Contracting for Expertise

Fritz L. Laux*

Instituto Tecnológico Autónomo de México

School of Business and CIEN-F Working Paper

22 July 1999

ABSTRACT

This paper shows that the key considerations in contracting for expertise are the potential for competition between experts and the potential for strategic use of the decision-maker's own private information. Although the problem may look like a straightforward application of mechanism design with correlated information, this is not the case. This is because, although we would expect expert opinion regarding an optimal policy decision to be correlated, we would not necessarily expect the same for the private incentives of experts. Intuitive contractual arrangements that ensure full information revelation, and are also robust to collusion among agent-experts, are developed and discussed.

Keywords: Agency, correlated information, collusion, cheap talk

JEL classification nos.: D82, L22, D73

Author address: Fritz L. Laux
ITAM, Departamento Académico de Administración
Río Hondo No. 1
Col. Tizapán - San Angel
01000 México, D.F.
MÉXICO

Phone: 52·5·662·4000, Ext. 3416
E-mail: fritz@eniac.rhon.itam.mx

* I am grateful to Larry Ausubel, Roberto Burguet, Thomas Jeitschko, Jean-Jacques Laffont, César Martinelli, Arijit Sen, and Tridib Sharma for helpful advice and suggestions.

1. Introduction

Consider the problem of a manager who must allocate budget between research and development projects. Should she make these decisions herself, using only her own personal information and limited insight, or should she solicit advice from potentially self-interested subordinates who may have more intimate knowledge of the issues involved? She has some idea regarding the likely profitability of each project but she could generally make better decisions if she could also use the information of her subordinates. Her problem is that she does not know whether or not any guidance given to her will be honest. Furthermore, waiting until project outcomes are observed (if they are ever observed within the tenure of her relationship with her subordinates) before rewarding any guidance given may not be practical.

In this scenario, we can think of her subordinates as being the engineers who do or will work on the various long-term projects. These engineers may have a special interest or bias toward one project over another. Perhaps they favor one project not because they expect it will be more profitable but because it will employ more interesting technology. Perhaps one project looks likely to provide better training and experience for them. Alternatively, in a procurement context, we can think of these experts as being self-interested consultants or outside vendors. Although procurement contests are often decided on the basis of cost competition (and are thus well modeled in a standard auction framework), many focus more on the solicitation of design proposals. This is especially true for the initiation phases of new project activity. The company (principal) that would like to solicit design proposals from outside vendors (agents) must consider the possibility that these vendors may have private interests. A consulting firm may prefer to recommend solution approaches that, if adopted, would further its entry into new lines of business. A manufacturer may prefer to recommend product designs that fit with its broader offerings or business development strategy. Will consultants or vendors suggest solutions or designs that are in the best interest of the purchasing company?[1]

[1] This can also be viewed from a broader perspective. Even if the interests of the consulting firm or outside vendor are taken into account, will decision making maximize social welfare or will the inefficiencies

Agency models generally describe situations where a principal knows what she wants done but for some reason must delegate actual task performance. The objective of this paper is to present a model of how a principal, who does not know what she wants to do, can use incentive contracting to elicit guidance from expert agents. As discussed above, the key consideration in such a contracting problem is that agent-experts will presumably have private interests. To the extent that the principal might act on any guidance they provide, they will have incentives to influence the principal's decision toward their personal interests and possibly away from the interests of the principal.

A key feature of the modeling of this problem, as it is done in this paper, is that the information provided by the experts is assumed to be "soft" or unverifiable (it is not based on hard evidence). Thus no agent can ever be known, with certainty, to have lied. This is consistent with contexts where experts are asked to provide simple opinions. It is also consistent with contexts of complicated decisions where different evidence can legitimately be presented with different emphasis, to support different conclusions. To abstract away from the problem of moral hazard, or shirking by agents in their research and advice-giving role, it is also assumed that their expertise is costless to the agents and accrues to them simply as a result of their position in an organization. Furthermore, and importantly, the bias or direction of the agents' private incentives is unknown to the principal and uncorrelated with any observable work product of the agents. From the perspective of the principal, an expert agent is just as likely to have private preferences for Project A over Project B as vice versa. Lastly, the comparative profitability of the project selection choice, if ever known by the principal, is assumed not to be known within the timeframe of her relationship with the agents. She cannot simply link incentive payments to project outcomes.

The analysis shows that there are two keys to resolving this contracting problem. First, the principal should solicit advice from more than one expert. Second, the principal needs to make strategic use of her own private information. By soliciting advice from more than one expert, the principal can use a comparison of expert reports to improve incentives for truthful reporting. This is because, since expert opinion should, by definition, be correlated,

---

of agency prevent this?

contradictory advice is subject to doubt. Contracts that reward agreement between experts and punish disagreement can thus tend, in expectation, to reward truthfulness. This well-known logic was developed by Crémer and McLean (1985, 1988) and was subsequently generalized to mechanism design contexts by McAfee and Reny (1992). In contexts where a manager can commit to contract terms, this correlation provides a means for incentive contracting in this information provision game.

The limitation of this comparison approach, however, is that it invites the possibility of gaming between experts and only works in contexts where experts believe that their peers will report truthfully. Contracts that reward agreement encourage agents to seek out alternate equilibrium strategies and to use pre-play communication by which they can arrange agreement in their reports. If agreement in reported opinion can be reached through coordination, collusion, or degenerate reporting strategies, then a comparison system is defeated.

More technically, the failure of the Crémer and McLean logic of contracting for the revelation of correlated information in this context results from the possibility that agents may have some private information that is not correlated. Although we would expect expert information regarding the policy choice for a decision maker to be correlated (both across experts and with the interests of the principal), we would not necessarily expect the private interests of experts over this policy choice to be correlated in either of these senses. This element of uncorrelated information complicates the problem in a way that prevents the standard Crémer and McLean solution from working. A reliable incentive contract for this context must cope with strategies that can vary (be conditioned) not only over the expert opinion of the agents but also over their uncorrelated private interests. The solution to be demonstrated is for the principal to make use of her private information, and the correlation of this information with the expertise of the agents, to eliminate the potential noncommunicative alternative equilibria that result from these richer strategic possibilities. Importantly, this logic can also be extended to show how the principal can design incentive contracts that are robust to the possibility of pre-play communication and collusion among expert agents.[2]

[2] Laffont and Martimort (1998) argue that, indeed, it is the potential for collusion that limits the applicability of mechanisms or arrangements, such as this one, that rely on competition between agents. Clearly, the robustness of any incentive system that relies on playing one agent off against another depends on some

With reference to the theory literature, in addition to its connection to the Crémer and McLean results, this paper also has application to what is often called the "cheap talk" problem. The cheap talk literature describes the problem of soliciting advice from experts when the payoffs of these experts are not directly dependent on the content of the advice they give — expert payoffs are only indirectly affected by the influence their advice may have on a policy maker's decision.[3] In the model of this paper, a single expert, asked to give advice, would want to reveal his information honestly only if, in the expectation of that expert, the information revealed would tend to influence policy toward the direction of his private interests (the game degenerates to a cheap talk problem). The analysis shows, however, that given the decision maker is able to commit to incentive contracting, the key to resolving such a problem is to solicit advice from more than one expert. Furthermore, unlike a cheap talk game, this agency formulation permits the decision maker to ensure full information revelation.

The applied theory literature most closely related to this present work consists of two papers, Gibbons (1988), and Prendergast (1993). Gibbons models the problem of an arbitrator who would like to impose a settlement, between management and a union, that as closely and fairly as possible reflects the state of their employment relationship. Although the arbitrator has some *ex ante* idea of what would be a fair settlement, he could make better decisions if he had access to the more expert knowledge of the disputants. Gibbons shows that there do exist equilibria in conventional arbitration under which, by using his private information strategically, the arbitrator can obtain useful information from the disputants. Unlike the present paper, however, Gibbons' results depend on common knowledge of the preferences of each of the disputants (labor wants to maximize and management wants to minimize the wage settlement).[4] Also unlike this present paper, although Gibbons finds existence results for full information revelation, he does not find conditions that ensure full

ability to prevent coordinated play by these agents.

[3]A fundamental result of this cheap talk literature is that these games generally do not yield fully revealing equilibria (Crawford and Sobel, 1982). This literature is surveyed by Farrell and Rabin (1996).

[4]Furthermore, and because of this, the arbitrator does not have to commit to bonus terms over expert reports in order to induce full information revelation.

information revelation from the disputants.

In contrast to the Gibbons paper, which does not assume the arbitrator can *ex ante* commit to "contract" terms (punishment and reward allocations that depend directly on disputant reports), Prendergast (1993) describes a game formulated entirely within the context of incentive contracting. He models the problem of a principal, P, who delegates project evaluation activities to two staff members. These subordinates can either perform their research with costly effort or shirk. To discourage shirking, P would like to condition incentive contracts on project evaluation information of her own. The difficulty is, however, that P's information is not entirely private; it is partially observed by the agents. Thus, attempts by P to provide incentives for diligent research effort will lead to adverse incentives for the agents to act as "yes men." Unlike Prendergast (1993), the present paper describes a context where information is costless to the experts and incentives for bias are unknown to P. Rather than explore the incentives that agents might have to shirk in their research activities, this present paper focuses on the cheap-talk-related problem of the incentives agents may have to manipulate policy decisions. The concern in this paper is not that, because of shirking, agents will wish to falsely agree with the principal but rather that, because of private interests in policy, they will wish to falsely influence the principal.

