Theoretical Economics 16 (2021), 1313–1350 1555-7561/20211313

Monitoring experts

Yaron Azrieli
Department of Economics, The Ohio State University

We study the design of contracts that incentivize experts to collect information and truthfully report it to a decision maker. We depart from most of the previous literature by assuming that the transfers cannot depend on the realized state or on the ex post payoff of the decision maker. The contract thus has to induce the experts to "monitor each other" by making the transfers contingent on the entire vector of reports. We characterize the least costly contract that implements any given vector of efforts and derive the cost function for the decision maker. We then study properties of optimal contracts by comparing the value of information and its cost.

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

JEL classification. D82, D86.

1. Introduction

Policy makers, managers, journal editors, and individual consumers, to name just a few, often rely on the advice of experts to make critical decisions. In cases where the expert is not directly affected by the eventual choice of the decision maker (DM, henceforth), and where the quality of information is not verifiable, a moral hazard problem naturally arises: The expert would prefer not to incur the cost of collecting information; the DM in turn needs to design the contract in a way that motivates the expert to exert effort, as well as to truthfully report his findings.

If the environment is such that the payment to the expert can depend on the ex-post realized state, or, at least, on the eventual payoff of the DM, then the situation is similar to the classic principal-agent problem. Indeed, incentives to provide high-quality information can be created by paying high rewards in cases where the expert's recommendation turned out to be a good one, and low rewards after a bad advice.[1] However, there are many cases in which contracting on the realized state or payoff is infeasible. One such scenario occurs when the state is observed only in the far future, e.g., when an expert is asked to evaluate the effectiveness of a certain proposal to reduce global

Yaron Azrieli: azrieli.2@osu.edu
I thank comments from James Best, Aislinn Bohren, Ozan Candogan, Rahul Deb, Yoram Halevy, PJ Healy, Teddy Kim, Dan Levin, Mallesh Pai, Jim Peck, Marcin Peski, Bruno Salcedo, Edward Schlee, Eran Shmaya, Richard van Weelden, Alex Wolitzky, and from participants at several seminars and conferences where this work has been presented. I am also grateful to three anonymous referees of *Theoretical Economics* for their helpful suggestions.

[1] See Section 4.2.2 for references to these types of models.

 https://doi.org/10.3982/TE4017

warming. Second, it may be difficult or even impossible to verify the true state, e.g., in the case where a political candidate hires a pollster to estimate public sentiment on a certain issue, or when a journal editor asks a referee to evaluate the quality of a paper. Designing performance-based contracts in this type of environments is a challenging task.

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.[2]

This method of peer-monitoring works not only in theory—we are aware of several real-world instances where similar methods have been applied. The Sensors and Sensing Systems Program of the National Science Foundation (NSF) ran a pilot review process in 2013 where each submitted grant proposal was evaluated by several PIs who also submitted proposals to the same program. The instructions state that “Each individual PI’s rankings will be compared to the global ranking, and the PI’s ranking will be adjusted in accordance with the degree to which his/her ranking matches the global ranking. This adjustment provides an incentive to each PI to make an honest and thorough assessment of the proposals to which they are assigned as failure to do so results in the PI placing himself/herself at a disadvantage compared to others in the group.”[3]

A second example comes from crowdsourcing platforms such as Amazon Mechanical Turk. Requestors (i.e., people who post jobs) on these platforms have the option to assign the same task to multiple workers, and to pay a worker only if his/her answer matches that of other workers assigned to the same task. Platforms explicitly market this tool as a way to ensure that tasks are completed in a high-quality manner.[4] A related area where similar methods have been used is experimental economics: One well-known example is the work of Krupka and Weber (2013) on social norms, where subjects were asked to reveal their view on what is appropriate behavior in various versions of the “dictator game”; subjects were paid a bonus if their response matched the modal response in

[2]Similar ideas have been studied before in various contexts; see the related literature section below.

[3]For a complete description of the review process see https://www.nsf.gov/pubs/2013/nsf13096/nsf13096.jsp. This method was based on a proposal of Merrifield and Saari (2009). An interesting discussion of this proposal by Rakesh Vohra appears in “The Theory of the Leisure Class” blog at https://theoryclass.wordpress.com/2013/06/06/a-mechanism-design-approach-to-peer-review. I thank Bruno Salcedo for pointing me to this example.

[4]See for example the explanation at https://blog.mturk.com/cooking-tip-5-ask-multiple-workers-to-complete-a-hit-ec21c9fc0734.

their session.[5] A similar elicitation method was used in Xiao and Houser (2005) to incentivize subjects tasked with categorizing natural language messages of previous subjects who played ultimatum games.

The main contribution of this paper is to provide insights into properties of optimal contracts in the above setup. To gain tractability, we consider a simple model with binary state-space and signal-space for each expert. The probability that an expert receives the "correct" signal increases in that expert's (continuous) effort. There is a common prior belief shared by the DM and all the experts, and signals of different experts are independent conditional on the true state. We focus on the case of risk-neutral parties and assume that experts are protected by limited liability.[6]

Our analysis of optimal contracts is based on the Grossman and Hart (1983) approach: As a first step, for each vector of efforts 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 effort vectors.

For a given effort vector, the minimization problem describing the least costly contract implementing that vector boils down to a linear program with three constraints. One is a standard first-order condition guaranteeing that there is no profitable deviation from the required effort level, *assuming truthful reporting*. But this is not sufficient, since an expert may nevertheless find it profitable to reduce his effort and misreport his signal. The two other constraints make sure that is not the case.[7] We point out that, except for the case of a uniform prior, one of these two adverse selection constraints is binding. Thus, the cost to the DM is strictly higher than in the case of pure moral hazard where signals are verifiable; see Section 4.2.1 for details.

Our first main result is that the least costly incentive compatible contract has a particularly simple form: For any effort vector, the experts are paid 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. Note that this contract punishes the entire group if one expert misreports (or is unlucky), similar to the classic Holmstrom (1982) mechanism; however, the event that triggers the punishment here (disagreement of reports) is of a different nature than in Holmstrom (low output). The solution allows us to derive an explicit formula for the DM's cost function.

We then move on to study the value of information for the DM. It is important to point out that this part is independent of the contracting problem and is typically missing in traditional models of moral hazard or adverse selection. In our setting of contracting for information, the output of the experts' work is an input for the DM choice of

[5]The cost of collecting information in this case can be thought of as the time and effort involved in contemplating what to choose.

[6]While this is a highly stylized setup, the binary-binary information structure is common in the literature and allows for a clean characterization of the cost-minimizing contract. In Section 4.2, we discuss several possible extensions, including the case in which each expert obtains information about a different dimension of the (multidimensional) state, the case of non-common priors, the case in which experts are heterogeneous in their cost or productivity, and the case of risk-averse experts.

[7]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 (2005, p. 148)).

action. This implies certain restrictions on the mapping from effort vectors to DM's expected utility. In particular, we prove the following: If two sets of experts have the same average accuracy[8] of signals, and in one of these sets the spread of accuracies is larger than the other, then the former gives higher utility to the DM *in every decision problem*. This is equivalent to saying that the former is more informative than the latter in the sense of Blackwell (1953). For example, ignoring the cost, in every decision problem two experts with respective accuracies $\frac{7}{8}$ and $\frac{5}{8}$ generate higher expected utility to the DM than two experts each with accuracy of $\frac{3}{4}$, and the latter generates higher expected utility than 3 experts with accuracy of $\frac{2}{3}$ each.

This latter result expresses a certain kind of convexity in the value of information,[9] and has important implications for the optimal contract problem. Technically, it implies that first-order conditions for effort vectors are not sufficient for optimality, even if the cost function derived in the first step is convex. But it also implies economically meaningful necessary conditions for optimality. For example, we show that even in our symmetric environment an optimal contract typically requires uneven compensations to the experts. The intuition is that a given total effort generates the least amount of information when divided equally between the experts.[10] Thus, discrimination between experts naturally follows from optimality considerations and need not be the result of prejudice or bias. A similar point in a very different context is made in Winter (2004).

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. See Section 6 for details, as well as for additional properties of optimal contracts.

An obvious shortcoming of the optimal contracts we derive is that they do not implement the desired equilibrium uniquely; after all, the experts can all report the same signal without collecting any information.[11] However, for the experts to be able to match their reports they need to coordinate on one of the signals, and if they cannot communicate with one another then this may be hard to achieve. This is especially true if none of the two signals is ex ante focal. Collecting information creates a focal point (truthfully reporting one's signal) even when a priori 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 (1960) suggests, then we should expect agents to be willing to incur a cost in order to create a focal point. In Section 7, we further discuss the multiplicity problem and show that the DM can utilize indirect mechanisms to reduce its severity.

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

[9]The classic result of Radner and Stiglitz (1984) expresses a different kind of nonconcavity in the value of information; see also Chade and Schlee (2002).

[10]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.

[11]The work of Prendergast (1993) emphasizes the problem of experts second-guessing each other in a similar setup to ours. However, in his model each expert observes a noisy signal of the information held by other experts in addition to his own private signal about the state.

2. 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 solve for the optimal contract when a DM uses peer monitoring of information providers: In Bohren and Kravitz (2016), 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 (2007) 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.[12]

While the method of incentivizing effort is similar to ours, 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. Our richer environment uncovers properties such as the nonmonotonicity 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 properties of optimal contracts that are satisfied uniformly across all problems.

The first paper to explicitly suggest that effort and truthful reporting may be induced by comparing the reports of multiple agents is Miller et al. (2005), whose mechanism elicits honest feedback from raters who experience a certain product.[13] Their mechanism assigns to each rater a "reference rater" and applies a proper scoring rule to the pair of reports. This paper has been extended by many authors, including Witkowski and Parkes (2012) who relax the common prior assumption; Dasgupta and Ghosh (2013) who showed how to eliminate bad equilibria when workers engage in multiple tasks; and Jurca and Faltings (2005, 2009) who studied the implications of potential collusion among raters. See also Friedman et al. (2007) for a handbook explanation of the Miller et al. mechanism. This literature does not explicitly model the value of information, nor the cost of obtaining it, and is therefore silent about contract optimality.

