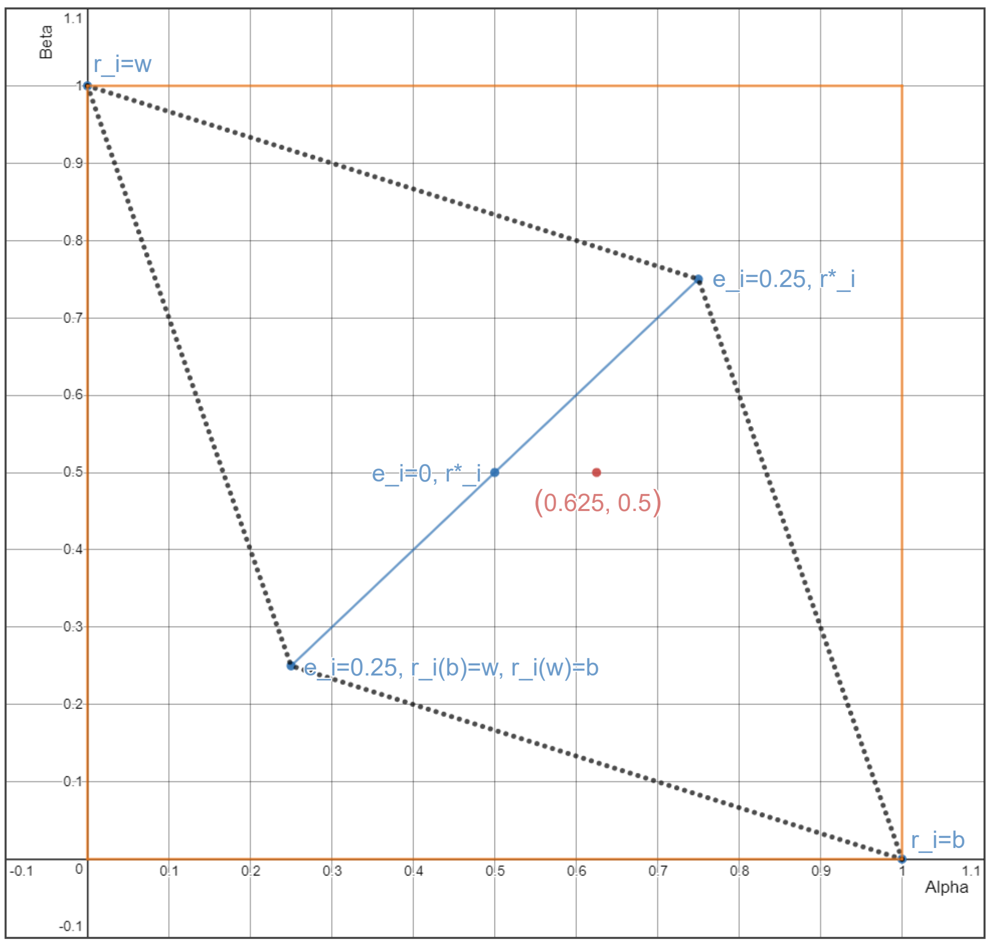

Figure 2: The set of information structures i can generate by playing pure strategies in the binary-binary model (with effort bounded by 0.25) is the blue interval and two corner dots. The black dashed line shows the boundary of the convex hull of this set – this is the set of information structures i can generate by playing pure or mixed strategies.

Proposition 6. Consider the binary-binary model of Section 2 with a uniform prior $\gamma = 0.5$, and suppose that either (1) c' is strictly concave, or (2) c' is strictly convex, $c'(0) = 0$ and $n \in \{2, 3\}$. Let $e = (e_1, \dots, e_n)$ be such that $0 < e_i < 0.5$ for every i. Then in the game induced by the optimal contract x^* of Theorem 1, (e, r^*) is the only equilibrium with truthful reporting and positive levels of effort.

Proof. Fix an interior e and let x^* be the least costly contract that implements e derived in Theorem 1. Consider the game in which each expert i only chooses effort level e_i and the realized signal is truthfully reported to the mechanism. Then it is immediate to verify that this is a supermodular game in the sense of Milgrom and Roberts [27] (see Theorem 4 in that paper). By Topkis [38, Theorem 4.2.1], the set of equilibria is a lattice. It follows that if e' and e'' are both equilibria then their coordinate-wise maximum is an equilibrium as well, so if there are two equilibria with positive efforts then there are two equilibria with positive efforts with one dominating the other. However, we now show that under the conditions of the proposition there can't be two equilibria with positive efforts in which one dominates the other; since e is one equilibrium, the result follows.

Suppose by contradiction that $0 < e'_i \leq e''_i$ for all i, that $e' \neq e''$, and that both are equilibria under x^*. Then for each i the first-order condition with respect to effort must hold at both e' and e''. This gives

$$\left[e'\left(N_{-i}\right) - \bar{e}'\left(N_{-i}\right)\right] \frac{x_i^*\left(\underline{b}\right) + x_i^*\left(\underline{w}\right)}{2} = c'(e'_i) \quad \text{and} \quad \left[e''\left(N_{-i}\right) - \bar{e}''\left(N_{-i}\right)\right] \frac{x_i^*\left(\underline{b}\right) + x_i^*\left(\underline{w}\right)}{2} = c'(e''_i).$$

Therefore, for each expert i we have

$$\frac{e'\left(N_{-i}\right) - \bar{e}'\left(N_{-i}\right)}{e''\left(N_{-i}\right) - \bar{e}''\left(N_{-i}\right)} = \frac{c'(e'_i)}{c'(e''_i)}.$$

Now, suppose that c' is strictly concave. Then $\frac{c'(x)}{x}$ is strictly decreasing on [0,0.5], which implies that $\frac{c'(e'_i)}{c'(e''_i)} \geq \frac{e'_i}{e''_i}$ for all i, with strict inequality whenever $e'_i < e''_i$. Thus,

$$\frac{e'\left(N_{-i}\right) - \bar{e}'\left(N_{-i}\right)}{e''\left(N_{-i}\right) - \bar{e}''\left(N_{-i}\right)} \geq \frac{e'_i}{e''_i}$$

for all i with strict inequality for at least one expert (recall that $e' \neq e''$). Cross-multiplying and summing-up these n inequalities gives

$$\sum_{i=1}^{n} e''_i \left[e'\left(N_{-i}\right) - \bar{e}'\left(N_{-i}\right)\right] > \sum_{i=1}^{n} e'_i \left[e''\left(N_{-i}\right) - \bar{e}''\left(N_{-i}\right)\right].$$

However, it is not hard to check that this last inequality is inconsistent with $0 < e'_i \leq e''_i$ for all i, hence the desired contradiction.