Another body of literature that is related to this present paper is that associated with the seminal paper of Milgrom and Roberts (1986). These papers, describing scenarios that can be thought of as analogous to those of evidence production in a court of law, showed how competition between interested experts (advocates) can be used to induce full information revelation. Their scenario, however, requires that all reports be based on verifiable, hard evidence. The present paper shows how, through the use of incentive contracting, the problem can be solved for the case of soft, unverifiable information.[5]

The remainder of this paper includes a brief description of the model to be used for anal-

[5] The article from this hard-evidence literature perhaps most closely related to the present paper is Shin (1998). He models the trade-off between adversarial versus inquisitorial procedures in arbitration. In contrast to Shin, this present paper provides insight into the potential for combining adversarial and inquisitorial procedures. Furthermore, unlike Shin, the present paper does this in a context of soft evidence (where, additionally, the biases of disputants are unknown).

ysis (Section 2); the development of existence (Section 3) and sufficiency results (Section 4) for the full information revelation mechanism; a discussion of the extensions to these results necessary to allow for the possibility of pre-play communication and collusion between agents (Section 5); and finally a summary of implications and conclusions (Section 6).

2. Model

The game has three players, a receiver or principal (P) and two senders or agents (A_i, $i = 1, 2$). P must choose the design, $\widehat{x} \in \{0, 1\}$, for a new product ($\widehat{x}$ can also represent project selection, facility location, and so forth). The optimal product design for the principal, $x^* \in \{0, 1\}$, is selected by nature according to a Bernoulli distribution, $F(x; q)$.[6] No players know the value of x^* but all have common prior beliefs that the parameter q, equal to the expected value of x^*, has been drawn from the uniform distribution on $[0, 1]$. In addition to these common priors on x^*, each player privately observes independent draws from this same $F(x; q)$ distribution, denoted $\theta_0, \theta_i, \theta_j$ (where 0 indexes the principal's draw and $i, j \in \{1, 2\}$ index agent draws). Each agent can share his expertise (observed draw, θ_i) with P by means of a report, $r_i \in \{0, 1\}$, that can be either truthful, $r_i = \theta_i$, or false, $r_i \neq \theta_i$.

Each agent has a preference for one of the two outcomes, $\widehat{x} = 0$ or $\widehat{x} = 1$. I refer to this preference as the agent's type, $x_i \in \{0, 1\}$, where knowledge of his type is private for each agent. These types define each agent's private interest in P's policy decision and it is common knowledge that they are independently drawn so that the probability that each $x_i = 1$ is $\frac{1}{2}$. In order to motivate honest reporting by the agents, P offers them bonus contracts, w, symmetrical across agents. Since P is assumed not to be able to commit to honest revelation of her information, θ_0, these bonus contracts must be conditioned on agent reports, r_i and r_j, and the ultimate policy decision made by P, $\widehat{x}$.[7] Thus a bonus contract, $w(r_i, r_j, \widehat{x})$, denotes the payment to be made by P to A_i when she receives report r_i from that agent, report r_j from the other agent, and makes policy choice $\widehat{x}$. All players are risk

[6] In other words, so that $x^* = 1$ with probability q.

[7] The importance of this requirement, of incentive compatibility for P in her use of private information, is demonstrated by Jost (1996).

neutral.

The utility functions for each player can be written as follows:

$$U_P = \Psi_P - |x^* - \widehat{x}| - w(r_1, r_2, \widehat{x}) - w(r_2, r_1, \widehat{x}) \tag{2.1}$$

$$U_i = w_i(r_i, r_j, \widehat{x}) - \lambda\, |x_i - \widehat{x}| - \Psi_A \tag{2.2}$$

where Ψ_P, a normalization constant, will henceforth always be assumed to be large enough to ensure $U^P > 0$. The value each agent places on a policy decision that is in agreement with (rather than against) the interests of his privately-known type is denoted λ.[8] The constant, Ψ_A, normalizes participation utilities for the agents and can be thought of as representing the outside opportunity or reservation wage of the agent-experts.[9]

The timing of the game follows the following 6 steps:

1. Nature draws q from the uniform distribution on $[0, 1]$.
2. P defines symmetric bonus terms, $w(r_i, r_j, \widehat{x})$, such that these terms are always acceptable to the agents (satisfy their participation constraints).
3. P, A_i, and A_j observe signals, $\theta_0, \theta_i, \theta_j$, and the agents observe their types, x_i, x_j.[10]
4. Agents simultaneously make their reports, r_i and r_j.
5. P chooses $\widehat{x}$ and pays bonuses, $w(r_i, r_j, \widehat{x})$, to both agents as promised.
6. Only after bonuses have been paid are the optimal policy, x^*, revealed by nature and the principal's utility realized.

[8] By this I mean that λ is the positive difference in A_i utility between the cases where $\widehat{x} = x_i$ and the cases where $\widehat{x} \neq x_i$.

[9] Assuming experts obtain their information costlessly, as a by-product of their ongoing work activity, this Ψ_A parameter could take a negative value, such as $-\frac{1}{2}\lambda$.

[10] The principal-agent relationship (staffing decisions and the development of suppliers) predates the initiation of the project.

Discussion. This game set-up, where player information on P's decision consists of Bernoulli draws from a prior distribution, allows a simple characterization of information where agent reports are soft, in the sense that they are incontrovertible. As mentioned in the introduction, this distinguishes the present paper from a large body of work in the sender-receiver game literature that assumes reports either consist of hard evidence or are verifiable (as in Milgrom and Roberts, 1986).

A necessary feature for this formulation is that x^* cannot be contracted on. As explicitly laid out in the game sequencing, this inability to contract on x^* could result from conditions where x^* is not expected to be realized until after contracts are implemented (or the agency relationship is terminated). An alternative interpretation would be that x^* is observable at an earlier phase of the game (perhaps Step 5) but that it is unverifiable, *i.e.*, a reported value of x^* can be manipulated by the principal. A third and perhaps more intuitive interpretation is that x^* is never observed. Although not explicitly modeled, one can conceptualize a scenario where P never learns the relative payoffs of policy alternatives (she only learns the payoff of the chosen policy). More explicit modeling of this "road not taken" interpretation might be adapted by introducing an additional noise term and interpreting U_P as an expected value.

The key to understanding this game is to realize that P would like to know the value of q. The best she can do on this score, however, is to generate a Bayesian estimate for Eq. The actual value of q is unobservable and any report, r_i, can be consistent with any observation, θ_0, and outcome, x^*. Given all players hold common prior beliefs that q is drawn from the uniform distribution on [0,1], the formula for this Bayesian estimate, for a random sampling of n Bernoulli draws indexed by i, is[11]

$$E(q|\theta_0, \theta_1, \ldots, \theta_n) = \frac{\sum \theta_i + 1}{n+2} \tag{2.3}$$

Thus, in a fully revealing equilibrium ($r_i = \theta_i$ for both agents), P's best estimate of the probability that $x^* = 1$ will be $\frac{r_1+r_2+\theta_0+1}{3+2}$.[12]

The full state space for this game can be identified as $\mathbf{X} \times \Theta$, where $\mathbf{X} = \{0,1\} \times \{0,1\} \times \{0,1\}$ represents the space of possible draws (x^*, x_i, x_j) over player preferences or

[11]For a derivation of this result, see Mood *et al.* (1974), p. 342, or Degroot (1970), Theorem 1, p. 160.

[12]Although agent types, x_i, are not revealed, this has no bearing on the principal's optimal policy choice.

types and $\Theta = \{0,1\} \times \{0,1\} \times \{0,1\}$ represents the space of informational draws that can collectively be observed by the principal and agents. For purposes of analysis, however, one can focus simply on the state space of player information relevant to Ex^*, which is Θ. Using equation (2.3), one can observe that the prior probability for each of the two possible outcomes in which all three draws agree is $\frac{1}{4}$. Likewise, by symmetry, the probability of each of the six possible outcomes where only two draws agree is $\frac{1}{12}$. All of these probabilities will be common knowledge to the players of this game. The problem for each player will thus be to maximize expected utility over this state space.

3. Existence of a Fully Revealing Equilibrium

This section develops the initial result of this paper — existence of a fully revealing equilibrium. In doing show, it illustrates how the mechanism design results of Crémer and McLean can be adapted to this scenario of contracting for expertise. Starting from the end of the game (Step 5) and moving backward, the first consideration for analysis is the principal's choice of $\widehat{x}$. This choice, in turn, depends on P's beliefs regarding the truthfulness of agents' reporting. We must find an equilibrium such that P believes agent reporting will be truthful and, in turn, agents believe that P's policy choice will reflect the information of truthful agent reporting. Assuming truthful reporting by the agents, application of formula (2.3) thus implies the following conditions for the principal's policy choice:

1. Whenever $r_i = r_j$, P prefers a policy choice, $\widehat{x}$, that agrees with agent reports, regardless of θ_0.

2. Whenever $r_i \neq r_j$, P prefers a policy of $\widehat{x} = \theta_0$.

To assure $\widehat{x}$ is in accordance with these conditions, it must be the case that $U_P(\widehat{x}) \geq U_P(x')$, for $x' \neq \widehat{x}$. By definition of U_P, this can be expressed as

$$\begin{aligned} &E\left[-\left|x^* - \widehat{x}\right| - w(r_i, r_j, \widehat{x}) - w(r_j, r_i, \widehat{x})\right] \geq \\ &\qquad E\left[-\left|x^* - x'\right| - w(r_i, r_j, x') - w(r_j, r_i, x')\right], \\ &\qquad\qquad \text{for } x' \neq \widehat{x}, \text{ and all truthful } r_i, r_j \end{aligned}$$

This implies eight incentive compatibility constraints, one for each of the eight states of nature observable by P in a truth-telling equilibrium, ($r_i = \theta_i \in \{0,1\}$, $r_j = \theta_j \in \{0,1\}$, $\theta_0 \in \{0,1\}$).