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

[13] A similar mechanism was suggested at about the same time by Prelec (2004) in his "truth serum" paper. This paper's main concern is inducing truthful reporting, not incentivizing effort.

In Pesendorfer and Wolinsky (2003), a consumer sequentially sample experts who need to exert costly and unobservable effort in order to determine the appropriate service. In equilibrium, the consumer sometimes searches until the recommendations of two experts coincide, and then buys the service from one of them. This way the experts are incentivized to exert effort. The focus of this paper is on the effect that the moral hazard problem has on price competition in the market for experts' advice.

Rahman (2012) emphasized the ability of a principal to monitor workers by secretly recommending actions and base compensations on reported signals as well as on these recommendations.[14] 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 alternative solution for the principal that we study here 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.[15]

It has been known since the early mechanism design literature that correlation in agents' types facilitates the extraction of private information. A classic reference is the work of Crémer and McLean (1988) on full extraction of surplus in auctions. Note that, in contrast to Crémer and McLean (1988), we assume limited liability, and hence that the experts typically obtain positive expected utility in the optimal contract. In addition, in our model the experts *choose* the information structure rather than it being exogenously given. A recent work of Bikhchandani and Obara (2017) extends the results of Crémer and McLean to environments where agents can acquire information. In Azrieli (2019), we build on the ideas of this literature, and in particular on the work of Rahman (2011), to characterize the vectors of information structures that the DM can implement in general environments that include the binary framework of the current paper as a special case.

Finally, 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 (2001), Battaglini (2002), and Gentzkow and Kamenica (2017). 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 (2003), Martinelli (2006), Gerardi and Yariv (2008), and Gershkov and Szentes (2009). In these papers, incentives to collect information are provided through the voting rule and not through transfers.

[14] See also Strausz (2012) on the connection between Rahman's paper and the classic mechanism design framework of Myerson (1982).

[15] In our formal model, we do not consider contracts that condition payments on secret recommendations to the experts. Allowing for such contracts significantly expands the possibilities for the DM: An expert can be monitored by asking another expert to exert effort only with small probability, with payments made only when monitoring takes place. This requires of course that payments become arbitrarily large as the probability of monitoring goes to zero, which makes such contracts less practical and extremely sensitive to experts' risk attitudes. As we show, experts can effectively monitor each other even when effort recommendations are public.

3. The contracting environment

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 $e_i \in [0, \frac{1}{2}]$. The cost of effort is described by the function $c : [0, \frac{1}{2}] \to \mathbb{R}$. We assume that c is strictly increasing, strictly convex, twice 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 following stochastic matrix, where each row corresponds to the distribution over signals conditional on a state:

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

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 cannot 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 nonnegative, 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 $e_i \in [0, \frac{1}{2}]$ 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

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

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

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

It will be convenient to introduce the following notation. For every subset of experts $A \subseteq N$ and every effort 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 $N^b_s = \{j : s_j = b\}$ and $N^w_s = N \setminus N^b_s = \{j : s_j = w\}$. Finally, let $r^* = (r^*_1, \dots, r^*_n)$ denote the vector of truthful reporting strategies. Using this notation, we have that

$$\mathbb{P}_{(e,r^*)}(s) = \gamma e(N^b_s)\bar{e}(N^w_s) + (1-\gamma)e(N^w_s)\bar{e}(N^b_s).$$

Say that a contract x implements the vector of efforts $e = (e_1, \dots, e_n)$ if (e, r^*) is an equilibrium of the game induced by x with payoff functions as in (1); using Myerson's (1982) 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.[16] Effort vector e is implementable if there exists a contract x that implements it.

4. Cost of information

4.1 *The cost-minimizing contract*

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

$$\psi_i(e) = \min_{x_i} \mathbb{E}_{(e,r^*)}\big[x_i(s)\big] \tag{COST}$$
$$\text{s.t. } (e_i, r^*_i) \in \arg\max_{(e'_i, r'_i)} \big\{U_i\big((e'_i, r'_i), (e_{-i}, r^*_{-i}); x_i\big)\big\} \text{ and } x_i(s) \geq 0 \ \forall s \in S.$$

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 $N_{-ij} = N \setminus \{i, j\}$.

[16]We sometimes refer to x as being an incentive compatible contract in this case.

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

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

Furthermore, the cost function for the DM is given by

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

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

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:*[17]

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

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

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

The constraints (2)–(4), together with the objective $\mathbb{E}_{(e,r^*)}[x_i(s)]$ and the nonnegativity constraints define a linear program that, given Lemma 1, is equivalent to

[17]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}.

(COST). We refer to this auxiliary program as (AUX):

$$\min_{x_i} \mathbb{E}_{(e,r^*)}\big[x_i(s)\big] \tag{AUX}$$
$$\text{subject to (2)–(4) and } x_i(s) \geq 0\ \forall s \in S.$$

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 nonnegative. Second, plugging x_i^* to the constraints (2)–(4) gives

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

and

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

respectively. We leave it for the interested reader to verify that these all indeed are satisfied.[18] 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:

$$\max_{z_1,z_2,z_3} \big\{c'(e_i)z_1 + c(e_i)(z_2 + z_3)\big\} \tag{DUAL}$$
$$\text{s.t. } z_2, z_3 \geq 0, \text{ and for every } s_{-i} \in S_{-i}$$
$$\big[\gamma e\big(s_{-i}^b\big)\bar{e}\big(s_{-i}^w\big) - (1-\gamma)e\big(s_{-i}^w\big)\bar{e}\big(s_{-i}^b\big)\big]z_1 - \mathbb{P}_{(e,r^*)}(w, s_{-i})z_2 + \mathbb{P}_{(e,r^*)}(b, s_{-i})z_3$$
$$\leq \mathbb{P}_{(e,r^*)}(b, s_{-i}) \tag{5}$$

and

$$\big[\gamma e\big(s_{-i}^b\big)\bar{e}\big(s_{-i}^w\big) - (1-\gamma)e\big(s_{-i}^w\big)\bar{e}\big(s_{-i}^b\big)\big]z_1 - \mathbb{P}_{(e,r^*)}(w, s_{-i})z_2 + \mathbb{P}_{(e,r^*)}(b, s_{-i})z_3$$
$$\geq -\mathbb{P}_{(e,r^*)}(w, s_{-i}). \tag{6}$$

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 $\underline{b}_{-i}$.

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

LEMMA 2. *Let*

$$z_1^* = \frac{\big(\gamma\bar{e}(N) + (1-\gamma)e(N)\big)\big(\gamma e(N_{-i}) + (1-\gamma)\bar{e}(N_{-i})\big)}{\gamma(1-\gamma)\big[e(N_{-i})^2 - \bar{e}(N_{-i})^2\big]},$$

$$z_2^* = \frac{(2\gamma-1)e(N_{-i})\bar{e}(N_{-i})}{\gamma(1-\gamma)\big[e(N_{-i})^2 - \bar{e}(N_{-i})^2\big]} \quad \textit{and} \quad z_3^* = 0.$$

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

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

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. □

REMARK 1. 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.

REMARK 2. If the cost function c satisfies $c'(0.5) < +\infty$ then the theorem remains true for any vector e with $0 < e_i \le 0.5$, 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.

Let $\psi(e) = \sum_{i=1}^n \psi_i(e)$. Thus, ψ is the DM's cost function, describing the total expected cost of the least costly contract that implements e.

EXAMPLE 1. Suppose that the DM hires $n = 2$ experts, that the prior is uniform $\gamma = 0.5$, and that the cost of effort is given by $c(e_i) = \frac{e_i^2}{2}$. Theorem 1 implies that the cost function for the DM is given by

$$\psi(e_1, e_2) = \psi_1(e_1, e_2) + \psi_2(e_1, e_2) = \left(\frac{1}{2} + 2e_1e_2\right)\left(\frac{e_1}{2e_2} + \frac{e_2}{2e_1}\right). \qquad \Diamond$$

EXAMPLE 2. Consider the same setup as in the previous example but with cost of effort given by $c(e_i) = \frac{e_i^2}{2} + e_i$. The cost function for the DM is then

$$\psi(e_1, e_2) = \psi_1(e_1, e_2) + \psi_2(e_1, e_2) = \left(\frac{1}{2} + 2e_1e_2\right)\left(\frac{e_1+1}{2e_2} + \frac{e_2+1}{2e_1}\right). \qquad \Diamond$$

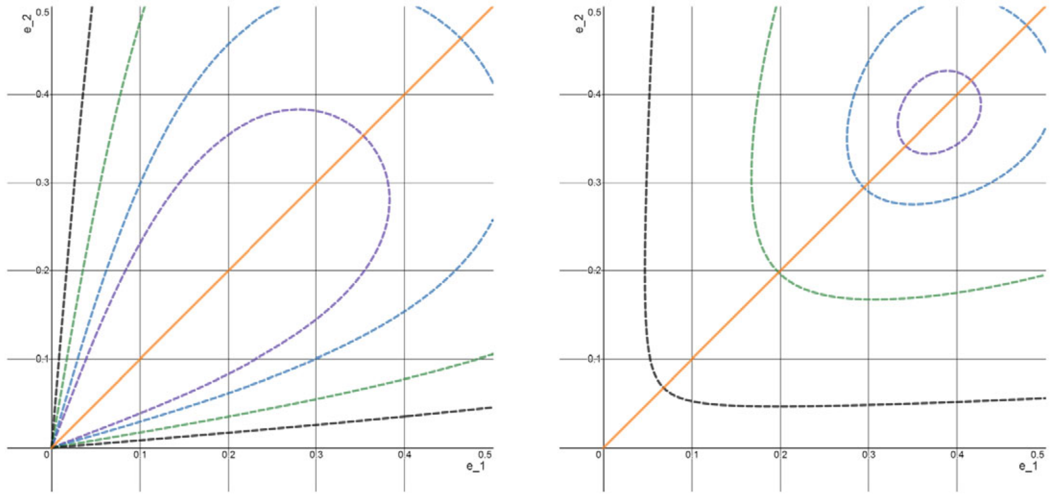

FIGURE 1. The left and right panels show several isocost curves (dashed lines) of ψ when the cost of effort is $c(e_i) = \frac{e_i^2}{2}$ and $c(e_i) = \frac{e_i^2}{2} + e_i$, respectively. Curves closer to the axes correspond to higher costs. The solid line is the main diagonal $e_1 = e_2$.