In the other case where c' is strictly convex and $c'(0) = 0$ we have that $\frac{c'(x)}{x}$ is strictly increasing, so

$$\frac{\left[e'\left(N_{-i}\right) - \bar{e}'\left(N_{-i}\right)\right]}{\left[e''\left(N_{-i}\right) - \bar{e}''\left(N_{-i}\right)\right]} = \frac{c'(e'_i)}{c'(e''_i)} \leq \frac{e'_i}{e''_i}$$

holds for every i with strict inequality at least once. If the number of experts is either $n = 2$ or $n = 3$ then similarly to the previous paragraph we get a contradiction. $\square$

As mentioned in the last proof, under the optimal contract x^* of Theorem 1 (and assuming truthful reporting), the game of effort choices is supermodular. It follows that the game has a largest equilibrium that dominates all other equilibria. Furthermore, it is immediate to check that the payoff of each expert is increasing in other experts' efforts, so by Theorem 7 in Milgrom and Roberts [27] the largest equilibrium Pareto-dominates all other equilibria. Thus, in environments with multiple positive-efforts equilibria, if the DM wants to implement an equilibrium other than the largest one then she should be more concerned about over-investment than about under-investment.[22]

Theorem 1 derives an explicit formula for the cost function in the binary-binary model. Initial attempts to extend this derivation to environments with more states and/or signals suggest that the problem may be significantly more difficult. Even if the precise form of the cost-minimizing contract is not tractable, it would be interesting to know if it has similar properties to the contract in the binary-binary case, e.g. whether it creates a relatively simple coordination game between the experts. Another potentially tractable model to consider is one in which signals are normally distributed with precision that increases in effort.

There is also much still to be done within the binary-binary framework. We have abstracted from many important features of the moral hazard literature: The experts (and DM) are risk neutral, the role of outside options is not considered,[23] and so does the role of reputation/career concerns. We leave these interesting issues for future work.

Finally, as mentioned before, Theorem 2 expresses convexity in the value of information and is therefore related to the works of Radner and Stiglitz [31] and Chade and Schlee [11]. An equivalent statement of this Theorem is that if e, e' are two effort vectors such that observing the signal of *one* randomly (uniformly) chosen expert from e is more informative than observing the signal of *one* randomly chosen expert from e', then observing the signals of *all* experts in e is more informative then observing the signals of *all* experts in e'. One may think that this is true for more general vectors of information structures, but this is not the case.[24] Still, non-concavity is likely to pose serious challenges for solving problems involving maximization of the value of information in other environments as well.

References

[1] A.A. Alchian and H. Demsetz (1972), Production, information costs, and economic organization, *American Economic Review* **62**, 777-795.

[22]Of course, the 'collusive' equilibrium with no effort and constant reporting Pareto-dominates all other equilibria, so it is crucial that the experts cannot communicate with each other and that no report can be easily identified as focal.

[23]Note that the experts make strictly positive rents in the optimal contract, so individual rationality is satisfied whenever outside options are sufficiently low.

[24]I thank Eran Shmaya for suggesting a simple counter example.

[2] Y. Azrieli and E. Lehrer (2008), The value of a stochastic information structure, *Games and Economic Behavior* **63**, 679-693.

[3] M. Battaglini (2002), Multiple referrals and multidimensional cheap talk, *Econometrica* **70**, 1379-1401.

[4] S. Bikhchandani and I. Obara (2017), Mechanism design with information acquisition, *Economic Theory* **63**, 783-812.

[5] D. Blackwell (1953), Equivalent comparisons of experiments, *The annals of mathematical statistics* **24**, 265-272.

[6] D. Blackwell and M.A. Girshick (1954), *Theory of games and statistical decisions*, John Wiley & Sons, Inc.

[7] J.A. Bohren and T. Kravitz (2016), Optimal contracting with costly state verification, with an application to crowdsourcing, working paper.

[8] P. Bolton and M. Dewatripont (2005), *Contract theory*, The MIT Press.

[9] G. Carroll (2017), Robust incentives for information acquisition, working paper.

[10] H. Chade and N. Kovrijnykh (2016) Delegated information acquisition with moral hazard, *Journal of Economic Theory* **162**, 55-92.

[11] H. Chade and E. Schlee (2002) Another look at the Radner-Stiglitz non-concavity in the value of information, *Journal of Economic Theory* **107**, 421-452.

[12] A. Clark (2017), Contracts for information acquisition, working paper.

[13] J. Crémer and R. McLean (1988), Full extraction of the surplus in Bayesian and dominant strategy auctions, *Econometrica* **56**, 1247-1257.

[14] M. Dewatripont and J. Tirole (1999), Advocates, *Journal of Political Economy* **107**, 1-39.

[15] M. Gentzkow and E. Kamenica (2017), Competition in persuasion, *Review of Economic Studies* **84**, 300-322.

[16] D. Gerardi and L. Yariv (2008), Information acquisition in committees, *Games and Economic Behavior* **62**, 436-459.

[17] A. Gershkov and B. Szentes (2009), Optimal voting schemes with costly information acquisition, *Journal of Economic Theory* **144**, 36-68.

[18] D. Gromb and D. Martimort (2007), Collusion and the organization of delegated expertise, *Journal of Economic Theory* **137**, 271-299.

[19] S.J. Grossman and O.D. Hart (1983), An analysis of the principal-agent problem, *Econometrica* **51**, 7-45.

[20] S. Häfner and C.R. Taylor (2018), Contracting for research: Moral hazard and the incentive to overstate significance, working paper.

[21] B. Holmstrom (1982), Moral hazard in teams, *The Bell Journal of Economics* **13**, 324-340.