Note first that, by this constraint, bonus terms $w(1,1,0)$ and $w(0,0,1)$ should never be paid in equilibrium (condition 1, above). To ensure this, it is immediate that P can establish her incentive compatibility for the two states of nature in which $r_i = r_j \neq \theta_0$ by simply setting $w(1,1,0)$ and $w(0,0,1)$ equal to some large number, M. This will also ensure incentive compatibility for the two $r_i = r_j = \theta_0$ states of nature. The remaining four states of nature require that, when agent reports disagree, $\widehat{x}$ agrees with θ_0 (this satisfies condition 2). The assumption that bonus terms are symmetrical across agents, because it ensures that P has no incentive to bias policy toward the reports of one agent against the reports of the other, ensures that two of these constraints are satisfied. The remaining two are satisfied by an additional symmetry assumption, that bonus contracts are symmetrical across policy alternatives. Thus $w(1,0,1)$ will be assumed equal to $w(0,1,0)$, $w(1,0,0) = w(0,1,1)$, and so forth. By use of these symmetry assumptions, together with the observation that P can costlessly establish bonus terms such that $w(1,1,0) = w(0,0,1) = M$, large, the issue of incentive compatibility for P can be suppressed.[13] For the remainder of this paper, I will maintain these symmetry assumptions and thereby assume that wage contracts are symmetrical both across agents and also across policy alternatives.

Agent incentives. Since agents are *ex ante* symmetrical, it is sufficient to perform an analysis of incentives for one agent only and then to apply these results for both agents. By risk neutrality, the utility associated with any pair of draws, x_i, θ_i, for agent A_i, and any hypothetical values of $r_i, r_2 \in \{0,1\}$, can be written as

$$U_i(x_i, \theta_i) = E[w(r_i, r_j, \widehat{x}[r_i, r_j, \theta_0]) \mid \theta_i] - \lambda \left| x_i - E[\widehat{x}(r_i, r_j, \theta_0) \mid \theta_i)] \right| - \Psi_A$$

Given that A_j reports honestly, A_i's evaluation of this expectation will be over the state space, $\{\Theta|\theta_i\} = \Theta_j \times \Theta_0 = \{0,1\} \times \{0,1\}$. The conditional expected likelihood of each of these states of nature occurring is developed as follows. If $\theta_i = 1$, then A_i's conditional expectation

[13] It is, of course, assumed there are no side arrangements between P and any A_i (at the expense of $\mathrm{A}_{j\neq i}$).

for θ_j, $E(\theta_j|\theta_i = 1)$, equals $\frac{2}{3}$.[14] Further, the conditional expectation, $E(\theta_0|\theta_j = \theta_i = 1)$, is $\frac{3}{4}$. Letting $p(\theta_0, \theta_j|\theta_i)$ denote the joint probability of signals θ_0 and θ_j, given $\theta_i = 1$, it follows that $p(1,1|1) = \frac{1}{2}$, $p(1,0|1) = p(0,1|1) = \frac{1}{6}$, and $p(0,0|1) = \frac{1}{6}$. Thus, taking the case of $\theta_i = 1$, a simple weighting by the likelihood of each state of nature implies

$$\begin{aligned} EU_i(x_i, \theta_i &= 1) = \max_{r_i} \frac{1}{2} w(r_i, 1, \widehat{x}(r_i, 1, 1)) + \frac{1}{6} w(r_i, 1, \widehat{x}(r_i, 1, 0)) \\ &+ \frac{1}{6} w(r_i, 0, \widehat{x}(r_i, 0, 1)) + \frac{1}{6} w(r_i, 0, \widehat{x}(r_i, 0, 0)) \\ &- \lambda \left[\frac{1}{2} |x_i - 1| + \frac{1}{3} |x_i - r_i| + \frac{1}{6} |x_i - 0| \right] - \Psi_A \end{aligned}$$

For truth telling to be incentive compatible it must thus be the case that, when $x_i = 0$ and $\theta_i = 1$,

$$\begin{aligned} &4w(1,1,1) + w(1,0,1) + w(1,0,0) \\ &- 3w(0,1,1) - w(0,1,0) - 2w(0,0,0) \geq 2\lambda \end{aligned}$$

By symmetry, the case of $x_i = 1$ and $\theta_1 = 0$ will be identical to this and the two constraints for the cases where $x_i = \theta_i$ will be slack. Applying the assumed symmetry in bonus terms across policy alternatives, this constraint can be simplified to

$$w(1,1,1) - w(0,1,1) \geq \lambda \tag{3.1}$$

Directly following from the above (and applying symmetry across policy alternatives), the participation constraint for the agent is

$$\frac{4}{6} w(1,1,1) + \frac{1}{6} w(1,0,1) + \frac{1}{6} w(0,1,1) - \frac{\lambda}{2} - \Psi_A \geq 0 \tag{3.2}$$

Thus, in defining $w(\cdot,\cdot,\cdot)$, the principal solves the following optimization problem:

$$\begin{aligned} \max_{w(\cdot,\cdot,\cdot)} U_P &= E\left[\Psi_P - |x^* - \widehat{x}| - w(r_i, r_j, \widehat{x}) - w(r_j, r_i, \widehat{x})\right] \\ s.t. &: \text{(3.1) and (3.2)} \end{aligned}$$

Employing the symmetry assumptions and evaluating expectations, this can be simplified to the following

$$\begin{aligned} \min_{w(\cdot,\cdot,\cdot)} z &= \frac{4}{6} w(1,1,1) + \frac{1}{6} w(1,0,1) + \frac{1}{6} w(0,1,1) \\ s.t. &: \text{(3.1) and (3.2)} \end{aligned}$$

[14] This and subsequent results regarding conditional expectations follow from formula (2.3).

Since this linear programming problem involves two constraints in three variables, there is a continuum of optimal solutions such that (3.1) and (3.2) are satisfied.

Proposition 1. *As long as the expected value to P of the agents' information is high enough to cover the costs of agent participation, there will exist a fully revealing equilibrium for this game.*

Proof. This result follows immediately from the above discussion and is proven by construction. Let $v \equiv w(1,1,1) = w(0,0,0)$, $v' \equiv w(1,0,1) = w(0,1,0)$, and $v'' \equiv w(0,1,1) = w(1,0,0)$. Assume P offers any set of bonus terms such that $v = v'' + \lambda$. Then, in order to satisfy participation constraints for the agents, P will have to set $v' = 6\Psi_A - 5v'' - \lambda$. Since such bonus terms will satisfy constraints (3.1) and (3.2), they will support fully revealing strategies by the agents. The expected cost of these bonus terms to P (summed over both agents) is simply the cost of participation for both agents, $\lambda + 2\Psi_A$. The expected benefit to P of the agents' information is $\frac{1}{30}$ (calculated by multiplying the probability that agent information alters P's decision times the probability that the affected change in policy benefits P).[15] Thus, whenever $\lambda + 2\Psi_A \leq \frac{1}{30}$,[16] the expected costs of these bonus terms do not exceed their expected benefit to P and thus payment of these terms is not dominated for P. Furthermore, since the expense of these bonus terms is minimal for agent participation, P has no incentive to deviate from these terms. Therefore, since these terms satisfy incentive compatibility for the agents, they support a fully revealing equilibrium. □

Note that, as is standard in the Crémer and McLean (CM) approach, the principal is simply using competition between agents, together with the correlation of agent information, to eliminate the agents' informational rents. The above result thus follows directly from P's

[15]Agent reports influence policy only when they both disagree with the principal's information. The *ex ante* likelihood of $\theta_i = \theta_j \neq \theta_0$ is $p(\theta_i = 0|\theta_0 = 1)p(\theta_j = 0|\theta_0 = 1, \theta_i = 0) = \left(\frac{1}{3}\right)\left(\frac{1}{2}\right) = \frac{1}{6}$. The likelihood the policy decision is improved by deference to agent reports under these circumstances is equal to the probability that $x^* = \theta_i = \theta_j \neq \theta_0$ minus the probability that $x^* = \theta_0 \neq \theta_i = \theta_j$. The likelihood that $x^* = 1$, given $\theta_i = \theta_j = 1$ and $\theta_0 = 0$, is $\frac{3}{5}$. Thus, the probability that $x^* = 1$ minus the probability that $x^* = 0$, given that $\theta_i = \theta_j = 1 \neq \theta_0$, is $\frac{1}{5}$. Multiplying this probability times the above probability that $\theta_i = \theta_j \neq \theta_0$ yields $\frac{1}{30}$.

[16]Recall that, if expertise accrues to the agents costlessly, the value of Ψ_A equals $-\frac{1}{2}\lambda$ (see Section 2).

assumed ability to contract over agent reports. One difference is that, unlike in the CM set-ups, in this case the principal has private information. For the existence result demonstrated above, however, this private information is redundant. The existence of a fully revealing equilibrium depends entirely on the costs of agent participation. If the expected value of the information gained by the principal exceeds the expected costs of hiring two agents, then a fully revealing equilibrium will exist.[17]

The key distinction between the set-up of this paper and the standard contracting with correlated information problem is that, in the game of this paper, the agents have not only correlated information over the likely profitability of the principal's policy choice (their expertise) but also uncorrelated private interests regarding that choice. The complexities introduced by these considerations, in terms of alternative noncollusive play, and the role of the principal's private information in resolving these complexities are explored in the next section. The further considerations introduced by the possibility of collusion are discussed in Section 5.

4. The Assurance of Full Information Revelation

The above proposition simply claims existence, not uniqueness. In addition to the above equilibria, there are two types of alternative equilibria. First, there are the trivial degenerate equilibria, standard to any communication game, that result from confusion over language.

[17]A standard criticism of this CM approach of using correlated information to extract the agents' informational rents is that it relies too strongly on reward and punishment strategies. In this game, since the correlation of information between agents is relatively strong, the use of these rewards and punishments may seem less objectionable. The difference between a reward payoff and a punishment payoff is simply λ. The assumption that limited liability constraints might bind in this environment is explored in a working paper version of this paper. This analysis shows that as limited liability binds, participation constraints become slack and wage terms $v' = w(1,0,1) = w(0,1,0)$ will be set equal to wage terms $v'' = w(0,1,1) = w(1,0,0)$. Following from this, the cost to P of inducing agent participation now becomes a function of agency rents (defined by the limited liability constraint).

To eliminate these I will assume clear language conventions, where "yes" means "yes."[18] Second, there is the potential for combinations of nontruthful reporting strategies by the agents that support alternative nontrivial degenerate and nondegenerate equilibria. These are the focus of this section.