Figure 1 shows four of the isocost curves of ψ in Examples 1 (left panel) and 2 (right panel). Notice that the isocosts are sometimes upward sloping and are ordered in an unusual way, reflecting nonmonotonicity of ψ. The reason underlying this nonmonotonicity is discussed in Section 6.2 below.

4.2 *Variations and extensions*

Theorem 1 derives the cost-minimizing contract in our basic setup. We now discuss whether and how this result would change for various changes in the underlying contracting environment.

4.2.1 *Observable signals* Consider a variant of our model in which the DM directly observes the experts' signals, or, alternatively, that experts' reports are freely verifiable. Without adverse selection, incentive compatibility is characterized by the first-order condition (2) alone. The set of feasible contracts is therefore the intersection of the hyperplane defined by (2) with the nonnegative orthant, similar to the frontier of the budget set of a consumer.[19] Since the objective is linear, there is always an optimal contract at one of the extreme points of the feasible set, i.e., an optimal contract in which the expert is paid at only one realized signal vector s. This s is determined by minimizing the ratio of the coefficients of $x_i(s)$ in the objective and in the constraint, just like when

[19]Note that for some vectors s the coefficient of $x_i(s)$ in (2) is negative, so the feasible set is unbounded; but for such s an optimal contract would have $x_i(s) = 0$, so we may ignore those for the current discussion.

solving for an optimal bundle when the goods are perfect substitutes for the consumer (though here we look for a minimum rather than a maximum).[20]

When $\gamma > 0.5$, it is not hard to check that the *unique* solution is at $s = \underline{b}$, that is, experts are paid only if they all obtain the ex ante more likely signal. Specifically, the 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 $s \in S$. The expected payment to expert i comes out to be

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

Notice that this contract is *not* feasible in the original problem with unobservable signals—experts are only paid when they report b so exerting no effort and reporting b is a dominant strategy. The expert therefore has additional rents due to the exclusive access he has to the signal. More generally, no contract that pays at just one s can satisfy both (3) and (4). This implies that none of the extreme points of the feasible set of contracts with observable signals remains feasible with unobservable signals.

When $\gamma = 0.5$, both "corners" corresponding to $s = \underline{b}$ and $s = \underline{w}$ are optimal, and, therefore, so is any of their convex combinations. In particular, the optimal contract derived in Theorem 1 is optimal in the program without the adverse selection constraints (3) and (4).[21] Thus, the expected payment to i is the same with observable or unobservable signals and equals

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

Finally, it is important to point out that with observable signals it is possible to provide incentives even without hiring multiple experts, so long as the prior is not uniform. Indeed, an expert exerting effort e_i gets the b signal with probability $\gamma(0.5 + e_i) + (1 - \gamma)(0.5 - e_i)$. If $\gamma > 0.5$, then this is an increasing function of e_i, so by conditioning the payment on the realized signal the DM can induce any level of effort. With a uniform prior the distribution of signals is independent of effort so hiring additional experts is necessary.

Chade and Kovrijnykh (2016, Section 3) analyzed essentially the same contracting problem as ours (with observable signals), but with a single expert and a particular decision problem.[22] They apply the first-order approach to find the optimal contract. As already mentioned, when signals are not observable the first-order condition on effort is not sufficient for incentive compatibility.

[20]It is well known that in the moral hazard problem with a risk-neutral agent protected by limited liability the least costly way to implement a given action involves a binary contract with just two payoff levels. See, e.g., Demougin and Fluet (1998). This continues to hold in our setup with multiple agents.

[21]We note that with a uniform prior program (COST) has multiple solutions.

[22]The decision problem they consider is essentially the one in Example 4 below.

4.2.2 *Contractible state* Suppose that, in addition to the vector of reports s, the payment to an expert can also depend on the realized state of nature. Formally, a contract with expert i is a mapping $x_i : \{B, W\} \times S \to \mathbb{R}_+$. Since i's collected information concerns the state, and since the reports of other experts s_{-i} are noisy estimates of the state, it can be expected that the DM has nothing to gain by conditioning the payment to i on s_{-i}. In other words, the state is "sufficient statistic" for incentivizing i. The following corollary confirms this intuition.

Corollary 1. *Fix $e = (e_1, \dots, e_n)$ and consider the variant of program (COST) in which x_i may depend on the state (as well as on s). Then there is an optimal solution in which x_i is constant in s_{-i}, i.e., the payment to i depends only on the state and on his own report. The expected payment under the optimal contract is* $(0.5 + e_i)c'(e_i)$.

Proof. Fix $e = (e_1, \dots, e_n)$ and an expert i. Note that the state being contractible is equivalent to having another expert $n + 1$ with $e_{n+1} = 0.5$ (assuming this fictitious expert is honest and obedient). Therefore, we can find the cost of an optimal contract by plugging in $(e_1, \dots, e_n, 0.5)$ to ψ_i of Theorem 1. Since $\bar{e}(N) = \bar{e}(N_{-i}) = 0$ in this case, $\psi_i(e_1, \dots, e_n, 0.5)$ boils down to $(0.5 + e_i)c'(e_i)$. However, the same cost can be achieved by a contract that depends only on s_i and on the state. Indeed, $\psi_i(e_i, 0.5) = (0.5 + e_i)c'(e_i)$ as well. □

It is interesting to note that the optimal contract x_i^* given in Theorem 1 depends on the entire vector of reports, even when $e_j = 0.5$ for some $j \neq i$. It follows that the optimal contract is not unique in such cases (whenever $n > 2$).

There are quite a few papers analyzing the contracting problem when the ex post realized state is contractible, e.g., Osband (1989), Zermeño (2011), Rappoport and Somma (2015), Carroll (2017), Clark (2017), and Häfner and Taylor (2018), among others. Since there is no need to hire multiple experts to generate incentives, these papers all consider a single expert.

4.2.3 *Risk averse experts* It is instructive to compare our setup with risk-neutral experts to the more classical case in which the experts are risk averse. Suppose that expert i's payoff given contract x_i and strategy profile (e, r) is given by $\mathbb{E}_{(e,r)}[u(x_i(s))] - c(e_i)$, where $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly increasing and concave. To facilitate the comparison to the basic model, we maintain the limited liability constraint $x_i \geq 0$.

For a given vector e, incentive compatibility is still characterized by the constraints (2)–(4), but with $u \circ x_i$ replacing x_i everywhere. Geometrically, this implies that the set of incentive compatible contracts is no longer polyhedral as in the risk-neutral case; it is not even convex.

We can however still show that a cost-minimizing contract x_i has the following property: There is no s_{-i} for which both $x_i(b, s_{-i})$ and $x_i(w, s_{-i})$ are positive.[23] In particular,

[23] Indeed, suppose that $x_i(b, s_{-i}) > 0$ and $x_i(w, s_{-i}) > 0$. Take $\epsilon, \delta > 0$ such that $x_i(b, s_{-i}) - \epsilon > 0$, $x_i(w, s_{-i}) - \delta > 0$, and $u(x_i(b, s_{-i})) - u(x_i(w, s_{-i})) = u(x_i(b, s_{-i}) - \epsilon) - u(x_i(w, s_{-i}) - \delta)$. Then it is easy to verify that this modified contract still implements e, and it is clearly cheaper for the DM.

with $n = 2$ experts the optimal contract still pays positive amounts only when the reports match, just as in the risk-neutral case. With more than two experts this is no longer the case, as for any report of i (b or w) there would typically be multiple s_{-i}'s at which the payment to i is positive. Therefore, roughly speaking, an optimal contract insures an expert against wrong signals obtained by other experts but not against his own bad luck.

4.2.4 *Multidimensional states* It is often the case that multiple experts are being consulted because each of them has expertise in a different aspect of the problem at hand. In our basic setup, all experts collect information about the same (one-dimensional) state, but we now sketch an extension in which the state is multidimensional and different experts are hired to collect information on different dimensions. As long as there is correlation between the dimensions the report of one expert can be used to monitor others just as in the one-dimensional case, though as we now show the cost for the DM would increase as the correlation decreases.[24]

Suppose that the state space is the product $\{B_1, W_1\} \times \{B_2, W_2\}$. The prior is given by the following matrix, where $\rho \in [\frac{1}{4}, \frac{1}{2}]$:

	B_2	W_2
B_1	ρ	$0.5-\rho$
W_1	$0.5-\rho$	ρ

Thus, if $\rho = \frac{1}{4}$ then the dimensions are independent, and if $\rho = \frac{1}{2}$ then there is perfect correlation and we are essentially back in the basic setup; higher ρ implies higher (positive) correlation.

The DM hires two experts to collect information, each on a separate dimension. The mapping from effort to information is just as in the original model, but for expert i the distribution over signals depends only on whether the state has B_i or W_i. All other ingredients are unchanged relative to the model of Section 3.