[22] N. Khanna, A. Peivandi, and M. Schroder (2015), Rewarding disagreement for optimal decisions, working paper.

[23] V. Krishna and J. Morgan (2001), A model of expertise, *Quarterly Journal of Economics* **116**, 747-775.

[24] R. Laraki (2004), On the regularity of the convexification operator on a compact set, *Journal of Convex Analysis* **11**, 209-234.

[25] A.W. Marshal, I. Olkin and B.C. Arnold (2011), *Inequalities: Theory of majorization and its applications*, Second Edition, Springer.

[26] C. Martinelli (2006), Would rational voters acquire costly information?, *Journal of Economic Theory* **129**, 225-251.

[27] P. Milgrom and J. Roberts (1990), Rationalizability, learning, and equilibrium in games with strategic complementarities, *Econometrica* **58** 1255-1277.

[28] R.B. Myerson (1982), Optimal coordination mechanisms in generalized principal-agent problems, *Journal of Mathematical Economics* **10**, 67-81.

[29] K. Osband (1989), Optimal forecasting incentives, *Journal of Political Economy* **97**, 1091-1112.

[30] N. Persico (2003), Committee design with endogenous information, *Review of Economic Studies* **70**, 1-27.

[31] R. Radner and J. Stiglitz (1984), A non-concavity in the value of information, in: *Bayesian models in economic theory*, edited by M. Boyer and R. Kihlstrom, Elsevier Science Publishers.

[32] D. Rahman (2011), Detecting profitable deviations, working paper.

[33] D. Rahman (2012), But who will monitor the monitor?, *American Economic Review* **102**, 2767-2797.

[34] D. Rappoport and V. Somma (2015), Incentivizing information design, working paper.

[35] R.T. Rockafellar (1970), *Convex analysis*, Princeton University Press.

[36] T.C. Schelling (1960), *The stratgey of conflict*, Harvard University Press.

[37] R. Strausz (2012), Mediated contracts and mechanism design, *Journal of Economic Theory* **147**, 1280-1290.

[38] D.M. Topkis (1998), *Supermodularity and complementarity*, Princeton University Press.

[39] E. Winter (2004), Incentives and discrimination, *American Economic Review* **94**, 764-773.

[40] L. Zermeño (2011), A principal-expert model and the value of menus, working paper.

A Missing proofs

Proof of Lemma 1:

If x_i is feasible for (COST) then clearly it must satisfy constraints (2)-(4), so it is feasible for (AUX) as well. Conversely, suppose that x_i satisfies (2)-(4). Then the first-order condition (2) guarantees that deviations (e'_i, r^*_i) (i.e., deviations only from the required effort level e_i without misreporting of observed signals) are not profitable. Indeed, convexity of c implies that $U_i\left((e'_i, r^*_i), (e_{-i}, r^*_{-i});\ x_i\right)$ is concave in e'_i, so the first-order condition is both necessary and sufficient for optimality.

Next, consider deviations (e'_i, r_i) with $r_i \equiv b$, i.e., i reports b regardless of his signal. If there exists such a profitable deviation then the deviation to $(0, r_i)$ is profitable as well, since it gives i the same expected transfer as (e'_i, r_i) at a minimal cost. But inequality (3) says that $(0, r_i)$ is not profitable, so (e'_i, r_i) is not profitable as well. A similar argument applies for deviations (e'_i, r_i) with $r_i \equiv w$.

Finally, consider deviations (e'_i, r_i) with $r_i(b) = w$ and $r_i(w) = b$ (i.e., the report is opposite from the observed signal). Then

$$\mathbb{P}_{\left((0,r^*_i),(e_{-i},r^*_{-i})\right)} \equiv \frac{e_i}{e_i + e'_i}\mathbb{P}_{\left((e'_i,r_i),(e_{-i},r^*_{-i})\right)} + \frac{e'_i}{e_i + e'_i}\mathbb{P}_{(e,r^*)},$$

that is, the distribution of reported vectors of signals when i exerts zero effort and reports truthfully is a convex combination of the distributions when i is honest and obedient and when he plays the proposed deviation (assuming all others are honest and obedient). This implies that

$$\mathbb{E}_{\left((0,r^*_i),(e_{-i},r^*_{-i})\right)}[x_i(s)] = \frac{e_i}{e_i + e'_i}\mathbb{E}_{\left((e'_i,r_i),(e_{-i},r^*_{-i})\right)}[x_i(s)] + \frac{e'_i}{e_i + e'_i}\mathbb{E}_{(e,r^*)}[x_i(s)]. \tag{14}$$

Therefore,

$$\begin{aligned} U_i\left((e, r^*);\ x_i\right) \quad \geq \quad & U_i\left((0, r^*_i), (e_{-i}, r^*_{-i});\ x_i\right) = \mathbb{E}_{\left((0,r^*_i),(e_{-i},r^*_{-i})\right)}[x_i(s)] = \\ & \frac{e_i}{e_i + e'_i}\mathbb{E}_{\left((e'_i,r_i),(e_{-i},r^*_{-i})\right)}[x_i(s)] + \frac{e'_i}{e_i + e'_i}\mathbb{E}_{(e,r^*)}[x_i(s)] \geq \\ & \frac{e_i}{e_i + e'_i}U_i\left((e'_i, r_i), (e_{-i}, r^*_{-i});\ x_i\right) + \frac{e'_i}{e_i + e'_i}U_i\left((e, r^*);\ x_i\right), \end{aligned}$$

where the first inequality is by the first paragraph of this proof, the first equality follows from $c(0) = 0$, the next equality is by (14), and the last inequality is by non-negativity of the cost function. It follows that $U_i((e, r^*);\ x_i) \geq U_i\left((e'_i, r_i), (e_{-i}, r^*_{-i});\ x_i\right)$, so (e'_i, r_i) is not a profitable deviation. $\square$

Proof of Lemma 2:

