

A Principal-Agent Model of Sequential Testing*

Dino Gerardi[†] Lucas Maestri[‡]

July 2011

Abstract

This paper analyzes the optimal provision of incentives in a dynamic information acquisition process. In every period, the agent can acquire costly information that is relevant to the principal's decision. Each signal may or may not provide definitive evidence in favor of the good state. Neither the agent's effort nor the realizations of his signals are observable. First, we assume that the agent has no private information at the time of contracting. Under the optimal mechanism, the agent is rewarded only when his messages are consistent with the state. The payments that the agent receives when he correctly announces the good state increase over time. We then characterize the optimal mechanisms when the agent has superior information about the state at the outset of the relationship. The principal prefers to offer different contracts if and only if the agent types are sufficiently diverse. Finally, all agent types benefit from their initial private information.

KEYWORDS: Dynamic Mechanism Design, Information Acquisition, Sequential Testing.

JEL CLASSIFICATION: C72, D82, D83.

*We are grateful to Bart Lipman, two anonymous referees, Luca Anderlini, Dirk Bergemann, James Choy, Johannes Hörner, Rick Mansfield, Ben Polak, Larry Samuelson, Juuso Välimäki, and numerous seminar audiences for helpful comments.

[†]Collegio Carlo Alberto, Via Real Collegio 30, 10024 Moncalieri, Italy. dino.gerardi@carloalberto.org.

[‡]Toulouse School of Economics, Manufacture des Tabacs, Aile Jean-Jacques Laffont MF 52421, allée de Brienne 31000, Toulouse, France. maestri.lucas@gmail.com.

1 Introduction

In many situations the power to make a decision and the ability to acquire relevant information do not reside in the same place. Firms, and more generally, decision makers, routinely consult experts who spend time and energy to determine the best course of actions.

Consider a pharmaceutical company developing a new drug for a certain disease. Before starting to sell the drug, the company wants to know if the drug leads to a side-effect which makes it inappropriate for patients with some condition (for example, high-blood pressure). Producing the drug is always advantageous, but informing the public about side-effects is desirable to avoid legal complaints. The company signs a contract with a team of scientists to perform experiments and assess the possibility of this side-effect. Because of physical constraints, the scientists cannot test more than a fixed number of patients at a time. When they report evidence of a side-effect, the company stops the testing procedure and starts selling the drug, warning the patients about the possible consequences of taking it. On the other hand, if the scientists do not find evidence of a side-effect after a certain number of tests, the company sells the drug with no warning.

A number of features are important in this example. First, the information acquisition process is dynamic. Physical or technological constraints impose limits on the number of patients that a laboratory can test in a given period, say a week. Hence, it may take several months before the scientists either find evidence of a side-effect or they recommend to quit the testing process. Second, the pharmaceutical company may lack the means or knowledge to monitor the team in performing the tests. In particular, the scientists may choose to save the cost of performing the tests and announce that the drug does (or does not) lead to a side-effect. Finally, the pharmaceutical company and the team of scientists may not be equally informed at the beginning of their relationship. Because of their past experience conducting similar experiments, the scientists may know facts that the company ignores or they may interpret the public information in a more sophisticated way. Therefore, it is reasonable to assume that they start the testing process with superior information.

In spite of the presence of various forms of informational asymmetries, a principal can still motivate an agent to invest in information acquisition and share his discoveries if after the principal makes a decision, some information becomes publicly available. In the above example, the presence or absence of side-effects will become evident after the drug is on the market for some time.

The goal of this paper is to study how the principal can use this information to overcome the problems of moral hazard and adverse selection present in this context. In particular, we analyze a dynamic mechanism design problem and characterize the principal's optimal contract. We also investigate how the different sources of private information affect the

agent's ability to extract a rent from the principal.

In the benchmark model, a principal who has to make a risky decision hires an agent to acquire information about the unknown state of the world. The state can be either good or bad. The principal and the agent are equally informed at the outset of the relationship. The agent can complete at most one test in every period and each test generates an informative (binary) signal about the state. One realization of the signal can be observed only when the state is good, while the other realization is possible under both states. Neither the agent's effort nor the signals are observable (or verifiable). After the principal makes his decision, the state is revealed.

The principal has the ability to commit to a contract which specifies both the length of the relationship and all the possible payments to the agent. The payments depend on the messages that the agent sends to the principal and on the state of the world. The agent is protected by limited liability and cannot make transfers to the principal. The principal chooses the contract to maximize his expected payoff. His goal is to determine the optimal length of the relationship and to offer the cheapest contract that induces the agent to acquire the signal and reveal it truthfully in every period until the deadline or until the agent finds definitive evidence in favor of the good state, whichever comes first.

We solve the principal's problem in two steps. First, we fix the deadline and construct the optimal contract with a given length. We then study the optimal length of the relationship. An incentive compatible contract must prevent different types of deviations. In particular, the agent may lie about the realizations of his signals. By controlling the release of information, the agent therefore decides when to terminate the relationship with the principal. If later payments are sufficiently generous, the agent may decide to delay the announcement of a major finding. Furthermore, the agent can choose the pace of his testing in the sense that he can shirk in one or several periods. This leads to asymmetric beliefs about the state. To see this, consider the following deviation. Suppose that in a certain period the agent shirks and claims to have evidence in favor of the bad state. Being unaware that the agent's message carries no informational content, the principal updates his belief. Consequently, his belief that the state is good is lower than the agent's.

We characterize the optimal contract with a fixed length. The agent is rewarded only when his messages are consistent with the state. The payments that the agent receives when he correctly reports the good state increase over time. As time passes, the agent becomes more pessimistic about the arrival of evidence in favor of the good state. Consequently, larger payments are necessary to motivate him to exert effort. At the same time, the *discounted* values of these payments are decreasing over time. This induces the agent to announce the good state as soon as he finds definitive evidence in its favor.

We show that the agent's information rent can be decomposed into two components.

The first component is due to the presence of moral hazard, i.e. the fact that the principal cannot monitor the agent’s effort. The second one is due to the presence of hidden information, i.e. the fact that the results of the tests are unobservable.¹ We investigate how the various parameters of the problem affect the two rents. In particular, the moral hazard rent is increasing in the quality of the signal, while the hidden information rent is decreasing. When the signal becomes more precise the agent’s belief that the state is good deteriorates more quickly over time, leading to larger moral hazard rents. On the other hand, when the precision of the signal is high it is very risky for the agent to guess that the state is good, leading to lower hidden information rents. Consequently, the model predicts a non-monotonic relationship between information rents and signal precision.

Next, we endogenize the length of the relationship and allow the principal to choose the deadline. We show that the optimal deadline is (generically) unique and that agency problems shorten the information acquisition process.

In the second part of the paper, we analyze the case in which the agent has superior information about the state at the time of contracting. For tractability, the agent has one of two types. The agent’s type represents his belief that the state is good and can be either high or low. We also assume that the length of the relationship is exogenous and equal across types. We derive the optimal contract and disentangle the agent’s information rents into three components due to moral hazard, hidden information and adverse selection.

In contrast to many models of adverse selection in which the principal is able to extract all the rents from a certain type (see, among others, Mussa and Rosen (1978), and Baron and Myerson (1982)), in our model both types strictly benefit from the fact that their initial type is private information. The additional rents of the two types are of different nature. The low type obtains a larger payment when the state is bad while the payments in the good state coincide with the payments of the optimal contract in the benchmark model. On the other hand, the payments to the high type are front-loaded and distorted upwards when the state is good.

Our paper has elements in common with both the literature on learning in dynamic agency and the literature on delegated expertise. Within the former literature, Bergemann and Hege (1998, 2005), and Hörner and Samuelson (2010) are particularly related to our work. These papers study the dynamic provision of venture capital when the quality of a project is unknown to the entrepreneur and to the venture capitalist. The successful completion of the project depends both on its quality and on the financing it receives. The

¹We follow Holmstrom and Milgrom (1987) and use the term “hidden information” to emphasize the fact that the informational asymmetry about the realizations of the signals arises after the contract is signed. In contrast, in standard models of adverse selection the agent is privately informed at the time of contracting.

information structure is similar to ours. Either the project is completed or the parties become more pessimistic about it. Our paper is different in two important aspects which reflect the nature of delegated expertise problems. First, we assume that the realization of the signal acquired by the agent is private information. Hence, the principal must provide incentives to the agent to reveal his information. Furthermore, we investigate the possibility that the agent has superior information at the time of contracting.

Within the literature on delegated expertise, Laffont and Tirole (1986) analyze a static delegation model in which the agent has superior information at the time of contracting. Lewis and Sappington (1997), and Cremer, Khalil, and Rochet (1998) consider costly information acquisition. Gromb and Martimort (2007) derive many insights by studying optimal contracting with many agents when they can collude with one another. Our paper differs from these papers because we focus on the interplay between information acquisition and dynamics, which is not central to these models. Lewis and Ottaviani (2008) study a dynamic delegation model in which the principal offers short-term contracts to motivate the agent to search for innovations. Their main insights are related to the ability of the agent to conceal discoveries, slowing down the learning process. In our model, the agent also wants to manipulate the belief evolution of the principal, but concealing discoveries does not play an important role.

Our paper is also related to the literature on dynamic expertise. Olszewski and Peski (2010) and Klein and Mylovanov (2010) analyze models in which the expert has private information about the precision of his signal. The principal's goal is to keep the experts with high precision and release the other ones. Our work differs from theirs since we focus our attention on costly information acquisition and abstract from informational asymmetries about the expert's precision. Manso (2010) studies how to optimally motivate an agent to work on tasks of unknown payoffs. His work focuses on the tension between exploitation and exploration (absent in our model), while we focus on dynamic incentives for information acquisition.

The paper is organized as follows. In Section 2, we describe the benchmark model and characterize the optimal mechanism with a fixed length. In Section 3, we investigate the optimal length of the relationship. In Section 4, we extend the analysis to the case in which the agent has private information at the outset of the relationship. Section 5 concludes. All the proofs are relegated to the Appendix.

2 The Model with Symmetric Initial Information

A risk neutral principal has to choose one of two risky actions: $A = B, G$. The payoff $U(A, \omega)$ of each action depends on the binary state of the world $\omega \in \{B, G\}$. The principal's

preferred action in state $\omega = B, G$ is $A = \omega : U(G, G) > U(B, G)$ and $U(B, B) > U(G, B)$. The prior probability that the state is $\omega = G$ is $p_0 \in (0, 1)$.

In period 0 the principal hires an agent to perform a number of tests. In this section, we assume that the agent has no private information at the outset of the relationship.

The agent can perform at most one test in each period. Performing a test is costly and we let $c > 0$ denote the cost of a single test. Every test generates an informative but noisy signal s about the state. The signal takes the value $s = B, G$ and has the following distribution:

$$\begin{aligned}\Pr(s = G|\omega = G) &= \alpha, \\ \Pr(s = G|\omega = B) &= 0,\end{aligned}$$

where $\alpha \in (0, 1)$ denotes the quality of the signal. Thus, signal G provides definitive evidence in favor of state G . Conditional on the state, the signals are independent across periods.

For every $t = 0, 1, \dots$, we denote by

$$p_t = \frac{p_0(1-\alpha)^t}{p_0(1-\alpha)^t + 1 - p_0} \tag{1}$$

the agent's belief that the state is $\omega = G$ if he observes t signals equal to B .

The two actions of the agent are denoted by e (acquiring the signal) and ne (not acquiring the signal). The agent's effort decision (whether he chooses e or ne) and the realization of the signal are not observable. The state of the world is revealed after the principal makes his decision. The principal can commit to a long term contract (or mechanism) w , specifying sequential payments to the agent that are contingent on the agent's messages and on the state of the world (of course, a payment may depend on the state only if it is made after the principal chooses an action $A = B, G$ and the state is observed).

In this section, we assume that the length of the contract $T \geq 2$ is fixed,² and we analyze the optimal length of the relationship in Section 3. The objective of the principal is to design the optimal mechanism, i.e. the cheapest contract w that induces the agent to exert effort and reveal the realization of the signal in every period $t = 0, \dots, T - 1$, until he finds definitive evidence $s = G$ in favor of state G . In this case we say that the contract is incentive compatible.

The agent is risk neutral and has limited liability in the sense that the principal's payments must be non-negative. For simplicity, we assume that the agent has zero reservation utility. Both the principal and the agent have the same discount factor $\delta \in (0, 1]$.

²We consider the case $T = 1$ in Section 3.

We now turn to a formal definition of a contract. To keep the notation as simple as possible, we begin with a few observations. First, recall that the contract is designed to give incentives to the agent to exert effort in every period (until he observes signal G). It is therefore without loss of generality to assume that the set of messages available to the agent in every period is $\{B, G\}$.³

Second, we say that a contract is *evidence-based* if the principal makes a payment to the agent only when the relationship ends and the state is observed. As already mentioned, this happens as soon as the agent announces message G or in period $T - 1$ if the agent announces message B in every period. In other words, in an evidence-based mechanism all the intermediate payments (i.e., the payments that take place before the principal makes a decision) are equal to zero. Of course, the final payments of an evidence-based mechanism may depend on the realized state.

In our model, the principal and the agent are risk neutral and share the same discount factor. Given these assumptions, we can always construct an optimal mechanism that is evidence-based. Intuitively, consider a mechanism that specifies an intermediate payment $x > 0$ in period $t < T - 1$. Suppose now that the principal sets the intermediate payment equal to zero and increases all the final payments in period $t' = t + 1, \dots, T - 1$, by the amount $(1/\delta)^{t'-t} x$. That is, he takes away the intermediate payment from the agent and gives it back with the interests at the end of the relationship. Clearly, both the principal and the agent are indifferent between the two mechanisms.

