Should Principals Inform Agents?:

Information Management in Agency Problems

by

Tracy R. Lewis*

and

David E. M. Sappington**

Revised March 1991

* University of California, Davis.
** University of Florida.

We are grateful to the National Science Foundation, the University of California Energy Research Group, and the Public Utilities Research Center at the University of Florida for financial support. We wish to thank Jacques Cremer, David Levine, Preston McAfee, Tom Palfrey, Gary Ramey, Richard Romano, Tom Ross, Jean Tirole, Ralph Winter and seminar participants at Barcelona, Cal Tech, U. C. Davis, Florida, Toronto, and Toulouse for helpful comments.

ABSTRACT

We examine the optimal management of the information structure in a class of agency problems. The principal's expected utility is shown to be a convex function of the accuracy of the agent's private information. Therefore, when she has unmitigated control of the environment, the principal will always prefer one of two extreme information structures: (1) where the agent is endowed with the best available knowledge of his environment; or (2) where the agent has no private information. Intermediate structures provide lower expected utility for the principal, but not necessarily for the agent.

1. Introduction.

Most models in the agency literature specify the agent's information exogenously, presuming it cannot be altered by either the principal or the agent. The repairman in Holmstrom's (1979) model and the regulated firm in Baron and Myerson's (1982) model are both endowed with perfect knowledge of their costs of operation. Similarly, in Maskin and Riley's (1984) analysis, each consumer is fully aware of his tastes for the monopolist's product.

In practice, a principal will often have some control over the information available to her agent. To illustrate, a monopoly seller can improve her customers' knowledge of their tastes for the monopolist's product by undertaking informative advertising, conducting sales demonstrations, and permitting hands-on experimentation with the product. Similarly, a sales force manager can improve the information of her sales force by collecting and disseminating detailed statistics on product performance and on the activities of competitors. Alternatively, an employer can limit the information of her subordinates about their environment, by, for example, failing to keep detailed records of accidents in the workplace, or by concealing the operating schedules and capabilities of monitors.[1]

In this paper, we explicitly consider the possibility that the principal and/or the agent may have control over the information available to the agent. We focus on the principal's preferences regarding the agent's information, characterizing the optimal information structure from the principal's point of view. A fundamental trade-off is present. When more accurate information is available to the agent, the agent's activity level (e.g., production) can be better tailored to his environment (e.g., realized production costs). This increases the principal's expected welfare, *ceteris paribus*. However, when the accuracy of the agent's private information is improved, the rents the agent commands generally increase. These rents come at the principal's expense.

We find that in the class of problems under consideration, this fundamental tradeoff is optimally resolved from the principal's perspective with one of two extreme information structures. The risk-neutral principal either prefers that the agent be endowed with the best available knowledge of his environment or that the risk-neutral agent have no private information. Either extreme can be strictly preferred by the principal, but intermediate information structures, modeled as linear combinations of the best available information and the principal's information, yield her lower expected utility. The principal's preference for extreme information structures is one manifestation of a more general conclusion: the principal's expected utility is a convex function of the accuracy of the agent's private information. The marginal benefits of improved accuracy increase with the level of accuracy, while the corresponding marginal costs decline, leaving the principal with a systematic preference for extreme information structures.

We are also able to show that the principal's expected utility always declines as the accuracy of the agent's private signal increases slightly above the accuracy of the principal's own information. As in Radner and Stiglitz' (1984) analysis of a single-person decision problem, the first-order gains from improved information are nil. However, in the setting analyzed here, there is a first-order increase in the agent's rents as his private information improves.

In very common settings, the principal's losses from improved information for the agent continue to outweigh the gains over the entire range of possible information structures. To illustrate, in Baron and Myerson's (1982) classic model, the regulator will always prefer the regulated firm to have no private information about its constant marginal cost of production whenever this cost is drawn from a uniform distribution and consumers' surplus is a quadratic function of the output of the regulated firm. Thus, to the extent that a principal has any control over the agent's knowledge of operating costs in models of this sort, the common assumption that the agent is endowed with perfect cost information is called into question. On the other hand, there are settings where the principal will prefer the agent to have the best available information about his environment. Two factors enhance this preference: (1) when the optimal level of activity for the agent is sensitive to the realized state of the world; and (2) when the rents the agent is able to command from superior knowledge are limited. Thus, a monopoly seller will prefer her customers to have good pre-purchase information about their valuation of the monopolist's product when their final demand for the product varies significantly with the realized valuation. The monopolist gains from the price discrimination she can undertake when consumers have precise information about their preferences. Similarly, in a procurement setting, the buyer will prefer the seller to have accurate cost information when the buyer's demand for the good is very elastic, but will prefer to deal with an uninformed seller when the buyer's demand is inelastic.

In many settings, the agent may also have some control over the accuracy of his private information. For example, a manufacturer may be able to undertake costly research and experimentation to obtain a more accurate assessment of his production costs before production begins. In such settings, the principal will alter the incentive scheme she offers the agent to influence his choice of information structure. The manner in which the incentive scheme is optimally altered for this purpose is described in section 5. First, though, the elements of the basic model under consideration are described in section 2. Our general findings are recorded in section 3. Section 4 presents additional conclusions drawn from three special cases of the general model. Conclusions are summarized in section 6. The Appendix contains formal proofs that are not central to the analysis.

Before proceeding, we briefly discuss the four papers in the literature that seem most closely related to our own. Sobel (1983) asks whether a principal prefers her risk-averse agent to share the principal's imperfect knowledge of the state of the world, or to know the state perfectly. Perfect knowledge of the state effectively reduces the agent's costs of production, but also requires that the agent be promised his opportunity wage for each realization of the state, rather than in expectation. Sobel (1983) finds that with a binary outcome space, a risk-neutral principal always prefers an informed agent to an uninformed one. With a risk-averse principal or many possible outcomes, either ranking can emerge. Aside from differences in how information affects the production technology and in the agent's aversion to risk, our approach differs from Sobel's in that we permit a continuum of information structures rather than just two.

Clay, Sibley, and Srinagesh (CSS) (1991) ask whether a profit maximizing firm will prefer to have customers commit to a payment structure before or after uncertainty about their demand is resolved. Given some restrictions on the nature of the uncertainty and on the policy instruments available to the firm, CSS show the firm will prefer to contract with customers before the uncertainty is resolved, provided the magnitude of the uncertainty is sufficiently small.

The other two related papers are those of Cremer and Khalil (CK) (1990a,b).[2] CK allow the risk-neutral agent to acquire a signal about the true state of nature either before or after the principal proposes a contract. This information is not valuable for planning purposes, because (in contrast to our analysis) the agent ultimately observes the realization of the state before acting. Thus, better information about the state only increases the agent's rents. Consequently, the principal always prefers that the agent *not* acquire the *ex ante* information about the state. CK's focus is on how the principal can mitigate the agent's incentive to acquire the information. Thus, of all our findings, those in section 5 are related most directly to CK's findings.