Suppose that the experts choose effort levels $(e_1, e_2) > 0$ and truthfully report their signals. Then a direct calculation shows that the distribution over pairs of reports is given by

$$\mathbb{P}_{(e,r^*)}(b, b) = \mathbb{P}_{(e,r^*)}(w, w) = 0.5 + (4\rho - 1)e_1 e_2,$$

and

$$\mathbb{P}_{(e,r^*)}(b, w) = \mathbb{P}_{(e,r^*)}(w, b) = 0.5 - (4\rho - 1)e_1 e_2.$$

From this, we can derive the analogue of program (COST) to find the cost-minimizing contract that implements (e_1, e_2) with, say, expert 1: The objective is

$$\big[0.5 + (4\rho - 1)e_1 e_2\big]\big[x_1(b, b) + x_1(w, w)\big] + \big[0.5 - (4\rho - 1)e_1 e_2\big]\big[x_1(b, w) + x_1(w, b)\big];$$

[24] In Dewatripont and Tirole (1999), two "advocates" collect information on two different dimensions of the state. Since the dimensions are independent, it is impossible to generate incentives by comparing messages. Instead, Dewatripont and Tirole assumed that information comes as "hard evidence" that can only be found if effort is exerted.

the first-order condition for effort (2) is now replaced by

$$e_2(4\rho-1)\big[x_1(b,b)+x_1(w,w)-x_1(b,w)-x_1(w,b)\big]=c'(e_1);$$

the misreporting constraints (3) and (4) are replaced by

$$\begin{aligned}&\big[0.5+(4\rho-1)e_1e_2\big]\big[x_1(b,b)+x_1(w,w)\big]\\ &\quad+\big[0.5-(4\rho-1)e_1e_2\big]\big[x_1(b,w)+x_1(w,b)\big]-c(e_1)\\ &\geq 0.5\big[x_1(b,b)+x_1(b,w)\big], 0.5\big[x_1(w,w)+x_1(w,b)\big].\end{aligned}$$

When $\rho=\frac{1}{4}$ (independent dimensions), there is clearly no feasible contract. For every $\rho>\frac{1}{4}$, the solution has $x_1(b,w)=x_1(w,b)=0$ as before (this is a consequence of the positive correlation). We can then show that $x_1(b,b)=x_1(w,w)=\frac{c'(e_1)}{2e_2(4\rho-1)}$ is optimal, and that the expected cost for the DM is $c'(e_1)[\frac{1}{2e_2(4\rho-1)}+e_1]$. Thus, the cost is decreasing in ρ and becomes arbitrarily large as $\rho\downarrow\frac{1}{4}$.

The above example is particularly simple since it involves only two dimensions/ experts and a prior with uniform marginals. We expect however that qualitatively the results would be similar with more dimensions and other priors. A more complete analysis is left for future work.

4.2.5 *(Non)common priors* The common prior assumption, while standard in the literature, may be problematic in applications if experts have prior knowledge of the issues before contracting takes place.[25] Consider a generalization of our setup that allows for heterogeneous priors, with γ_i being the ex ante probability assigned to state B by expert i, $i=1,\dots,n$. Note that, from the point of view of expert i, the distribution of the signal vector s depends only on γ_i and on the effort vector e, not on the priors of other experts $j\neq i$ (assuming everyone is honest and obedient). This implies that the set of incentive compatible contracts x_i is exactly the same as in the common prior case, with γ_i taking the role of γ.

On the other hand, disagreement between the priors of an expert and the DM does have consequences for the cost-minimizing contract. Denote by γ_0 the prior of the DM. The objective function of (COST) is now calculated using γ_0, while the feasible set is determined by γ_i. If γ_0 is sufficiently different from γ_i, then the optimal contract may be located at a different extreme point of the feasible set than in the case where $\gamma_0=\gamma_i$. For a concrete example, suppose that $n=2$, fix (e_1,e_2), and let γ_1 (the prior of expert 1) be only slightly above 0.5. If the DM has the same prior, $\gamma_0=\gamma_1$, then recall that constraint (3) binds at the optimal contract x_1^*, and it is not hard to check that $x_1^*(b,b)>x_1^*(w,w)$. If, on the other hand, γ_0 is close to 1 then the DM believes that reports (b,b) are much more likely than (w,w), so would find such a contract too costly. The optimal contract would now pay a lower sum at (b,b) and a higher sum at (w,w). Furthermore, the binding constraint would be (4) rather than (3).

[25]The effectiveness of a policy to reduce global warming mentioned in the Introduction is a good example.

4.2.6 *Heterogeneous experts* Our model assumes complete symmetry between the experts in terms of their cost of effort and their production function. This is unrealistic in many contexts. However, we now argue that allowing heterogeneity along these dimensions does not qualitatively change Theorem 1.

Consider first the case in which experts differ in their cost of effort, namely that expert's i cost of exerting effort e_i is $c_i(e_i)$. Since an expert does not directly care about the cost of effort of other experts, the optimal contract in this case would be the same as in Theorem 1 with c_i replacing c.

Second, suppose that when i exerts effort e_i he gets the correct signal with probability $\frac{1}{2} + g_i(e_i)$ and the wrong signal with the complementary probability $\frac{1}{2} - g_i(e_i)$ (the same probabilities in both states), where g_i is a strictly increasing function with $g_i(0) = 0$ and $g_i(0.5) = 0.5$. This generalization captures heterogeneity in the productivity of experts. Here, too, the result of Theorem 1 remains essentially unchanged, provided that the composite function $c \circ g_i^{-1}$ is convex. Indeed, by the change of variable $t_i = g_i(e_i)$ one gets back to the original setup with $c_i(t_i) := c(g_i^{-1}(t_i))$ being the cost of attaining accuracy t_i.

5. Value of information

Finding the cost of obtaining information, we now consider the value of information for the DM. A decision problem is a triplet (γ, A, u), where as before γ is the prior over the states $\{B, W\}$, A is the set of actions available to the DM, and $u : \{B, W\} \times A \to \mathbb{R}$ describes the utility of the DM for each state-action pair. Let $\Delta(\{B, W\})$ be the set of probability distributions over the states, and identify each distribution in this set with the probability $q \in [0, 1]$ that the state is B. Any decision problem has an induced value function $v : \Delta(\{B, W\}) \to \mathbb{R}$ given by

$$v(q) = \max_{a \in A} \{qu(B, a) + (1 - q)u(W, a)\},$$

i.e., $v(q)$ is the maximal achievable expected utility of the DM given that her belief is q.[26] The function v is the pointwise 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 (2008)). It will be convenient to work directly with v rather than explicitly modeling decision problems. Let $\mathcal{V}$ be the set of all convex and continuous functions $v : [0, 1] \to \mathbb{R}$.

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 $m(e)$, then the value of information in decision problem $v \in \mathcal{V}$ is[27]

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

[26] Existence of a maximum is guaranteed whenever A is compact and u is continuous on A.

[27] The distribution over posteriors as well as the value of information obviously depend on the prior γ. For expositional reasons, we omit γ from the notation whenever no confusion may arise.

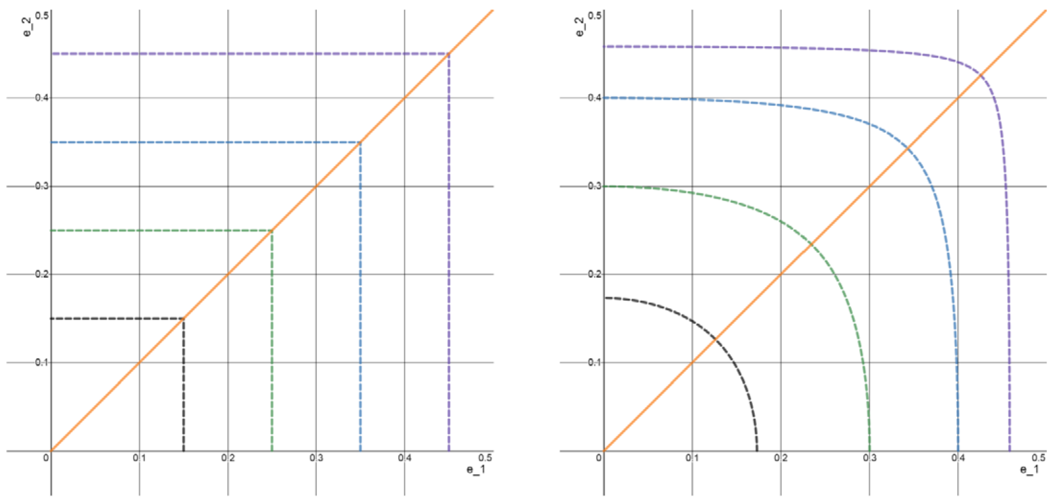

Figure 2. The left panel shows several indifference curves for the DM when the decision problem is as in Examples 3 and 4. The right panel shows indifference curves for Example 5. Indifference curves further away from the origin correspond to higher levels of expected utility.

Example 3. Suppose that the set of available alternatives is $A = \{B, W\}$ (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 $v(q) = \max\{q, 1 - q\}$. Suppose that the prior is uniform $\gamma = 0.5$ and that the DM hires two experts $n = 2$. Then the optimal alternative for the DM is independent of the signal reported by the expert exerting the lower of the two efforts. It is therefore easy to check that the value of information is given by $V_v(e_1, e_2) = 0.5 + \max\{e_1, e_2\}$. ◊

Example 4. 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 $0 \le q \le 0.5$ and $v(q) = 2q - 1$ for $0.5 < q \le 1$. With a uniform prior and two experts, the value of information is given by $V_v(e_1, e_2) = \max\{e_1, e_2\}$. ◊

Example 5. Let the set of alternatives be the unit interval $A = [0, 1]$, 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 $v(q) = -q(1 - q)$. If $\gamma = 0.5$ and $n = 2$, then a tedious but straightforward calculation gives $V_v(e_1, e_2) = -\frac{(1-4e_1^2)(1-4e_2^2)}{4(1+4e_1e_2)(1-4e_1e_2)}$. ◊

Figure 2 illustrates several indifference curves for the DM in each of the Examples 3, 4 (left panel), and 5 (right panel). Notice that in all these examples the upper contour sets are not convex. As we shall now see, this is no coincidence.