Finally, suppose that the agent announces message G and the state turns out to be B . This event can occur only if the agent deviates and lies about the realization of his signal. Thus, it is without loss of generality to inflict the hardest punishment possible on the agent. Since the agent is protected by limited liability, this corresponds to setting the corresponding payment equal to zero. We say that a contract is *extreme* if all the payments that take place when the state is B and the agent reports message G are equal to zero.

In what follows we restrict attention to the class of evidence-based and extreme contracts.⁴ To simplify the exposition, we refer to them simply as contracts (or mechanisms). Formally, a contract w is the collection of the following payments:

$$w = \left((w(t))_{t=0}^{T-1}, w(G), w(B) \right).$$

Recall that the contract reaches period $t \geq 1$ only if the agent announces message B in

³Consider an arbitrary incentive compatible mechanism and let $M_t \supseteq \{B, G\}$ denote the set of available messages in period $t = 0, \dots, T - 1$. Consider now a modified mechanism with the set of messages $\{B, G\}$ in every period and with payments equal to the corresponding ones of the original mechanism. Clearly, the modified mechanism is also incentive compatible. This is simply because the set of deviations is smaller under the new mechanism than under the original one.

⁴We will briefly come back to this point at the end of Section 2.1.

every period $t' < t$. For $t = 0, \dots, T - 1$, $w(t)$ is the payment that the agent receives in period t if he announces message G and the state is G . For $\omega = B, G$, $w(\omega)$ is the payment that the agent receives in period $T - 1$ if he announces message B and the state is equal to ω .

To sum up, we analyze the following game. In period 0 the principal offers a contract w . Because the agent has zero reservation utility and limited liability, the participation constraints are automatically satisfied and the agent accepts the contract. Then, in each period the agent decides whether to exert effort (action e) or not (action ne) and sends a message from the set $\{B, G\}$. The game ends either when the agent announces message G or in period $T - 1$ if the agent always reports message B .

We can now define the agent's strategies. In every period, the agent observes the private history of acquisition decisions and signal realizations as well as the public history of reports. Clearly, there is only one public history that is relevant in period t . This is the history in which the agent announces message B in every period $t' < t$. We can therefore restrict attention to private histories (and ignore public histories).

Consider an arbitrary period. If the agent exerts effort then he can either observe signal B or signal G . On the other hand, if the agent shirks we say that he observes ne (his decision). Thus, for any $t > 0$, $H^t = \{ne, B, G\}^t$ is the set of private histories at the beginning of period t (or, equivalently, at the end of period $t - 1$). We set H^0 equal to the empty set.

We let σ denote an arbitrary (pure) strategy. Formally, $\sigma = (\sigma_t^A, \sigma_t^M)_{t=0}^{T-1}$, where

$$\begin{aligned}\sigma_t^A &: H^t \rightarrow \{e, ne\}, \\ \sigma_t^M &: H^{t+1} \rightarrow \{B, G\}.\end{aligned}$$

A strategy has two components: the action strategy and the message strategy. The first component $(\sigma_t^A)_{t=0}^{T-1}$ specifies the agent's decisions at the information acquisition stage. The second component $(\sigma_t^M)_{t=0}^{T-1}$ maps the private histories of the agents into reports to the principal. We let Σ denote the set of strategies available to the agents.

We denote by Σ^* the set of strategies under which the agent acquires the signal and reveals it truthfully in every period (on path). Formally, Σ^* is the set of all strategies σ such that: (i) $\sigma_0^A = e$; (ii) for every $t > 0$

$$\sigma_t^A(B, \dots, B) = e;$$

and (iii) for every $t \geq 0$ and every $s = B, G$

$$\sigma_t^M(B, \dots, B, s) = s.$$

Given a contract w , we let $u(\sigma, p_0; w)$ denote the agent's expected utility in period 0 if he follows strategy σ .⁵ Clearly, if σ and σ' are two strategies in Σ^* , then $u(\sigma, p_0; w) = u(\sigma', p_0; w)$. With a slight abuse of notation, we let $u(w) = u(\sigma, p_0; w)$ with $\sigma \in \Sigma^*$.

Then the principal's (linear programming) problem is given by

$$\begin{aligned} & \min_{w \geq 0} u(w) \\ \text{s.t. } & u(w) \geq u(\sigma, p_0; w) \text{ for every } \sigma \in \Sigma. \end{aligned} \tag{2}$$

A contract is incentive compatible if it satisfies all the constraints in (2). The optimal contract is the solution to the above problem.

2.1 The Optimal Mechanism with a Fixed Length

We start the analysis by characterizing the class of contracts that induce the agent to acquire and reveal the signal in every period. In principle, an incentive compatible mechanism has to satisfy a large number of constraints since the agent may shirk and lie in one or several periods. The next lemma simplifies the analysis dramatically. Lemma 1 below identifies a much smaller set of constraints which are necessary and sufficient to guarantee incentive compatibility.

In order to state our next result, we need to introduce some additional notation. Fix a contract w . Consider period $t = 0, \dots, T-1$, and suppose that the agent's belief (that the state is $\omega = G$) is equal to $p \in [0, 1]$. With another minor abuse of notation, we let $u(t, p; w)$ denote the agent's expected utility, computed in period t , when he acquires and reveals the signal in every period $t' \geq t$:

$$\begin{aligned} u(t, p; w) = & p\alpha w(t) + p(1-\alpha)\alpha\delta w(t+1) + \dots + p(1-\alpha)^{T-1-t}\alpha\delta^{T-1-t}w(T-1) + \\ & p(1-\alpha)^{T-t}\delta^{T-1-t}w(G) + (1-p)\delta^{T-1-t}w(B) - \\ & c \left[1 + \delta(p(1-\alpha) + 1-p) + \dots + \delta^{T-1-t} \left(p(1-\alpha)^{T-1-t} + 1-p \right) \right]. \end{aligned}$$

Notice that $u(0, p_0; w) = u(w)$. We also let $u(T, p; w)$, $p \in [0, 1]$, be given by:

$$u(T, p; w) = \frac{1}{\delta} [pw(G) + (1-p)w(B)]. \tag{3}$$

For notational simplicity, we drop the argument w in $u(t, p; w)$ and $u(\sigma, p_0; w)$ when there is no ambiguity.

⁵In the next section we consider the case in which the agent's initial belief is private information. For this reason it is convenient to make explicit the dependence of the agent's utility on the prior p_0 .

Lemma 1 *A contract w is incentive compatible if and only if it satisfies the following constraints:*

$$u(0, p_0; w) \geq p_0 w(0); \quad (4)$$

$$u(t, p_t; w) \geq \delta u(t+1, p_t; w), \quad t = 0, \dots, T-1; \quad (5)$$

$$w(t) \geq \delta w(t+1), \quad t = 0, \dots, T-2. \quad (6)$$

We say that the agent guesses state G in a certain period t if in that period he shirks and announces message G . Constraint (4) guarantees that the agent does not guess state G in period 0. Constraint (5) considers “one-period deviations.” The agent cannot find it profitable to shirk and announce message B in a single period (in all the remaining periods the agent acquires and reveals the signal). Finally, because of constraint (6), the agent does not want to delay the announcement of message G once he discovers that the true state is indeed G . Not surprisingly, an incentive compatible mechanism must necessarily satisfy constraints (4)-(6).

It is more interesting to see why constraints (4)-(6) provide sufficient conditions for incentive compatibility. First, constraints (6) together with constraint (5) for $t = T-1$ guarantee that it is optimal for the agent to tell the truth as soon as he discovers that the state is good. This, in turn, implies that it cannot be optimal to invest in information acquisition and lie after observing signal B . Intuitively, such a strategy is dominated by a strategy under which the agent sends message G without acquiring the signal.

Next, notice that under a “guessing” strategy the agent receives a positive payment if and only if the state is good. The fact that the discounted sequence of payments $\{\delta^t w(t)\}_{t=0}^{T-1}$ is (weakly) decreasing implies that if it is not profitable to guess state G in period 0, then it cannot be profitable to guess it in any other period $t > 0$.

We are therefore left with strategies under which the agent can do two things in every period. He can either acquire the signal and reveal it truthfully, or he can shirk and send message B . Constraint (5) is enough to prevent deviations under which the agent shirks in one or *several* periods.

To give some intuition, let us consider the agent in period t . Among the remaining discounted payments that the agent can receive in state G , $w(t)$ is the largest one. However, the agent can get this payment only if he acquires the signal in period t (recall that we have restricted attention to strategies under which the agent announces G only if he observes signal G). Consider two different scenarios. In the first scenario, the agent has acquired the signal in every period $t' < t$. Constraint (5), evaluated at t , implies that given the belief p_t the agent is willing to pay the cost c to have a chance to receive the payment $w(t)$. In the second scenario, the agent has shirked at least once before t . In period t , his belief is larger than p_t . Therefore, the agent is more optimistic that he will receive the large payment $w(t)$

than under the first scenario. In other words, in the second scenario the agent has a strict incentive to exert effort. Thus, we conclude that if it is not profitable for the agent to shirk once, then, *a fortiori*, it will not be profitable to shirk several times.

Given Lemma 1 we can restate the principal's problem as

$$\begin{aligned} & \min_{w \geq 0} u(w) \\ & \text{s.t. (4)-(6).} \end{aligned}$$

We let $w^*(p_0) = \left((w^*(t; p_0))_{t=0}^{T-1}, w^*(G; p_0), w^*(B; p_0) \right)$ denote the optimal mechanism when the prior is p_0 . Again, to simplify the notation we drop the argument p_0 in $w^*(p)$ when there is no ambiguity.

The next proposition characterizes the optimal contract. In the proposition and in the rest of the paper we adopt the following convention: if the lower bound of a summation is strictly larger than the upper bound then the summation is equal to zero.

Proposition 1 *The optimal contract is*

$$\begin{aligned} w^*(t) &= \frac{c}{\alpha p_t} + c \sum_{t'=t+1}^{T-1} \delta^{t'-t} \left(\frac{1}{p_{t'}} - 1 \right), \quad t = 0, \dots, T-1; \\ w^*(G) &= 0; \\ w^*(B) &= \frac{c}{\alpha(1-p_0)\delta^{T-1}} + c \sum_{t'=0}^{T-2} \delta^{-t'}. \end{aligned} \tag{7}$$

In the proof of Proposition 1, we solve a relaxed problem in which we ignore the non-negativity constraints and constraint (6). We show that the remaining constraints (4) and (5) must be binding and that it is optimal to set $w^*(G) = 0$. Therefore, the solution to the relaxed problem is the solution to a linear system with $T + 2$ equations (the constraints (4) and (5) and the equation $w(G) = 0$) and $T + 2$ unknowns ($w(0), \dots, w(T-1), w(G), w(B)$). The unique solution to the system is the contract w^* defined in equation (7) which satisfies the non-negativity constraints and, furthermore, has the following feature:

$$w^*(t) = \delta w^*(t+1) + \frac{c(1-\delta)}{p_t \alpha}$$

for $t = 0, \dots, T-2$. We conclude that the contract w^* satisfies constraint (6) and is optimal.

The optimal contract presents a number of interesting properties. The agent receives a positive payment only when his reports match the state (notice that $w^*(G) = 0$). Intuitively, the cheapest way to motivate the agent to acquire information is to reward him

in the states that support his announcements. The sequence of payments $\{w^*(t)\}_{t=0}^{T-1}$ that the agent receives when he correctly announces the good state is (weakly) increasing over time.⁶ As time passes without observing a good signal, the agent becomes more pessimistic about the arrival of evidence in favor of the good state. Consequently, the principal has to pay larger payments to motivate agent to invest in information acquisition. At the same time, the sequence of discounted payments $\{\delta^t w^*(t)\}_{t=0}^{T-1}$ is (weakly) decreasing over time. Hence, the optimal contract rewards earlier discoveries. This is required to prevent the agent from delaying the announcement of evidence in favor of the good state. Finally, the agent is tempted to make no investment in information acquisition and guess the good state immediately. To preclude this deviation, the principal promises a large payment at the end of the relationship if the state is bad (and the agent reports the bad signal in every period).

So far we have restricted attention to contracts which are both evidence-based and extreme. Do there exist optimal contracts without these properties? The answer depends on which property of the contract we consider. First, optimal contracts must be evidence-based. To see this, suppose that there exists an optimal contract which specifies an intermediate payment $x > 0$ in period $t < T - 1$. Without loss of generality, we may assume that the agent receives no payment if he reports message G and the state is B (i.e., the contract is extreme). Then we must be able to construct an optimal evidence-based mechanism w by increasing all the final payments in period $t' = t + 1, \dots, T - 1$, by the amount $(1/\delta)^{t'-t} x$. By construction, we have $w(G) > 0$. But recall that $w^*(G) = 0$. Thus, this contradicts the fact that w^* is the unique optimal mechanism in the class of evidence-based and extreme contracts.

Second, the optimal contract is not necessarily unique within the class of *all* contracts. Indeed, it is easy to construct examples of optimal mechanisms which are not extreme. We summarize our discussion in the following remark.

Remark 1 *If a contract is optimal then it is evidence-based. The optimal contract is not necessarily unique. Among the optimal contracts, one and only one is extreme.*

2.2 The Agent's Information Rent

Recall that the optimal contract w^* satisfies constraint (4) with equality. Thus, the agent's expected payoff under w^* can be easily computed:

$$u(w^*) = p_0 w^*(0) = \frac{c}{\alpha} + c(1 - p_0) \sum_{t=1}^{T-1} \left(\frac{\delta}{1 - \alpha} \right)^t.$$

⁶More precisely, the sequence is strictly increasing if $\delta < 1$ and constant if $\delta = 1$.