To simplify the notation, we write $\mathbb{P}$ instead of $\mathbb{P}_{(e,r^*)}$ when no confusion may arise. Also, it will be convenient to write $\mathbb{P}(s_{-i}|B) = e(s^b_{-i})\bar{e}(s^w_{-i})$ and $\mathbb{P}(s_{-i}|W) = \bar{e}(s^b_{-i})e(s^w_{-i})$ for the conditional probability of s_{-i} given each state of nature. Using this notation, constraint (5) for $s_{-i} = \underline{b}_{-i}$ at z^* becomes

$$\left[\gamma\mathbb{P}(\underline{b}_{-i}|B) - (1-\gamma)\mathbb{P}(\underline{b}_{-i}|W)\right] z^*_1 - \mathbb{P}(w, \underline{b}_{-i})z^*_2 \geq \mathbb{P}(\underline{b}),$$

or, more explicitly,

$$[\gamma e(N_{-i}) - (1-\gamma)\bar{e}(N_{-i})]\, z^*_1 - [\gamma(0.5-e_i)e(N_{-i}) + (1-\gamma)(0.5+e_i)\bar{e}(N_{-i})]\, z^*_2 \leq \gamma e(N) + (1-\gamma)\bar{e}(N).$$

It is tedious but straightforward to verify that this constraint holds with equality.

Now, consider constraint (5) at z^* for some other s_{-i}. After a slight rearrangement it becomes

$$\frac{\gamma\mathbb{P}(s_{-i}|B) - (1-\gamma)\mathbb{P}(s_{-i}|W)}{\mathbb{P}(b, s_{-i})} z^*_1 - \frac{\mathbb{P}(w, s_{-i})}{\mathbb{P}(b, s_{-i})} z^*_2 \leq 1.$$

Thus, to establish this inequality it is sufficient to show that $\frac{\gamma\mathbb{P}(s_{-i}|B)-(1-\gamma)\mathbb{P}(s_{-i}|W)}{\mathbb{P}(b,s_{-i})}$ is maximized at $s_{-i} = \underline{b}_{-i}$ and that $\frac{\mathbb{P}(w,s_{-i})}{\mathbb{P}(b,s_{-i})}$ is minimized at $s_{-i} = \underline{b}_{-i}$ (note that $z^*_1, z^*_2 \geq 0$). For the coefficient of z^*_1 we have

$$\begin{aligned}\frac{\gamma\mathbb{P}(s_{-i}|B) - (1-\gamma)\mathbb{P}(s_{-i}|W)}{\mathbb{P}(b, s_{-i})} &= \frac{\gamma\mathbb{P}(s_{-i}|B) - (1-\gamma)\mathbb{P}(s_{-i}|W)}{\gamma(0.5+e_i)\mathbb{P}(s_{-i}|B) + (1-\gamma)(0.5-e_i)\mathbb{P}(s_{-i}|W)} = \\ &\frac{1}{(0.5+e_i) + \frac{1-\gamma}{\gamma}(0.5-e_i)\frac{\mathbb{P}(s_{-i}|W)}{\mathbb{P}(s_{-i}|B)}} - \frac{1}{\frac{\gamma}{1-\gamma}(0.5+e_i)\frac{\mathbb{P}(s_{-i}|B)}{\mathbb{P}(s_{-i}|W)} + (0.5-e_i)},\end{aligned}$$

which clearly increases in the likelihood ratio $\frac{\mathbb{P}(s_{-i}|B)}{\mathbb{P}(s_{-i}|W)}$, and is hence maximal at $s_{-i} = \underline{b}_{-i}$. And for the coefficient of z^*_2 we have

$$\frac{\mathbb{P}(w, s_{-i})}{\mathbb{P}(b, s_{-i})} = \frac{\gamma(0.5-e_i)\mathbb{P}(s_{-i}|B) + (1-\gamma)(0.5+e_i)\mathbb{P}(s_{-i}|W)}{\gamma(0.5+e_i)\mathbb{P}(s_{-i}|B) + (1-\gamma)(0.5-e_i)\mathbb{P}(s_{-i}|W)} = \frac{\gamma(0.5-e_i)\frac{\mathbb{P}(s_{-i}|B)}{\mathbb{P}(s_{-i}|W)} + (1-\gamma)(0.5+e_i)}{\gamma(0.5+e_i)\frac{\mathbb{P}(s_{-i}|B)}{\mathbb{P}(s_{-i}|W)} + (1-\gamma)(0.5-e_i)},$$

which decreases in $\frac{\mathbb{P}(s_{-i}|B)}{\mathbb{P}(s_{-i}|W)}$ and so minimized at $s_{-i} = \underline{b}_{-i}$. This proves that (5) holds at z^* for every for every $s_{-i} \in S_{-i}$.

The proof that constraints (6) hold at z^* is similar. First, it is not hard to check that for $s_{-i} = \underline{w}_{-i}$ constraint (6) is satisfied with equality at z^*. Next, for any other s_{-i} we can rewrite the constraint as

$$\frac{\gamma\mathbb{P}(s_{-i}|B) - (1-\gamma)\mathbb{P}(s_{-i}|W)}{\mathbb{P}(w, s_{-i})} z^*_1 - z^*_2 \geq -1.$$

The coefficient of z_1^* is decreasing in the likelihood ratio $\frac{\mathbb{P}(s_{-i}|W)}{\mathbb{P}(s_{-i}|B)}$ and hence minimized at $s_{-i} = \underline{w}_{-i}$. It follows that (6) is satisfied for every s_{-i} at z^*. This completes the proof. $\square$

Proof of Corollary 1:

First,

$$\frac{\partial \psi_i}{\partial e_i} = \frac{\big[\gamma e(N_{-i}) + (1-\gamma)\bar{e}(N_{-i})\big]\big[-\gamma\bar{e}(N_{-i}) + (1-\gamma)e(N_{-i})\big] + (2\gamma-1)e(N_{-i})\bar{e}(N_{-i})}{\gamma(1-\gamma)\big[e(N_{-i})^2 - \bar{e}(N_{-i})^2\big]} c'(e_i)$$
$$+ \frac{\big[\gamma\bar{e}(N) + (1-\gamma)e(N)\big]\big[\gamma e(N_{-i}) + (1-\gamma)\bar{e}(N_{-i})\big]}{\gamma(1-\gamma)\big[e(N_{-i})^2 - \bar{e}(N_{-i})^2\big]} c''(e_i).$$

The first term simplifies to just $c'(e_i) > 0$, and the second term is clearly positive, which proves that $\frac{\partial\psi_i}{\partial e_i} > 0$.