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

this will later allow us to draw general conclusions about which effort vectors may be optimal for the DM. As is well-known since Blackwell (1953), 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 effort vectors $e = (e_1 \geq e_2 \geq \cdots \geq e_n)$ and $e' = (e'_1 \geq e'_2 \geq \cdots \geq e'_m)$. Say that *e dominates e′* if $e_i \geq e'_i$ for every $i = 1, \dots, \max\{m, n\}$, and that *e weakly majorizes e′* if $\sum_{i=1}^k e_i \geq \sum_{i=1}^k e'_i$ for every $k = 1, \dots, \max\{m, n\}$, where in case $m \neq n$, the shorter of the two vectors is appended with zeroes.[28] 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 Girshick (1954, Theorem 12.3.1 on p. 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 v, V_v is nondecreasing in each expert's effort. The next Theorem 2, which may be of independent interest, strengthen this result by showing that if e weakly majorizes e' then $m(e)$ is more informative than $m(e')$. In other words, for every v, the value of information V_v is a Schur-convex function of effort vectors.[29] In particular, except for pathological examples, the value of information is not quasi-concave in efforts. We note that a different kind of nonconcavity in the value of information has been shown by Radner and Stiglitz (1984) (see also Chade and Schlee (2002)).

To gain some intuition, consider the case of two experts exerting equal efforts $e_1 = e_2 = 1/4$. This means that each of their signals has accuracy of $3/4$. Now, suppose that expert 1 increases his effort to $e_1 = 1/4 + \delta$, while expert 2 decreases his effort by the same amount to $e_2 = 1/4 - \delta$. When $\delta = 1/4$, the first expert learns the state for sure and expert 2's signal is completely uninformative, which is clearly better than the original effort vector in any decision problem. Theorem 2 implies that informativeness is not only higher at $\delta = 1/4$ than at $\delta = 0$, rather it is monotonically increasing throughout the interval $\delta \in [0, 1/4]$; moreover, higher spread of efforts continues to be more informative with any number of experts.

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

[28] 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.

[29] Roughly speaking, a function is Schur-convex if it is (1) symmetric and (2) convex in a restricted set of directions.

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. (2011, Proposition A.9 on p. 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 $z_1, \dots, z_L$ such that $z_1 = z$, $z_L = z'$, and z_l is obtained from z_{l-1} by a PD transfer for every $l = 2, \dots, L$. It is well known (see, e.g., Marshal et al. (2011, Proposition A.1 on p. 155)) that if z (exactly) majorizes z' then z' can be obtained from z by a sequence of PD transfers.[30]

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 $m(e'_1, e'_2)$.* □

6. 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 prior γ, the cost of effort function $c \in \mathcal{C}$, and the value function $v \in \mathcal{V}$. For any e, let $\pi_{c,v,\gamma}(e) = V_{v,\gamma}(e) - \psi_{c,\gamma}(e)$ be the net expected utility of the DM given effort vector e, where, by convention, $\pi_{c,v,\gamma}(e) = -\infty$ when e is not implementable. We sometimes omit the subscripts c, v, γ when no confusion may arise. Even though the arguments of π are effort 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.

Given the shapes of the cost and value of information functions (see Figures 1 and 2), a characterization of optimal contracts is too much to ask for. In particular, even when these functions are smooth, first-order conditions are generally not sufficient for optimality since by Theorem 2 the value function is not concave. Nevertheless, in what follows we prove several properties of optimal contracts that hold for large classes of environments.

6.1 *Asymmetry*

Given the symmetry in the cost of effort and the technology for collecting information across experts, one may expect that optimal contracts typically involve equal levels of

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

effort. Our first result is that quite the opposite is true—the optimal contract often involves discriminating between the experts. The intuition for this result comes directly from Theorem 2: Getting two signals of the same accuracy 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 typically lower than the cost of the latter, so we cannot immediately conclude that equal efforts are not optimal. Nevertheless, in the next proposition we prove that near any decision problem v there is another decision problem $\tilde{v}$ such that, for every cost function c, equal efforts are not optimal.

Proposition 1. *Fix $v \in \mathcal{V}$, a prior $\gamma \in [0.5, 1)$, 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 $c \in \mathcal{C}$ and every even $n \geq 2$, if $e = (e_1 \dots, e_n) > 0$ satisfies $e_1 = \cdots = e_n$ then there is $e' = (e'_1 \dots, e'_n) > 0$ such that $\pi_{c,\tilde{v},\gamma}(e) < \pi_{c,\tilde{v},\gamma}(e')$.

Proof. Given $v \in \mathcal{V}$, $\gamma \in [0.5, 1)$, and $\epsilon > 0$, let $v' \in \mathcal{V}$ be given by[31] $v'(q) = \epsilon \max\{0, q - \gamma\}$, and let $\tilde{v} = v + v'$. 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 $q \in [0, 1]$.

We first consider the case of two experts $n = 2$. Fix some $c \in \mathcal{C}$ and an effort vector e with $0 < e_1 = e_2 < 0.5$. Define $e(\delta) = (e_1 + \delta, e_2 - \delta)$, where $|\delta|$ is small enough so that these efforts remain between 0 and 0.5. Note that $e(0) = e$. We show that $\pi_{c,\tilde{v},\gamma}(e(\delta))$ has a strict local minimum at $\delta = 0$, from which part (ii) of the proposition follows.

First, from Theorem 2 we know that for all δ,

$$V_{v,\gamma}\big(e(\delta)\big) \geq V_{v,\gamma}(e). \tag{7}$$

Second, a direct calculation gives

$$V_{v',\gamma}\big(e(\delta)\big) = \begin{cases} 2\epsilon\gamma(1-\gamma)(e_1 + \delta) & \text{if } 0 \leq \delta \leq \bar{\delta}, \\ 2\epsilon\gamma(1-\gamma)(e_2 - \delta) & \text{if } -\bar{\delta} \leq \delta \leq 0. \end{cases} \tag{8}$$

Also, since $\psi_{c,\gamma}$ is symmetric and differentiable, the derivative $\frac{d\psi_{c,\gamma}(e(\delta))}{d\delta}$ is zero at $\delta = 0$. It follows that the right derivative at $\delta = 0$ of the difference

$$V_{v',\gamma}\big(e(\delta)\big) - \psi_{c,\gamma}\big(e(\delta)\big)$$

is $+2\epsilon\gamma(1-\gamma)$ and the left derivative of this difference at $\delta = 0$ is $-2\epsilon\gamma(1-\gamma)$. Thus, for all $\delta \neq 0$ sufficiently close to zero,

$$V_{v',\gamma}\big(e(\delta)\big) - \psi_{c,\gamma}\big(e(\delta)\big) > V_{v',\gamma}(e) - \psi_{c,\gamma}(e). \tag{9}$$

[31] Recall that v' is obtained from a decision problem with two alternatives similar to Example 4.

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

$$\begin{aligned}\pi_{c,\tilde{v},\gamma}\big(e(\delta)\big) &= V_{v,\gamma}\big(e(\delta)\big) + V_{v',\gamma}\big(e(\delta)\big) - \psi_{c,\gamma}\big(e(\delta)\big)\\ &\geq V_{v,\gamma}(e) + V_{v',\gamma}\big(e(\delta)\big) - \psi_{c,\gamma}\big(e(\delta)\big)\\ &> V_{v,\gamma}(e) + V_{v',\gamma}(e) - \psi_{c,\gamma}(e) = \pi_{c,\tilde{v},\gamma}(e),\end{aligned}$$

where the first equality is by linearity of expectation, the weak inequality is by (7), the strict inequality is by (9), and the last equality is again by linearity of expectation.

Consider now the case of even $n > 2$ and fix e with $0 < e_1 = \cdots = e_n < 0.5$. Similar to the $n = 2$ case, define $e(\delta) = (e_1 + \delta, e_2 - \delta, e_3, \dots, e_n)$. As before, we have that $V_{v,\gamma}(e(\delta)) \geq V_{v,\gamma}(e)$ and that $\frac{d\psi_{c,\gamma}(e(\delta))}{d\delta} = 0$ at $\delta = 0$, so to complete the proof it is enough to show that $V_{v',\gamma}(e(\delta))$ has a strictly positive right derivative and a strictly negative left derivative at $\delta = 0$.

We can calculate $V_{v',\gamma}(e(\delta))$ by conditioning on the realized number $k \in \{0, 1, \dots, n-2\}$ of b signals received by experts $3, \dots, n$. Let $p(k)$ be the probability of exactly k b signals among these experts, and let $q(k)$ be the posterior of state B conditional on this event. Note that $p(k)$ and $q(k)$ are independent of δ. We have

$$V_{v',\gamma}\big(e(\delta)\big) = \sum_{k=0}^{n-2} p(k) V_{v',q(k)}(e_1 + \delta, e_2 - \delta).$$

For every k the right derivative of $V_{v',q(k)}(e_1 + \delta, e_2 - \delta)$ at $\delta = 0$ is nonnegative by Theorem 2, and for $k = \frac{n-2}{2}$ we have $q(k) = \gamma$, so the right derivative is strictly positive by (8). For the same reason, the left-derivative is strictly negative at $\delta = 0$.

Finally, $e_1 = \cdots = e_n = 0.5$ is not optimal (when implementable) since the DM can learn the state at a lower cost by choosing $e_1 = 0.5$ and $e_i < 0.5$ for each $i = 2, \dots, n$ (this is easy to verify directly). □

6.2 *The cost and benefit of monitoring*

Our next observation is that a necessary condition for optimality can be obtained by studying the cost and benefit of improved monitoring. We start with the following definition.

Definition 1. Fix c, γ and n. Let $e_{-i} = (e_1, \dots, e_{i-1}, e_{i+1}, \dots, e_n)$ be an effort vector of all experts except i. We say that $e_i \in (0, 0.5)$ is a dominated effort level given e_{-i} if there is $e'_i > e_i$ such that $\psi_{c,\gamma}(e_i, e_{-i}) > \psi_{c,\gamma}(e'_i, e_{-i})$.