To recognize the potential for these alternative equilibria, notice that the above development of incentive compatibility constraints for agent i was all done under the assumption that agent j was fully revealing. Alternative constraints exist for the cases of other opposing-agent strategies. That is to say, a full treatment of potential equilibrium outcomes must consider the incentive compatibility of truthful reporting by A_i given that A_j reports falsely in some states of nature. For example, it is possible that A_i believes that A_j will report truthfully whenever $\theta_j = x_j$ but will report falsely whenever $\theta_j \neq x_j$. Such beliefs effect the incentive compatibility constraints for A_i.

For the traditional Crémer and McLean approach to mechanism design with correlated information, these considerations of alternative strategies for opposing agents are not necessary. This is because the traditional set-up is defined so that agents will have dominant strategies. With contracting for expertise, as modeled here, this is not necessarily the case. This is because the agents' private interests regarding P's policy decision are independently distributed private information.

The key to eliminating these alternative equilibria is for the principal to make use of her private information. Regardless of opponent agent strategies, there will still be a correlation between the information of each agent, θ_i, and the information of the principal, θ_0. P can take advantage of this to define bonus terms that, in expectation, align the private interests of the agents (aggregated as the sum of expected bonus receipts and self-interest over the principal's policy decision) with her own.

The following development proceeds first by providing some intuition for how these bonus terms are developed, and then presents a formal proposition and proof. At Step 4 of the

[18]The issue of language assumptions is standard in the communication-game, especially cheap-talk, literature. For a discussion of these issues see Farrell and Rabin (1996). Since the possibility that a principal might interpret an agent's recommendation for Policy A as being a recommendation for Policy B adds little insight to this analysis, it will hereafter be ignored.

game, when agents make their reports, they each have information sets, $\Theta_i \times X_i$, with four elements.[19] I denote the interim state of nature where $\theta_i = 1$, $x_i = 0$ as partition 1, the case where $\theta_i = x_i = 1$ as partition 2, the case where $\theta_i = x_i = 0$ as partition 3, and the case where $\theta_i = 0$, $x_i = 1$ as partition 4. Taking the perspective of A_j, for each of these four partitions in his information set, he can choose either to report honestly, $r_j \in \{0, 1 \mid r_j = \theta_j\}$, or to lie, $r_j \in \{0, 1 \mid r_j \neq \theta_j\}$. Thus, for the continuation game that begins after P has defined bonus terms, A_j has a strategy space with 2^4 or 16 elements in pure strategies. To label these strategies, I use the notation, 1234, to denote the strategy of telling the truth for all four of these information partitions, 123, to denote the strategy of telling the truth in only the first three partitions, *null*, to denote the strategy of always lying, and so forth.

Clearly, the choice of strategy by A_j affects the expected value of r_j and, in turn, the expected payoff for alternative strategies for A_i. For example, consider the A_j strategy denoted 234. Let $p(r_i, r_j, \theta_0)$ refer to the joint probability of the three outcomes in the parentheses. If A_j has adopted strategy 234, then honest reporting by A_i will lead to the following distribution of outcomes:

$$\begin{array}{ll} p(1,1,1) = \frac{3}{24} & \qquad p(0,0,0) = \frac{7}{24} \\ p(1,1,0) = \frac{1}{24} & \qquad p(0,0,1) = \frac{3}{24} \\ p(1,0,1) = \frac{5}{24} & \qquad p(0,1,0) = \frac{1}{24} \\ p(1,0,0) = \frac{3}{24} & \qquad p(0,1,1) = \frac{1}{24} \end{array}$$

In contrast, if A_j adopts a policy, 1234, of always reporting honestly, then honest reporting by A_i will yield a distribution of outcomes: $p(1, 1, 1) = p(0, 0, 0) = \frac{6}{24}$, $p(1, 1, 0) = \ldots = p(0, 1, 1) = \frac{2}{24}$.

Incentive compatibility for A_i to report truthfully that $r_i = 1$ when $x_i = 0$ in this first case requires (recalling the symmetry assumptions) that

$$-2w(1, 1, 1) + 2w(1, 0, 1) \geq 3\lambda$$

Incentive compatibility in the second case, where A_j never lies, requires simply that

$$w(1, 1, 1) - w(0, 1, 1) \geq \lambda$$

[19] $\Theta_i = \{0, 1\}$ and $X_i = \{0, 1\}$ are the spaces of possible θ_i and x_i draws, respectively.

As above, to simplify notation, let the variables v, v', and v'' represent the symmetrical bonus terms $w(1,1,1)$, $w(1,0,1)$, and $w(0,1,1)$, respectively.

Note that, in the calculation of A_i's expected payoffs, a lie by A_j when $\theta_j = x_j = 1$ (partition 2) is equivalent to a lie by A_j when $\theta_j = 1 \neq x_j$ (partition 1). Likewise, lies by A_j in partitions 3 and 4 are equivalent to A_i. Because of this, for the calculation of A_i payoffs, the strategy space for A_j can be reduced to only 9 elements. A complete listing of incentive compatibility (IC) constraints for A_i, for the case where he may wish to misreport $r_i = x_i = 0$ whenever $\theta_i = 1$ (his partition 1), can thus be developed as follows (Table A-1, Column 3).

Table A-1

Case	A_j Strat.	IC against A_i lying in his part. 1 ($r_i = x_i = 0$ when $\theta_i = 1$)	IC against A_i lying in his part. 2 ($r_i = 0$ when $\theta_i = x_i = 1$)
i	1234	$v - v'' \geq \lambda$	$v - v'' \geq -\lambda$
ii	234	$4v - v' - 3v'' \geq 2\lambda$	$4v - v' - 3v'' \geq -2\lambda$
iii	123	$-2v + 2v' \geq 3\lambda$	$-2v + 2v' \geq -3\lambda$
iv	12	$3v - v' - 2v'' \geq \lambda$	$3v - v' - 2v'' \geq -\lambda$
v	34	$-3v + 2v' + v'' \geq 2\lambda$	$-3v + 2v' + v'' \geq -2\lambda$
vi	23	$v' - v'' \geq 3\lambda$	$v' - v'' \geq -3\lambda$
vii	2	$2v - 2v'' \geq 3\lambda$	$2v - 2v'' \geq -3\lambda$
viii	3	$-4v + 3v' + v'' \geq 4\lambda$	$-4v + 3v' + v'' \geq -4\lambda$
ix	null	$-v + v' \geq 2\lambda$	$-v + v' \geq -2\lambda$

By symmetry, the incentive compatibility constraints against a temptation for A_i to misreport that $r_i = x_i = 1$ when $\theta_i = 0$ (i.e., in partition 4) will be identical to these except that the constraint for A_j strategy 234 will become the constraint for A_j strategy 123, and so forth. Incentive compatibility against a temptation by A_i to lie that $r_i = 0$ when $x_i = \theta_i = 1$ (partition 2) will be the same as above except that the right hand side values will be negative. For example, for the A_j strategy 234, it must be the case that $4v - v' - 3v'' \geq -2\lambda$. A full listing of these constraints for the case when A_i is in partition 2 is provided in Column 4 of Table A-1.

These above results thus enumerate the incentive compatibility constraints that ensure truthful reporting by A_i against all possible A_j strategies. For an agent's reporting strategy to be sustainable in equilibrium, it cannot be dominated in any partition. To align agent incentives, the principal thus simply needs to find bonus terms that satisfy these constraints. These are provided in the following proposition. As will be shown in the proof, the alignment of incentives through these bonus terms relies on a logic of iterated elimination of strictly dominated strategies.

Proposition 2. *Suppose the principal and the agents agree on language conventions. Then, for any* $\varepsilon > 0$*, the following bonus terms will be sufficient to rule out misreporting in any Bayesian perfect continuation of this game:* $v = \left(\frac{1}{2}\lambda + \Psi_A\right)(1+\varepsilon)$, $v' = (3\lambda + \Psi_A)(1+\varepsilon)$, *and* $v'' = -(\lambda - \Psi_A)(1+\varepsilon)$. *Note that, as* $\varepsilon \to 0$*, the expected cost of these wage terms for* P *will approach the participation costs of the agents* ($\lambda + 2\Psi_A$)*, and full extraction of agent rents will be realized.*

Proof. Following the above logic, Table A-1 presents a full listing, accounting for all possible strategies of the opposing agent, of the incentive compatibility conditions against lying by an agent for each of his four information partitions. Under the proposed bonus terms, a strategy that requires A_i to misreport with any positive probability in either partition 2 or 3 is strictly dominated for him. Observe that, by symmetry, the same will be true for A_j. Elimination of these strictly dominated strategies from the above tables allows for the removal of all remaining (partition 1) incentive compatibility constraints except those derived from cases *i*, *ii*, *iii*, and *vi* on the table. By the hypothesized bonus terms, an A_i strategy of lying with any positive probability is strictly dominated for all of these remaining cases. By symmetry, lying by A_j is also strictly dominated. The only surviving strategies are full revelation by both agents.

The binding constraints that determine the hypothesized bonus terms are from cases *iii* and *vi* for partition 1 from the table and participation for the agents. Since the agent participation constraint is binding, the full extraction of informational rents is verified. □

Sensitivity analysis. To further explore the implications of this above proposition, consider the parameter λ, which represents the strength of the private interests of the agents in this game. The logic implies that, for any $\lambda > 0$, P will not be able to ensure a fully revealing equilibrium without using her private information (v' cannot equal v''; it must be greater than or equal to $v'' + 4\lambda$). In order words, if the players in this game believe that the expert-agents have any, even the smallest amount, of independently distributed private interests, then strategic use of private information by the principal will be necessary to ensure a fully revealing outcome. Thus, in games such as this one, where agents' incentives are unknown (e.g., where agents may just as well have private preferences in the direction of either Project B as in the direction of Project A), the extension to the Crémer and McLean concept that is presented above is important.