As for the other derivative $\frac{\partial\psi_i}{\partial e_j}$ with $j \neq i$, note first that the denominator $\gamma(1-\gamma)\big[e(N_{-i})^2 - \bar{e}(N_{-i})^2\big]$ of ψ_i is increasing in e_j, and that the second term $(2\gamma-1)e(N_{-i})\bar{e}(N_{-i})c(e_i)$ in the numerator of ψ_i is decreasing in e_j. To prove that the derivative is negative it is therefore enough to prove that the ratio

$$\frac{\big[\gamma\bar{e}(N) + (1-\gamma)e(N)\big]\big[\gamma e(N_{-i}) + (1-\gamma)\bar{e}(N_{-i})\big]}{e(N_{-i})^2 - \bar{e}(N_{-i})^2}$$

decreases in e_j. After some rearranging, the numerator of the derivative of this ratio with respect to e_j becomes

$$\begin{aligned}
&\Big\{2\gamma(1-\gamma)\big[e(N_{-i})e(N_{-j}) - \bar{e}(N_{-i})\bar{e}(N_{-j})\big] - 2e_j e(N_{-ij})\bar{e}(N_{-ij})\big[\gamma^2(0.5-e_i) + (1-\gamma)^2(0.5+e_i)\big]\Big\} \times \\
&\Big\{e(N_{-i})^2 - \bar{e}(N_{-i})^2\Big\} - \\
&2\Big\{e(N_{-i})e(N_{-ij}) + \bar{e}(N_{-i})\bar{e}(N_{-ij})\Big\} \times \\
&\Big\{\gamma^2\bar{e}(N)e(N_{-i}) + (1-\gamma)^2 e(N)\bar{e}(N_{-i}) + \gamma(1-\gamma)\big[e(N)e(N_{-i}) + \bar{e}(N)\bar{e}(N_{-i})\big]\Big\}.
\end{aligned}$$

Eliminating some of the clearly negative terms, we get that this expression is bounded above by

$$\begin{aligned}
&2\gamma(1-\gamma)\big[(e(N_{-i})e(N_{-j}) - \bar{e}(N_{-i})\bar{e}(N_{-j}))\big]\big[e(N_{-i})^2 - \bar{e}(N_{-i})^2\big] - \\
&2\gamma(1-\gamma)\big[e(N_{-i})e(N_{-ij}) + \bar{e}(N_{-i})\bar{e}(N_{-ij})\big]\big[(e(N)e(N_{-i}) + \bar{e}(N)\bar{e}(N_{-i}))\big].
\end{aligned}$$

It is immediate to verify that this last expression is negative, which completes the proof. $\square$

Proof of Lemma 3:

Fix $e_1 \geq e_2$. The set of possible signals in the information structures $m(e_1, e_2)$ can be identified with $\{\emptyset,\ 1,\ 2,\ 12\}$, corresponding to the coalition of experts who got signal b. For each signal A in this set denote by $p_e(A) = \frac{1}{2}\left[e(A)\bar{e}(A^c) + \bar{e}(A)e(A^c)\right]$ the probability that signal A is observed, and by $q_e(A) = \frac{\frac{1}{2}e(A)\bar{e}(A^c)}{p_e(A)}$ the posterior probability that the state is B after signal A is observed (assuming a uniform prior). We view the posterior of state B as a $[0,1]$-valued random variable which takes the values $\{q_e(A)\}$ with corresponding probabilities $\{p_e(A)\}$. The cumulative distribution function (cdf) of this variable is

$$F_e(t) = \sum_{\{A:\ q_\alpha(A) \leq t\}} p_e(A).$$

Let $e'_1 \geq e'_2$ be obtained from (e_1, e_2) by a PD transfer, i.e., $e_1 + e_2 = e'_1 + e'_2$ and $e_1 \geq e'_1$. The probabilities $p_{e'}(A)$ and $q_{e'}(A)$, and the cdf $F_{e'}(t)$ are defined in an analogous way to the above definitions. By Blackwell and Girshick [6, Theorem 12.4.1 on page 332], $m(e)$ is more informative than $m(e')$ if and only if

$$\int_0^x F_e(t)dt \geq \int_0^x F_{e'}(t)dt \tag{15}$$

holds for every $x \in [0,1]$. To complete the proof we now show that (15) holds at the four atoms of F_e, i.e. at the points $x = q_e(\emptyset), q_e(1), q_e(2)$, and $q_e(12)$. Since F_e and $F_{e'}$ are non-decreasing step-functions this would imply that (15) holds for every $x \in [0,1]$. Indeed, if $\int_0^x F_e(t)dt < \int_0^x F_{e'}(t)dt$ at some $x \in [0,1]$, then the same must be true at one of the jumps of F_e adjacent to x.

We will need the following simple observations, whose proofs can be found at the end of this proof:

(a) $q_e(\emptyset) \leq q_e(2) \leq \frac{1}{2} \leq q_e(1) \leq q_e(12)$.

(b) $q_{e'}(\emptyset) \leq q_{e'}(2) \leq \frac{1}{2} \leq q_{e'}(1) \leq q_{e'}(12)$.

(c) $q_e(\emptyset) \leq q_{e'}(\emptyset)$, $q_e(2) \leq q_{e'}(2)$, $q_{e'}(1) \leq q_e(1)$, and $q_{e'}(12) \leq q_e(12)$.

(d) $F_e(t) = 1 - F_e(1-t)$ and $F_{e'}(t) = 1 - F_{e'}(1-t)$ for every $t \in [0,1]$.

1. $\underline{x = q_e(\emptyset)}$:

From observations (b) and (c) it immediately follows that $q_e(\emptyset)$ is smaller than the four possible posteriors under e'. Thus, $F_{e'}(t) = 0$ for every $t \in [0, q_e(\emptyset)]$, which implies $\int_0^{q_e(\emptyset)} F_{e'}(t)dt = 0$. Inequality (15) at $x = q_e(\emptyset)$ follows.