2. The Basic Model.

There are two risk-neutral actors in the model: the principal and the agent. The principal's utility is $V(Q) - T$. T is a transfer payment from the principal to the agent. Q is the realized level of activity, such as the level of production or consumption by the agent. The agent's utility is $U(Q,\theta) + T$. $\theta \in [\underline{\theta}, \overline{\theta}]$ is the realized state of nature. The agent's utility is assumed to decrease with θ i.e., $U_2(Q,\theta) < 0 \;\; \forall Q > 0$. Thus, in the procurement interpretation of our model where the agent produces output Q for the principal, θ will index the agent's cost of production. Higher costs of production lower the agent's utility, *ceteris paribus*. In the monopoly seller interpretation of our model where the agent purchases a good of quality Q from the principal, θ will be inversely related to the agent's valuation of quality. Higher values of θ correspond to lower valuations of quality, and thus to lower levels of utility for any given level of quality. No further restrictions on $V(\cdot)$ or $U(\cdot)$ are necessary for our main results.

The central concern in our analysis is the information about θ that is available to the agent. Initially, the agent shares the principal's beliefs about θ. These beliefs are captured by the density function $f(\theta) > 0 \quad \forall\, \theta \in [\underline{\theta}, \overline{\theta}]$ and the corresponding distribution function $F(\theta)$. Before a contract is signed, however, the agent may have an opportunity to update his prior beliefs about θ, perhaps through experimentation and careful study of the environment in which he is to work. The outcome of this process is a signal $s \in [\underline{s}, \overline{s}]$ about θ that the agent observes privately. Initially, we presume the principal has complete control over the accuracy of the agent's private signal. For example, the principal may control the agent's access to the workplace or the release of detailed product specifications. Formally, this control is modeled by allowing the principal to choose the probability $\gamma \in [0, 1]$ that the agent's signal is drawn from an informative distribution, $G(s)$. With probability $1 - \gamma$, the signal is drawn from an uninformative distribution, $G^0(s)$. We will refer to γ as the accuracy of the agent's private signal. Although the principal's choice of γ is common knowledge in the model, neither the principal nor the agent can be certain whether the realized signal is informative when $\gamma \in (0, 1)$. Furthermore, the agent's assessment of γ is assumed to be unaffected by his private observation of s, implying $G(s) = G^0(s)$ $\forall s \in [\underline{s}, \bar{s}]$.

When s is informative (i.e., when s is drawn from G(s)), the conditional distributions of θ, $F(\theta|s)$, differ in the sense of strict first-order stochastic dominance according to the realization of s. Formally, for each $s \in [\underline{s}, \bar{s}]$, $F_2(\theta|s) \leqq 0$ $\forall$ $\theta \in [\underline{\theta}, \bar{\theta}]$, with strict inequality for some values of θ. When s is uninformative (i.e., when s is drawn from $G^0(s)$), the relevant conditional distribution of θ is $F(\theta)$, which is invariant to the realization of s. Given this set-up, the conditional distribution of θ given signal s and accuracy γ is simply γ $F(\theta|s) + [1 - \gamma]$ $F(\theta)$, a linear combination of the most informative and the least informative information structures.[3]

A special case of the most informative information structure provides the agent with perfect knowledge of the realized state of nature. In this case, $F(\theta|s) = 0$ $\forall \theta < s$ and $F(\theta|s) = 1$ $\forall \theta \geqq s$, $\forall$ $\theta, s \in [\underline{\theta}, \bar{\theta}]$. As noted in the introduction, the agency literature commonly assumes that this information structure is available, and that $\gamma = 1$. Our immediate task is to determine the principal's optimal choice of γ for any given $F(\theta|s)$ distribution characterized by strict stochastic dominance. By limiting the accuracy of the agent's private signal, the principal can limit the rents the agent commands from his private information. However, she also limits the information about the environment that is available before actions (e.g., production or purchase decisions) are undertaken. On the other hand, by choosing a value of γ close to unity, the principal provides the agent with good information for planning purposes at the cost of increased rents for the agent.[4]

The principal's choice of γ is observable and irreversible. This choice is the first action in the model. Next, the principal commits to an incentive scheme $\{Q(\cdot), T(\cdot)\}$ which specifies a menu of possible activity levels, Q, and associated transfer payments, T, from the principal to the agent. The activity and payment pairs are indexed by the agent's report of his privately observed signal, s. After he observes the realization of s, the agent chooses his most preferred $\{Q(\cdot), T(\cdot)\}$ pair, formally by reporting his private information. The specified activity is then carried out, and the promised transfer payment is made. We assume that after he observes s, the agent can refuse all the options offered by the principal. If he does so, the relationship between the principal and agent is terminated, and the agent earns his reservation utility level, which is normalized to zero.

To examine the principal's optimal strategy in this environment, we analyze the $\{Q(\cdot), T(\cdot)\}$ menu most preferred by the principal for each possible value of $\gamma \in [0, 1]$. We then derive her preferred value of γ. Formally, the principal's problem in this setting for given γ, [P-γ], is the following:

$$\underset{Q(s),\, T(s)}{\text{Maximize}} \quad \int_{\underline{s}}^{\bar{s}} [V(Q(s)) - T(s)]\, dG(s) \tag{2.1}$$

subject to $\forall\, s, \hat{s} \in [\underline{s}, \bar{s}]$:

$$\pi(s\,; s \,|\, \gamma) \geqq 0\,, \quad \text{and} \tag{2.4}$$

$$\pi(s\,; s \,|\, \gamma) \geqq \pi(\hat{s}; s \,|\, \gamma)\ ; \tag{2.3}$$

$$\text{where} \quad \pi(\hat{s}\,; s \,|\, \gamma) - T(\hat{s}) + \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s), \theta)\, [\gamma dF(\theta \,|\, s) + (1-\gamma) dF(\theta)]. \tag{2.4}$$

The individual rationality constraints (2.2) ensure the agent's net expected utility is nonnegative for every realization of his private signal. The incentive compatibility constraints (2.3) guarantee the agent truthfully reports his private information. The principal's objective (2.1) is to maximize the expected value of the utility she receives from Q less transfer payments to the agent.

For future reference, three definitions are useful. $Q_\gamma(s)$ will denote the induced activity level in the solution to [P-γ] when signal s is realized. $V^*(\gamma, Q(s))$ will represent the principal's net expected utility given γ and any incentive compatible contract, $Q(s)$. Finally, $V^*(\gamma)$ will denote the net expected utility of the principal in the solution to [P-γ]. The incentive compatibility constraints require $Q'(s) \leqq 0$ a.e.. Therefore, $V^*(\gamma) = \underset{Q(s)}{\text{maximium}} \{V^*(\gamma, Q(s))\}$ subject to $Q'(s) \leqq 0$ a.e..

3. General Findings.

In this section, we characterize the level of accuracy for the agent's private signal that the principal prefers. We begin by analyzing the principal's problem [P-γ] in more detail.