Since it is so important that P be informed in this game, another question to ask is what if P is assumed to be somewhat less informed than the agents. Suppose, for example, that P only randomly observed θ_0 (suppose, with probability η, she observes θ_0 and that otherwise she observes nothing). This will force an increase in the coefficients of λ for the above contract terms. As $\eta \to 0$ this will require that bonus terms v' become larger and bonus (punishment) terms v'' become smaller.[20] A standard criticism of the Crémer and McLean approach is that, as the correlation of agent information approaches zero, the reward and punishment payments necessary to induce information revelation become impracticably large (approach infinity). This model illustrates an analogous result. As the principal becomes less and less informed, the reward and punishment payments necessary to defeat alternative nonrevealing equilibria also approach infinity. Thus, for information revelation to be assured

[20] This result seems intuitive. The less informed the principal is, the more vulnerable she becomes to manipulation by her agent-experts. Thus, to maintain an alignment of agent incentives, she must use larger incentive payments. The resulting behavior may appear erratic or arbitrary. A relatively uninformed customer may go to a bad auto mechanic for years before noticing a missed diagnosis or self-serving recommendation. Once she believes she may have observed bad service, however, she will likely switch mechanics immediately. Indeed, by herd-like behavior, she may switch on the basis of a rumor that her friend got bad service. A better informed customer, since she will accumulate evidence more quickly, may be willing to wait for stronger evidence before making a switch.

by this approach, it is important that the private information of the principal be within some rough parity with the information of the agents.

5. The Potential for Collusion

Since the above results depend on the ability of the principal to play one agent off against the other, it is natural to ask the extent to which they are robust to the potential for collusion. To investigate this question, two general approaches seem to be most appealing. First, one could focus exclusively on the potential for pre-play communication between agents and consider the possibility of Bayesian or generalized correlated equilibrium strategies by the agents such that they would coordinate on when and how they misrepresent their information.[21] I refer to an equilibrium that is robust to this type of generalized correlated play between agents as "communication-proof." Second, to consider robustness to the broadest possible scope for collusive activity, one could extend this communication-proofness logic to allow for transferable utility between agents and enforceable side contracting. A logic for this stronger refinement is presented by Laffont and Martimort (1998), in their concept of "collusion-proofness." We proceed by considering both of these concepts of collusion, in turn.

Communication-proofness. Looking first at the refinement of robustness to generalized correlated equilibria, this concept is used to represent the broadest possibility of communication between players in a game. Since a full characterization of the potential extensive forms communication processes might take would be impractical, the correlated equilibrium concept makes use of a third party coordinator or mediator, who maximizes the interests of the parties to a communication process. This artificial construct is referred to as mediated communication. As applied to this game, the concept is modeled by allowing each agent to fully share his private information with the mediator who then gives advice to the agents on what to report to the principal. Clearly, for the concept to have any appeal, it is required

[21] Myerson (1991, pp. 249-263) uses the terminology *communication equilibrium* and *generalized correlated equilibrium* to describe the generalization of the concept of correlated equilibrium as it would be applied to a Bayesian game.

that providing information to the mediator and following the advice of the mediator must be incentive compatible for each agent.

The application of this concept to the game of this paper is somewhat complicated. Thus space limitations prevent a complete formal development. However, to illustrate its implications, suppose the third-party mediator proposed that agents revealed their private information to it and that, whenever $x_i = x_j \neq \theta_i = \theta_j$, it would suggest that reports $r_i = r_j \neq \theta_i = \theta_j$ be coordinated. In cases other than the $x_i = x_j \neq \theta_i = \theta_j$ state of nature, it would suggest that agents report honestly. Would such an arrangement for coordinated strategies prevent P from inducing full information revelation?

Results of a full analysis of potential generalized correlated equilibria strategies show that it would not. P can design incentive contracts that ensure full information revelation and are robust to collusion via *any* generalized correlated equilibrium in the continuation play of the agents.[22] Furthermore, these contracts do not require payment of any informational rents. The logic of this result is very much analogous to that of Proposition 2; a slight increase in v', coupled with slight decreases in v and v'', are enough to lure defection from any suggested correlated reporting scheme.[23]

What generalized correlated equilibria strategies do for the agents is logically equivalent to allowing each agent to observe a binary partition of the state space that he believes his opponent is playing from. The signal or advice of the mediator simply identifies which of the two elements of this partition his opponent occupies. Although this tightens the restriction on incentive compatibility for honest play by each agent, it does not prevent P from being able to define bonus terms that ensure truth-telling will be a dominant strategy. A formal development of a revised game that is adapted to this communication-proofness context and the proof of these results are in Appendix 2.

Collusion-Proofness. The concept of collusion-proofness is used to assess the robustness of an equilibrium to the broadest possible scope of arrangements for communication and

[22] The resulting bonus terms are $v = (\frac{1}{3}\lambda + \Psi_A)(1+\varepsilon)$, $v' = (\frac{10}{3}\lambda + \Psi_A)(1+\varepsilon)$, and $v'' = (-\frac{5}{3}\lambda + \Psi_A)(1+\varepsilon)$.

[23] P cannot defeat all strategies whereby agents swap reports (A_i reports θ_j and A_j reports θ_i), but such strategies are outcome equivalent to truth telling by the agents.

side contracting between players in a game. For the analysis of this game, the key distinction between collusion-proofness and communication-proofness is that agents can *commit* to playing according to the recommendations of the mediator. This has the effect of weakening the requirement of incentive compatibility for agent participation in collusive arrangements. Whereas the generalized correlated equilibrium concept requires that agent participation in collusive arrangements must be *ex interim* incentive compatible, under collusion-proofness it only has to be *ex ante* incentive compatible.

As in the correlated equilibrium concept, since a full characterization of the potential extensive forms communication and bargaining processes might take would be impractical, this collusion-proofness concept makes use of a third party, who maximizes the interests of the parties to a collusion process. Since this equilibrium concept involves commitment by the agents (enforceable side contracting), this third party is best thought of as a mechanism designer. As above, the concept is modeled by allowing each agent to fully share his private information with this third party mechanism designer, who then, in exchange for whatever side-payment is specified by the mechanism, specifies how the agents will make their reports. For incentive compatibility, given whatever rules the mechanism specifies, it is required that the agents be willing to participate honestly.[24]

To illustrate the implications of the collusion-proofness requirement, suppose the third-party mechanism designer proposed a modified (binding) version of the above communication-proofness example. The agents reveal their private information to the third party who then requires that whenever $x_i = x_j \neq \theta_i = \theta_j$, the agents report $r_i = r_j \neq \theta_i = \theta_j$. Transfers to the agents would be zero and, in cases other than the $x_i = x_j \neq \theta_i = \theta_j$ state of nature, agents would report honestly. If one accepts that it is reasonable for the agents to be able to enter into this kind of binding arrangement, then these collusive strategies will indeed defeat the principal's strategies from the above Proposition 2 (and also the revised communication-proof form of these strategies).

How can the principal respond to this potential for collusion? There are several options.

---

[24]Laffont and Martimort also apply restrictions to the way agents update beliefs regarding the types of their opponents when collusive arrangements are rejected. Specifically, they require the appealing refinement that updating in reponse to out-of-equilibrium behavior be passive.

Perhaps most straightforwardly, the principal can randomly employ additional experts. Although likely increasing payroll or participation costs, this would increase the difficulty of forming collusive arrangements and greatly reduce the scope for gains from collusion for the agents. The likelihood that $\theta_1 = \theta_2 = \theta_3 \neq x_1 = x_2 = x_3$ is substantially smaller than the likelihood that $\theta_1 = \theta_2 \neq x_1 = x_2$.

Alternatively, within the framework of the existing two-agent model, P can respond by randomizing over the introduction a fourth bonus payment, $w(0, 0, 1)$. In other words, the principal can establish a policy such that, when the reports of the agents both contradict her private information, she will randomly choose between the policy indicated by her prior, without benefit of agent advice ($\widehat{x} = \theta_0 \neq r_1 = r_2$), and a policy of following agent advice ($\widehat{x} = r_1 = r_2 \neq \theta_0$). The credibility of such a mixing strategy would require that P set this $w(0, 0, 1)$ bonus term in such a way that she was indifferent between accepting or rejecting agent reports. As long as the private interests of the agents are not too strong relative to the interests of the principal, this is possible.[25] The choice between these solution approaches would thus depend on the value of λ and the relative costs of hiring additional experts versus the costs of randomizing the principal's policy choice in a way that occasionally ignores expert opinion. For a formal development of the revised game that is adapted to collusion-proofness and the proof of this result, see Appendix 1.

6. Conclusion

This paper has developed a model of incentive contracting for the provision of information from interested experts. It shows that key considerations in such situations are the potential for competition between experts and the potential for the principal to use her own private information. At first glance, the problem looks like a straightforward application of the

[25] By the requirement of incentive compatibility for the principal (at Step 5 of the game) discussed toward the beginning of Section 3, P must do this in a way that prevents degeneracy and preserves agent incentives for truth-telling. It turns out that, for values of $\lambda < \frac{1}{5}$, this is possible. Note that only P engages in mixed strategies. The agents continue to use pure, truth-telling, strategies. The values of bonus terms for this mixed strategy approach depend, nonlinearly, on the value of λ.

logic of mechanism design with correlated information. As was shown, however, this is not the case. The problem is that the private information of experts is likely to have two dimensions. Although we would expect expert opinion regarding an optimal policy decision to be correlated, we would not necessarily expect the same for the private incentives of experts. In the scenarios of this paper, even the smallest amount of independently distributed private interest, on the parts of the experts, is enough to upset the standard contracting for correlated information result.

Because of this, the mechanism design problem faced by the principal in contracting for expertise is somewhat complicated. The expected payoffs to experts for alternate reporting strategies will depend on each expert's beliefs regarding the strategy of his opponent, across his opponent's multidimensional type space. Because of this, multiple equilibria can exist. The solution for a principal, to ensure full information revelation, is to use her private information regarding the policy choice to design incentive contracts that reshape the private incentives of the experts. By doing so she can eliminate these multiple equilibria and induce full information revelation by agent-experts as a dominant strategy equilibrium. Without use of the principal's private information, this is not possible.