In words, e_i is dominated (given e_{-i}) if the DM can save costs by inducing i to exert more effort. The motivation for this definition stems from the obvious fact that if e_i is dominated given e_{-i} then the entire effort vector $e = (e_i, e_{-i})$ cannot be optimal in any decision problem. Indeed, since higher effort implies a more informative signal, the DM would have both higher value and lower cost at (e'_i, e_{-i}) than at e. Put differently, a necessary condition for optimality of e in *some* decision problem is that none of its coordinates e_i is dominated given e_{-i}.

How can an increase in the effort of one of the experts lead to cost savings for the DM? The answer is that an increase in e_i has opposite effects on the expected costs of the contracts with expert i and with the rest of the experts $j \neq i$. This is formalized in the following proposition.

Proposition 2. *For any interior e and two experts $i \neq j$, it holds that $\frac{\partial \psi_i(e)}{\partial e_i} > 0$ and $\frac{\partial \psi_j(e)}{\partial e_i} < 0$.*

It is not surprising that an increase in expert's i effort requires higher expected payment to that expert. But higher e_i also means that i's report is more correlated with the state and, therefore, that it allows better monitoring of the efforts and reports of other experts. For example, recall from Corollary 1 that in the extreme case $e_i = 0.5$ the report of i always matches the state, and there exist optimal contracts with other experts j that only rely on i's report for monitoring. More generally, the improved monitoring allows the DM to reduce the payments of all other experts.

The overall change in the cost for the DM when e_i increases depends on which of the above two effects is stronger. Namely, 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, so the overall marginal cost of inducing more effort may be either positive or negative. In the latter case, e is not optimal in any decision problem as explained above.

We now demonstrate that Definition 1 is not vacuous by describing two scenarios in which dominated effort levels do exist. In both cases, sufficiently low effort levels are dominated, implying that they are never part of an optimal contract.

Corollary 2. *Fix c, γ and n. If $c'(0) = c''(0) = 0$, then for every e_{-i} there is a number $f(e_{-i}) \in (0, 0.5)$ such that any $e_i < f(e_{-i})$ is dominated given e_{-i}.*

Proof. Fix e_{-i}. Since $c'(0) = c''(0) = 0$, it follows from the proof of Proposition 2 that

$$\lim_{e_i \downarrow 0} \frac{\partial \psi_i(e_i, e_{-i})}{\partial e_i} = 0.$$

In addition, it follows from the same proof that

$$\lim_{e_i \downarrow 0} \frac{\partial \psi_j(e_i, e_{-i})}{\partial e_i} < 0.$$

Therefore, the total cost $\psi_{c,\gamma}(e_i, e_{-i})$ is initially decreasing in e_i, so any small enough e_i is dominated given e_{-i}. □

If there are just two experts then the conclusion of Corollary 2 holds without any additional assumption on the cost function.

Corollary 3. *Fix c and γ, and suppose that $n = 2$. Then for every e_1 there is a number $f(e_1) \in (0, 0.5)$ such that any $e_2 < f(e_1)$ is dominated given e_1. Similarly, any $e_1 < f(e_2)$ is dominated given e_2.*

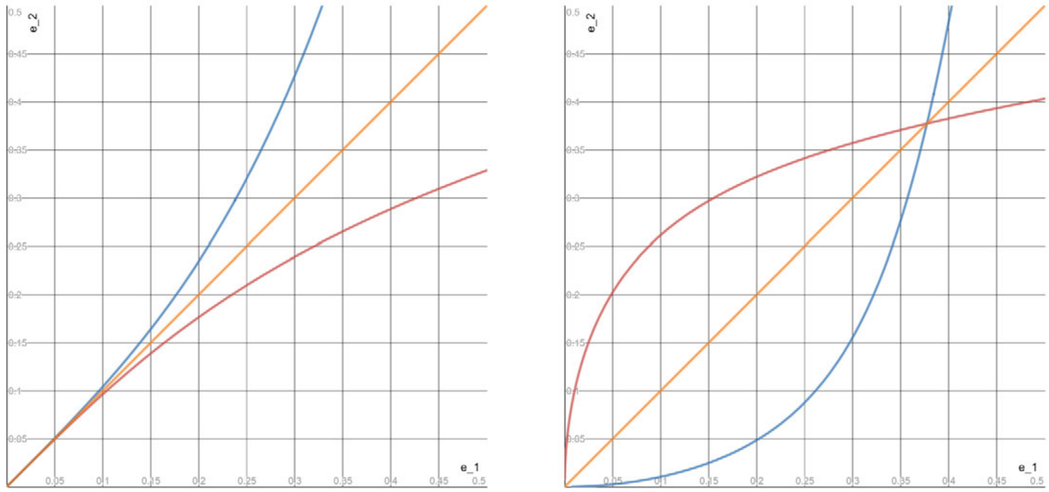

Figure 3. The left and right panels show the bounds f for the cost functions $c(e_i) = \frac{e_i^2}{2}$ and $c(e_i) = \frac{e_i^2}{2} + e_i$, respectively. In each panel, the concave curve is the bound $f(e_1)$ and the convex is $f(e_2)$. At any point (e_1, e_2) below the concave curve e_2 is dominated given e_1, and at any point (e_1, e_2) to the left of the convex curve e_1 is dominated given e_2. Thus, an optimal contract is never (for no decision problem) below the concave curve or to the left of the convex curve.

Proof. Fixing e_1, it follows from the expression for $\psi_i(e)$ in Theorem 1 that $\lim_{e_2 \downarrow 0} \psi_1(e_1, e_2) = +\infty$ and that $\lim_{e_2 \downarrow 0} \psi_2(e_1, e_2)$ is finite. Therefore, the total cost $\psi_{c,\gamma}(e_1, e_2)$ diverges to $+\infty$ as $e_2 \downarrow 0$,[32] implying that any small enough e_2 is dominated given e_1. The same argument applies when e_2 is fixed and $e_1 \downarrow 0$. □

To illustrate the bounds implied by Corollary 3, we revisit Examples 1 and 2 of Section 4. For a fixed e_1, the cost function $\psi_{c,\gamma}(e_1, e_2)$ in these examples is convex in e_2, first decreasing and then increasing. Denoting by $f(e_1)$ the effort level e_2 at which $\psi_{c,\gamma}(e_1, e_2)$ is minimal, we have that any $e_2 < f(e_1)$ is dominated given e_1. By symmetry, the same f also applies for e_1 when e_2 is held fixed. Figure 3 shows the resulting bounds for the two examples.

6.3 *Quality over quantity*

Our last observation is that under certain conditions optimal contracts never involve many low-effort experts. The reason is that such contracts are less informative and more costly than contracts with few experts each of which exerting high effort. To formalize this, for each $\bar{t} \in (0, 0.5)$ and $\bar{n} \geq 2$ denote by $e(\bar{t}, \bar{n})$ the constant effort vector with $\bar{n}$ experts each exerting effort $\bar{t}$.

Proposition 3. *Fix γ and c with $c'(0) > 0$. Let $e = (e_1, \dots, e_n) > 0$. If $\bar{t}, \bar{n}$ satisfy:*

(i) $e_i \leq \bar{t}$ for all i;

[32] This is not surprising given that $(e_1, 0)$ is not implementable. See Remark 1.

(ii) $\sum_{i=1}^{n} e_i \leq \bar{t}\bar{n}$*; and*

(iii) $n \geq \frac{2\psi_{c,\gamma}(e(\bar{t},\bar{n}))}{c'(0)}$,

then $e(\bar{t}, \bar{n})$ *is more informative and less costly than* e*. In particular, if such* $\bar{t}, \bar{n}$ *exist then* e *is not optimal for any decision problem* $v \in \mathcal{V}$*.*

Proof. First, for every $e = (e_1, \ldots, e_n) > 0$ and every expert i we have that $\psi_i(e) \geq (0.5 + e_i)c'(e_i)$, since this is the expected payment to i when the state itself is contractible (recall the discussion in the previous subsection). It follows that the total cost of implementing e satisfies

$$\psi_{c,\gamma}(e) = \sum_{i=1}^{n} \psi_i(e) \geq \frac{1}{2} \sum_{i=1}^{n} c'(e_i) \geq \frac{n}{2} c'(0),$$

where the last inequality is by convexity of c. Therefore, if $n \geq \frac{2\psi_c(e(\bar{t},\bar{n}))}{c'(0)}$ (condition (iii) of the proposition) then $\psi_{c,\gamma}(e) > \psi_{c,\gamma}(e(\bar{t}, \bar{n}))$.

Second, if $\bar{t}, \bar{n}$ satisfy conditions (i) and (ii) of the proposition then e is weakly majorized by $e(\bar{t}, \bar{n})$, so $e(\bar{t}, \bar{n})$ is more informative than e by Theorem 2. $\square$

To illustrate how Proposition 3 may be applied, consider the case where $\gamma = 0.5$ and $c(e_i) = \frac{e_i^2}{2} + e_i$. For any $\bar{t}, \bar{n}$ we have that

$$\psi_{c,\gamma}\big(e(\bar{t}, \bar{n})\big) = \bar{n} \frac{(0.5 + \bar{t})^{\bar{n}} + (0.5 - \bar{t})^{\bar{n}}}{(0.5 + \bar{t})^{\bar{n}-1} - (0.5 - \bar{t})^{\bar{n}-1}} (\bar{t} + 1).$$

Plugging in $\bar{t} = \frac{1}{3}$ and $\bar{n} = 2$, we get $\psi_{c,\gamma}(\frac{1}{3}, \frac{1}{3}) = \frac{26}{9}$. Therefore, any $e = (e_1, \ldots, e_n)$ with (i) $e_i \leq \frac{1}{3}$ for all i, (ii) $\sum_{i=1}^{n} e_i \leq \frac{2}{3}$, and (iii) $n \geq \frac{2 * \frac{26}{9}}{c'(0)} = 5.77$ is dominated by $(\frac{1}{3}, \frac{1}{3})$ and, therefore, never optimal.

In the case where fully learning the state is implementable, we can obtain a uniform bound on the number of experts in an optimal contract.