This is the information rent that the principal has to pay to motivate the agent to work and be honest. In our model, the agent has two sources of private information since both his action and his signal are unobservable. Thus, the information rent can be divided into two components: a moral hazard component and a hidden information component.

To see this, consider a variant of our model in which the signal is verifiable. Recall that the signal structure is such that either the agent finds definitive evidence in favor of state G or he does not. It is therefore natural to assume that if the agent shirks then he and the principal observe signal B with probability one. Only the agent who invests in information acquisition can discover the evidence. The new model is one with moral hazard but no hidden information. That is, the contract must induce the agent to exert effort in every period. However, incentive compatibility does not impose a truthtelling requirement since the outcome (B or G) generated by the agent's action and the state is publicly observed.

It is straightforward to show that in the new model the optimal contract is identical to w^* except that the payment $w(B)$ is equal to zero.⁷ In contrast, in our original model $w^*(B) > 0$ and the agent receives it in period $T - 1$ with probability $1 - p_0$ (if the state is B). Thus, the discounted expected utility of $w^*(B)$ represents the hidden information component of the agent's information rent and is equal to

$$\delta^{T-1} (1 - p_0) w^*(B) = \frac{c}{\alpha} + c(1 - p_0) \sum_{t=1}^{T-1} \delta^t.$$

The difference between $u(w^*)$ and the expression above

$$u(w^*) - \delta^{T-1} (1 - p_0) w^*(B) = c(1 - p_0) \sum_{t=1}^{T-1} \left(\frac{\alpha \delta}{1 - \alpha} \right)^t$$

represents the moral hazard component of the information rent.⁸

The information rent $u(w^*)$ is a U-shaped function of the signal quality α , and goes to infinity both when α is close to zero and when α is close to one. This reflects the combined effect that the signal quality has on the two components.

⁷When the realization of the signal is verifiable, a contract is incentive compatible if and only if it satisfies constraint (5). Notice that $w(B)$ does not appear in the constraints. Also, if $w(G)$ decreases, it becomes easier to satisfy all the constraints. To minimize the cost of the contract, it is therefore optimal to set $w(B) = 0$, $w(G) = 0$ and to satisfy all the constraints with equality. The unique solution to the system is the vector of payments $(w^*(0), \dots, w^*(T - 1))$.

⁸In principle, one could consider a different way to decompose the rent. In particular, one could compare the benchmark model with a model in which the agent's effort is observable but his signal is not. However, in this case the optimal contract is such that the agent receives the payment c if and only if he exerts effort. Thus, the agent's rent is equal to zero. It is then impossible to evaluate the impact that the two different forms of informational asymmetry have on the rent in our original model.

The moral hazard rent is increasing in α . To give some intuition, suppose for a moment that the agent's effort and his signal are verifiable. Consider a certain period t and let p_t (defined in equation (1)) denote the common belief. In this case it is enough to pay $c/\alpha p_t$ upon observing signal G to motivate the agent to work. The ratio p_t/p_{t+1} is increasing in α . Thus, as α grows the ratio between the payment in $t+1$ and the payment in t increases. Consider the payments $c/\alpha p_t$ and $c/\alpha p_{t+1}$ but now assume that the effort of the agent is not observable. Clearly, when α is higher the agent is more tempted to shirk in t in order to get the larger payment in $t+1$. Therefore, the principal must give a higher information rent to the agent when the quality of the signal improves.

On the other hand, the hidden information rent is decreasing in α . Under the optimal contract w^* , the agent is rewarded only if his messages and the state coincide. When the quality of the signal is high it is risky to guess state G in the first period. The agent can pay the cost c and find out, with high probability, the correct state. Thus when α is large a low value of $w^*(B)$ is sufficient to prevent the agent from deviating to a guessing strategy.

The comparative statics with respect to the remaining parameters of the model coincide for the two components of the information rent. The ratio p_t/p_{t+1} is decreasing in p_0 . Thus, if we start with the payments $c/\alpha p_t$ and $c/\alpha p_{t+1}$ of the model with observable effort, the agent's incentives to shirk in t (in order to get the reward in $t+1$) become stronger when state G becomes less likely (p_0 decreases). Therefore, the payments $(w^*(t))_{t=0}^{T-1}$ and the moral hazard rent are decreasing in p_0 . This, in turn, makes it more profitable for the agent to guess state G when that state is less likely. As a consequence, the hidden information rent is also decreasing in p_0 .

Finally, the information rent is increasing in c and δ . By shirking in a certain period t the agent saves the cost c but, at the same time, he eliminates the possibility of getting a reward in t . Clearly, shirking is more profitable when the test is particularly costly or when the agent becomes more patient. When c and δ are high the principal must reward the agents with larger payments $(w^*(t))_{t=0}^{T-1}$ and a larger moral hazard rent. In turn, the larger the payments $(w^*(t))_{t=0}^{T-1}$, the stronger the incentives to use a guessing strategy. It follows that c and δ also have a positive impact on the hidden information rent.

3 The Optimal Length of the Contract

As anticipated in Section 2, we now investigate the case in which the length of the contract is endogenous. Hence, the principal chooses the length T to maximize his expected payoff. To be able to compare all possible lengths, we first need to describe the optimal

contract when $T = 1$. One can easily verify that

$$w^*(0) = \frac{c}{p\alpha}, \quad w^*(G) = 0, \quad w^*(B) = \frac{c}{(1-p)\alpha}$$

and the agent's information rent is $u(w^*) = c/\alpha$.

For every $T = 0, 1, \dots$, let $V(T)$ denote the principal's discounted expected utility from the decision when he can observe at most T signals. Also, let $C(T)$ denote the discounted expected cost of inducing the agent to acquire T signals. We have $C(0) = \tilde{C}(0) = 0$ and for every $T \geq 1$

$$\begin{aligned} C(T) &= \tilde{C}(T) + \frac{c}{\alpha} + c(1-p_0) \sum_{t=1}^{T-1} \left(\frac{\delta}{1-\alpha}\right)^t, \\ \tilde{C}(T) &= c + c \sum_{t=1}^{T-1} \delta^t [p_0(1-\alpha)^t + 1 - p_0], \end{aligned} \tag{8}$$

where $\tilde{C}(T)$ denotes the expected cost of T signals when the principal has direct access to them (or, equivalently, when the agent's effort is verifiable).

The optimal length T^* maximizes $V(T) - C(T)$. We let \tilde{T} denote the efficient length, i.e. the length that maximizes $V(T) - \tilde{C}(T)$. While closed-form solutions of T^* and \tilde{T} are not readily available, it is simple to compare them.

Fact 1 *The optimal length T^* is (weakly) smaller than the efficient length \tilde{T} .*

The above result follows from the fact that the agent's rent is increasing in the length of the contract T . As one would expect, the presence of informational asymmetries yields a suboptimal outcome. In particular, testing is stopped too early and the gains of additional signals are not realized.

So far we have assumed that the duration of the relationship is deterministic. However, one could imagine a more general class of mechanisms in which the principal randomizes among contracts of different lengths. Are mechanisms with random duration optimal? The answer is no. More precisely, there is always an optimal contract with deterministic length. Furthermore, for generic values of the parameters of the model, the optimal length of the contract is unique.

Fix a positive integer K and consider a random mechanism under which the relationship lasts (at most) T_k periods with probability γ_k , $k = 1, \dots, K$. At the beginning of period 0, the principal randomly chooses a length in the set $\{T_1, \dots, T_K\}$ according to the probability distribution $(\gamma_1, \dots, \gamma_K)$. The principal does not inform the agent about his choice. If the selected length is T_k , the relationship ends as soon as the agent announces the good signal or in period T_{k-1} if the agent reports the bad signal in period $t = 0, \dots, T_{k-1}$. As usual,

the payments to the agent depend on his messages and on the state of the world. The cost of the optimal mechanism with such a random length is equal to:⁹

$$\gamma_1 C(T_1) + \dots + \gamma_K C(T_K),$$

where $C(T_k)$ is defined in equation (8) and represents the cost of the optimal contract with length T_k . This immediately implies that the principal is willing to randomize among two or more lengths if and only if he is completely indifferent among all of them.

4 Asymmetric Initial Information

In Section 2, we assume that the principal and the agent share the same information about the state at the outset of the relationship. This is a restrictive assumption if the agent is an expert who has been exposed to similar problems in the past. In such cases it seems natural to assume that the principal and the agent enter their relationship with different levels of information. For example, suppose that the state is identically and independently distributed across problems according to an unknown probability distribution. Consider the relationship between an agent who has already consulted for many different principals and a new principal. Since the agent has observed the realization of the state in sufficiently many problems, it is reasonable to assume that he knows the true distribution. On the other hand, the principal does not have access to past information and is, therefore, uninformed.

The goal of this section is to analyze how the principal motivates an informed agent to carry out his task. In particular, we investigate how the additional source of private information affects the optimal mechanism and the agent's information rent.

To allow the expert to possess initial information, we modify the model presented in Section 2 and let the agent have a private type at the beginning of period 0. For tractability, we assume that there are two possible types, low or high. The agent's type is correlated with the state which is equal to G with probability p_0 and B with probability $1 - p_0$. We find it convenient to denote the two types by their beliefs, p_0^ℓ and p_0^h , that the state is G . Thus, if we let $\Pr(p_0^k|\omega)$ denote the probability of type p_0^k , $k = h, \ell$, in state $\omega = B, G$, we

⁹For brevity, we only provide a sketch of the proof and omit the details. We consider a relaxed problem in which only two types of deviations are admitted. The agent can either guess state G in period 0 or he can shirk and announce message B in one and only one period. These deviations correspond to constraints (4) and (5) in Lemma 1. The solution to the relaxed problem satisfies all the constraints with equality and the payments are positive only if the agent's messages match the state. The cost of the optimal contract of the relaxed problem is equal to $\gamma_1 C(T_1) + \dots + \gamma_K C(T_K)$. Finally, it is easy to check that the optimal contract of the relaxed problem is incentive compatible.

have:

$$p_0^k = \frac{p_0 \Pr(p_0^k|G)}{p_0 \Pr(p_0^k|G) + (1 - p_0) \Pr(p_0^k|B)}.$$

We assume $1 > \Pr(p_0^h|G) > \Pr(p_0^h|B) > 0$ and, thus, $1 > p_0^h > p_0^\ell > 0$. We also let

$$\rho = p_0 \Pr(p_0^h|G) + (1 - p_0) \Pr(p_0^h|B)$$

denote the (unconditional) probability of the high type p_0^h . Thus, the principal believes that the agent's type is high with probability ρ and low with probability $1 - \rho$.

We assume that the number of available signals is $T \geq 2$.¹⁰ For $k = h, \ell$ and $t = 0, \dots, T$, we let

$$p_t^k = \frac{p_0^k (1 - \alpha)^t}{p_0^k (1 - \alpha)^t + 1 - p_0^k}$$

denote the agent's belief that the state is G if his type is p_0^k and he observes t signals equal to B .

As in Section 2, the principal tries to induce the agent to acquire and reveal the signal in every period $t = 0, \dots, T - 1$ (until he observes signal G).¹¹ Since the agent has private information about his type, a mechanism (w^h, w^ℓ) consists of a pair of contracts, one for each type. In this section, we focus on evidence-based and extreme mechanisms (for simplicity, we refer to them simply as mechanisms) and, therefore, $w^k = \left((w^k(t))_{t=0}^{T-1}, w^k(G), w^k(B) \right)$, for $k = h, \ell$. This is without loss of generality in the sense that there is an optimal mechanism in the class of evidence-based and extreme mechanisms (we briefly return to this point at the end of the section).

Thus, the game between the principal and the agent is as follows. In period 0, the principal offers a pair of contracts (w^h, w^ℓ) and the agent chooses one. In every period $t = 0, \dots, T - 1$, the agent decides whether to exert effort or not and sends a message from the set $\{B, G\}$. The game ends as soon the agent announces message G (or in period $T - 1$ if he reports message B in every period). The agent receives the payment specified by the contract that he chose.

¹⁰We rule out the simplest case $T = 1$ because the notation developed for the general case $T \geq 2$ should be slightly modified when $T = 1$. However, it is straightforward to extend all the results developed in this section to the special case $T = 1$.

¹¹To gain some intuition on how the agent's initial information affects the shape of the mechanism, we develop a simple model in which the length of the contract is the same for both types of the agent. This is a reasonable assumption if there is an exogenous deadline T by which the principal has to make a decision and the additional value of each signal is large compared to the cost of effort c . Of course, it is easy to imagine situations in which the principal prefers to induce the two types to acquire different numbers of signals. While some of our results in this section (Lemmata 2 and 3, the second part of Lemma 4, and Proposition 2) easily extend to the case of contracts with different lengths, we do not have a general characterization of the optimal mechanism.

If a mechanism (w^h, w^ℓ) is incentive compatible, then it is optimal for the type p_0^k , $k = h, \ell$, to choose the contract w^k and to acquire and reveal the signal in every period. The principal's problem is to find the cheapest incentive compatible mechanism.

We say that a contract $w = \left((w(t))_{t=0}^{T-1}, w(G), w(B) \right)$ is *suitable* for type p_0^k if it satisfies constraints (4)-(6) when the prior is p_0^k . Clearly, if (w^h, w^ℓ) is an incentive compatible mechanism, then for $k = h, \ell$, the contract w^k must be suitable for type p_0^k .

An incentive compatible mechanism must also satisfy a set of constraints which prevent the agent from lying about his initial type. In principle, we have one constraint for each strategy $\sigma \in \Sigma$ since a type who lies can then choose any contingent plan of actions and messages. However, we now show that it is without loss of generality to ignore many of these constraints.