To assess the robustness of this solution approach to the potential for collusion between experts, two equilibrium refinements were evaluated. The first, generalized correlated equilibrium, was used to assess robustness to any form of communication between expert-agents (communication-proofness). It showed that, by a slight revision to incentive contracting terms, the principal can continue to ensure a fully informative equilibrium (with no informational rents). The second, collusion-proofness, showed that even if agents are allowed to enforceably side contract over collusive arrangements, given that the private interests of the agents are small enough relative to the interests of the principal, she can still ensure full information revelation (at the cost of some sacrifice of efficiency in the use of expert information).

An appeal of this logic of contracting for expertise is that the contracting arrangements described seem to be relatively realistic. Surely managers facing important decisions in situations where there is a potential bias in expert opinion do solicit advice from more than one expert. Large engineering firms maintain staffs of project engineers who compete with each other in this advice-giving role. Matrix and network-form organizations provide additional examples of these kinds of arrangements. Furthermore, the form of bonus terms described would seem to arise in a somewhat natural way. When project or product design advice is solicited from two project engineers, it is likely that one of these engineers will be rewarded with an assignment of project leadership. To win such an assignment represents a substantial reward in prestige, organizational power, and, likely, monetary compensation. When the advice of competing expert project engineers is in agreement, corresponding to the mid-range bonus case, the likelihood of winning the assignment is significant but diluted. When advice of the engineers disagrees, corresponding to the larger bonus case, that engineer who agrees with the manager benefits from an undiluted chance at winning the assignment (and likely higher prestige from being selected as a clear winner). The engineer who loses the competition will suffer a corresponding larger loss of prestige. Lastly, regarding the principal's use of her private information, this seems entirely consistent with the importance we place on knowledge and ability in upper management. It also seems to fit with the age-old management tactic of "holding one's cards close to one's chest."

As for the broader application of these kinds of arrangements, in addition to the intraorganizational contracting for expertise contexts motivating this paper, the concept also has some application in procurement scenarios. For example, although procurement contests for military equipment are often decided primarily on the basis of cost competition (and thus are modeled well in a standard auction framework), most large procurements are decided through submission of design as well as cost proposals. Indeed, since both the military procurement authority and its bidders typically share common estimates of budget availability and since the incentive of the procurement authority may well be to spend all of its available budget, the design competition is often considered to be paramount. Thus, it would seem that an important function of these contests is to elicit expert information on optimal design. Bidder incentives in advocating alternative design concepts are not clear and may be biased in various directions independent of those that are in the interest of the procurement authority.

Another area for application is in the monitoring or hierarchical agency problems of Tirole (1986, 1992) and McAfee and McMillan (1995). This line of research explores the incentives of supervisors who have been delegated the responsibility to monitor and direct agents in a hierarchy. As applied to this literature, the results of this paper show how the formulation of these hierarchical agency problems can be generalized to allow for soft, unverifiable information in supervisory monitoring. Also, as stressed in the modeling of Laffont (1990), this paper shows how these models can be adapted to situations where supervisors (referred to as agent-experts in the present paper) may have opportunities to report falsely in more than one direction. An implication of this paper for the monitoring and hierarchical agency literature is to emphasize the importance of parallelism or redundancy in monitoring. Parallel or overlapping systems of oversight will create situations where the private information of expert-agent monitors is correlated. The principal can take advantage of this correlation to reduce agency rents and to make otherwise problematic arrangements workable. Furthermore, to implement such arrangements, the principal should not neglect the strategic use of her own private information.

Appendix 1: Collusion-Proofness

The modeling of collusion-proofness is, as much as possible, in exact parallel to the modeling of the main text. It follows the framework laid out by Laffont and Martimort (1998). First a third-party mechanism designer, whose purpose it is to facilitate collusion, is introduced into the game. After the principal has offered the agents her grand contract (step 2 of the original game) and agents have observed their private information (step 3 of the original game), this third party mechanism designer proposes a binding, incentive-compatible, side mechanism by which the agents can collude. If agents accept this side mechanism, they report their private information to the third party who then manipulates reports in such a way as to maximize the joint welfare of the two agents. Replacing step 4 of the original game, the third party passes these manipulated reports to the principal. If one or both of the agents refuse the side mechanism, then reports are passed by the third party to the principal without manipulation. The full timing of this revised game is detailed below:[1]

1. Nature draws q from the uniform distribution, $G(q)$, on $[0,1]$.

2. P defines bonus terms, $w(r_i, r_j, \widehat{x})$, such that these terms are always acceptable to the agents. These bonus terms are assumed to be symmetrical both across agents and across policy alternatives.

3. P, A_1, and A_2 observe signals, $\theta_0, \theta_1, \theta_2$, and the agents observe their types, x_1, x_2.

4. A third-party proposes a side-mechanism mapping $(r_1, r_2) \rightarrow (m_1, m_2)$ and including side payments, $y_i(r_1, r_2, x_1, x_2)$, paid by the third party to the agents in return for their reports, such that $\sum_i y_i(r_1, r_2, x_1, x_2) = 0$, $\forall\ (r_1, r_2, x_1, x_2)$.

5. If both agents accept the side mechanism, the third party reports (m_1, m_2) to P and side contracts $y_i(\theta_1, \theta_2, x_1, x_2)$ are implemented. If either agent refuses the side mechanism, the third party passes through the unmanipulated reports of the agents and no side contracts are implemented.

6. P chooses $\widehat{x}$ and pays bonuses, $w(m_i, m_j, \widehat{x})$, to both agents as promised.

[1] Also, Laffont and Martimort (1998, p. 8) apply restrictions to the way agents update beliefs regarding the types of their opponents when collusive arrangements are rejected. Specifically, they require the appealing refinement that updating in reponse to out-of-equilibrium behavior be passive.

7. Only after bonuses have been paid are the optimal policy, x^*, revealed by nature and the principal's utility realized.

As discussed in the body of the paper, such third-party mechanisms for implementing collusion defeat the Proposition 2 incentive contracting approach of the principal. One way in which the principal can respond to this potential for collusion is by introduction of a fourth bonus payment, $w(0,0,1)$. Using this fourth bonus-payment alternative, the principal can establish a policy such that, when the reports of the agents both contradict her private information, she will randomly choose between the policy indicated by her prior, without benefit of agent advice ($\widehat{x} = \theta_0 \neq m_1 = m_2$), and a policy of following agent advice ($\widehat{x} = m_1 = m_2 \neq \theta_0$). The credibility of such a mixing strategy requires that P set this $w(0,0,1)$ bonus term in such a way that she is indifferent between accepting or rejecting agent reports. This in turn limits the scope for such mixing strategies to scenarios in which $\lambda < \frac{1}{5}$. The formalization of this result is provided in the following proposition. For simplicity, I denote symmetrical bonus terms as follows: $v = w(1,1,1)$, $v' = w(1,0,1)$, $v'' = w(0,1,1)$, and $v''' = w(0,0,1)$.

Proposition 3. *For values of $\lambda < \frac{1}{5}$, there will exist randomizations over the use of this fourth wage term, v''', that are, by the definition of Laffont and Martimort (1998), collusion-proof. Using these randomizations, P can induce full information revelation by the agents.*

Proof. To prove these results, first, consider the conditions that must hold for P to be indifferent between setting $\widehat{x} = \theta_1 = \theta_2 \neq \theta_0$ versus setting $\widehat{x} = \theta_0 \neq \theta_1 = \theta_2$.

- The probability that $\theta_1 = \theta_2 = \theta_0 \neq x_1 = x_2$ is $\left(\frac{1}{2}\right)\left(\frac{1}{4}\right) = \frac{1}{8}$.
- The probability that $\theta_1 = \theta_2 \neq \theta_0 = x_1 = x_2$ is $\left(\frac{1}{6}\right)\left(\frac{1}{4}\right) = \frac{1}{24}$.
- The probability that $\theta_1 = \theta_2 \neq \theta_0$ and x_1 and/or $x_2 \neq \theta_0$ is $\left(\frac{1}{6}\right)\left(\frac{3}{4}\right) = \frac{1}{8}$.
- Assume that, when $\theta_1 = \theta_2 \neq \theta_0$, $\widehat{x} = \theta_0$ with probability p and, when $\theta_1 = \theta_2 \neq x_1 = x_2$, agents lie with probability q. In all other states of nature agents tell the truth.
- This implies that the probability that P observes $m_1 = m_2 \neq \theta_0$ is equal to prob($\theta_1 = \theta_2 \neq \theta_0$ and $x_1$and/or $x_2 \neq \theta_0$) + $(1-q)\cdot$prob($\theta_1 = \theta_2 \neq \theta_0 = x_1 = x_2$) + $q\cdot$prob($\theta_1 = \theta_2 = \theta_0 \neq x_1 = x_2$) $= \frac{1}{8} + (1-q)\frac{1}{24} + q\frac{1}{8} = \frac{4}{24} + q\frac{2}{24}$. Thus, given that $m_1 = m_2 \neq \theta_0$, the probability that $\theta_1 = \theta_2 = \theta_0$ is equal to $= \frac{3q}{4+2q}$.

- This implies that for payoffs to be such that P is indifferent, it must be the case that the payoff to P from setting $\widehat{x} = \theta_0 \neq m_1 = m_2$, $-2v''' + \frac{3}{5}\frac{3q}{4+2q} + \frac{2}{5}\left(1 - \frac{3q}{4+2q}\right)$, equals the gain from setting $\widehat{x} = m_1 = m_2$, $-2v + \frac{3}{5}\left(1 - \frac{3q}{4+2q}\right) + \frac{2}{5}\frac{3q}{4+2q}$

$$\Rightarrow q = \frac{1 - 10\left(v - v'''\right)}{1 + 5\left(v - v'''\right)}$$

So, for P to credibly mix in an equilibrium where $q = 0$, it must be the case that $v - v''' = \frac{1}{10}$. Any solution that requires that the difference between v and v''' be greater than $\frac{1}{10}$ would violate the incentive compatibility of mixing for P and could thus lead to degenerate play.