2. $\underline{x = q_e(2)}$:

From observation (a) we have that $\int_0^{q_e(2)} F_e(t)dt = [q_e(2) - q_e(\emptyset)]p_e(\emptyset)$, and from observations (b) and (c) we have that either $\int_0^{q_e(2)} F_{e'}(t)dt = [q_e(2) - q_{e'}(\emptyset)]p_{e'}(\emptyset)$ or $\int_0^{q_e(2)} F_{e'}(t)dt = 0$. In the latter case there is nothing to prove, so suppose the former is true. We therefore need to show that

$$[q_e(2) - q_e(\emptyset)]p_e(\emptyset) \geq [q_e(2) - q_{e'}(\emptyset)]p_{e'}(\emptyset),$$

or equivalently that

$$q_e(2)[p_e(\emptyset) - p_{e'}(\emptyset)] \geq q_e(\emptyset)p_e(\emptyset) - q_{e'}(\emptyset)p_{e'}(\emptyset). \tag{16}$$

Using the equality $e_1 + e_2 = e'_1 + e'_2$, simple algebra gives that the right-hand side of (16) is equal to $\frac{1}{2}(e_1e_2 - e'_1e'_2)$. Also, it is easy to verify that $p_e(\emptyset) - p_{e'}(\emptyset) = e_1e_2 - e'_1e'_2$, so (16) becomes

$$q_e(2)(e_1e_2 - e'_1e'_2) \geq \frac{1}{2}(e_1e_2 - e'_1e'_2).$$

Since the area of a rectangle with a given perimeter decreases in the difference between its length and its width, we have that $e_1e_2 - e_1'e_2' \leq 0$, and by observation (a) we have that $q_e(2) \leq \frac{1}{2}$. This proves (16).

3. $\underline{x = q_e(1)}$:

This inequality is the "mirror image" of the inequality of the previous case. Indeed, using the symmetry of F_e around 0.5 (observation (d)) and a simple change of variables we get that

$$\int_0^{q_e(1)} F_e(t)dt = q_e(1) - \int_0^1 F_e(t)dt + \int_0^{1-q_e(1)} F_e(t)dt,$$

and similarly that

$$\int_0^{q_e(1)} F_{e'}(t)dt = q_e(1) - \int_0^1 F_{e'}(t)dt + \int_0^{1-q_e(1)} F_{e'}(t)dt.$$

Now, since the expected posterior is equal to the prior, we have that $\int_0^1 F_e(t)dt = \int_0^1 F_{e'}(t)dt$. Thus, inequality (15) at $x = q_e(1)$ is equivalent to $\int_0^{1-q_e(1)} F_e(t)dt \geq \int_0^{1-q_e(1)} F_{e'}(t)dt$. But notice that $1 - q_e(1) = q_e(2)$, so the last inequality is the same as the one proved for $x = q_e(2)$.

4. $\underline{x = q_e(12)}$:

As in the previous case, it is simple to show that inequality (15) at $x = q_e(12)$ is equivalent to the inequality at $x = q_e(\emptyset)$ proven above. We omit the details.

Proofs of observations (a)-(d):

(a): The posterior probability of state B is clearly nondecreasing (with respect to set inclusion) in the coalition of experts who obtained signal b. Thus, to prove observation (a) we only need to check that $q_e(2) \leq \frac{1}{2} \leq q_e(1)$. The latter inequality immediately follows from $e_1 \geq e_2$, since

$$q_e(2) = \frac{1}{1 + \frac{(0.5+e_1)(0.5-e_2)}{(0.5+e_2)(0.5-e_1)}} \quad \textit{and} \quad q_e(1) = \frac{1}{1 + \frac{(0.5+e_2)(0.5-e_1)}{(0.5+e_1)(0.5-e_2)}}.$$

(b): The proof is identical to that of observation (a) (recall that $e_1' \geq e_2'$).

(c): We have

$$q_e(\emptyset) = \frac{1}{1 + \frac{(0.5+e_1)(0.5+e_2)}{(0.5-e_1)(0.5-e_2)}} \quad \textit{and} \quad q_{e'}(\emptyset) = \frac{1}{1 + \frac{(0.5+e_1')(0.5+e_2')}{(0.5-e_1')(0.5-e_2')}},$$

so we need to show that $\frac{(0.5+e_1)(0.5+e_2)}{(0.5-e_1)(0.5-e_2)} \geq \frac{(0.5+e_1')(0.5+e_2')}{(0.5-e_1')(0.5-e_2')}$. The latter is equivalent to $(0.5 + e_1)(0.5 + e_2)(e_1'e_2' - e_1e_2) \geq (0.5 - e_1)(0.5 - e_2)(e_1'e_2' - e_1e_2)$, which follows from $e_1'e_2' \geq e_1e_2$.

Next,

$$q_e(2) = \frac{1}{1 + \frac{(0.5+e_1)(0.5-e_2)}{(0.5-e_1)(0.5+e_2)}} \quad \textit{and} \quad q_{e'}(2) = \frac{1}{1 + \frac{(0.5+e_1')(0.5-e_2')}{(0.5-e_1')(0.5+e_2')}},$$

so $q_e(2) \leq q_{e'}(2)$ is equivalent to $\frac{(0.5+e_1)(0.5-e_2)}{(0.5-e_1)(0.5+e_2)} \geq \frac{(0.5+e_1')(0.5-e_2')}{(0.5-e_1')(0.5+e_2')}$, which follows from $e_1 \geq e_1'$ and $e_2 \leq e_2'$. The rest of the inequalities are proved in a similar fashion, the details are omitted.