Using standard techniques (e.g., Baron and Myerson (1982)), the local incentive compatibility constraints are readily shown to imply that at the solution to [P-γ], the agent's expected utility decreases with the realized value of s at the rate $\gamma \int_{\underline{\theta}}^{\bar{\theta}} U(Q_\gamma(s), \theta)\, dF_2(\theta | s) < 0$. Therefore, the individual rationality constraints bind only for the highest realizations of the signal. Furthermore, the agent's equilibrium expected utility when he sees signal s is

$$\pi(s | \gamma) = -\gamma \int_{s}^{\bar{s}} \int_{\underline{\theta}}^{\bar{\theta}} U(Q_\gamma(\tilde{s}), \theta)\, dF_2(\theta | \tilde{s})\, d\tilde{s}. \tag{3.1}$$

Lemma 1 follows immediately from (3.1), using integration by parts in (2.1).

Lemma 1.

$$V^*(\gamma, Q(s)) = \gamma \int_{\underline{s}}^{\bar{s}} \{V(Q(s)) + \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s), \theta)[dF(\theta | s) + dF_2(\theta | s)\frac{G(s)}{g(s)}]\}\, dG(s)$$

$$+ [1-\gamma] \int_{\underline{s}}^{\bar{s}} \{V(Q(s)) + \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s), \theta)\, dF(\theta)\}\, dG(s). \tag{3.2}$$

Proposition 1. $V^*(\gamma)$ is a convex function of γ.

Proof. Let $\gamma'' = \delta\,\gamma + [1 - \delta]\,\gamma'$ where $\delta \in (0,1)$. Then because (3.2) is linear in δ:

$$V^*(\gamma'', Q(s)) = \delta V^*(\gamma, Q(s)) + [1-\delta] V^*(\gamma', Q(s)). \tag{3.3}$$

Now, let $Q(s\,|\,\gamma)$ denote the output schedule that maximizes the right hand side of the expression in (3.2). Then

$$V^*(\gamma'') = \delta\, V^*(\gamma, Q(s\,|\,\gamma'')) + [1-\delta] V^*(\gamma', Q(s\,|\,\gamma''))$$

$$\leq \delta\, V^*(\gamma, Q(s\,|\,\gamma)) + [1-\delta]\, V^*(\gamma', Q(s\,|\,\gamma'))$$

$$= \delta\, V^*(\gamma) + [1-\delta]\, V^*(\gamma').$$

The first equality follows from (3.3). The inequality follows from the definition of $Q(s\,|\,\gamma)$. The last equality follows from the definition of $V^*(\gamma)$.[5] ■

The mathematical intuition for Proposition 1 is the following. With risk neutral parties and with updated beliefs about θ characterized as linear combinations of the most and least informative information structures, the principal's expected net utility from a given activity schedule, $Q(s)$, turns out to a linear function of γ. Therefore, holding the activity schedule constant, the principal's expected utility will vary with γ at a linear rate. However, by adjusting the activity schedule as γ is varied, the principal can ensure her expected utility increases more than linearly or decreases less than linearly with γ, leading to the identified convexity of $V^*(\gamma)$.

An immediate consequence of Proposition 1 is:

Corollary 1. $V^*(\gamma) \leq$ maximum $\{V^*(0), V^*(1)\}$ $\forall\, \gamma \in [0, 1]$, i.e., the principal's net expected utility in the solution to [P-γ] is highest either when the agent's private signal reflects the best available information structure ($\gamma = 1$) or when the signal provides no additional information about γ ($\gamma = 0$).

Thus, among all possible information structures in the class under consideration, the structure most preferred by the principal involves either the highest ($\gamma = 1$) or the lowest ($\gamma = 0$) accuracy for the agent's signal. Intermediate values of γ never provide a higher expected return to the principal. An intuitive explanation of this result is provided most readily by considering the procurement interpretation of the model, wherein the agent produces output Q at constant marginal cost θ. There are two effects that arise when the accuracy of the agent's private signal is increased in this setting. The *high signal effect* is detrimental to the principal. When the agent observes a high realization of s, he is more certain that production costs will be high the more accurate his signal. Consequently, he requires greater compensation from the principal to produce any specified level of output. The countervailing *low signal effect* is beneficial from the principal's perspective. A low realization of s signals lower expected costs to the agent the more accurate the signal. Consequently, the agent requires less compensation for a given level of production to achieve the same level of expected utility.

The principal's anticipated losses from increasing γ due to the high signal effect are proportional to the induced output levels for high realizations of s. The corresponding gains due to the low signal effect are proportional to the induced output levels for low s realizations. Most importantly, the more accurate the agent's private signal, the smaller the induced output for the large realizations of s and the greater the induced output for the small realizations of s under the optimal incentive scheme.[6] Consequently, as γ increases, the marginal expected losses from the high signal effect are diminished and the marginal expected gains from the low signal effect are increased. Therefore, the principal's expected welfare is a convex function of γ in this setting, and by analogy, in a variety of other settings.

Before turning to examine this procurement setting and other special cases of the model in more detail, one additional general conclusion is recorded in Proposition 2. The Proposition states that starting from the point of symmetric information where $\gamma = 0$, the agent gains but the principal loses as the accuracy of the agent's private signal is improved slightly. A marginal increase in γ above zero does not improve the accuracy of the agent's private information sufficiently to provide any useful planning information. (See Radner and Stiglitz (1984).) However, the improved accuracy of his private information does enable the agent to command higher rents.[7] In the statement of Proposition 2, $U^*(\gamma) \equiv \int_{\underline{s}}^{\bar{s}} \pi(s\,|\,\gamma)\, dG(s)$ denotes the agent's expected utility in the solution to [P-γ] before he observes his private signal.

Proposition 2. $V^{*\prime}(\gamma)|_{\gamma=0} \leqq 0$ and $U^{*\prime}(\gamma)|_{\gamma=0} \geqq 0$, with strict inequalities if

$$Q_o \equiv \underset{Q}{\text{argmax}} \left\{ V(Q) + \int_{\underline{\theta}}^{\bar{\theta}} U(Q, \theta) dF(\theta) \right\} > 0.$$

Proof. From Lemma 1, the envelope theorem provides

$$V^{*\prime}(\gamma) = \int_{\underline{s}}^{\bar{s}} \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s), \theta) \left[dF(\theta\,|\,s) - dF(\theta) + dF_2(\theta\,|\,s) \frac{G(s)}{g(s)} \right] dG(s).$$

Therefore, since $Q(s) = Q_o \quad \forall\ s \in [\underline{s}, \bar{s}]$ in the solution to [P-0],

$$V^{*\prime}(\gamma)|_{\gamma=0} = \int_{\underline{s}}^{\bar{s}} \int_{\underline{\theta}}^{\bar{\theta}} U(Q_0, \theta)\, dF_2(\theta\,|\,s)\, G(s)\, ds \leq 0, \quad \text{with strict inequality if } Q_0 > 0.$$

From (3.1) $U^{*\prime}(\gamma) = -\gamma \int_{\underline{s}}^{\bar{s}} \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s), \theta)\, dF_2(\theta\,|\,s)\, G(s)\, ds.$ Therefore,

$$U^{*\prime}(\gamma)|_{\gamma=0} = - \int_{\underline{s}}^{\bar{s}} \int_{\underline{\theta}}^{\bar{\theta}} U(Q_0, \theta)\, dF_2(\theta\,|\,s)\, G(s)\, ds \geqq 0, \quad \text{with strict inequality if } Q_0 > 0.$$

■

4. Applications.

In this section, we present three examples which provide additional insight regarding the principal's choice of an information structure. The first example depicts a procurement setting, and demonstrates that the principal will prefer the agent to have access to the most informative information structure when the agent's private information is particularly valuable for planning purposes and when the rents the agent can command from this information are relatively small. The second example illustrates that in procurement settings commonly analyzed in the agency literature, the principal may always prefer the agent to share the principal's imperfect knowledge of the environment. The third example illustrates how these insights extend to a setting where the principal is a monopolist and there is uncertainty about consumers' valuation of product quality or variety. For simplicity, in all three examples we assume that the most informative information structure provides the agent with perfect information regarding the realized state of nature.