Consider a pair of contracts (w^h, w^ℓ) with w^k suitable for type p_0^k , $k = h, \ell$. Consider an arbitrary type p_0^k and suppose that he chooses the contract $w^{k'}$, $k' \neq k$, designed for the other type. Since $w^{k'}$ is suitable for $p_0^{k'}$ it is easy to see that it is optimal for type p_0^k to reveal the signal truthfully. Furthermore, any strategy in which type p_0^k guesses state G in some period t yields a payoff weakly smaller than $p_0^k w^{k'}(0)$ (for both claims, see the discussion following Lemma 1 and its proof).

Let us now restrict attention to the strategies under which the agent reveals the signal truthfully and reports message B when he shirks. We denote this set of strategies by Σ' . Formally, a strategy σ belongs to Σ' if for every $t = 0, \dots, T-1$, every $h^t \in H^t$, and every $s = B, G$,

$$\begin{aligned}\sigma_t^M(h^t, s) &= s, \\ \sigma_t^M(h^t, ne) &= B.\end{aligned}$$

Suppose that the high type p_0^h faces the contract w^ℓ and that this contract is suitable for p_0^ℓ . It is easy to see that the optimal strategy in Σ' for p_0^h is to acquire the signal in every period. Consider period t and suppose that the agent acquires and reveals the signal in every period $t' > t$. The fact that w^ℓ is suitable for p_0^ℓ implies that in period t the agent strictly prefers to work and be honest if his belief is strictly larger than p_t^ℓ . Clearly, under any strategy in Σ' , the high type's belief in period t must be strictly larger than p_t^ℓ . Thus, any strategy in Σ' under which the agent shirks in at least one period (on-path) is strictly dominated by the strategy under which the agent exerts effort in every period.

Suppose now that the low type p_0^ℓ faces the contract w^h . Knowing that w^h is suitable for p_0^h is not enough to pin down the low type's optimal strategy in Σ' . The easiest way to see this is to consider a contract w^h that satisfies constraint (5) for $T-1$ with strict inequality. Suppose that the low type's belief in period $T-1$ is smaller than p_{T-1}^h . Given this limited amount of information, it is clearly impossible to determine whether the agent prefers to

acquire and reveal the signal or to shirk and announce message B .

We summarize the discussion above about incentive compatibility in the following lemma. Recall that given a contract w , $u(t, p; w)$ denotes the agent's expected utility, computed in period t , when his belief is p and he acquires and reveals the signal in every period $t' \geq t$. Also, $u(\sigma, p; w)$ denotes the agent's expected utility in period 0 if he follows the strategy σ .

Lemma 2 *A mechanism (w^h, w^ℓ) is incentive compatible if and only if w^k is suitable for p_0^k , $k = h, \ell$, and it satisfies the following inequalities:*

$$u(0, p_0^k; w^k) \geq p_0^k w^{k'}(0), \quad k = h, \ell, \quad k' \neq k,$$

$$u(0, p_0^h; w^h) \geq u(0, p_0^h; w^\ell),$$

$$u(0, p_0^\ell; w^\ell) \geq u(\sigma, p_0^\ell; w^h), \quad \sigma \in \Sigma'.$$

The next step of our analysis is to characterize the optimal mechanisms.

4.1 The Optimal Mechanism with Asymmetric Initial Information

The principal's problem is given by

$$\begin{aligned} \min_{w^h \geq 0, w^\ell \geq 0} \quad & \rho u(0, p_0^h; w^h) + (1 - \rho) u(0, p_0^\ell; w^\ell) \\ \text{s.t.} \quad & (w^h, w^\ell) \text{ is incentive compatible.} \end{aligned} \tag{9}$$

The set of incentive compatible mechanisms is non-empty. To see this, notice that the principal can always offer the following contract \bar{w} to both types. The contract \bar{w} is identical to $w^*(p_0^\ell)$, the optimal contract offered to the agent when he has no private information and the prior is p_0^ℓ , except that $\bar{w}(B)$ is such that the high type is indifferent between acquiring and revealing the signal in every period and guessing state G in period 0 (thus, $\bar{w}(B) > w^*(B; p_0^\ell)$). Clearly, the contract \bar{w} is suitable both for p_0^h and p_0^ℓ . In fact, as we will see below, when the two beliefs p_0^h and p_0^ℓ are sufficiently close to each other, it is indeed optimal to offer only the contract \bar{w} . However, when the difference between the beliefs is sufficiently large, the principal prefers to offer two different contracts.

In general, there are multiple solutions to the principal's problem. However, there are some features of the contracts that are common to all optimal mechanisms. In particular, the payments to the low type in state G coincide with the payments of the contract $w^*(p_0^\ell)$.

Proposition 2 *If (w^h, w^ℓ) is an optimal mechanism then*

$$w^\ell(t) = w^*(t; p_0^\ell)$$

for every $t = 0, \dots, T - 1$, and

$$w^\ell(G) = w^*(G; p_0^\ell).$$

Intuitively, to screen the types of the agent, it is convenient to provide each of them with larger (smaller) payments in the state that he considers relatively more (less) likely. Hence, the principal should provide low payments to the low type when he correctly announces the good state. Among all the contracts that are suitable for p_0^ℓ , $w^*(p_0^\ell)$ specifies the lowest payments in state G . We conclude that there is no distortion on these payments in the contract of the low type.

Consider an incentive compatible mechanism (w^h, w^ℓ) and suppose that w^ℓ and $w^*(p_0^\ell)$ do not have the same payments in state G . Suppose now that the principal lowers the payments $(w^\ell(t))_{t=0}^{T-1}$ and $w^\ell(G)$ to make them equal to the payments of $w^*(p_0^\ell)$. At the same time the principal increases the payment $w^\ell(B)$ so that the low type is indifferent between the old contract w^ℓ and the new contract, which we call \hat{w}^ℓ , when he acquires and reveals the signal in every period.

Let us now evaluate how the change from w^ℓ to \hat{w}^ℓ affects the utility of the high type when he lies about his type and acquires and reveals the signal in every period. Compared to the low type, the high type assigns higher probabilities to the payments in state G (which are lower in \hat{w}^ℓ than in w^ℓ) and lower probability to the payment in state B (which is higher in \hat{w}^ℓ than in w^ℓ). Clearly, if the low type is indifferent between the two contracts, the high type must strictly prefer w^ℓ to \hat{w}^ℓ . Thus, under the mechanism (w^h, \hat{w}^ℓ) , the high type has a strict incentive to choose the contract w^h . But then the principal can lower some of the payments of w^h without making it profitable for the high type to imitate the low type. Therefore, the original mechanism (w^h, w^ℓ) is not optimal.

Proposition 2 shows that the payments to the low type in state G are not distorted from the optimal mechanism $w^*(p_0^\ell)$. However, the fact that the agent's initial belief is private information does have an impact on the contract of the low type.

Lemma 3 *If (w^h, w^ℓ) is an incentive compatible mechanism then*

$$u(0, p_0^\ell; w^\ell) > p_0^\ell w^\ell(0).$$

Suppose that the mechanism (w^h, w^ℓ) is incentive compatible. The high type weakly prefers to acquire and reveal the signal with the contract w^h rather than guess state G

in the first period with the contract w^ℓ . But then the low type strictly prefers the first alternative to the second one. This is because the guessing strategy is more tempting for the high type who is more optimistic about the good state. We conclude that choosing the contract w^ℓ and guessing state G in the first period is not the most profitable deviation for the low type. Thus, the incentive $u(0, p_0^\ell; w^\ell) \geq p_0^\ell w^\ell(0)$ does not bind.

Taken together, Proposition 2 and Lemma 3 imply that the information rent of the low type under an optimal mechanism (w^h, w^ℓ) satisfies

$$u(0, p_0^\ell; w^\ell) > p_0^\ell w^\ell(0) = p_0^\ell w^*(0; p_0^\ell) = u(0, p_0^\ell; w^*(p_0^\ell)).$$

Compared to the benchmark case (i.e., no initial private information) with prior p_0^ℓ , the low type obtains a higher information rent. The additional rent comes in the form of a payment $w^\ell(B)$ which is strictly larger than $w^*(B; p_0^\ell)$.

So far we have considered the low type. We now look for general properties of the optimal contract of the high type. The next lemma shows that in the first period the high type must receive the same payment as the low type. Furthermore, the high type must be indifferent between exerting effort in every period and guessing state G immediately.

Lemma 4 *If (w^h, w^ℓ) is an optimal mechanism then*

$$w^h(0) = w^\ell(0) = w^*(0; p_0^\ell),$$

and

$$u(0, p_0^h; w^h) = p_0^h w^h(0).$$

First, we provide some intuition for the second result. Suppose that $u(0, p_0^h; w^h) > p_0^h w^h(0)$. Also, consider the case $u(0, p_0^\ell; w^\ell) > p_0^\ell w^h(0)$, i.e. the constraint in which the low type chooses w^h and guesses state G immediately is not binding.¹² Then the principal can increase the value of $w^h(0)$ and decrease the value of $w^h(B)$ so that the high type is indifferent between the old contract w^h and the new contract, which we call \hat{w}^h , when he acquires and reveals the signal in every period. However, given any strategy in Σ' , the low type is strictly worse off with the contract \hat{w}^h than with the contract w^h . The logic is similar to that of Proposition 2. The principal makes the deviations of a certain type more costly by decreasing (increasing) the payments that the type deems relatively more (less) likely. Finally, the principal can also lower the payment $w^\ell(B)$ by a small amount and the

¹²In the proof we show that if $u(0, p_0^\ell; w^\ell) = p_0^\ell w^h(0)$, then the principal can lower the payment $w^h(B)$ by a small amount and all the incentive constraints are still satisfied.

new mechanism is incentive compatible. This shows that the original mechanism (w^h, w^ℓ) is not optimal.

Since $u(0, p_0^h; w^h) = p_0^h w^h(0)$, $w^h(0)$ must be weakly greater than $w^\ell(0)$ otherwise the high type would have an incentive to choose the contract w^ℓ and guess state G immediately. In the proof we rule out the case $w^h(0) > w^\ell(0)$ by showing that the principal can lower $w^h(0)$ and some other payments of the contract w^h without violating the incentive constraints.

Lemma 4 has a number of important implications. First, notice that the information rent of the high type is $p_0^h w^*(0; p_0^\ell)$ and that this is larger than $p_0^h w^*(0; p_0^h)$, the rent that he would obtain in the benchmark model with prior p_0^h . Thus, both types benefit from the fact that their initial belief is private information. This is in contrast to many models of adverse selection in which the principal is able to extract all the rents from a certain type.

Second, the contract offered to the low type is the same among all optimal mechanisms. Suppose that (w^h, w^ℓ) and $(\hat{w}^h, \hat{w}^\ell)$ are two optimal mechanisms. Then $u(0, p_0^\ell; w^\ell)$ and $u(0, p_0^\ell; \hat{w}^\ell)$ must coincide, since $u(0, p_0^h; w^h)$ and $u(0, p_0^h; \hat{w}^h)$ coincide. However, the contracts w^ℓ and \hat{w}^ℓ have the same payments in state G . The low type can be indifferent among w^ℓ and \hat{w}^ℓ (when he works in every period) if and only if $w^\ell(B) = \hat{w}^\ell(B)$. Therefore, w^ℓ must be equal to \hat{w}^ℓ .

Third, the solution to the principal's problem does not depend on the probability distribution of the two types (i.e., the parameter ρ). Clearly, the set of incentive compatible mechanisms does not vary with ρ . Lemma 4 guarantees that for every ρ the high type receives the same utility $p_0^h w^*(0; p_0^\ell)$ under an optimal mechanism. Thus, the utility of the low type must also be the same for all values of ρ .

Proposition 2, Lemma 3 and Lemma 4 identify a part of the contracts that is common to all the optimal mechanisms. The remaining part may vary among optimal mechanisms. Moreover, it is also affected by the distance between the two beliefs p_0^h and p_0^ℓ .

Recall that \bar{w} denotes the contract that is identical to $w^*(p_0^\ell)$, except for the value of $\bar{w}(B)$, which is such that

$$u(0, p_0^h; \bar{w}) = p_0^h \bar{w}(0) = p_0^h w^*(0; p_0^\ell).$$

As mentioned above, the mechanism (\bar{w}, \bar{w}) under which the principal offers the same contract \bar{w} to both types is incentive compatible. The next proposition identifies necessary and sufficient conditions for the optimality of such a mechanism.

Proposition 3 *The mechanism (\bar{w}, \bar{w}) is optimal if and only if $p_0^\ell \in [p_1^h, p_0^h]$. Furthermore, if $p_0^\ell \in (p_1^h, p_0^h)$, then (\bar{w}, \bar{w}) is the unique optimal mechanism.*

Recall from Lemma 4 that a mechanism (w^h, w^ℓ) can be optimal only if w^h gives to the high type the same rent as the contract \bar{w} . Furthermore, we must have $w^h(0) = \bar{w}(0)$. We also know from Proposition 2 that the contract w^ℓ is such that the low type is indifferent between exerting effort in every period and exerting effort *only after* the first one.

Let us now investigate how the principal could improve upon the mechanism (\bar{w}, \bar{w}) . To do that the principal must offer to the high type a contract w^h which is less attractive than \bar{w} for the low type *both* when he expends effort in every period *and* when he does so only in period $t = 1, \dots, T - 1$. Otherwise, the low type would have an incentive to accept the contract of the high type and adopt one of the two strategies described above.