(d): $F_e(t)$ is the probability that the posterior of state B is less or equal to t, while $1 - F_e(1 - t)$ is the probability that the posterior of state W is less or equal to t. Since the prior and the information structure are symmetric between the two states, these two probabilities must be equal. The same argument holds for $F_{e'}$. $\square$

Proof of Lemma 4:

Suppose that $e_1 \geq e_2$. Then the posterior probability of state B is greater or equal to $q = 0.5$ when expert 1 sends signal b and less or equal to $q = 0.5$ when he sends signal w. Thus, the optimal choice for the DM is independent of expert 2's signal, which implies that $V(e_1, e_2) = V(e_1)$. Finally, the distribution over posteriors induced by $m_1(e_1)$ has a mass of 0.5 at the posterior $q = 0.5 + e_1$ and a mass of 0.5 at $q = 0.5 - e_1$. Therefore,

$$V(e_1) = 0.5v(0.5 + e_1) + 0.5v(0.5 - e_1).$$

Since v is the maximum of two linear functions that cross at $q = 0.5$ there are numbers $\alpha_1, \beta_1, \alpha_2, \beta_2$, with $\alpha_1 > \alpha_2$ such that

$$v(0.5 + e_1) = \alpha_1(0.5 + e_1) + \beta_1$$

and

$$v(0.5 - e_1) = \alpha_2(0.5 - e_1) + \beta_2.$$

Thus,

$$V(e_1) = 0.5(\alpha_1(0.5 + e_1) + \beta_1) + 0.5(\alpha_2(0.5 - e_1) + \beta_2) = \frac{\alpha_1 - \alpha_2}{2} e_1 + \frac{0.5\alpha_1 + \beta_1 + 0.5\alpha_2 + \beta_2}{2}.$$

The case $e_2 > e_1$ is similar. $\square$

Proof of Lemma 5:

The mapping $m_i \mapsto \mathbb{P}_{(m_i, m^*_{-i})}$ is affine, and therefore $\mathbb{E}_{(m_i, m^*_{-i})}[x_i(s)]$ is affine in m_i for any fixed x_i. Since C_i is convex (assumption (A2)) it follows that $f(m_i, x_i)$ is concave in its first argument, and since C_i is continuous (assumption (A3)) it follows that $f(m_i, x_i)$ is continuous in its first argument. Finally, f is clearly affine and continuous in x_i for any given m_i. $\square$

Proof of Lemma 6:

Let $g : \mathbb{R}^p \to \mathbb{R}^q$ be a non-constant linear function and define the set $K = Ker(g)^{\perp} \bigcap P$, where $Ker(g)^{\perp}$ is the orthogonal complement of $Ker(g)$ (the kernel of g) and P is the unit sphere of $\mathbb{R}^p$. Let $\alpha = \min_{z \in K} \|g(z)\|_q$, where the minimum is attained due to compactness and continuity. Moreover, $\alpha > 0$ by construction. Fix some $y \in \mathbb{R}^p$ and let $\bar{y}$ be the projection of y to $Ker(g)$. We claim that $\|g(y)\|_q \geq \alpha \|y - \bar{y}\|_p$. Indeed, the inequality is trivially satisfied when $y \in Ker(g)$, and if $y \notin Ker(g)$ then

$$\|g(y)\|_q = \|g(y) - g(\bar{y})\|_q = \|g(y - \bar{y})\|_q = \|y - \bar{y}\|_p \left\| g\left(\frac{y - \bar{y}}{\|y - \bar{y}\|_p}\right) \right\|_q \geq \alpha \|y - \bar{y}\|_p,$$

where the first equality follows from $\bar{y} \in Ker(g)$, the second and third follow from the linearity of g, and the inequality is by the definition of α.

Now, the above conclusion remains valid if g is affine rather than linear, so we may apply this to the function T of the lemma. Thus, denoting by $\tilde{m}_i$ the projection of m_i to $Ker(T)$, there is $\alpha > 0$ such that $\|T(m_i)\|_S \geq \alpha \|m_i - \tilde{m}_i\|$ for every $m_i \in M_i$. Finally, since M_i is polyhedral (assumption (A1)), there is $\alpha' > 0$ such that $\|m_i - \tilde{m}_i\| \geq \alpha' \|m_i - \bar{m}_i\|$ for all $m_i \in M_i$. The product $\beta' = \alpha\alpha'$ then satisfies the required inequality. □

B An alternative general framework

The purpose of this appendix is to show how the model of Section 3 can be derived from a more primitive framework in which experts choose how much effort to exert and what to report to the DM (as in the binary-binary model of Section 2). In particular, we show that assumptions (A1)-(A3) of Section 3 are consequences of standard assumptions in this basic framework; and that the notion of implementability of information structures m^* in Definition 1 is equivalent to implementability of strategies that induce m^* in this basic framework.

As in the model of Section 3, the set of possible states of nature is $\Omega = \{\omega_1, \ldots, \omega_K\}$, and the common prior over Ω is $\gamma = (\gamma_1, \ldots, \gamma_K)$. The set of experts is $N = \{1, \ldots, n\}$, and S_i is the set of signals that expert $i \in N$ may observe. Each expert i chooses an effort level $e_i \in [0, \bar{e}]$. Every e_i determines a mapping $m_i(e_i) : \Omega \to \Delta(S_i)$, which we identify with a stochastic matrix with K rows and $|S_i|$ columns. Here $m_i(e_i)(k, s^i)$ is the probability that signal $s^i \in S^i$ realizes conditional on the state being $\omega_k \in \Omega$, given effort level e_i.

We assume that informativeness increases with effort: If $e_i > e_i'$ then $m_i(e_i)$ is strictly more informative than $m_i(e_i')$ in the sense of Blackwell [5]. We will also need to assume that the technology m_i is continuous on $[0, \bar{e}]$. The cost of exerting effort e_i is $c_i(e_i)$, where $c_i : [0, \bar{e}] \to \mathbb{R}_+$ is strictly increasing.

Let $S = S_1 \times \ldots \times S_n$. A contract is a list $x = (x_1, \ldots, x_n)$ with $x_i : S \to \mathbb{R}_+$ for each $i \in N$. A contract x defines a game between the experts: A pure strategy for expert i is a pair (e_i, r_i), where e_i is i's effort level and $r_i : S_i \to S_i$ determines the report that i sends to the DM as a function of the signal he observes. It is convenient to think of r_i as a (stochastic) matrix of dimensions $|S_i| \times |S_i|$, with $r_i(s_i', s_i)$ equals 1 if $r_i(s_i') = s_i$ and equals zero otherwise. Notice that the information that the DM receives from i under the strategy (e_i, r_i) is a garbling of the information that i privately observes, and that the stochastic matrix that describes this information structure is the product $b_i(e_i)r_i$.