<u>Example 1.</u> $$V(Q) = \frac{1}{e} Q^{e} \quad \forall Q \geqq 0; \tag{4.1}$$

$$U(Q, \theta) = - \theta Q \quad \forall Q \geqq 0; \text{ and} \tag{4.2}$$

$$F(\theta) = \left[\frac{\theta}{\bar{\theta}}\right]^{\alpha} \quad \forall \theta \in [0, \bar{\theta}] \quad \text{where } \alpha \geq 0. \tag{4.3}$$

Example 1 identifies a procurement setting in which the seller (the agent) produces output Q at constant marginal cost of production, θ. The buyer's (the principal's) valuation of output reflected in (4.1) corresponds to a demand function ($Q(p) = p^{-\eta}$) for the agent's output with constant elasticity, $\eta > 0$, where $e = \frac{\eta - 1}{\eta}$. Thus, since $\frac{d}{d\eta}(\frac{\eta - 1}{\eta}) > 0$, higher values of e correspond to higher demand elasticities. The higher the demand elasticity, the more sensitive to the realization of θ is the output level that maximizes the difference between the benefits and the costs of production.[8] The agent's private information about θ is more valuable for planning purposes the higher is e in this sense.

The distribution function in (4.3) allows for convenient measurement of the principal's difficulty in controlling the rents of an agent who is privately informed about the realization of θ. Increases in α shift probability mass from the lower to the higher realizations of $\theta \in [0, \bar{\theta}]$,[9] thereby complicating the principal's problem. The equilibrium rents of the privately-informed agent in this Example increase as θ declines at a rate proportional to the induced output level. Thus, since output increases as θ falls under any incentive compatible policy, the rents that accrue to the agent for the lower θ realizations are reduced by limiting the agent's output for the higher realizations of θ. The principal is willing to implement these output distortions when the higher θ realizations are deemed unlikely. But as α increases and the higher θ realizations become relatively more probable, the reduction in total surplus associated with the output reductions weighs more heavily in the principal's calculus, leading to higher equilibrium output levels and, consequently, more rents for the agent.

These considerations suggest that in Example 1, the principal will prefer: (1) no new private information for the agent for the lower values of e and the higher values of α; and (2) perfect information for the agent for the higher values of e and the lower values of α, *ceteris paribus*. In fact, this is precisely what simulations reveal. To illustrate, consider the representative setting where $\bar{\theta} = 10$ and $\alpha = 1$. Then for $e < 0.33$ (corresponding to a demand elasticity of $\eta \in (0, 1.5)$), the principal prefers the agent to be uninformed (i.e., $V^*(0) > V^*(1)$). She prefers the agent to be perfectly informed (i.e., $V^*(1) > V^*(0)$) when her demand is more elastic, i.e., for $e > 0.33$ or $\eta > 1.5$. The principal is indifferent between the two extreme information structures when $e = 0.33$.

Next, fixing the principal's demand elasticity at this point of indifference ($e = 0.33$ or $\eta = 1.5$), simulations reveal that the principal prefers the agent to be perfectly informed if $\alpha < 1$ and to be uninformed if $\alpha > 1$.[10]

<u>Example 2.</u> $V(Q) = aQ - \frac{1}{2}bQ^2 \quad \forall\, Q \geqq 0$ (4.4)

where $b > 0$ and $a > 2\bar{c}$ are constants;

$U(Q, \theta) = -\,\theta Q \quad \forall\, Q \geqq 0$; and (4.5)

$F(\theta) = \frac{\theta}{\bar{\theta}} \quad \forall\, \theta \in [0, \bar{\theta}]$. (4.6)

Example 2 again depicts a procurement setting, this time with uniform cost uncertainty, (4.6), and a quadratic valuation function, (4.4), corresponding to a linear inverse demand function ($P(Q) = a - bQ$). The restriction $a > 2\bar{c}$ ensures that positive output levels are always induced in the solution to [P-γ] for all $\gamma \in [0, 1]$ in the setting of Example 2. Straightforward calculations reveal that in this setting, the principal always prefers her agent to be uninformed (i.e., $V^*(0) > V^*(1)$).[11,12] The agent, on the other hand, prefers to be perfectly informed about his costs (i.e., $U^*(1) \geqq U^*(\gamma) \;\; \forall\, \gamma \in [0, 1]$) provided $a > \frac{13}{6}\bar{c}$. However, when the principal's valuation of the agent's performance is less pronounced (i.e., when $a \in (2\bar{c}, \frac{13}{6}\bar{c})$, so that the intercept of the principal's demand curve is relatively small), the agent prefers to be less than perfectly informed.[13] This preference arises because as the total potential surplus shrinks with a, the principal is more willing to implement output distortions to limit the rents of the informed agent. These distortions become particularly severe as the accuracy of the agent's information (γ) increases towards unity. Hence, to limit the distortions and the associated reduction in expected rents, the agent prefers his private signal to reflect his actual costs imperfectly. Thus, the principal and agent do not necessarily have preferences that are diametrically opposed regarding the informational environment in which they operate.

Example 3. $V(Q) = -cQ \quad \forall Q \geqq 0$, where $c > 0$; (4.7)

$$U(Q, t) = \frac{t}{\delta} Q^{\delta} \quad \forall Q \geqq 0, \text{ where } \delta \neq 0 \text{ and } \delta < 1\,; \tag{4.8}$$

$$F(t) = \frac{1}{\bar{t} - \underline{t}} \quad \forall t \in [\underline{t}, \bar{t}]\,. \tag{4.9}$$

Example 3 considers a setting where the principal is a monopolist and the agent is a consumer. The monopolist's production costs are given in (4.7). The consumer with taste parameter t values Q units of the monopolist's product according to (4.8). $\epsilon = \frac{1}{1-\delta}$ is the elasticity of demand, so higher values of δ correspond to more elastic demand. The consumer's taste parameter, t, follows a uniform distribution, as reflected in (4.9). Higher realizations of t reflect higher valuations of Q. The change of variables from θ to t is convenient for expositional purposes. Example 3 is consistent with the general formulation of the model if, for example, $t = \frac{1}{\theta}$. One can interpret Q as the number of units of the monopolist's product that the consumer purchases. Alternatively, the consumer might be assumed to always purchase one unit of the product, and Q could be viewed as the level of quality imbedded in the product. We adopt the latter interpretation in the ensuing discussion. To emphasize the range of applications of our model, we also assume there is a continuum of consumers, with individual taste parameters distributed as in (4.9).