First, let us consider the strategy under which the low type acquires and reveals the signal in every period. At the beginning of period 0, he is less optimistic than the high type (about the good state). Therefore, w^h is less attractive than \bar{w} if it promises less when the agent correctly anticipates the bad state (i.e., $w^h(B) < \bar{w}(B)$). Of course, this implies that w^h is more generous than \bar{w} when the agent announces the good state in period $t \geq 1$.

Consider now the strategy under which the low type shirks in period 0 and exerts effort in any other period. Let us compare now the two types at the beginning of $t = 1$, after the high type has observed a bad signal (and the low type has shirked). If $p_0^\ell < p_1^h$, the low type is still less optimistic than the high type. In this case too, the low type prefers \bar{w} to the new contract w^h if $w^h(B) < \bar{w}(B)$. In other words, when $p_0^\ell < p_1^h$, the changes to the contract of the high type that prevent the two deviations of the low type go in the same direction. Suppose now that $p_0^\ell > p_1^h$. In this case, at the beginning of $t = 1$ the low agent is more optimistic than the high type. To make the new contract w^h less attractive than \bar{w} for the low type, the principal must increase the payment $w^h(B)$. But this change is just the opposite of what is needed to prevent the first deviation. When the initial beliefs of the two types are sufficiently close, there is no room to change the contract of the high type and prevent the two deviations of the low type. This is because the two deviations lead to opposite results in terms of the comparison between the types' beliefs.

Finally, the case $p_0^\ell = p_1^h$ is special in the sense that the change that prevents the first deviation (i.e., $w^h(B) < \bar{w}(B)$) leaves the low type indifferent in terms of the second deviation. While it is not possible to improve upon (\bar{w}, \bar{w}) , this mechanism is not necessarily the unique optimal one.

We now turn to the case $p_0^\ell \leq p_1^h$. In this case, there are multiple optimal mechanisms. We describe one in the next proposition and then address the issue of multiplicity.

Proposition 4 *Suppose that $p_0^\ell \leq p_1^h$. There exists an optimal mechanism (w^h, w^ℓ) which satisfies*

$$w^h(B) \in [w^*(B; p_0^h), w^\ell(B))$$

and one of the following two conditions:

(i) there exists $\hat{t} \in \{1, \dots, T-1\}$ such that

$$\begin{aligned} w^h(t) &= \frac{w^*(0; p_0^\ell)}{\delta^t}, & t < \hat{t} \\ w^h(\hat{t}) &\in \left(w^*(\hat{t}; p_0^\ell), \frac{w^h(\hat{t}-1)}{\delta} \right] \\ w^h(t) &= w^*(t; p_0^\ell), & t > \hat{t} \\ w^h(G) &= 0 \end{aligned}$$

(ii)

$$\begin{aligned} w^h(t) &= \frac{w^*(0; p_0^\ell)}{\delta^t}, & t = 0, \dots, T-1 \\ w^h(G) &\in \left(0, w^h(T-1) - \frac{c}{\alpha p_{T-1}^h} \right]. \end{aligned}$$

In the proof we start with an arbitrary optimal mechanism $(\hat{w}^h, \hat{w}^\ell)$. We let $\Delta \geq 0$ denote

$$\Delta = u(0, p_0^h; \hat{w}^h) - \delta u(1, p_0^h; \hat{w}^h).$$

This and the fact that \hat{w}^h must satisfy the conditions in Lemma 4 immediately give us the value of $\hat{w}^h(B)$, which is increasing in Δ and coincides with $w^*(B; p_0^h)$ when $\Delta = 0$ (see equation (29)).

Then we construct the mechanism (w^h, w^ℓ) described in Proposition 4. Of course, $w^\ell = \hat{w}^\ell$ and $w^h(0) = \hat{w}^h(0)$. We also let $w^h(B) = \hat{w}^h(B)$. The remaining payments of w^h are determined using the following algorithm. In the first step, we let w^h have the same payments as w^ℓ after period one and choose $w^h(1)$ to satisfy

$$u(0, p_0^h; w^h) - \delta u(1, p_0^h; w^h) = \Delta. \quad (10)$$

If the solution $w^h(1)$ is between $w^\ell(1)$ and $\frac{w^\ell(0)}{\delta}$ we stop. Otherwise we let $w^h(1) = \frac{w^\ell(0)}{\delta}$ and move to the second step. In step t , $t = 2, \dots, T$, we let the payment $w^h(t')$, $t' < t$, be equal to $\frac{w^\ell(0)}{\delta^{t'}}$. We also let w^h have the same payments as w^ℓ after period t and choose the remaining payment $w^h(t)$ (this is $w^h(G)$ if we are in step T) to solve equation (10).

Finally, we show that the resulting mechanism (w^h, w^ℓ) is incentive compatible and has the same expected cost as the original mechanism $(\hat{w}^h, \hat{w}^\ell)$.

Of course, the value of $w^h(\hat{t})$ (or the value of $w^h(G)$ if (w^h, w^ℓ) satisfies condition (ii)) depends on Δ which is endogenous. For the special case in which $p_0^\ell = \frac{p_0^h(1-\alpha)^t}{p_0^h(1-\alpha)^{t+1} - p_0^h}$ for some $t = 1, 2, \dots$, it is possible to show that there exists an optimal mechanism $(\hat{w}^h, \hat{w}^\ell)$

such that $u(0, p_0^h; \hat{w}^h) = \delta u(1, p_0^h; \hat{w}^h)$. In other words, we can assume that Δ is equal to zero. However, this is not true in general. There are examples in which Δ is strictly positive for any optimal mechanism.

The mechanism described in Proposition 4 has very intuitive properties. After some period \hat{t} , the two contracts w^h and w^ℓ specify the same payments in state G . However, before \hat{t} the principal sets the payment of w^h at their highest possible levels.¹³ This is useful to separate the two types and prevent the low type from choosing the contract w^h . In fact, in the initial periods the low type is much less optimistic that he will receive a payment in state G . Even if these initial payments are large, he will not find it profitable to choose w^h and exert effort. As time goes on and the high type observes more signals equal to B , his posterior gets closer to the initial belief of the low type. If the later payments of w^h are large (and sufficient to motivate the high type) then the low type could find it profitable to choose w^h and start to acquire the signal after a few periods.

Suppose that the low type chooses the contract w^h . Since the payments of w^h after \hat{t} are the same as the payments of $w^*(p_0^\ell)$ (the optimal contract in the benchmark model with prior p_0^ℓ) and $w^h(\hat{t}) \geq w^*(\hat{t}, p_0^\ell)$, the low type does not have an incentive to shirk in $\hat{t}, \dots, T-1$ (under any strategy the belief of the low type in period t must be at least p_t^ℓ).

Before \hat{t} the low type may prefer to shirk. However, the decision to shirk should not be delayed. If $t < \hat{t} - 1$, then the agent is indifferent between the payment $w^h(t)$ in t and the payment $w^h(t+1)$ in $t+1$. By definition, they have the same discounted value. However, the agent prefers to pay the cost c in $t+1$ rather than in t (this preference is strict if $\delta < 1$).

To sum up, given w^h it is optimal for the low type to adopt the following strategy. He shirks in the first t periods (for some $t < \hat{t}$) and then acquires and reveals the signal in $t+1, \dots, T-1$. Given this, it is easy to see why the solution to the principal's problem is not unique. Consider the optimal mechanism described in Proposition 4. For example, suppose that given w^h all the strategies under which the low type works in period 1 and/or period 2 are strictly dominated. Suppose now that the principal increases $w^h(1)$ by a small amount and decreases $w^h(2)$ to keep constant the rent of the high type. As far as the low type is concerned, this change affects only strategies that are strictly dominated. Since the original contract w^h satisfies all the constraints in (5) with strict inequality, the new contract of the high type is still suitable for p_0^h . We have therefore constructed a new optimal mechanism.

So far we have restricted attention to evidence-based and extreme mechanisms. Allowing for intermediate payments has no consequences for the contract of the low type. Recall

¹³Recall that $w^h(t)$ must be weakly smaller than $\frac{w^h(0)}{\delta^t}$ otherwise the high type could delay the announcement of the signal G . Also, notice that if the optimal mechanism (w^h, w^ℓ) satisfies condition (ii) in Proposition 4, then all the payments $w^h(1), \dots, w^h(T-1)$ are set at their highest possible levels.

that in any optimal mechanism (w^h, w^ℓ) the payment $w^\ell(G)$ is equal to zero. The contract of the low type must be evidence-based.

Moreover, when $p_0^\ell > p_1^h$, (\bar{w}, \bar{w}) remains the unique optimal mechanism in the class of extreme mechanisms. However, when $p_0^\ell \leq p_1^h$ there are optimal evidence-based mechanisms (w^h, w^ℓ) under which all the payments of w^h are strictly positive. In this case it is possible to modify the contract w^h to allow for intermediate payments. In other words, there are optimal mechanisms under which the contract of the high type is not evidence-based.

Finally, it is easy to construct examples of optimal mechanisms which are not extreme. This is true both when $p_0^\ell > p_1^h$ and when $p_0^\ell \leq p_1^h$.

The results in this section show that the introduction of asymmetric information at the time of contracting leads to a number of novel predictions. First, all types benefit from their initial private information. Compared to the optimal contract of the benchmark model, each type receives a larger payment when the state is bad. In addition, the high type receives larger payments when he correctly announces the good state. These payments are also front-loaded if the initial beliefs of the two types are sufficiently diverse.

5 Conclusions

This paper analyzes the optimal provision of incentives in a dynamic information acquisition process. In every period the agent can acquire costly information that is relevant to the principal's decision. The agent's effort and the realizations of his signals are unobservable. The principal commits to a long-term contract that specifies the payments to the agent. The optimal contract induces the agent to perform the test and reveal its outcome truthfully in every period and minimizes his expected utility.

First, we assume that the agent has no private information at the outset of the relationship. Under the optimal contract, the agent is rewarded only when his reports are supported by the state. The payments that the agent receives when he announces the good state increase over time. We show that agency problems shorten the information acquisition process. We then extend the analysis to the case in which the agent has superior information at the time of contracting. We characterize the optimal mechanisms and show that the contract offered to the low type is minimally distorted. The principal prefers to offer different contracts if and only if the types' beliefs are sufficiently diverse. Finally, all the types benefit from their initial private information.

In our model, the state of the world is observed no matter what decision the principal makes. However, one can imagine situations in which the state is revealed only if the principal chooses certain actions. For instance, consider an oil company hiring a team of geologists to perform preliminary investigations and give a recommendation about the

profitability of a new site. If the company decides not to invest in the site, its profitability remains unknown. Consider a variant of the model in which the principal does not learn the state if the agent reports the bad signal in every period. If the agent has no private information at the time of contracting, the analysis is similar to the one in Sections 2 and 3.¹⁴ A mechanism is incentive compatible if and only if it satisfies constraints (4)-(6), and constraints (4) and (5) are binding under the optimal contract. Compared to our benchmark model, the agent is able to extract larger information rents since the principal has less instruments to monitor his effort. As in our model, the information acquisition problem is shorter than the efficient one because of agency problems.

Under our information structure, the agent's beliefs evolve in a simple way. Either the agent becomes certain that the state is good or his belief that the state is bad increases. Although this information structure is commonly used in models of dynamic agency (see, among others, Bergemann and Hege 1998, 2005 and Hörner and Samuelson 2010), it is natural to consider more general information structures under which all the realizations of the signal contain some noise. Our results are robust to small perturbations. Consider a variant of the benchmark model in which the probability of observing the good signal when the state is bad is $\nu > 0$. Consider the optimal contract that induces the agent to acquire and reveal the signal in every period $t = 0, \dots, T - 1$ until he observes a good signal. If ν is sufficiently small, the optimal contract is determined by the same set of binding constraints as in Section 2. As ν converges to zero, the optimal contract approaches the contract described in Proposition 1.

Preliminary investigation also suggests that a number of properties of our optimal mechanism extend to the case in which the signals do not provide extreme evidence in favor of the states. First, moral-hazard, hidden-information and adverse-selection rents are present in general environments. Second, it is possible to show that information rents are non-monotonic in the precision of the signal. Finally, the rents are increasing in the length of the information acquisition process which is, therefore, shorter than the efficient one. Obtaining a closed-form solution for the optimal contract in general settings seems difficult since the binding constraints may vary with the parameters of the model. We leave this challenging task for future research.

¹⁴Extending the analysis when the agent is privately informed is less immediate and beyond the scope of the paper.

Appendix

Proof of Lemma 1.

We have only to prove that if a mechanism w satisfies constraints (4)-(6) then it is incentive compatible. Notice that if σ and σ' are two strategies which induce the same on-path behavior, then $u(\sigma, p_0; w) = u(\sigma', p_0; w)$.

Let $\sigma \in \Sigma \setminus \Sigma^*$ be a strategy such that $\sigma_t^M(h^t, s) \neq s$ for some on-path history (h^t, s) (i.e., on path the agent acquires the signal and lies about its realization). We claim that there exists another strategy σ' which yields a weakly larger payoff: $u(\sigma', p_0; w) \geq u(\sigma, p_0; w)$. First, suppose $\sigma_t^M(h^t, G) = B$. The agent's continuation payoff after observing signal G is bounded above by

$$\max \{ \delta w(t+1), \dots, \delta^{T-1-t} w(T-1), \delta^{T-1-t} w(G) \}.$$

If instead the agent reports message G , his continuation payoff is equal to $w(t)$. From constraint (5) at $T-1$,

$$\begin{aligned} 0 \leq u(T-1, p_{T-1}) - \delta u(T, p_{T-1}) = \\ -c + p_{T-1} \alpha [w(T-1) - w(G)], \end{aligned} \tag{11}$$

we obtain

$$w(T-1) > w(G). \tag{12}$$

Let σ' be identical to σ except that we set $\sigma_t^M(h^t, G) = G$. Inequality (12) and constraint (6) immediately imply $u(\sigma', p_0; w) \geq u(\sigma, p_0; w)$.