Now, to define the constraints for collusion-proofness for the agents against a bonus term system that uses payoff $v''' = w(0,0,1)$, whenever $\theta_1 = \theta_2 \neq \theta_0$ with probability p, we first define the baseline payoffs to the agents assuming their information is revealed truthfully. This is:

$$EU_A = \frac{4-p}{6}v + \frac{1}{6}v' + \frac{1}{6}v'' + \frac{p}{6}v''' - \frac{1}{2}\lambda - \Psi_A$$

Using symmetry, the number of possible states of nature over θ_1, θ_2, x_1, and x_2 draws can be reduced to eight. For each of these states of nature, deviations from truthful reporting can be defined as follows (in state of nature 7, there are two possible coordinated deviations that could yield gains for the agents).

(1) If $\theta_i = x_i = x_j = \theta_j$, assume, if other ICs are met, deviation is dominated.

(2) If $\theta_i = x_i = x_j \neq \theta_j$, assume $m_i = m_j = x_i$.

(3) If $\theta_i = x_i \neq x_j \neq \theta_j$, assume $m_i = x_i \neq m_j = x_j$.

(4) If $\theta_i = x_i \neq x_j = \theta_j$, assume $m_i = x_i = \theta_i = m_j \neq \theta_j$, with prob .5, else $m_i = \theta_j$.

(5) If $\theta_i \neq x_i \neq x_j = \theta_j$, assume $m_i = x_i \neq m_j = x$.

(6) If $\theta_i \neq x_i \neq x_j \neq \theta_j$, assume $m_i = x_i = \theta_i = m_j \neq \theta_j$, with prob .5, else $m_i = \theta_j$.

(7a) If $\theta_i \neq x_i = x_j \neq \theta_j$, assume $m_i = m_j \neq \theta_i = \theta_j$.

(7b) If $\theta_i \neq x_i = x_j \neq \theta_j$, assume $m_i = \theta_i = \theta_j \neq m_j$, with prob .5., else $m_i = x_i$

(8) If $\theta_i \neq x_i = x_j = \theta_j$, assume $m_i = m_j = x_i$.

Now, for each of these fundamental deviations, I define the impact on agent payoffs in terms of its difference from the baseline truth-telling case (including the impact on both agents of one's deviation).

(1) If Proposition 2 ICs are met, deviation is dominated.

(2) $\frac{1}{3}v - \frac{1}{6}vp + \frac{1}{6}pv''' - \frac{1}{6}p\lambda - \frac{1}{6}v' - \frac{1}{6}v'' + \frac{1}{6}\lambda$

(3) $\frac{1}{3}v' + \frac{1}{3}v'' - \frac{2}{3}v + \frac{1}{6}vp - \frac{1}{6}pv'''$

(4) $\frac{1}{3}v - \frac{1}{6}vp + \frac{1}{6}pv''' - \frac{1}{6}v' - \frac{1}{6}v''$

(5) $\frac{1}{3}v' + \frac{1}{3}v'' - \frac{2}{3}v + \frac{1}{6}vp - \frac{1}{6}pv'''$

(6) $\frac{1}{3}v - \frac{1}{6}vp + \frac{1}{6}pv''' - \frac{1}{6}v' - \frac{1}{6}v''$

(7a) $-\frac{1}{3}vp + \frac{1}{3}pv''' - \frac{1}{3}p\lambda + \frac{1}{2}\lambda$

(7b) $\frac{1}{3}v' + \frac{1}{3}v'' - \frac{2}{3}v + \frac{1}{6}vp - \frac{1}{6}pv''' + \frac{1}{6}p\lambda$

(8) $\frac{1}{3}v - \frac{1}{6}vp + \frac{1}{6}pv''' - \frac{1}{6}p\lambda - \frac{1}{6}v' - \frac{1}{6}v'' + \frac{1}{6}\lambda$

Combining these constraints (while removing redundancies) with the objective function for the principal, as defined for this problem, yields the following nonlinear programming problem.

$$
\begin{array}{lll}
\min & 2\left(\left(4-p\right)v + v' + v'' + pv'''\right) - \left(\dfrac{p}{5}\right)\left(\dfrac{1}{6}\right) & \text{(obj)} \\
s.t. & \dfrac{4-p}{6}v + \dfrac{1}{6}v' + \dfrac{1}{6}v'' + \dfrac{p}{6}v''' - \dfrac{1}{2}\lambda \geq 0 & \text{(IR)} \\
 & \dfrac{1}{3}v - \dfrac{1}{6}vp + \dfrac{1}{6}pv''' - \dfrac{1}{6}p\lambda - \dfrac{1}{6}v' - \dfrac{1}{6}v'' + \dfrac{1}{6}\lambda \leq 0 & (2) \\
 & \dfrac{1}{3}v' + \dfrac{1}{3}v'' - \dfrac{2}{3}v + \dfrac{1}{6}vp - \dfrac{1}{6}pv''' \leq 0 & (3) \\
 & \dfrac{1}{3}v - \dfrac{1}{6}vp + \dfrac{1}{6}pv''' - \dfrac{1}{6}v' - \dfrac{1}{6}v'' \leq 0 & (4) \\
 & -\dfrac{1}{3}vp + \dfrac{1}{3}pv''' - \dfrac{1}{3}p\lambda + \dfrac{1}{2}\lambda \leq 0 & \text{(7a)} \\
 & \dfrac{1}{3}v' + \dfrac{1}{3}v'' - \dfrac{2}{3}v + \dfrac{1}{6}vp - \dfrac{1}{6}pv''' + \dfrac{1}{6}p\lambda \leq 0 & \text{(7b)} \\
 & v - v''' = \dfrac{1}{10} & \text{(IC for } P\text{)} \\
 & \lambda = \text{ constant} &
\end{array}
$$

The solution to this program depends on the value of λ. For values of $\lambda < \frac{1}{5}$, there does exist a feasible solution. Now since these constraints define incentive compatibility of each individual state of nature that could be observable by a mechanism designer in this game, the satisfaction of these constraints will also ensure collusion-proofness against any combinations of these deviations (combination deviations will be additive). Furthermore, it can be verified that the wage terms that satisfy these constraints will also ensure truthful revelation against non-collusive agent strategies.

Finally, note that, for any $p > 0$, the outcome of the resulting game will not be first-best efficient. Indeed, as $p \to 1$ (which occurs as $\lambda \to \frac{1}{5}$), the benefit of agent information to P goes to zero. $\square$

Appendix 2: Communication-Proofness

Considerations of the weaker refinement of a generalized correlated equilibrium are very similar to the Appendix 1 collusion-proofness analysis. The difference is that it is no longer assumed that agents can make binding commitments to play by the mechanism of the third party mediator.[2] In terms of revisions to the game, these are the same as for the collusion-proofness set-up except that Steps 4 and 5 are revised as follows:

4. A third-party proposes a side-mechanism whereby each agent can report his private information, (θ_i, x_i), to the mediator and the mediator will then give advice, m_i, to each agent regarding what he should report to the principal.

5. If the side mechanism is incentive compatible, then both agents report their private information to the mediator honestly and follow his advice regarding their reporting to the supervisor. If the side mechanism is not incentive compatible, then it is ignored by the agents.

As demonstrated by the following Proposition 4, P can defeat collusion against this sort of correlated equilibrium in continuation play without payment of informational rents and without recourse to mixed strategies equilibria. The logic of this proposition is very much analogous to that of Proposition 2; P can define bonus terms, v, v', and v'', such that the aggregate interests of each agent are, in expectation, aligned with those of the principal. What generalized correlated equilibrium strategies do for the agents is essentially provide each agent with knowledge of a partition of the state space of his opponent agent. The signal or advice of the mediator can be interpreted as telling each agent, A_i, what partition his opponent, A_j, is playing from. Although this tightens the restriction on incentive compatibility for honest play by each agent, it does not prevent P from being able to define incentive compatible bonus terms.

Proposition 4. *For any $\varepsilon > 0$, the following bonus terms will be sufficient to rule out misreporting for any nondegenerate continuation of this game that is consistent with generalized correlated equilibria strategies for the agents: $v = \left(\frac{1}{3}\lambda + \Psi_A\right)(1+\varepsilon)$, $v' = \left(\frac{10}{3}\lambda + \Psi_A\right)(1+\varepsilon)$, and $v'' = -\left(\frac{5}{3}\lambda - \Psi_A\right)(1+\varepsilon)$. Note that, as $\varepsilon \to 0$, the expected cost of these wage terms for P will approach $\lambda + 2\Psi_A$ and full extraction of agent rents will be realized.*

[2] Also, no transfers are allowed, but these are not required to implement the above collusive strategies.

Proof. Without loss of generality, the impact of the mediator's signal or advice to players can be represented as a simple partition (into two parts) of each agent's beliefs regarding the state space from which his opponent agent is playing. In other words, the mediator's signal can be modeled as a piece of information that tells A_i which of two partitions for A_j's possible states of nature A_j is in. Note that there are four possible states of nature (SON) for each player and that a state space with four elements can be bifurcated in eight possible ways. Since the signal of the mediator will tell A_i which partition A_j is occupying, a full enumeration of the possible state spaces A_i may believe his opponent is playing from would consider 15 possibilities ($16 - 1$, because A_j cannot occupy the null space). Since, for the logic of the proof, A_i is indifferent regarding the value of x_j, this enumeration can be reduced to six elements, as illustrated below.

Potential partitions are:

Case	θ_2	x_2
1	1	0
2	1	1
3	0	0
4	0	1

$\Rightarrow$

null	1234
1	234
2	134
3	124
4	123
12	34
13	24
14	23

$\Rightarrow$ simplification $\Rightarrow$

partition of θ_2
1 or 2 or 12
3 or 4 or 34
14 or 13 or 24 or 23
134 or 234
124 or 123
1234

Within each of these partitions, we consider the incentive compatibility constraints for truth-telling by A_1 against all possible strategies for A_2 (the constraints for the full state space, 1234, are enumerated in the proof to Proposition 2).