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

Proof. Let $\bar{t} = 0.5$ and $\bar{n} = 2$. The cost of $e(\bar{t}, \bar{n})$ is given by $\bar{n}c'(\bar{t}) = 2c'(0.5)$. Thus, by the first part of the proof of the last proposition, if e has n experts and $n \geq \frac{2\psi_{c,\gamma}(e(0.5,2))}{c'(0)} = \frac{4c'(0.5)}{c'(0)}$ then $\psi_{c,\gamma}(e) > \psi_{c,\gamma}(e(\bar{t}, \bar{n}))$. Also, $e(\bar{t}, \bar{n})$ is clearly at least as informative as any e. $\square$

Considering again the example where $\gamma = 0.5$ and $c(e_i) = \frac{e_i^2}{2} + e_i$, we have $c'(0) = 1$ and $c'(0.5) = 1.5$, so $\frac{4c'(0.5)}{c'(0)} = 6$. It follows that it is never (for no decision problem) optimal to hire more than 6 experts.

7. Collusion

7.1 *No-effort equilibria*

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, under the optimal contract x^* derived in Theorem 1 there would typically be a no-effort equilibrium which Pareto-dominates the intended honest-obedient equilibrium: If the experts could communicate with each other and coordinate on one of the announcements b or w then they would secure the highest possible payoff without exerting any effort.[33]

If the DM is worried about the possibility of such "collusion," then she may prefer another contract that, while increasing the cost for her, implements the desired vector of efforts e uniquely. Unfortunately, it is easy to see that for any contract there always exists a no-effort equilibrium. Indeed, fix some contract and suppose that all experts other than i choose a strategy with zero effort. Then the reported messages of these experts are independent of the state, so even if i exerts positive effort his message would be independent of the messages of the other experts. But this implies that i can achieve the same expected payment without incurring the cost of effort. Hence, i's best-response is to exert zero effort as well. Once we restrict attention to no-effort strategies the experts play a finite game, so Nash's theorem guarantees existence of equilibrium, possibly only in mixed strategies. Note that the same argument applies also for indirect (finite) mechanisms where the set of messages may be different than the set of signals.[34]

In the case of $n = 2$ experts, the problem with no-effort equilibria is more severe in the following sense: If $x = (x_1, x_2)$ is *any* contract that implements $e = (e_1, e_2) > 0$, then in the game induced by x both strategy profiles (b, b) and (w, w) (with zero efforts) are equilibria; moreover, for each expert at least one of these equilibria gives a higher payoff than under honesty and obedience (though there need not be one equilibrium that gives higher payoffs to both experts). However, we now illustrate via an example how indirect mechanisms can be used to make honesty and obedience the Pareto-dominant equilibrium.

Suppose that the prior is uniform $\gamma = 0.5$, that the cost of effort is $c(e_i) = 10e_i^2$, and that the DM would like to implement $e = (0.25, 0.25)$. It is straightforward to verify that a cost-minimizing contract that implements e is given by $x_1(\underline{b}) = x_2(\underline{b}) = x_1(\underline{w}) = x_2(\underline{w}) = 10$ and $x_i(s) = 0$ otherwise. Thus, without efforts the experts play the following matrix

---

[33] It may be that $x_i^*(\underline{b}) < x_i^*(\underline{w})$ for some experts, while $x_i^*(\underline{b}) > x_i^*(\underline{w})$ for the others, so different experts would prefer different zero-effort equilibria. But if the prior is uniform (or sufficiently close to uniform) then $\underline{b}$ gives the highest possible payoff to all experts.

[34] No-effort equilibrium exists even if the signal that the experts receive when they exert no effort is informative. Indeed, if all experts except i ignore their signal, then it is a best response for i to ignore his signal as well.

game:

	b	w
b	10, 10	0, 0
w	0, 0	10, 10

By coordinating their announcements, the experts obtain a payoff of 10, while the payoff under honesty and obedience is $10 * (\frac{3}{4} * \frac{3}{4} + \frac{1}{4} * \frac{1}{4}) - 10 * (\frac{1}{4})^2 = \frac{45}{8}$. Suppose that the DM adds another possible message, denoted ϕ, and modifies the payoffs as follows:

	b	w	ϕ
b	$10 + \lambda, 10 + \lambda$	λ, λ	$0, 0$
w	λ, λ	$10 + \lambda, 10 + \lambda$	$0, 10 + \delta$
ϕ	$10 + \delta, 0$	$0, 0$	$1, 1$

Here, λ and δ are small positive constants satisfying $\frac{5}{8}\delta < \lambda < \delta$. Note first that, since $\lambda < \delta$, the strategy profiles (b, b) and (w, w) are no longer equilibria in this modified game. In fact, the no-effort equilibrium with the highest payoff is the one where both experts choose b or w with equal probabilities, yielding an expected payoff of $5 + \lambda$ for each expert. Furthermore, this game still implements the desired effort levels of $e_1 = e_2 = 0.25$ in the sense that choosing these efforts and truthfully reporting the observed signal (i.e., choosing the action that corresponds to the observed signal) is an equilibrium.[35] The payoff to each expert under this desired equilibrium is $\frac{45}{8} + \lambda$, higher than in any no-effort equilibrium. Since λ can be made arbitrarily small, the cost to the DM is virtually the same as in the original contract. Thus, while it is impossible to completely eliminate no-effort equilibria in our environment, indirect mechanisms have the potential to make them less attractive.

In an environment with a contractible state, Khanna et al. (2015) proposed a mechanism that induces a group of experts to acquire information and truthfully reveal their signals. Their mechanism rewards the experts for being correct about the state, but also for *disagreeing* with other experts. They show that even if agents can communicate (using cheap-talk messages) before submitting their reports the mechanism implements the desired equilibrium uniquely. As explained above, this is not possible in our setup with noncontractible state.

7.2 *Multiple equilibria with effort*

Even if we restrict attention to equilibria in which the experts exert positive levels of effort and truthfully report their signals, it is still not necessarily the case that honesty and obedience is the only equilibrium. However, as the following proposition shows, at least in some environments we do get uniqueness.

[35] To see why, note first that when both experts choose b or w their payoffs are the same as in the original contract plus the constant λ, implying that, restricted to those strategies, incentives are the same as before. The condition $\frac{5}{8}\delta < \lambda$ guarantees that there is no profitable deviation to strategies involving the action ϕ whenever λ, δ are sufficiently small.

Proposition 4. *Suppose that the prior is uniform $\gamma = 0.5$ and 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 (1990) (see Theorem 4 in that paper). By Topkis (1998, Theorem 4.2.1), the set of equilibria is a lattice. It follows that if e' and e'' are both equilibria then their coordinatewise 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 cannot be two equilibria with positive efforts in which one dominates the other; since e is one equilibrium, the result follows.

Suppose by contradiction that $0 < e'_i \le e''_i$ 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

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

Therefore, for each expert i we have

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

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

$$\frac{e'(N_{-i}) - \bar{e}'(N_{-i})}{e''(N_{-i}) - \bar{e}''(N_{-i})} \ge \frac{e'_i}{e''_i}$$

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

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

However, it is not hard to check that this last inequality is inconsistent with $0 < e'_i \le e''_i$ 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

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

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

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 (1990) 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 underinvestment.

8. Final comments

In Theorem 1, we explicitly derive the cost function for the binary-binary model of the paper. As shown in Section 4.2, the key feature that experts get paid only when all the reports match continues to hold in several variants of the basic setup. Furthermore, one can show that this property holds even if the symmetry between the two states is relaxed, i.e., if the probability of observing the correct signal differs between states B and W. The key is that the vectors of signals s_{-i} at which the likelihood ratio $\frac{\mathbb{P}(s_{-i}|B)}{\mathbb{P}(s_{-i}|W)}$ is maximal and minimal are $s_{-i}=\underline{b}_{-i}$ and $s_{-i}=\underline{w}_{-i}$, respectively. It would be interesting to know whether this generalizes to environments with more states and signals.

The binary-binary framework also allows to derive the strong nonconcavity result of Theorem 2. This result is of independent interest and may be useful in other applications. There are (at least) two natural ways to extend Theorem 2 to more general information structures: First, 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'. This property does not extend to more general vectors of information structures.[36] Second, the theorem is equivalent to the claim that $m(e_1,e_2)$ is more informative than $m(\lambda e_1+(1-\lambda)e_2,(1-\lambda)e_1+\lambda e_2)$ for any $\lambda\in[0,1]$. This property also does not generalize beyond symmetric binary information structures.[37] While Theorem 2 seems to be specific to the type of information structures we consider, it suggests that deriving the demand for information is likely to pose serious challenges in many cases.

[36] I thank Eran Shmaya for suggesting a simple counter example.
[37] This potential generalization was suggested to me by Marcin Peski.