Consider now the case in which (h^t, B) and (h^t, G) are two on-path histories and $\sigma_t^M(h^t, B) = \sigma_t^M(h^t, G) = G$. Clearly, the agent strictly prefers to save the cost c and send the uninformed message G . Thus, let σ' be identical to σ , except that we set $\sigma_t^A(h^t) = ne$ and $\sigma_t^M(h^t, ne) = G$. Then we have $u(\sigma', p_0; w) > u(\sigma, p_0; w)$.

Given these initial observations, in what follows, we restrict attention to strategies σ under which the agent reveals truthfully all the signals that he acquires: $\sigma_t^M(h^t, s) = s$ for all on-path histories (h^t, s) . Let σ' be a strategy that prescribes guessing state G (i.e., shirking and announcing message G) in the first period. Let $\hat{\sigma}$ be any strategy such that the agent guesses state G in some period $\hat{t} > 0$ (and this is part of the on-path behavior). We now show that the agent weakly prefers σ' to $\hat{\sigma}$. Let $\tau_1 < \dots < \tau_{\hat{k}} < \hat{t}$, for some $\hat{k} \leq \hat{t}$, denote the \hat{k} periods in which the agent acquires the signal under the strategy profile $\hat{\sigma}$.¹⁵

¹⁵Recall that we have already ruled out strategies under which the agent lies after observing the signal G . Thus, the agent acquires the signal in period τ_k , $k = 2, \dots, \hat{k}$, only if he has never observed the signal G before.

We have

$$\begin{aligned}
& u(\sigma', p_0) = p_0 w(0) \geq \\
& p_0 \alpha \delta^{\tau_1} w(\tau_1) + p_0 (1 - \alpha) \alpha \delta^{\tau_2} w(\tau_2) + \dots + p_0 (1 - \alpha)^{\hat{k}-1} \alpha \delta^{\tau_{\hat{k}}} w(\tau_{\hat{k}}) + p_0 (1 - \alpha)^{\hat{k}} \delta^{\hat{t}} w(\hat{t}) - \\
& c \left[\delta^{\tau_1} + \delta^{\tau_2} (p_0 (1 - \alpha) + 1 - p_0) + \dots + \delta^{\tau_{\hat{k}}} (p_0 (1 - \alpha)^{\hat{k}-1} + 1 - p_0) \right] = u(\hat{\sigma}, p_0),
\end{aligned}$$

where the inequality follows from $c > 0$ and constraint (6) (this constraint implies $w(0) \geq \delta^t w(t)$ for every t). Combining this result with constraint (4) we obtain $u(0, p_0) \geq u(\hat{\sigma}, p_0)$.

It remains to consider strategies σ under which the agent's on-path behavior is to tell the truth when he acquires the signal and to send message B when he shirks: $\sigma_t^M(h^t, s) = s$ and $\sigma_t^M(h^t, ne) = B$, for all on-path histories (h^t, s) , (h^t, ne) . The last step of the proof is to show that constraint (5) implies

$$u(t, p) > \delta u(t+1, p) \quad (13)$$

for every $t = 0, \dots, T-1$ and every $p > p_t$. This will be enough to conclude that the mechanism w is incentive compatible because it implies that any strategy in which the agent shirks more than once is strictly dominated. To see this, consider a strategy $\tilde{\sigma}$ under which the agent shirks in two or more periods. Let \tilde{t} denote the last period in which the agent shirks. Let \tilde{p} denote the agent's belief in period \tilde{t} . Because the agent has shirked at least once before \tilde{t} , we have $\tilde{p} > p_t$. Clearly, $u(\tilde{t}, \tilde{p}) > \delta u(\tilde{t}+1, \tilde{p})$ implies that the agent has a strict incentive to exert effort in period \tilde{t} (given that he will exert effort in all future periods). On the other hand, any strategy in which the agent shirks only once is weakly dominated by the strategies in Σ^* because of constraints (5). Thus, the mechanism w is incentive compatible.

We now prove inequality (13). One can immediately see from inequality (11) that inequality (13) holds for $t = T-1$.

Given the contract w , define $a(T-1; w)$ and $b(T-1; w)$ as follows:

$$\begin{aligned}
a(T-1; w) &= w(B) - c, \\
b(T-1; w) &= \alpha w(T-1) + (1 - \alpha) w(G) - w(B),
\end{aligned}$$

and for every $t = 0, \dots, T-2$ define recursively

$$\begin{aligned}
a(t; w) &= -c + \delta a(t+1; w), \\
b(t; w) &= \alpha w(t) - \delta \alpha a(t+1; w) + \delta (1 - \alpha) b(t+1; w).
\end{aligned} \quad (14)$$

Notice that for every t and every p

$$u(t, p) = -c + p\alpha w(t) + (1 - p\alpha) \delta u\left(t + 1, \frac{p(1 - \alpha)}{1 - p\alpha}\right).$$

Using an induction argument it is easy to check that for every $t = 0, \dots, T - 1$, and every $p \in [0, 1]$

$$u(t, p) = a(t) + b(t)p, \tag{15}$$

where we have dropped the argument w in $a(t; w)$ and $b(t; w)$ to simplify the notation.

Thus, we have

$$u(t, p) - \delta u(t + 1, p) = a(t) - \delta a(t + 1) + [b(t) - \delta b(t + 1)]p.$$

Constraint (5) implies that the above expression is non-negative when $p = p_t$. To conclude the proof of the lemma it is, therefore, sufficient to show that for $t = 0, \dots, T - 2$

$$b(t) - \delta b(t + 1) > 0,$$

which is equivalent to

$$w(t) - \delta a(t + 1) - \delta b(t + 1) > 0. \tag{16}$$

The above inequality is satisfied for $t = T - 2$. In fact

$$\begin{aligned} w(T - 2) - \delta a(T - 1) - \delta b(T - 1) &= \\ w(T - 2) - \delta\alpha w(T - 1) - \delta(1 - \alpha)w(G) + \delta c &= \\ [w(T - 2) - \delta w(T - 1)] + \delta(1 - \alpha)[w(T - 1) - w(G)] + \delta c &> 0, \end{aligned}$$

where the inequality follows from constraint (6), inequality (12), and, of course, $c > 0$.

We now proceed by induction. We assume that inequality (16) holds for $t' > t$ and show that it also holds at t . We have

$$\begin{aligned} w(t) - \delta a(t + 1) - \delta b(t + 1) &= \\ w(t) - \delta\alpha w(t + 1) - \delta^2(1 - \alpha)[a(t + 2) + b(t + 2)] + \delta c &= \\ [w(t) - \delta w(t + 1)] + \delta(1 - \alpha)[w(t + 1) - \delta a(t + 2) - \delta b(t + 2)] + \delta c &> 0, \end{aligned}$$

where, again, the inequality follows from constraint (6), the induction hypothesis, and $c > 0$.

Proof of Proposition 1.

We start by solving a relaxed problem in which we minimize $u(w)$ subject to the constraints (4) and (5). That is, we ignore the non-negativity constraints and constraint (6). Then we verify that the solution to the relaxed problem satisfies them.

Given a contract w , we define $\tilde{\varphi}(w), \varphi(0, w), \dots, \varphi(T-1, w)$ as follows:

$$\tilde{\varphi}(w) = u(0, p_0; w) - p_0 w(0) = p_0 \left[-(1-\alpha)w(0) + \alpha \sum_{t=1}^{T-1} (1-\alpha)^t \delta^t w(t) + (1-\alpha)^T \delta^{T-1} w(G) \right] + (1-p_0) \delta^{T-1} w(B) - \tilde{\psi},$$

and for every $t = 0, \dots, T-1$

$$\varphi(t, w) = u(t, p_t; w) - \delta u(t+1, p_t; w) = p_t \left[\alpha w(t) - \alpha^2 \sum_{t'=t+1}^{T-1} (1-\alpha)^{t'-t-1} \delta^{t'-t} w(t') - \alpha (1-\alpha)^{T-1-t} \delta^{T-1-t} w(G) \right] - \psi_t,$$

where $\tilde{\psi}, \psi_0, \dots, \psi_{T-1}$ are $T+1$ constants. $\tilde{\psi}$ is the agent's (discounted) expected cost of testing when he acquires and reveals the signal in every period. For every t , ψ_t represents the difference between the expected costs of two different strategies. Under the first strategy, the agent starts to acquire the signal in period t . Under the second strategy he shirks in period t and starts to acquire the signal in period $t+1$. For example, ψ_{T-1} is equal to c .

We now show that if \hat{w} is a solution to the relaxed problem, then

$$\begin{cases} \tilde{\varphi}(\hat{w}) = 0 \\ \varphi(t, \hat{w}) = 0, \quad t = 0, \dots, T-1 \\ \hat{w}(G) = 0. \end{cases} \quad (17)$$

Notice that $w(B)$ appears only in $\tilde{\varphi}(w)$ with a positive coefficient. If $\tilde{\varphi}(\hat{w}) > 0$ then we can lower $\hat{w}(B)$ by a small amount and all the constraints are still satisfied. Similarly, $w(0)$ appears only in $\tilde{\varphi}(w)$, with a negative coefficient, and in $\varphi(0, w)$, with a positive coefficient. Again, if $\varphi(0, \hat{w}) > 0$ we can lower $\hat{w}(0)$ by a small amount and all the constraints are still satisfied. Thus, $\tilde{\varphi}(\hat{w}) = \varphi(0, \hat{w}) = 0$ if \hat{w} is a solution to the relaxed problem.

Suppose that \hat{w} solves the relaxed problem and $\varphi(t, \hat{w}) > 0$ for some $t = 1, \dots, T-1$. Let \hat{t} denote the smallest integer for which the inequality holds. Then consider a new contract w' which is identical to \hat{w} except that we set

$$w'(\hat{t}) = \hat{w}(\hat{t}) - \varepsilon, \\ w'(0) = \hat{w}(0) - \alpha (1-\alpha)^{\hat{t}-1} \delta^{\hat{t}} \varepsilon$$

for some small positive ε . Notice that $\varphi(\hat{t}, w') > 0$ for ε sufficiently small. By construction, $\tilde{\varphi}(w') = \varphi(0, w') = 0$. For $t = 1, \dots, \hat{t} - 1$, $\varphi(t, w') > \varphi(t, \hat{w})$ since $\varphi(t, w)$ is decreasing in $w(\hat{t})$. Also, $\varphi(t, w') = \varphi(t, \hat{w})$ for $t > \hat{t}$. Thus, the contract w' , which is cheaper than \hat{w} , satisfies all the constraints of the relaxed problem.

Finally, suppose by contradiction that \hat{w} solves the relaxed problem and $\hat{w}(G) > 0$. Consider a new contract w' that is identical to \hat{w} except that we set

$$\begin{aligned} w'(G) &= 0, \\ w'(0) &= \hat{w}(0) - (1 - \alpha)^{T-1} \delta^{T-1} \hat{w}(G). \end{aligned}$$

We have $\tilde{\varphi}(w') = \varphi(0, w') = 0$ and $\varphi(t, w') > \varphi(t, \hat{w})$ for every $t > 0$. This contradicts the fact that \hat{w} is a solution to the relaxed problem.

Consider now the system of linear equations (17). It is easy to check that the (unique) solution w^* is given by equation (7) in Proposition 1.

Clearly, all the payments in w^* are non-negative. Also, notice that w^* satisfies all the constraints in (5) with equality. Consider $t = 0, \dots, T - 2$. We have

$$\begin{aligned} 0 &= u(t, p_t; w^*) - \delta u(t + 1, p_t; w^*) = \\ &= -c + p_t \alpha w^*(t) + (1 - p_t \alpha) \delta u(t + 1, p_{t+1}; w^*) + \\ &= \delta c - \delta p_t \alpha w^*(t + 1) - (1 - p_t \alpha) \delta^2 u(t + 2, p_{t+1}; w^*) = \\ &= -c + p_t \alpha w^*(t) + \delta c - \delta p_t \alpha w^*(t + 1), \end{aligned}$$

where the last equality follows from $u(t + 1, p_{t+1}; w^*) = \delta u(t + 2, p_{t+1}; w^*)$, constraint (5) at $t + 1$. Constraint (6) is therefore satisfied since for every $t = 0, \dots, T - 2$,

$$w^*(t) = \delta w^*(t + 1) + \frac{c(1 - \delta)}{p_t \alpha}.$$

Thus, the contract w^* defined in Proposition 1 solves the principal's problem.

Proof of Proposition 2.

Let (w^h, w^ℓ) be an optimal mechanism and suppose that $w^\ell(t) \neq w^*(t; p_0^\ell)$ for some t , and/or $w^\ell(G) \neq w^*(G; p_0^\ell)$, where $w^*(p_0^\ell)$ is the optimal contract defined in Proposition 1.

Let \hat{w}^ℓ denote the contract which is identical to $w^*(p_0^\ell)$ except that we set $\hat{w}^\ell(B)$ to satisfy

$$u(0, p_0^\ell; \hat{w}^\ell) = u(0, p_0^\ell; w^\ell).$$

Among all the contracts that are suitable for p_0^ℓ , $w^*(p_0^\ell)$ has the lowest payments in state G . Therefore, we have $\hat{w}^\ell(t) \leq w^\ell(t)$ for every t , with strict inequality at $t = 0$, $\hat{w}^\ell(G) \leq w^\ell(G)$, and $\hat{w}^\ell(B) > w^\ell(B)$. Clearly, the mechanism \hat{w}^ℓ is suitable for p_0^ℓ .