The central question here is how well informed about their own preferences for her product does the monopolist wish consumers to be. If consumers are very unsure about their personal valuations of quality (as when consumers are not certain how often they would substitute a new cooking utensil for other kitchen aids), the monopolist need not be concerned with high-valuation consumers pretending to be low-valuation consumers when choosing how much quality to purchase. Thus, the rents of high-valuation consumers can be limited. On the other hand, when consumers are well informed about their personal valuation of product quality (perhaps through in-home demonstrations, for example), the monopolist may be better able to tailor the price of its product to the imbedded quality level so as to capture some of the surplus that might otherwise flow to the high-valuation consumers.

Tedious calculations and simulations demonstrate that the monopolist prefers consumers to be perfectly informed about their tastes when their elasticity of demand is sufficiently high (i.e., $V^*(1) > V^*(0)$ for $\delta \in (0, 1)$), and to share the principal's imperfect knowledge of t when their demand elasticity is sufficiently low (i.e., $V^*(0) > V^*(1)$ for $\delta < 0$).[14] When consumer demand is sufficiently elastic, lucrative opportunities for price discrimination dominate the monopolist's calculus, and she prefers to sell to customers with perfect knowledge of their preferences. These findings suggest a monopolist may be more likely to implement a greater variety of product qualities and provide detailed product information the more elastic is consumer demand for product quality. An interesting welfare implication also emerges. The price discrimination that is implemented when consumers are certain of their personal valuations of quality (coupled with the provision of information to consumers) is Pareto improving. The monopolist is clearly better off when she chooses to provide information to consumers and price discriminate. Each consumer also gains because regardless of his personal taste for quality, he receives nonnegative net utility from his purchase decision. When they share the principal's imperfect knowledge of t, consumers receive only zero in expected net utility.

5. Information Control by the Agent.

We now turn to the possibility that the agent may, at some personal cost, be able to influence the information structure under which he operates. We assume that if the agent undertakes no effort to improve the accuracy of his signal, his private signal will convey no additional information about the environment (i.e., $\gamma = 0$). However, at personal cost $K(\gamma) > 0$, the agent can improve the accuracy of his private signal to $\gamma \in (0, 1]$. $K(\gamma)$ is assumed to increase with γ at a nondecreasing rate, i.e, $K'(\gamma) > 0$ and $K''(\gamma) \geqq 0 \ \forall \gamma > 0$.

Because the principal cannot observe the realized accuracy of the agent's private signal (or, equivalently, the costs incurred by the agent in increasing γ), she cannot simply dictate how accurate a signal the agent must select. However, the principal can influence the agent's self-interested choice of γ through her design of the incentive scheme. The focus in this section is on how this influence is reflected in the optimal incentive scheme.

The timing in this setting is as follows. First, the principal commits to the terms of the incentive scheme $\{Q(s), T(s)\}$. Second, the agent selects $\gamma \in [0, 1]$ by incurring personal cost $K(\gamma)$. Third, the agent observes the realization of his private signal, $s \in [\underline{s}, \overline{s}]$, and selects his most preferred $(Q(s), T(s))$ option by making a report about s to the principal. Finally, "production" occurs and the promised payment is delivered to the agent by the principal.

Formally, the principal's problem, [P-K], is the following:

$$\underset{Q(s),\, T(s),\, \gamma}{\text{Maximize}} \quad \int_{\underline{s}}^{\overline{s}} [V(Q(s)) - T(s)]dG(s)$$

subject to $\forall s, \hat{s} \in [\underline{s}, \overline{s}]$:

$$\pi(s|s;\gamma) \geqq 0\,, \tag{5.1}$$

$$\pi(s|s;\gamma) \geqq \pi(\hat{s}|s;\gamma)\,, \quad \text{and} \tag{5.2}$$

$$\gamma \in \underset{\gamma'}{\operatorname{argmax}} \left\{ \int_{\underline{s}}^{\bar{s}} \pi(\tilde{s}\,|\,\tilde{s};\gamma')\,dG(\tilde{s}) - K(\gamma') \right\}, \tag{5.3}$$

where $\pi(\hat{s}\,|s;\gamma) \equiv T(\hat{s}) + \int_{\underline{\theta}}^{\bar{\theta}} U(Q(\hat{s}),\theta)\,[\gamma dF(\theta\,|\,s) + (1-\gamma)dF(\theta)].$

Relations (5.1) and (5.2) are, respectively, the standard individual rationality and incentive compatibility constraints. Relation (5.3) is the accuracy selection constraint, reflecting the agent's self-interested choice of γ given the incentive scheme offered by the principal. In the ensuing analysis, we let λ denote the Lagrange multiplier associated with the accuracy selection constraint, (5.3). It can be verified that if λ is positive (negative) at the solution to [P-K], then the principal's expected net benefits would increase if the agent selected a slightly higher (lower) value of $\gamma \in (0, 1)$.

Proposition 3. Suppose the second-order conditions are satisfied.[15] Then at an interior solution to [P-K]:

$$\int_{\underline{s}}^{\bar{s}} \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s),\theta)[dF(\theta\,|s) - dF(\theta) + dF_2(\theta\,|\,s)\frac{G(s)}{g(s)}]\,ds - \lambda K''(\gamma) = 0; \text{ and} \tag{5.4}$$

$$V'(Q(s)) + \int_{\underline{\theta}}^{\bar{\theta}} U_1(Q(s),\theta)[\gamma dF(\theta\,|\,s) + (1-\gamma)dF(\theta)]$$
$$- [\lambda-\gamma] \int_{\underline{\theta}}^{\bar{\theta}} U_1(Q(s),\theta)dF_2(\theta\,|\,s)\,\frac{G(s)}{g(s)} \qquad \forall\, s \in [\underline{s}, \bar{s}]\,. \tag{5.5}$$

Equation (5.5) reveals how the principal will alter the incentive scheme to influence the agent's choice of γ.[16] If the principal wishes to limit the agent's incentive to improve the accuracy of his signal, the principal induces more severe distortions in the agent's activity level. The usual wedge between the marginal value of the activity and its expected marginal cost given signal s is increased by $|\lambda \int_{\underline{\theta}}^{\overline{\theta}} U_1(Q(s), \theta)\, dF_2(\theta | s) \frac{G(s)}{g(s)}|$. To best understand the role of the increased distortion, return to the procurement setting of Example 2 where the principal prefers the agent to share her imperfect cost information. In that setting with $\gamma = 0$, the agent will produce a fixed output level, $Q_o > 0$, in return for a fixed payment, T_o, that provides zero rents in expectation. Notice that although the agent's realized rents will be strictly positive when his marginal cost of production, θ, turns out to be sufficiently small, he will suffer strict losses for the higher realizations of θ. Consequently, the agent will have an incentive to acquire a more accurate signal about θ under this incentive scheme. When a high realization of the signal is observed, indicating that a high θ realization is likely, the agent can refuse to produce Q_o, thereby avoiding a large loss. To limit the value of a more accurate signal to the agent, the principal can reduce the output the agent is asked to produce when he reports a high realization of this signal. The smaller output reduces the difference in total production costs associated with any two distinct realizations of marginal cost, and thereby reduces the maximum loss the agent could suffer by agreeing to produce. In this way, the value to the agent of a more accurate signal is diminished.