Case	Partition	A_2 report true in SON	IC for A_1, $\theta_1 = 1, x_1 = 0$
1	1 and/or 2	12	$-v + \frac{1}{4}v' + \frac{3}{4}v'' + \frac{1}{4}\lambda \leq 0$
2		1 or 2	$-\frac{1}{2}v + \frac{1}{2}v'' + \frac{3}{8}\lambda \leq 0$
3		null	$v - \frac{3}{4}v' - \frac{1}{4}v'' + \frac{3}{4}\lambda \leq 0$
4	3 and/or 4	34	$v - \frac{1}{2}v' - \frac{1}{2}v'' + \frac{1}{2}\lambda \leq 0$
5		3 or 4	$\frac{1}{2}\lambda \leq 0$
6		null	$-v + \frac{1}{2}v' + \frac{1}{2}v'' + \frac{1}{2}\lambda \leq 0$
7	1,3 = 1,4 = 2,3 = 2,4	14	$-\frac{1}{3}v + \frac{1}{3}v'' + \frac{1}{3}\lambda \leq 0$
8	(treat all as 14)	1	$-v + \frac{1}{3}v' + \frac{2}{3}v'' + \frac{1}{3}\lambda \leq 0$
9		4	$v - \frac{2}{3}v' - \frac{1}{3}v'' + \frac{2}{3}\lambda \leq 0$
10		null	$\frac{1}{3}v - \frac{1}{3}v' + \frac{2}{3}\lambda \leq 0$
11	1,2,3 = 1,2,4	124	$-\frac{6}{10}v + \frac{1}{10}v' + \frac{5}{10}v'' + \frac{3}{10}\lambda \leq 0$
12	(treat all as 124)	12	$-v + \frac{3}{10}v' + \frac{7}{10}v'' + \frac{3}{10}\lambda \leq 0$
13		14 (= 24)	$\frac{2}{10}v - \frac{3}{10}v' + \frac{1}{10}v'' + \frac{5}{10}\lambda \leq 0$
14		4	$v - \frac{7}{10}v' - \frac{3}{10}v'' + \frac{7}{10}\lambda \leq 0$
15		1 (= 2)	$-\frac{2}{10}v - \frac{1}{10}v' + \frac{3}{10}v'' + \frac{5}{10}\lambda \leq 0$
16		null	$\frac{6}{10}v - \frac{5}{10}v' - \frac{1}{10}v'' + \frac{7}{10}\lambda \leq 0$
17	2,3,4 = 1,3,4	134	$-\frac{1}{8}v' + \frac{1}{8}v'' + \frac{3}{8}\lambda \leq 0$
18	(treat all as 134)	34	$v - \frac{5}{8}v' - \frac{3}{8}v'' + \frac{5}{8}\lambda \leq 0$
19		13 (= 14)	$-\frac{4}{8}v + \frac{1}{8}v' + \frac{3}{8}v'' + \frac{3}{8}\lambda \leq 0$
20		1	$-v + \frac{3}{8}v' + \frac{5}{8}v'' + \frac{3}{8}\lambda \leq 0$
21		3 (= 4)	$\frac{4}{8}v - \frac{3}{8}v' - \frac{1}{8}v'' + \frac{5}{8}\lambda \leq 0$
22		null	$-\frac{1}{8}v' + \frac{1}{8}v'' + \frac{5}{8}\lambda \leq 0$

Now, reviewing the above, two constraints can be eliminated. To see this, suppose that A_2 is in partition 3 and/or 4 and lies in SON 3 and/or 4. In these scenarios, incentive compatibility for A_1 not to lie from his SONs 1 and/or 2 will not be feasible. Note however, that such lies by A_1 and A_2 will cancel each other out. Thus, the aggregate report to P will be truthful and these constraints are redundant.

To proceed with the proof, hypothesize the following bonus terms, $v = \frac{1}{3}, v' = \frac{10}{3}, v'' = -\frac{5}{3}$. Given that these above cases are ruled out and given these hypothesized bonus terms, any misreporting by A_i when $\theta_i = x_i$ will be dominated. Thus, by symmetry, lying by any agent who is in SON 2 or 3 will be dominated. Eliminating these strategies for A_2 yields the following reduced table of constraints.

Case	Partition	Report true in SON	IC, $\theta_1 = 1, x_1 = 0$
1	1 and/or 2	12	$-v + \frac{1}{4}v' + \frac{3}{4}v'' + \frac{1}{4}\lambda \leq 0$
2		2	$-\frac{1}{2}v + \frac{1}{2}v'' + \frac{3}{8}\lambda \leq 0$
3		null	assumed A_2 strat. is dominated
4	3 and/or 4	34	$v - \frac{1}{2}v' - \frac{1}{2}v'' + \frac{1}{2}\lambda \leq 0$
5		3	$\frac{1}{2}\lambda \leq 0$ ruled out
6		null	assumed A_2 strat. is dominated
7	1,3 = 1,4 = 2,3 = 2,4	14	$-\frac{1}{3}v + \frac{1}{3}v'' + \frac{1}{3}\lambda \leq 0$
8	(treat all as 14)	1	$-v + \frac{1}{3}v' + \frac{2}{3}v'' + \frac{1}{3}\lambda \leq 0$
9		4	$v - \frac{2}{3}v' - \frac{1}{3}v'' + \frac{2}{3}\lambda \leq 0$
10		null	$\frac{1}{3}v - \frac{1}{3}v' + \frac{2}{3}\lambda \leq 0$
11	1,2,3 = 1,2,4	124	$-\frac{6}{10}v + \frac{1}{10}v' + \frac{5}{10}v'' + \frac{3}{10}\lambda \leq 0$
12	(treat all as 124)	12	$-v + \frac{3}{10}v' + \frac{7}{10}v'' + \frac{3}{10}\lambda \leq 0$
13		24	$\frac{2}{10}v - \frac{3}{10}v' + \frac{1}{10}v'' + \frac{5}{10}\lambda \leq 0$
14		4	assumed A_2 strat. is dominated
15		2	$-\frac{2}{10}v - \frac{1}{10}v' + \frac{3}{10}v'' + \frac{5}{10}\lambda \leq 0$
16		null	assumed A_2 strat. is dominated
17	2,3,4 = 1,3,4	134	$-\frac{1}{8}v' + \frac{1}{8}v'' + \frac{3}{8}\lambda \leq 0$
18	(treat all as 134)	34	$v - \frac{5}{8}v' - \frac{3}{8}v'' + \frac{5}{8}\lambda \leq 0$
19		13	$-\frac{4}{8}v + \frac{1}{8}v' + \frac{3}{8}v'' + \frac{3}{8}\lambda \leq 0$
20		1	assumed A_2 strat. is dominated
21		3	$\frac{4}{8}v - \frac{3}{8}v' - \frac{1}{8}v'' + \frac{5}{8}\lambda \leq 0$
22		null	assumed A_2 strat. is dominated

The resulting linear programming problem for P is feasible and yields the hypothesized bonus terms. Multiplying these bonus terms by $1 + \varepsilon$, for any $\varepsilon > 0$, ensures that full information revelation will be dominant for the agents in any Bayesian or generalized correlated equilibrium. By inspection, it can be seen that, as $\varepsilon \to 0$, the participation constraints for the agents will be binding and, thus, information rents for the agents will go to zero. □

References

Crawford, Vince and Joel Sobel, "Strategic Information Transmission," *Econometrica*, November 1982, *50* (6), 1431–51.

Crémer, Jacques and Richard P. McLean, "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, March 1985, *53* (2), 345–61.

___ **and** ___ , "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, November 1988, *56* (6), 1247–57.

DeGroot, Morris H., *Optimal Statistical Decisions*, McGraw-Hill, 1970.

Farrell, Joseph and Matthew Rabin, "Cheap Talk," *Journal of Economic Perspectives*, Summer 1996, *10* (3), 103–18.

Gibbons, Robert, "Learning in Equilibrium Models of Arbitration," *American Economic Review*, December 1988, *78* (5), 896–912.

Jost, Peter-Jurgen, "On the Role of Commitment in a Principal-Agent Relationship with an Informed Principal," *Journal of Economic Theory*, February 1996, *68* (2), 510–30.

Laffont, Jean-Jacques, "Analysis of Hidden Gaming in a Three-Level Hierarchy," *Journal of Law, Economics, and Organization*, Fall 1990, *6* (2), 301–24.

___ **and David Martimort**, "Mechanism Design with Collusion and Correlation," *mimeo*, June 1998.

McAfee, R. Preston and John McMillan, "Organizational Diseconomies of Scale," *Journal of Economics and Management Strategy*, Fall 1995, *4* (3), 399–426.

___ **and Philip J. Reny**, "Correlated Information and Mechanism Design," *Econometrica*, March 1992, *60* (2), 395–421.

Milgrom, Paul R. and John Roberts, "Relying on the Information of Interested Parties," *Rand Journal of Economics*, Spring 1986, *17* (1), 18–32.

Mood, Alexander M., Franklin A. Graybill, and Duane C. Boes, *Introduction to the Theory of Statistics*, third ed., McGraw-Hill, 1974.

Myerson, Roger B., *Game Theory: Analysis of Conflict*, Harvard University Press, 1991.

Prendergast, Canice, "A Theory of "Yes Men."," *American Economic Review*, September 1993, *83* (4), 757–70.

Shin, Hyun Song, "Adversarial and Inquisitorial Procedures in Arbitration," *RAND Journal of Economics*, Summer 1998, *29* (2), 378–405.

Tirole, Jean, "Hierarchies and Bureaucracies: On the Role of Collusion in Organizations," *Journal of Law, Economics, and Organization*, Fall 1986, *2* (2), 181–214.

——, "Collusion and the Theory of Organizations," in Jean-Jacques Laffont, ed., *Advances in Economic Theory: Proceedings from the Sixth World Congress of the Econometric Society, Vol. 2*, 1992, pp. 151–206.