We now show that

$$u(0, p_0^h; w^\ell) > u(0, p_0^h; \hat{w}^\ell). \quad (18)$$

In fact, we have

$$\begin{aligned} & u(0, p_0^h; w^\ell) - u(0, p_0^h; \hat{w}^\ell) = \\ & [u(0, p_0^h; w^\ell) - u(0, p_0^\ell; w^\ell)] - [u(0, p_0^h; \hat{w}^\ell) - u(0, p_0^\ell; \hat{w}^\ell)] = \\ & (p_0^h - p_0^\ell) \left[\alpha \sum_{t=0}^{T-1} (1-\alpha)^t \delta^t w^\ell(t) + (1-\alpha)^T \delta^{T-1} w^\ell(G) - \delta^{T-1} w^\ell(B) \right] - \\ & (p_0^h - p_0^\ell) \left[\alpha \sum_{t=0}^{T-1} (1-\alpha)^t \delta^t \hat{w}^\ell(t) + (1-\alpha)^T \delta^{T-1} \hat{w}^\ell(G) - \delta^{T-1} \hat{w}^\ell(B) \right] = \\ & (p_0^h - p_0^\ell) \left[\alpha \sum_{t=0}^{T-1} (1-\alpha)^t \delta^t (w^\ell(t) - \hat{w}^\ell(t)) + \right. \\ & \left. (1-\alpha)^T \delta^{T-1} (w^\ell(G) - \hat{w}^\ell(G)) - \delta^{T-1} (w^\ell(B) - \hat{w}^\ell(B)) \right] > 0. \end{aligned}$$

This and the fact that the mechanism (w^h, w^ℓ) is incentive compatible imply

$$u(0, p_0^h; w^h) \geq u(0, p_0^h; w^\ell) > u(0, p_0^h; \hat{w}^\ell). \quad (19)$$

Notice also that

$$u(0, p_0^h; w^h) > p_0^h \hat{w}^\ell(0), \quad (20)$$

since $\hat{w}^\ell(0) < w^\ell(0)$ and (w^h, w^ℓ) is incentive compatible.

We now construct a contract \hat{w}^h such that the mechanism $(\hat{w}^h, \hat{w}^\ell)$ is incentive compatible and

$$u(0, p_0^h; w^h) > u(0, p_0^h; \hat{w}^h).$$

Clearly, this means that the original mechanism (w^h, w^ℓ) is not optimal.

The exact form of the contract \hat{w}^h depends on the contract w^h . We need to distinguish among three different cases. In what follows, ε denotes a small positive number.

(i) First suppose that $w^h(t) = w^*(t; p_0^h)$ for every t and $w^h(G) = w^*(G; p_0^h) = 0$. Notice that

$$w^h(0) = w^*(0; p_0^h) < w^*(0; p_0^\ell) = \hat{w}^\ell(0) < w^\ell(0),$$

and, thus,

$$u(0, p_0^h; w^h) \geq p_0^h w^\ell(0) > p_0^h w^h(0), \quad (21)$$

where the first inequality follows from the fact that under (w^h, w^ℓ) the high type does not have an incentive to choose the contract w^ℓ and guess state G in the first period.

In this case we let \hat{w}^h be identical to w^h except that we set

$$\hat{w}^h(B) = w^h(B) - \varepsilon.$$

It follows from inequalities (19)-(21) that for ε sufficiently small the new mechanism $(\hat{w}^h, \hat{w}^\ell)$ is incentive compatible.

(ii) Suppose now that there exists a period t such that the constraint

$$u(t, p_t^h; w^h) \geq \delta u(t+1, p_t^h; w^h) \quad (22)$$

is satisfied with strict inequality. In this case let \hat{t} denote the largest integer for which the above constraint is not binding. If $\hat{t} > 0$, we let

$$\hat{w}^h(\hat{t}) = w^h(\hat{t}) - \varepsilon,$$

and

$$\hat{w}^h(t) = w^h(t) - \alpha \delta^{\hat{t}-t} \varepsilon$$

for $t = 0, \dots, \hat{t} - 1$. The remaining payments of \hat{w}^h are equal to the corresponding payments of w^h .

If $\hat{t} = 0$, we let \hat{w}^h be identical to w^h except that we set

$$\hat{w}^h(0) = w^h(0) - \varepsilon.$$

For ε sufficiently small, the contract \hat{w}^h is suitable for p_0^h . It follows from inequalities (19) and (20) that under the contract $(\hat{w}^h, \hat{w}^\ell)$ the high type does not have an incentive to choose the contract \hat{w}^ℓ (provided that ε is small enough). It is also obvious that the low type does not have an incentive to choose the contract \hat{w}^h (every payment of \hat{w}^h is weakly smaller than the corresponding payment of w^h). We conclude that for ε sufficiently small the mechanism $(\hat{w}^h, \hat{w}^\ell)$ is incentive compatible.

(iii) Finally, suppose that constraint (22) is binding in every period. In this case we must have $w^h(G) > 0$ (if $w^h(G) = 0$ and all the constraints in (22) are binding then we are in case (i)). We let

$$\hat{w}^h(G) = w^h(G) - \varepsilon,$$

and

$$\hat{w}^h(t) = w^h(t) - \delta^{T-1-t} \varepsilon$$

for $t = 0, \dots, T - 1$. Furthermore, $\hat{w}^h(B) = w^h(B)$. Again, for ε sufficiently small the mechanism $(\hat{w}^h, \hat{w}^\ell)$ is incentive compatible (the proof is identical to the proof of case (ii)).

Proof of Lemma 3.

Before proceeding with the proof of the lemma, we need to establish a preliminary result. Given a contract w , recall the definition of $a(0; w)$ in equation (14) in the proof of Lemma 1.

Claim 1 *Suppose the contract w is suitable for some $p_0 \in (0, 1)$. Then $a(0; w) > 0$.*

Proof.

Fix a contract w . Recall from equation (15) that for every $p \in [0, 1]$,

$$u(0, p; w) = a(0; w) + b(0; w)p.$$

Thus, $a(0; w)$ coincides with $u(0, 0; w)$, the agent's expected utility in period 0 when his belief is 0 and he acquires and reveals the signal in every period. Then we have

$$a(0; w) = u(0, 0; w) = -c(1 + \delta + \dots + \delta^{T-1}) + \delta^{T-1}w(B).$$

Fix $p_0 \in (0, 1)$ and consider the following problem:

$$\begin{aligned} & \min_{w \geq 0} w(B) \\ & \text{s.t. } w \text{ is suitable for } p_0. \end{aligned}$$

It is immediate to check that the optimal contract $w^*(p_0)$ solves the above problem. We conclude that if w is suitable for p_0 , then

$$a(0; w) \geq a(0; w^*(p_0)) = c \left(\frac{1}{\alpha(1-p_0)} - 1 \right) > 0.$$

This concludes the proof of the claim.

We now continue with the proof of Lemma 3.

Suppose that the mechanism (w^h, w^ℓ) is incentive compatible. It follows from $a(0; w^h) > 0$ (the contract w^h is suitable for p_0^h), and the fact that the high type does not have an incentive to choose w^ℓ and guess state G in the first period,

$$u(0, p_0^h; w^h) = a(0; w^h) + b(0; w^h)p_0^h \geq p_0^h w^\ell(0),$$

that

$$u(0, p_0^\ell; w^h) = a(0; w^h) + b(0; w^h)p_0^\ell > p_0^\ell w^\ell(0).$$

Since the mechanism (w^h, w^ℓ) is incentive compatible the low type does not have an incentive to choose w^h and acquire and reveal the signal in every period. Therefore, we have

$$u(0, p_0^\ell; w^\ell) \geq u(0, p_0^\ell; w^h) > p_0^\ell w^\ell(0).$$

Proof of Lemma 4.

We first show that if (w^h, w^ℓ) is an optimal mechanism, then

$$u(0, p_0^h; w^h) = p_0^h w^h(0). \quad (23)$$

Suppose that the mechanism is incentive compatible and $u(0, p_0^h; w^h) > p_0^h w^h(0)$. We distinguish between two cases.

First suppose that

$$u(0, p_0^\ell; w^\ell) = p_0^\ell w^\ell(0). \quad (24)$$

We know from Lemma 3 that $u(0, p_0^\ell; w^\ell) > p_0^\ell w^\ell(0)$. This and equality (24) imply $w^h(0) > w^\ell(0)$. Furthermore, it follows from $a(0; w^\ell) > 0$ (see Claim 1) and equality (24) that $u(0, p_0^h; w^\ell) < p_0^h w^h(0)$. Recall that $u(0, p_0^h; w^h) > p_0^h w^h(0)$. Thus, we have $u(0, p_0^h; w^h) > p_0^h w^\ell(0)$ and $u(0, p_0^h; w^h) > u(0, p_0^h; w^\ell)$. Clearly, we can lower the payment $w^h(B)$ by a small amount and the new mechanism is still incentive compatible. This shows that the original mechanism (w^h, w^ℓ) is not optimal.

Thus, let us assume that

$$u(0, p_0^\ell; w^\ell) > p_0^\ell w^\ell(0). \quad (25)$$

Consider now a new contract \hat{w}^h for the high type which is identical to w^h except that we set

$$\hat{w}^h(0) = w^h(0) + \varepsilon$$

for some small positive ε and choose $\hat{w}^h(B) < w^h(B)$ such that

$$u(0, p_0^h; \hat{w}^h) = u(0, p_0^h; w^h).$$

Clearly, for ε sufficiently small the contract \hat{w}^h is suitable for p_0^h . Also, inequality (25) implies that for ε small enough, $u(0, p_0^\ell; \hat{w}^\ell) > p_0^\ell \hat{w}^\ell(0)$. Finally, for every $\sigma \in \Sigma'$,

$$u(\sigma, p_0^\ell; w^h) > u(\sigma, p_0^\ell; \hat{w}^h).$$

The proof of this inequality is identical to the proof of inequality (18), so we omit the details.

Notice that the new mechanism (\hat{w}^h, w^ℓ) is incentive compatible for ε sufficiently small, and all the constraints in which the low type lies about his type are satisfied with strict inequality. Also recall from Lemma 3 that $u(0, p_0^\ell; w^\ell) > p_0^\ell w^\ell(0)$. Therefore, we can decrease the payment $w^\ell(B)$ by a small amount and the new mechanism is still incentive compatible. But then the original mechanism (w^h, w^ℓ) cannot be optimal (notice that, by construction, the mechanisms (w^h, w^ℓ) and (\hat{w}^h, w^ℓ) have the same expected cost).

Next, we show that if (w^h, w^ℓ) is an optimal mechanism, then $w^h(0) = w^\ell(0) = w^*(0; p_0^\ell)$.

Suppose that the mechanism (w^h, w^ℓ) is optimal. Equality (23) and the fact that (w^h, w^ℓ) is incentive compatible immediately imply $w^h(0) \geq w^\ell(0)$. By contradiction, suppose that $w^h(0) > w^\ell(0)$.

First, assume that $u(0, p_0^h; w^h) = u(0, p_0^h; w^\ell)$. Suppose that the principal offers the mechanism (w^ℓ, w^ℓ) , i.e. the same contract w^ℓ to both types. Since

$$u(0, p_0^h; w^\ell) = u(0, p_0^h; w^h) = p_0^h w^h(0) > p_0^h w^\ell(0),$$

the contract w^ℓ is suitable for both types. However, notice that under the contract w^ℓ each type strictly prefers to acquire and reveal the signal in every period rather than guess state G in period 0. Thus, if we lower the payment $w^\ell(B)$ by a small amount, the new contract will remain suitable for both types. By construction, the mechanisms (w^h, w^ℓ) and (w^ℓ, w^ℓ) have the same expected cost. But then the original mechanism (w^h, w^ℓ) cannot be optimal.

To conclude the proof, suppose that (w^h, w^ℓ) is optimal, $w^h(0) > w^\ell(0)$, and $u(0, p_0^h; w^h) > u(0, p_0^h; w^\ell)$. Notice that it cannot be the case that $w^h(G) = 0$ and w^h satisfies the constraint

$$u(t, p_t^h; w^h) \geq \delta u(t+1, p_t^h; w^h)$$

with equality in every period. If this were the case, then $w^h(0) = w^*(0; p_0^h) < w^*(0; p_0^\ell) = w^\ell(0)$.

If the above constraint is satisfied with strict inequality for some t , then we construct a new contract \hat{w}^h as in case (ii) in the proof of Proposition 2. Otherwise, if all the constraints are binding, then we construct a new contract \hat{w}^h as in case (iii) in the proof of Proposition 2.

The fact that $u(0, p_0^h; w^h) > u(0, p_0^h; w^\ell)$ implies that for ε sufficiently small, the new mechanism (\hat{w}^h, w^ℓ) is incentive compatible. Also, the expected cost of (\hat{w}^h, w^ℓ) is strictly smaller than the expected cost of the original mechanism (w^h, w^ℓ) which, therefore, cannot be optimal.

Proof of Proposition 3.