We denote a pure strategy profile by $(e, r) = ((e_1, \ldots, e_n),\ (r_1, \ldots, r_n))$. Each (e, r) induces a distribution over the vector of signals $s \in S$ that the DM observes, denoted $\mathbb{P}_{(e,r)}$:

$$\mathbb{P}_{(e,r)}(s) = \sum_{k=1}^{K} \gamma_k \prod_{i=1}^{n} (m_i(e_i)r_i)(k, s_i). \tag{17}$$

Note that (17) assumes that different experts' signals are independent conditional on the state. The payoff to expert i given pure strategy profile (e, r) is

$$\mathbb{E}_{(e,r)}[x_i(s)] - c_i(e_i). \tag{18}$$

Experts can also use mixed strategies. To avoid unnecessary technical issues we only consider finite-support distributions over pure strategies. When i plays the mixed strategy σ_i that assigns probability λ^l to the pure strategy (e_i^l, r_i^l) $(l = 1, \dots L,\ \sum_l \lambda_l = 1)$, the induced information structure is the convex combination $m_i(\sigma_i) := \sum_l \lambda^l m_i(e_i^l) r_i^l$. With abuse of notation we denote the expected cost of such a mixed strategy σ_i by $c_i(\sigma_i) = \sum_l \lambda^l c_i(e_i^l)$. The payoff in (18) is extended to profiles of mixed strategies as usual.

Let M_i be the set of all information structures (i.e., mappings from Ω to $\Delta(S_i)$) that can be induced by some (pure or mixed) strategy of i:

$$M_i = \{m_i(\sigma_i)\ :\ \sigma_i \textit{ is a strategy of expert } i\}. \tag{19}$$

For each $m_i \in M_i$ let $C_i(m_i)$ be the cost of the least costly way for i to induce m_i:

$$C_i(m_i) = \inf\{c_i(\sigma_i)\ :\ \sigma_i \textit{ induces } m_i\}. \tag{20}$$

Lemma 7. The set M_i defined in (19) satisfies assumption (A1) of Section 3. The cost function C_i defined in (20) satisfies assumption (A2) of Section 3.

Proof. First, every $m_i \in M_i$ is a convex combination of garblings of information structures from the image of m_i. Since informativeness increases with effort, and since the set of garblings of a given information structure is convex, it follows that every $m_i \in M_i$ is a garbling of $m_i(\bar{e})$. Conversely, every garbling of $m_i(\bar{e})$ can be induced as a mixture of pure strategies of the form $(\bar{e}, r_i^l)$ and is therefore in M_i. It follows that M_i is equal to the set of all garblings of $m_i(\bar{e})$, which proves that M_i is a closed and convex polyhedral set. Thus, (A1) is satisfied.

Next, for every $e_i \in [0, \bar{e}]$ let $M_i(e_i) = \{m_i(e_i)r_i\}_{r_i}$ be the set of information structures that can be induced by exerting effort e_i and then applying some (pure) reporting strategy. For m_i that can be induced by some pure strategy define $E_i(m_i) = \min\{e_i \in [0, \bar{e}]\ :\ m_i \in M_i(e_i)\}$, where the minimum is attained due to continuity of m_i. The composite function $c_i \circ E_i$, defined for all m_i's that can be induced by pure strategies, describes the cost required to induce each such m_i using pure strategies only. The cost function C_i is then the convexification of $c_i \circ E_i$ defined over the convex-hull M_i of the domain of $c_i \circ E_i$. In other words, C_i is the largest convex function on M_i that is point-wise below $c_i \circ E_i$. This implies convexity (A2).

□

Lemma 7 only concerns assumptions (A1) and (A2). In order to conclude that C_i is Lipschitz continuous on M_i (assumption (A3)) we need to impose further restrictions on the primitives of the environment. From the proof of the Lemma 7 we have that C_i is the convexification of the composite function $c_i \circ E_i$. Since M_i is polyhedral, it follows from Laraki [24] that a sufficient condition for C_i to be Lipschitz is that $c_i \circ E_i$ is Lipschitz. This latter condition is satisfied whenever (1) c_i is Lipschitz on $[0, \bar{e}]$; and (2) for every pure reporting strategy r_i, the mapping $e_i \mapsto m_i(e_i)r_i$ is either constant on the entire interval $[0, \bar{e}]$, or it is one-to-one and its inverse is Lipschitz on its image. We omit the proof, but will assume for the next result that $c_i \circ E_i$ is continuous.

Consider the game of Section 3 with M_i from (19) and C_i from (20). We have the following.

Lemma 8. Fix a contract x. Then $m^* \in M$ is an equilibrium of the game of Section 3 if and only if there is an equilibrium σ such that[25] $m(\sigma) = m^*$. In particular, m^* is implementable according to Definition 1 if and only if there is a contract x and an equilibrium σ of the game induced by x such that $m(\sigma) = m^*$.

Proof. First, note that since $c_i \circ E_i$ is continuous, it follows from standard arguments that the infimum in (20) is attained.

Suppose that σ is an equilibrium of the game induced by x and let $m^* = m(\sigma)$. Consider a deviation $m_i' \in M_i$ for expert i. Let σ_i' be a strategy for i such that $m_i(\sigma_i') = m_i'$ and $C_i(m_i') = c_i(\sigma_i')$ (existence of such σ_i' follows from the first sentence of this proof). Then the payoff to i by choosing m_i' in the game of Section 3 is the same as his payoff for choosing σ_i' in the game of this appendix, and his payoff by choosing m_i^* in the game of Section 3 is at least his payoff for choosing σ_i in the game of this appendix. Since σ_i is a best response to σ_{-i}, it follows that deviating to m_i' is not profitable, hence m^* is an equilibrium of the game of Section 3.

To prove the converse, start with an equilibrium m^* of the game of Section 3, and consider σ such that $m_i(\sigma_i) = m_i^*$ and $c_i(\sigma_i) = C_i(m_i)$ for every i. Then it is immediate to check in a similar way to the previous paragraph that σ is an equilibrium of the game of this appendix. This completes the proof. $\square$

[25] The notation $m(\sigma) = m^*$ means that $m_i(\sigma_i) = m_i^*$ for every $i \in N$.