One implication of this argument is that in equilibrium, the delivered output may vary with the realization of the agent's private signal even though this signal provides no information about the likely realization of θ (i.e., even though $\gamma = 0$). This possibility will emerge when: (i) $K(\gamma)$ is sufficiently small for small γ that the agent will choose $\gamma > 0$ under the (Q_o, T_o) contract; and yet (ii) the principal finds it optimal to induce $\gamma = 0$ in the solution to [P-K].[17] In this case, the principal must commit herself to an incentive scheme in which induced output declines with the realization of the agent's private signal in order to dissuade the agent from improving the accuracy of his private signal above zero.

To provide a bit more insight regarding the agent's choice of γ in the setting of Example 2, consider Corollary 2.

<u>Corollary 2.</u> In the setting of Example 2, suppose $\bar{\theta} = 1$, $K(\gamma) = \frac{1}{2}k\gamma^2$, and $k > \frac{a}{2b}$.

Then at the solution to [P-K]:

(i) $\gamma \in (0, 1)$;

(ii) $\lambda < 0$;

(iii) $\frac{d\gamma}{dk} < 0$;

(iv) $\frac{d\gamma}{da} > 0$; and

(v) $\frac{d\gamma}{db} < 0$.

Property (iii) of Corollary 2 reports that when it becomes less costly for the agent to improve the accuracy of his private signal, he will purchase a more accurate signal. Properties (iv) and (v) state that γ will also increase as the total potential surplus rises, either through an increase in the intercept (a) or a decrease in the slope (b) of the principal's demand curve. With higher levels of total surplus available, the agent's information becomes more valuable for planning purposes, and the rent he commands from his private information becomes relatively less problematic. Hence, the principal reduces the quantity distortions that would otherwise induce the agent to choose a smaller

γ. Property (i) states that provided the agent's marginal cost of increasing γ rises sufficiently rapidly with γ (i.e., provided $k > \frac{a}{2b}$), the agent's private information will not be perfect (i.e., $\gamma < 1$). Property (i) also reports that the agent will always improve the accuracy of his signal above zero, since the marginal cost of doing so is small ($K'(0) = 0$). Property (ii) indicates the principal would prefer a smaller γ at the solution to [P-K], as should be expected in the setting of Example 2.

More generally, there are also settings where, at the solution to [P-K], the principal's expected net benefits would increase if the agent selected a higher value of γ (i.e., where $\bar{\lambda} > 0$). In such settings, the principal *reduces* the usual distortions in the agent's activity level, as is evident from equation (5.5). The expanded output in the procurement setting, for example, makes the agent more willing to pay to improve the accuracy of his private signal, and thereby avoid producing large output levels for relatively meager compensation when costs are likely to be higher. It is even possible to have $\lambda \geqq \gamma$ at the solution to [P-K], so that (from equation (5.5)) the agent is induced to produce at or above the efficient level of output for all realizations of his private signal.[18] Thus, important qualitative differences in the optimal incentive scheme can arise when the agent has control over the accuracy of his private information.

6. Conclusions.

We have identified a class of problems in which a principal who has complete control over the informational environment will always choose an extreme one. She will either ensure that her agent has no additional information or that he has the best available knowledge about the random state of nature. The former choice is preferred when additional information about the state is not very important for planning purposes, and when the private information leads to significant increases in the agent's expected rents. The principal will prefer her agent to have the best available state information when the information is valuable for planning purposes and does not afford the agent significant rents.

When the agent has control over the informational environment, the principal will generally alter the incentive scheme she offers to the agent. To motivate the agent to choose a less accurate private signal, the principal will induce more severe distortions in the agent's activity level. The reduced activity lowers the agent's perceived gains from more accurate information, inducing him to spend less on increasing γ. In contrast, larger activity levels are induced when the principal finds it advantageous to have the agent acquire a more accurate private signal. In fact, the optimal incentive scheme may induce the agent to undertake activity in excess of its efficient level for all realizations of the agent's private signal.

The principal's preference for extreme information structures was shown to derive principally from two assumptions: (1) both the principal and agent are risk neutral; and (2) conditional beliefs about the state of nature under all feasible information structures are linear combinations of beliefs under the most and least informative information structures. The first assumption is standard in the agency literature. The second assumption is special, but does incorporate plausible settings, including Blackwell rankings of information structures. Investigation of the optimal management of the informational environment in alternative settings appears worthwhile.

The same is true of studies that explore alternative formulations of the relationships among the state of nature, the activity level, and the utility functions and policy instruments of the principal and agent. As just one illustration, suppose the agent is a regulated monopoly that may acquire private information about consumer demand for the product it sells. In this setting, where the critical information asymmetry concerns the agent's revenues rather than his costs, it is possible to show that the aforementioned extreme preferences of the principal persist under some conditions, but for different reasons. Lewis and Sappington (LS) (1988) have shown that in this "unknown demand" model, the agent cannot command any private rents from perfect knowledge of demand when marginal production costs increase with output. Therefore, the principal will always prefer that the agent acquire perfect demand information, since the information is valuable for planning purposes and requires no sacrifice of information rents.[19] On the other hand, LS have also shown that when marginal production costs decline with output, the optimal policy for the principal generally makes no use of the agent's perfect private knowledge of demand. Therefore, because private knowledge of demand does increase the agent's rents, the principal will prefer that the agent share her imperfect knowledge.[20]

One other direction for future research concerns the dimensionality of the information asymmetry. In the monopoly seller interpretation of our model, for example, we assumed the only possible information asymmetry concerned a single parameter that reflected a consumer's marginal valuation of quality or quantity. More generally, there may also be a fixed component to a consumer's valuation of quality or quantity, and the seller or consumers may be able to influence the consumers' knowledge of this component of tastes also. In such expanded settings, interesting questions arise concerning the *type* of information the seller would like her consumers to have.[21] A satisfactory answer to these questions could provide insight concerning the optimal advertising and promotion policies of firms.

FOOTNOTES

1. Some thoughts along these lines are presented in Demski and Sappington (1986).

2. Craswell (1988) examines the design of legal rules to motivate an agent to acquire the socially optimal amount of valuable planning information.

3. To prove this conclusion formally, let $h(s|\theta)$ denote the conditional probability of the informative signal. Note that $g^0(s)$, the conditional density of the uninformative signal, is independent of θ. Given our assumptions, Bayes formula reveals that $f(\theta|s, \gamma)$, the conditional distribution of θ given signal s and accuracy γ, is given by

$$\begin{aligned}
f(\theta|s, \gamma) &= [\gamma h(s|\theta) + (1-\gamma)g^0(s)]f(\theta)/[\gamma\int_{\underline{\theta}}^{\bar{\theta}} [h(s|\theta) + (1-\gamma)g^0(s)]dF(\theta) \\
&= [\gamma h(s|\theta) + (1-\gamma)g^0(s)]f(\theta)/[\gamma g(s) + (1-\gamma)g^0(s)] \\
&= [\gamma h(s|\theta) + (1-\gamma)g(s)]f(\theta)/g(s) \\
&= \gamma f(\theta|s) + (1-\gamma)f(\theta),
\end{aligned}$$

where $f(\cdot)$, $g(\cdot)$ and $g^0(\cdot)$ denote the density functions associated, respectively, with the distribution functions $F(\cdot)$, $G(\cdot)$ and $G^0(\cdot)$.