Appendix: 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 $U_i((e'_i, r^*_i), (e_{-i}, r^*_{-i}); x_i)$ 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

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

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

$$\begin{aligned}
&\mathbb{E}_{((0,r^*_i),(e_{-i},r^*_{-i}))}\big[x_i(s)\big] \\
&\quad = \frac{e_i}{e_i + e'_i}\mathbb{E}_{((e'_i,r_i),(e_{-i},r^*_{-i}))}\big[x_i(s)\big] + \frac{e'_i}{e_i + e'_i}\mathbb{E}_{(e,r^*)}\big[x_i(s)\big].
\end{aligned} \tag{10}$$

Therefore,

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

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 (10), and the last inequality is by non-negativity of the cost function. It follows that $U_i((e, r^*); x_i) \geq U_i((e'_i, r_i), (e_{-i}, r^*_{-i}); x_i)$, so (e'_i, r_i) is not a profitable deviation. □

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

$$\big[\gamma\mathbb{P}(\underline{b}_{-i}|B) - (1-\gamma)\mathbb{P}(\underline{b}_{-i}|W)\big]z_1^* - \mathbb{P}(w, \underline{b}_{-i})z_2^* \geq \mathbb{P}(\underline{b}),$$

or, more explicitly,

$$\begin{aligned}&\big[\gamma e(N_{-i}) - (1-\gamma)\bar{e}(N_{-i})\big]z_1^* - \big[\gamma(0.5-e_i)e(N_{-i}) + (1-\gamma)(0.5+e_i)\bar{e}(N_{-i})\big]z_2^* \\ &\quad \leq \gamma e(N) + (1-\gamma)\bar{e}(N).\end{aligned}$$

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

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

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

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

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. □

Proof of Proposition 2. First,

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

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 $\gamma(1-\gamma)[e(N_{-i})^2 - \bar{e}(N_{-i})^2]$ of ψ_i is increasing in e_j, and that the second term $(2\gamma - 1) \times e(N_{-i})\bar{e}(N_{-i})c(e_i)$ 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

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

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

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

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

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 $p_e(A) = \frac{1}{2}[e(A)\bar{e}(A^c) +$

$\bar{e}(A)e(A^c)]$ 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 (1954, Theorem 12.4.1 on page 332), $m(e)$ is more informative than $m(e')$ if and only if

$$\int_0^x F_e(t)\,dt \geq \int_0^x F_{e'}(t)\,dt \tag{11}$$

holds for every $x \in [0, 1]$. To complete the proof we now show that (11) 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 nondecreasing step-functions this would imply that (11) 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 (11) 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

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

or equivalently that

$$q_e(2)\big[p_e(\emptyset) - p_{e'}(\emptyset)\big] \geq q_e(\emptyset)p_e(\emptyset) - q_{e'}(\emptyset)p_{e'}(\emptyset). \tag{12}$$

Using the equality $e_1 + e_2 = e'_1 + e'_2$, simple algebra gives that the right-hand side of (12) 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 (12) becomes

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

Since the area of a rectangle with a given perimeter decreases in the difference between its length and its width, we have that $e_1e_2 - e'_1e'_2 \leq 0$, and by observation (a) we have that $q_e(2) \leq \frac{1}{2}$. This proves (12).

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 (11) 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 (11) 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

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

(b): The proof is identical to that of observation (a) (recall that $e'_1 \geq e'_2$).

(c): We have

$$q_e(\emptyset) = \frac{1}{1 + \dfrac{(0.5+e_1)(0.5+e_2)}{(0.5-e_1)(0.5-e_2)}} \quad \text{and} \quad q_{e'}(\emptyset) = \frac{1}{1 + \dfrac{\big(0.5+e'_1\big)\big(0.5+e'_2\big)}{\big(0.5-e'_1\big)\big(0.5-e'_2\big)}},$$

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,

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

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 $F_{e'}$. □

References

Azrieli, Y. (2019). "Delegated expertise: Implementability with peer-monitoring." Working paper. [1318]

Azrieli, Y. and E. Lehrer (2008), "The value of a stochastic information structure." *Games and Economic Behavior*, 63, 679–693. [1329]

Battaglini, M. (2002), "Multiple referrals and multidimensional cheap talk." *Econometrica*, 70, 1379–1401. [1318]

Bikhchandani, S. and I. Obara (2017), "Mechanism design with information acquisition." *Economic Theory*, 63, 783–812. [1318]

Blackwell, D. (1953), "Equivalent comparisons of experiments." *The Annals of Mathematical Statistics*, 24, 265–272. [1316, 1331]

Blackwell, D. and M. A. Girshick (1954), *Theory of Games and Statistical Decisions*. John Wiley & Sons, Inc. [1331, 1345]

Bohren, J. A. and T. Kravitz (2016), "Optimal contracting with costly state verification, with an application to crowdsourcing." Working paper. [1317]

Bolton, P. and M. Dewatripont (2005), *Contract Theory*. The MIT Press. [1315]

Carroll, G. (2017), "Robust incentives for information acquisition." Working paper. [1326]

Chade, H. and N. Kovrijnykh (2016), "Delegated information acquisition with moral hazard." *Journal of Economic Theory*, 162, 55–92. [1325]

Chade, H. and E. Schlee (2002), "Another look at the Radner–Stiglitz non-concavity in the value of information." *Journal of Economic Theory*, 107, 421–452. [1316, 1331]

Clark, A. (2017), "Contracts for information acquisition." Working paper. [1326]

Crémer, J. and R. McLean (1988), "Full extraction of the surplus in Bayesian and dominant strategy auctions." *Econometrica*, 56, 1247–1257. [1318]

Dasgupta, A. and A. Ghosh (2013), "Crowdsourced judgement elicitation with endogenous proficiency." In *Proceedings of the 22nd International Conference on World Wide Web*, 319–330. [1317]

Demougin, D. and C. Fluet (1998), "Mechanism sufficient statistic in the risk-neutral agency problem." *Journal of Institutional and Theoretical Economics*, 154, 622–639. [1325]

Dewatripont, M. and J. Tirole (1999) *Advocates, Journal of Political Economy*, 107, 1–39. [1327]

Friedman, E., P. Resnick, and R. Sami (2007), "Manipulation-resistant reputation systems." In *Algorithmic Game Theory* (N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, eds.). Cambridge University Press. [1317]

Gentzkow, M. and E. Kamenica (2017), "Competition in persuasion." *Review of Economic Studies*, 84, 300–322. [1318]

Gerardi, D. and L. Yariv (2008), "Information acquisition in committees." *Games and Economic Behavior*, 62, 436–459. [1318]

Gershkov, A. and B. Szentes (2009), "Optimal voting schemes with costly information acquisition." *Journal of Economic Theory*, 144, 36–68. [1318]

Gromb, D. and D. Martimort (2007), "Collusion and the organization of delegated expertise." *Journal of Economic Theory*, 137, 271–299. [1317]

Grossman, S. J. and O. D. Hart (1983), "An analysis of the principal-agent problem." *Econometrica*, 51, 7–45. [1315]

Häfner, S. and C. R. Taylor (2018), "Contracting for research: Moral hazard and the incentive to overstate significance." Working paper. [1326]

Holmstrom, B. (1982), "Moral hazard in teams." *The Bell Journal of Economics*, 13, 324–340. [1315]

Jurca, R. and B. Faltings (2005), *Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms*. International Workshop on Internet and Network Economics. Springer. [1317]

Jurca, R. and B. Faltings (2009), "Mechanisms for making crowds truthful." *Journal of Artificial Intelligence Research*, 34, 209–253. [1317]

Khanna, N., A. Peivandi, and M. Schroder (2015), "Rewarding disagreement for optimal decisions." Working paper. [1339]

Krishna, V. and J. Morgan (2001), "A model of expertise." *Quarterly Journal of Economics*, 116, 747–775. [1318]

Krupka, E. L. and R. A. Weber (2013), "Identifying social norms using coordination games: Why does dictator game sharing vary?" *Journal of the European Economic Association*, 11, 495–524. [1314]

Marshal, A. W., I. Olkin, and B. C. Arnold (2011), *Inequalities: Theory of Majorization and Its Applications*, second edition. Springer. [1332]

Martinelli, C. (2006), "Would rational voters acquire costly information?" *Journal of Economic Theory*, 129, 225–251. [1318]

Merrifield, M. R. and D. G. Saari (2009), "Telescope time without tears: A distributed approach to peer review." *Astronomy & Geophysics*, 50, 4.16–4.20. [1314]

Milgrom, P. and J. Roberts (1990), "Rationalizability, learning, and equilibrium in games with strategic complementarities." *Econometrica*, 58, 1255–1277. [1340, 1341]

Miller, N., P. Resnick, and R. Zeckhauser (2005), "Eliciting informative feedback: The peer-prediction method." *Management Science*, 51, 1359–1373. [1317]

Myerson, R. B. (1982), "Optimal coordination mechanisms in generalized principal-agent problems." *Journal of Mathematical Economics*, 10, 67–81. [1318, 1320]

Osband, K. (1989), "Optimal forecasting incentives." *Journal of Political Economy*, 97, 1091–1112. [1326]

Persico, N. (2003), "Committee design with endogenous information." *Review of Economic Studies*, 70, 1–27. [1318]

Pesendorfer, W. and A. Wolinsky (2003), "Second opinions and price competition: Inefficiency in the market for expert advice." *The Review of Economic Studies*, 70, 417–437. [1318]

Prelec, D. (2004), "A Bayesian truth serum for subjective data." *Science*, 306, 462–466. [1317]

Prendergast, C. (1993), "A theory of "Yes men"." *American Economic Review*, 83, 757–770. [1316]

Radner, R. and J. Stiglitz (1984), "A non-concavity in the value of information." In *Bayesian Models in Economic Theory* (M. Boyer and R. Kihlstrom, eds.). Elsevier Science Publishers. [1316, 1331]

Rahman, D. (2011), "Detecting profitable deviations." Working paper. [1318]

Rahman, D. (2012), "But who will monitor the monitor?" *American Economic Review*, 102, 2767–2797. [1318]

Rappoport, D. and V. Somma (2015), "Incentivizing information design." Working paper. [1326]

Schelling, T. C. (1960), *The Stratgey of Conflict*. Harvard University Press. [1316]

Strausz, R. (2012), "Mediated contracts and mechanism design." *Journal of Economic Theory*, 147, 1280–1290. [1318]

Topkis, D. M. (1998), *Supermodularity and Complementarity*. Princeton University Press. [1340]

Winter, E. (2004), "Incentives and discrimination." *American Economic Review*, 94, 764–773. [1316]

Witkowski, J. and D. C. Parkes (2012), "Peer prediction without a common prior." In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 964–981. [1317]

Xiao, E. and D. Houser (2005), "Emotion expression in human punishment behavior." *Proceedings of the National Academy of Sciences*, 102, 7398–7401. [1315]

Zermeño, L. (2011), "A principal-expert model and the value of menus." Working paper. [1326]

Co-editor Thomas Mariotti handled this manuscript.

Manuscript received 5 November, 2019; final version accepted 13 January, 2021; available online 22 January, 2021.