4. Although the implicit ranking of information structures we adopt is not without loss of generality, it does admit a convenient parameterization, and also coincides with other rankings of information structures in some special cases. For instance, when the unconditional distribution of the signal is uniform and the most informative information structure provides a perfect signal about θ, higher values of γ correspond to information structures that are more informative in the sense of Blackwell (1951).

5. As its proof suggests, Proposition 1 generalizes to the case where the agent's private signal is drawn from one of $N > 2$ distinct distributions.

6. In this procurement setting (see Examples 1 and 2 below for more details), $Q_\gamma(s)$ is determined by

$$V'(Q_\gamma(s)) = \gamma[\theta + \frac{F(\theta)}{f(\theta)}] + [1-\gamma] \int_{\underline{\theta}}^{\bar{\theta}} \theta dF(\theta).$$

Therefore,

$$\frac{dQ_\gamma(s)}{d\gamma} = \frac{\theta + \frac{F(\theta)}{f(\theta)} - \int_{\underline{\theta}}^{\bar{\theta}} \theta dF(\theta)}{V''(Q_\gamma(s))},$$

which is negative (positive) for the higher (lower) realizations of θ provided $V''(Q) < 0$ $\forall\, \theta$ and $\frac{d}{d\theta}\{\frac{F(\theta)}{f(\theta)}\} \geq 0$ $\forall\, \theta$. Notice, too, that an increase in γ in this setting raises the principal's realized welfare for given s at the rate

$$- Q_\gamma(s)[\theta + \frac{F(\theta)}{f(\theta)} - \int_{\underline{\theta}}^{\bar{\theta}} \theta dF(\theta)].$$

Thus, the principal's welfare is reduced when high signals are realized (the high signal effect) and increased when low signals are realized (the low signal effect).

7. Radner and Stiglitz (1981) include an exogenous cost of better information in their model of decision-making by a single individual. In our agency model, the cost of improved information is endogenous, arising from the increased rents the agent receives in equilibrium. Singh (1985) abstracts from cost considerations, but finds in a moral hazard model that the marginal value to the principal of imperfect information about the agent's effort level to be zero at the

point of no information.

8. The constant elasticity formulation for demand ensures that the elasticity of $Q^*(\theta)$ with respect to θ, $|Q^{*\prime}(\theta)\frac{\theta}{Q^*(\theta)}|$, increases linearly with η, where $Q^*(\theta) \equiv \underset{Q}{\operatorname{argmax}}\left\{V(Q) - \theta\ Q\right\}$.

9. To see this, notice that for $\alpha = 1$, θ follows the uniform distribution on $[0, \bar{\theta}]$, giving rise to a distribution function that increases linearly with θ. For $\alpha > 1$, the distribution function is an increasing, strictly convex function of θ, so less (more) weight is placed on the smaller (larger) realizations of θ. Conversely, the distribution function in (4.3) is an increasing, strictly concave function of θ for $\alpha \in (0, 1)$, reflecting more (less) weight on the smaller (larger) realizations of θ relative to the uniform distribution.

10. The simulations performed allowed α to vary between 0.01 and 25.0, in increments of .01.

11. Of course, $V^*(1) \geqq V^*(0)$ if $Q_o = 0$ in the setting of Example 2, and more generally. In words, suppose the principal would not induce any output from the agent when $\gamma = 0$ (since expected production costs are too high relative to the total expected surplus). Then the principal will prefer the agent to become perfectly informed because: (1) when high costs are realized, no output will be induced, and the principal and agent each achieve the same net benefits as when $\gamma = 0$; and (2) when particularly low costs are realized, Pareto gains may be possible from strictly positive production levels.

12. In fact, this conclusion generalizes to the setting where the agent's marginal cost of production follows a generalized uniform distribution (i.e., where relation (4.6) is replaced by (4.3)). To see this, first notice that pointwise maximization of (3.2) reveals

$$Q_\gamma(\theta) = \frac{1}{b}\left[a - \gamma\theta\frac{1+\alpha}{\alpha} - (1-\gamma)\frac{\alpha}{1+\alpha}\bar{\theta}\right].$$

Substituting this value back into (3.2) and differentiating with respect to γ reveals after some manipulation that

$$V^{*\prime}(\gamma) \stackrel{s}{=} -a + \bar{\theta}\left\{\gamma\frac{\alpha^2+3\alpha+1}{\alpha(2+\alpha)} + (1-\gamma)\frac{\alpha}{1+\alpha}\right\}.$$

This expression is negative provided $a > \bar{\theta}\,[\gamma\frac{1+\alpha}{\alpha} + (1-\gamma)\frac{\alpha}{1+\alpha}]$, which is necessary to ensure $Q_\gamma(\theta) > 0 \quad \forall\, \gamma, \theta$.

13. The value of γ that maximizes the agent's expected rents is readily shown to be $\frac{3}{10}\,[2a - \bar{\theta}]$ for $a \in [2\bar{\theta}, \frac{13}{6}\bar{\theta}]$.

14. Standard techniques reveal that for $\underline{t} = z\bar{t}$, $V^*(1) \gtreqless V^*(0)$ as $\frac{1 - (2z-1)^{\frac{2-\delta}{1-\delta}}}{1-z}\,[1-\gamma]/[2-\gamma] \lesseqgtr [\frac{1+z}{2}]^{\frac{1}{1-\delta}}$. To ensure $Q_\gamma(t) > 0 \quad \forall\, t \in [\underline{t}, \bar{t}]$, we require $z \in (.5, 1)$. Therefore, for $\delta < 0$, $(\frac{1+z}{2})^{\frac{1}{1-\delta}} < 1$, $\frac{1-\delta}{2-\delta} > \frac{1}{2}$, and $(2z-1)^{\frac{2-\delta}{1-\delta}} < 2z-1$. Together these facts reveal $V^*(0) > V^*(1) \quad \forall\, \delta < 0$. Next, letting $\rho = \underline{t}\,/\,\bar{t}$, simulations that allow ρ to vary between .5 and .9 in increments of .1 reveal $V^*(1) > V^*(0)$ for values of δ that vary between .1 and .9 in increments of .1.

15. The second order conditions referred to here are conditions sufficient to ensure: (1) the global incentive compatibility constraints are satisfied at the identified solution to [P-K] when only the local constraints are imposed; and (2) the accuracy selection constraint (5.3) is satisfied at the identified solution to [P-K] when (5.3) is replaced by the requirement that the

derivative of the expression in (5.3) with respect to γ be zero. An example of such sufficient conditions is presented in Corollary 2 below.

16. Equation (5.4) defines the agent's induced choice of γ at an interior solution. The left-hand side of (5.4) could exceed the right-hand side at a boundary solution where $\gamma = 1$.

17. Formally, suppose $\alpha = 1$, $\bar{\theta} = 10$ and $e < 0.33$ in the setting of Example 1, so that $V^*(0) > V^*(1)$. Also, suppose $K(\gamma) = k\gamma \quad \forall\ \gamma \in [0, 1]$, and define $\tilde{k} = \int_0^{\bar{\theta}} \tilde{Q}\, \frac{\theta}{\bar{\theta}}\, d\theta$ where $V'(\tilde{Q}) = \frac{\bar{c}}{2}$. If $k = \tilde{k} - \beta$, then for $\beta > 0$, (5.3) reveals that the agent will choose $\gamma = 1$ if $Q(\theta) = \tilde{Q} \quad \forall\ \theta \in [0, \bar{\theta}]$. However, since $V^*(0) > V^*(1)$, the presumed linear structure of $K(\cdot)$ ensures it will be optimal for the principal to induce $\gamma = 0$ if β is sufficiently small. She does so by implementing the quantity schedule, $\hat{Q}(\theta)$ for which $k = \int_0^{\bar{\theta}} \hat{Q}(\theta) \frac{\theta}{\bar{\theta}}$, and $V'(\hat{Q}(\theta)) = \frac{\bar{\theta}}{2} - \lambda \frac{F(\theta)}{f(\theta)} \quad \forall\ \theta \in [\underline{\theta}, \bar{\theta}]$.

18. To derive conditions which ensure $\lambda \geqq \gamma$ in the solution to [P-K], return to the setting of Example 1 where $\alpha = 1$, and let $K(\gamma) = k\gamma$. From (5.5), define $Q_\lambda(\theta)$ by $V'(Q_\lambda(\theta)) = \theta + [1 - \lambda] \frac{F(\theta)}{f(\theta)} \quad \forall\ \theta \in [\underline{\theta}, \bar{\theta}]$. It is straightforward to verify that for $Q(\theta) = Q_\lambda(\theta)$ and $\gamma = 1$, the left-hand side of (5.4) is equal to $\frac{1}{\bar{\theta}}[2-\lambda]^{-\eta} \left[\frac{\bar{\theta}\ \theta^{1-\eta}}{2[1-\eta]} - \frac{2\ \theta^{2-\eta}}{2-\eta} \right]_0^{\bar{\theta}}$. This expression is infinitely large for $\lambda \in (1.5, 2)$, so the solution to [P-K] will have $\gamma = 1$ (by complementary slackness) and $\lambda \geqq 1$ if $k \leq \left[[2-\lambda]^{\eta}\ [2-\eta]\ \bar{\theta}^{\eta-1} \right]^{-1}$ for the specified $\lambda \geqq 1$, using (5.3).

19. In fact, more accurate information may even reduce the agent's information rents in this setting. One can show that under certain conditions, the agent can command rents from a private signal that is more accurate than the principal's but is not perfect.

20. It remains to determine whether the regulator will optimally employ the firm's *imperfect* private knowledge of demand. If so, then the possibility remains that the principal might prefer the agent to acquire some imperfect private demand information.

21. The corresponding extension to the procurement setting would involve uncertainty about both the marginal and fixed costs of production. Since private knowledge of fixed costs will provide rents to the agent without providing information that is useful in planning the optimal level of production, a principal with elastic demand will likely prefer an improvement in the agent's knowledge of marginal rather than fixed production costs.

APPENDIX

Proof of Proposition 3.

Using (3.1), standard techniques which involve integration by parts ensure the that the solution to [P-K] is identified by maximizing the following Lagrangian function with respect to γ and (pointwise) with respect to Q(s).

$$L = \int_{\underline{s}}^{\bar{s}} \Big\{ V(Q(s)) + \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s), \theta)[\gamma dF(\theta|s) + (1-\gamma)dF(\theta) + \gamma dF_2(\theta|s)\frac{G(s)}{g(s)}] \Big\} dG(s)$$

$$- \lambda \left\{ \int_{\underline{s}}^{\bar{s}} \int_{\underline{\theta}}^{\bar{\theta}} U(Q(s), \theta) dF_2(\theta|s) G(s) ds - K'(\gamma) \right\}.$$

Equations (5.4) and (5.5) follow immediately from differentiation of L.

■

Proof of Corollary 2.

Substitution and manipulation of equations (5.4) and (5.5) in the identified setting reveal

$$\lambda = - \frac{2[6a - 3][6bk - 1]}{[12bk + 1][12bk + 5] + 48bk - 8},$$

$$\gamma = \frac{[6a - 3][12bk + 1]}{[12bk + 1][12bk + 5] + 48bk - 8}, \quad \text{and}$$

$$Q(c) = \frac{1}{b}\left[a - \frac{1}{2}(1 - \gamma) + [\lambda - 2\gamma]\theta\right].$$

Straightforward calculations then reveal that $\lambda < 0$ and $\gamma \in (0, 1)$ provided $bk > \frac{a}{2} > 1$. This condition also ensures $Q(\theta) > 0$ and $Q'(\theta) \leqq 0 \quad \forall\, \theta \in [0, 1]$. Finally, differentiation immediately provides properties (iii) - (v) in the Corollary.

■

REFERENCES

Baron, David and Roger Myerson, "Regulating a Monopolist with Unknown Costs," *Econometrica*, 50, 1982, 911-930.

Blackwell, David, "Comparison of Experiments," in J. Neyman (ed.), *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*. Berkeley, University of California, 1951, 93-102.

Clay, Karen, David Sibley, and Padmanabhan Srinagesh, "Ex Post vs. Ex Ante Pricing: Optional Calling Plans and Tapered Tariffs", Bellcore technical memorandum, January 1991.

Craswell, Richard, "Pre-Contractual Investigation as an Optimal Precaution Problem," *Journal of Legal Studies*, 17, June 1988, 401-436.

Cremer, Jacques, and Fahad Khalil, "Gathering Information Before Accepting a Contract," Virginia Polytechnic Institute mimeo, May 1990a.

Cremer, Jacques and Fahad Khalil, "Gathering Information Before the Contract is Offered," Virginia Polytechnic Institute mimeo, October 1990b.

Demski, Joel and David Sappington, "On the Timing of Information Release," *Information Economics and Policy*, 2, 1986, 307-316.

Holmstrom, Bengt, "Moral Hazard and Observability," *Bell Journal of Economics*, 10, Spring 1979, 74- 91.

Maskin, Eric and John Riley, "Monopoly with Incomplete Information," *Rand Journal of Economics*, 15, Summer 1984, 171-196.

Radner, Roy and Joseph Stiglitz, "A Nonconcavity in the Value of Information," in Marcel Boyer and Richard Kihlstrom (eds.), *Bayesian Models in Economic Theory*. Amsterdam: North-Holland, 1984.

Singh, Nirvikar, "Monitoring and Hierarchies: The Marginal Value of Information in a Principal-Agent Model," *Journal of Political Economy*, 93, June 1985, 599-610.

Sobel, Joel, "Information Control in the Principal-Agent Problem," University of California at San Diego mimeo, 1983.