Given a contract w^h and a prior $p \in (0, 1)$, we define

$$\begin{aligned} \Lambda(p, w^h) &= u(0, p; w^h) - u(0, p; \bar{w}) = \\ p &\left[\sum_{t=0}^{T-1} (1-\alpha)^t \alpha \delta^t (w^h(t) - \bar{w}(t)) + (1-\alpha)^T \delta^{T-1} (w^h(G) - \bar{w}(G)) \right] + \\ &(1-p) \delta^{T-1} (w^h(B) - \bar{w}(B)). \end{aligned}$$

We start with two preliminary observations. First, if a mechanism (w^h, w^ℓ) is optimal, then $w^h(B) \leq \bar{w}(B)$. The proof of this result is by contradiction. Notice that if (w^h, w^ℓ) is optimal, then $\Lambda(p_0^h, w^h) = 0$. This and $w^h(B) > \bar{w}(B)$ would imply $\Lambda(p_0^h, w^h) > 0$, and, thus,

$$u(0, p^\ell; w^\ell) \geq u(0, p^\ell; w^h) > u(0, p^\ell; \bar{w}).$$

We conclude that the mechanism (w^h, w^ℓ) is not optimal, since it is more expensive than the incentive compatible mechanism (\bar{w}, \bar{w}) .

Second, it is immediate to see that if (w^h, w^ℓ) is optimal and $w^h(B) = \bar{w}(B)$, then the mechanism (\bar{w}, \bar{w}) is also optimal.

We are now ready to prove that (\bar{w}, \bar{w}) is optimal when $p_0^\ell \in [p_1^h, p_0^h)$. It is enough to show that $u(0, p_0^\ell; w^\ell) \geq u(0, p_0^\ell; \bar{w})$ for any mechanism (w^h, w^ℓ) that satisfies: (i) $u(0, p_0^h; w^h) = u(0, p_0^h; \bar{w})$; (ii) $w^h(0) = \bar{w}(0)$; and (iii) $w^h(B) < \bar{w}(B)$.

Let σ^1 denote the strategy under which the agent shirks in the first period and acquires and reveals the signal in every other period $t > 0$. Using $w^h(0) = \bar{w}(0)$ we can rewrite $\Lambda(p_0^h, w^h) = 0$ as

$$\begin{aligned} p_0^h (1-\alpha) &\left[\sum_{t=1}^{T-1} (1-\alpha)^{t-1} \alpha \delta^t (w^h(t) - \bar{w}(t)) + (1-\alpha)^{T-1} \delta^{T-1} (w^h(G) - \bar{w}(G)) \right] + \\ &(1-p_0^h) \delta^{T-1} (w^h(B) - \bar{w}(B)) = 0. \end{aligned}$$

We divide both sides by $(1-p_0^h\alpha)$ and obtain

$$\begin{aligned} p_1^h &\left[\sum_{t=1}^{T-1} (1-\alpha)^{t-1} \alpha \delta^t (w^h(t) - \bar{w}(t)) + (1-\alpha)^{T-1} \delta^{T-1} (w^h(G) - \bar{w}(G)) \right] + \\ &(1-p_1^h) \delta^{T-1} (w^h(B) - \bar{w}(B)) = 0. \end{aligned}$$

Recall that $w^h(B) < \bar{w}(B)$ and $p_0^\ell \geq p_1^h$. This and the above equality imply

$$\begin{aligned} 0 \leq p_0^\ell &\left[\sum_{t=1}^{T-1} (1-\alpha)^{t-1} \alpha \delta^t (w^h(t) - \bar{w}(t)) + (1-\alpha)^{T-1} \delta^{T-1} (w^h(G) - \bar{w}(G)) \right] + \\ &(1-p_0^\ell) \delta^{T-1} (w^h(B) - \bar{w}(B)) = u(\sigma^1, p_0^\ell; w^h) - u(\sigma^1, p_0^\ell; \bar{w}). \end{aligned} \tag{26}$$

Finally,

$$u(0, p_0^\ell; \bar{w}) = u(\sigma^1, p_0^\ell; \bar{w}) \leq u(\sigma^1, p_0^\ell; w^h) \leq u(0, p_0^\ell; w^\ell),$$

where the equality follows from the fact constraint (5) is binding when the contract is \bar{w} , $t = 0$ and the prior is p_0^ℓ (notice that $u(\sigma^1, p_0^\ell; \bar{w}) = \delta u(1, p_0^\ell; \bar{w})$), while the last inequality holds because (w^h, w^ℓ) is incentive compatible.

We now turn to uniqueness. Suppose that $p_0^\ell \in (p_1^h, p_0^h)$. One can immediately check that in this case inequality (26) is strict. Thus, if (w^h, w^ℓ) is optimal, then $w^h(B) = \bar{w}(B)$. It is possible to show that

$$\max_{\sigma \in \Sigma'} u(\sigma, p_0^\ell; w^h) > u(0, p_0^\ell; \bar{w})$$

for any contract w^h that satisfies (i) $u(0, p_0^h; w^h) = u(0, p_0^h; \bar{w})$; (ii) $w^h(0) = \bar{w}(0)$; and (iii) $w^h(B) = \bar{w}(B)$. For brevity, we omit the proof of this claim. Of course, this shows that when $p_1^\ell > p_0^h$ there is no optimal mechanism other than (\bar{w}, \bar{w}) .

It remains to show that (\bar{w}, \bar{w}) is not optimal when $p_0^\ell < p_1^h$. Consider the mechanism (w^h, w^ℓ) , defined as follows. Let ε denote a small positive number. The contract w^h is identical to \bar{w} except that we set

$$\begin{aligned} w^h(1) &= \bar{w}(1) + \varepsilon, \\ w^h(B) &= \bar{w}(B) - \frac{p_0^h(1-\alpha)\alpha}{(1-p_0^h)\delta^{T-2}}\varepsilon. \end{aligned}$$

The contract w^ℓ is identical to \bar{w} except that we set

$$w^\ell(B) = \bar{w}(B) - \varepsilon.$$

It is easy to check that $u(0, p_0^h; w^h) = u(0, p_0^h; \bar{w})$ and that, for ε sufficiently small, w^h is suitable for p_0^h .

The fact that $p_0^\ell < p_1^h$ implies that for every $\sigma \in \Sigma'$,

$$u(\sigma, p_0^\ell; \bar{w}) > u(\sigma, p_0^\ell; w^h).$$

By definition,

$$u(0, p_0^\ell; \bar{w}) > p_0^\ell \bar{w}(0) = p_0^\ell w^\ell(0) = p_0^\ell w^h(0).$$

We conclude that for ε sufficiently small the mechanism (w^ℓ, w^h) is incentive compatible. Therefore, (\bar{w}, \bar{w}) is not optimal.

Proof of Proposition 4.

We start with a preliminary observation. If $(\hat{w}^h, \hat{w}^\ell)$ is an optimal mechanism, then

$$u(1, 1; \hat{w}^h) \geq u(1, 1; \hat{w}^\ell), \quad (27)$$

where one should recall that $u(t, 1; \hat{w}^k)$, $t = 0, \dots, T-1$, denotes the agent's expected utility, computed in period t , when his beliefs is one, he acquires and reveals the signal in every period $t' \geq t$, and the contract is \hat{w}^k . Inequality (27) follows from $\hat{w}^h(0) = \hat{w}^\ell(0) = w^*(0; p_0^\ell)$, and the fact that given the mechanism $(\hat{w}^h, \hat{w}^\ell)$, it is not profitable for the type p_0^k , $k = h, \ell$, to choose $\hat{w}^{k'}$, $k' \neq k$, and acquire and reveal the signal in every period.

Let $(\hat{w}^h, \hat{w}^\ell)$ be an optimal contract (recall that there exists a solution to the principal's problem) and let $\Delta \geq 0$ be equal to

$$\Delta = u(0, p_0^h; \hat{w}^h) - \delta u(1, p_0^h; \hat{w}^h).$$

This and the constraints

$$u(0, p_0^h; \hat{w}^h) = p_0^h \hat{w}^h(0) = p_0^h w^*(0; p_0^\ell) \quad (28)$$

have the following implications:

$$\hat{w}^h(B) = w^*(B; p_0^h) + \frac{(1-\alpha)\Delta}{(1-p_0^h)\alpha\delta^{T-1}}, \quad (29)$$

and

$$u(1, 1; \hat{w}^h) = \frac{w^*(0; p_0^\ell)}{\delta} - \frac{c}{p_0^h \alpha \delta} - \frac{\Delta}{p_0^h \alpha \delta}. \quad (30)$$

From equations (28)-(30) we can derive the values of $u(0, p_0^\ell; \hat{w}^h)$ and $\delta u(1, p_0^\ell; \hat{w}^h)$, which for brevity we refer to as v_0 and v_1 , respectively:

$$u(0, p_0^\ell; \hat{w}^h) = \left(\frac{p_0^h - p_0^\ell}{p_0^h} \right) \left(-c + \frac{c + (1-\alpha)\Delta}{(1-p_0^h)\alpha} \right) + p_0^\ell w^*(0; p_0^\ell) := v_0,$$

and

$$\delta u(1, p_0^\ell; \hat{w}^h) = p_0^\ell \left(w^*(0; p_0^\ell) - \frac{c}{p_0^h \alpha} - \frac{\Delta}{p_0^h \alpha} \right) + (1-p_0^\ell) \left(\frac{c + (1-\alpha)\Delta}{(1-p_0^h)\alpha} \right) := v_1.$$

Thus, v_t , $t = 0, 1$, denotes the utility of the low type when he chooses the contract \hat{w}^h and he acquires and reveals the signal in every period $t' \geq t$ (before t the agent shirks and announces message B).

Recall from the definition of $u(T, p; w)$ in equation (3), that $u(T, 1; \hat{w}^h)$ is equal to $\frac{\hat{w}^h(G)}{\delta}$. For every $t = 2, \dots, T$, we have

$$u(1, 1, \hat{w}^h) = \sum_{t'=1}^{t-1} (1-\alpha)^{t'-1} \delta^{t'-1} (\alpha \hat{w}^h(t') - c) + (1-\alpha)^{t-1} \delta^{t-1} u(t, 1; \hat{w}^h).$$

This and the fact that for every t , $\hat{w}^h(t) \leq \frac{w^*(0; p_0^\ell)}{\delta^t}$ imply

$$\delta^t u(t, 1; \hat{w}^h) \geq \frac{1}{(1-\alpha)^{t-1}} \left[w^*(0; p_0^\ell) \left(1 - \alpha \sum_{t'=1}^{t-1} (1-\alpha)^{t'-1} \right) - \frac{c}{p_0^h \alpha} + c \sum_{t'=1}^{t-1} (1-\alpha)^{t'-1} \delta^{t'} - \frac{\Delta}{p_0^h \alpha} \right].$$

Finally, using this inequality and the definition of $\hat{w}^h(B)$ in equation (29) we have that for every $t = 2, \dots, T$,

$$\begin{aligned} \delta^t u(t, p_0^\ell; \hat{w}^h) &= \delta^t p_0^\ell u(t, 1; \hat{w}^h) + \delta^t (1 - p_0^\ell) u(t, 0; \hat{w}^h) \geq \\ &\frac{p_0^\ell}{(1-\alpha)^{t-1}} \left[w^*(0; p_0^\ell) \left(1 - \alpha \sum_{t'=1}^{t-1} (1-\alpha)^{t'-1} \right) - \frac{c}{p_0^h \alpha} + c \sum_{t'=1}^{t-1} (1-\alpha)^{t'-1} \delta^{t'} - \frac{\Delta}{p_0^h \alpha} \right] + \\ &(1 - p_0^\ell) \left[-c (\delta^t + \dots + \delta^{T-1}) + \delta^{T-1} \left(w^*(B) + \frac{(1-\alpha)\Delta}{(1-p_0^h)\alpha\delta^{T-1}} \right) \right] := v_t. \end{aligned}$$

Therefore, for $t = 2, \dots, T$, v_t is a lower bound to the utility that the lower type can obtain when he chooses the contract \hat{w}^h and he starts to acquire and reveal the signal in period t (he shirks and sends message B before t).

We conclude that under the optimal mechanism $(\hat{w}^h, \hat{w}^\ell)$, the utility of the low type is bounded below by

$$u(0, p_0^\ell; \hat{w}^\ell) \geq \max\{v_0, \dots, v_T\}. \quad (31)$$

We are now ready to construct an optimal mechanism (w^h, w^ℓ) which satisfies the conditions in Proposition 4. We set $w^\ell = \hat{w}^\ell$, $w^h(0) = \hat{w}^h(0)$, and $w^h(B) = \hat{w}^h(B)$. The rest of the contract w^h is constructed using an algorithm that involves T steps.

In step 1, we set $w^h(t) = w^*(t; p_0^\ell)$, for $t = 2, \dots, T-1$, and $w^h(G) = w^*(G; p_0^\ell) = 0$. Also, we choose $w^h(1)$ to solve

$$u(1, 1; w^h) = \frac{w^*(0; p_0^\ell)}{\delta} - \frac{c}{p_0^h \alpha \delta} - \frac{\Delta}{p_0^h \alpha \delta} = u(1, 1; \hat{w}^h). \quad (32)$$

It follows from inequality (27) and the way the mechanism (w^h, w^ℓ) is designed in step 1 that the solution $w^h(1)$ to the equation (32) must be weakly greater than $w^*(1; p_0^\ell)$.¹⁶ If the solution $w^h(1)$ is weakly smaller than $\frac{w^*(0; p_0^\ell)}{\delta}$, then the algorithm stops at step 1. Otherwise, we set $w^h(1) = \frac{w^*(0; p_0^\ell)}{\delta}$ and move to step 2.

Next, we describe step $t = 2, \dots, T - 1$. We set $w^h(t') = \frac{w^*(0; p_0^\ell)}{\delta^{t'}}$ for $t' < t$, $w^h(t') = w^*(t'; p_0^\ell)$ for $t' = t + 1, \dots, T - 1$, and $w^h(G) = w^*(G; p_0^\ell) = 0$. Finally, we choose $w^h(t)$ to solve the equation (32). If the solution $w^h(t)$ is weakly smaller than $\frac{w^*(0; p_0^\ell)}{\delta^t}$, then the algorithm stops at step t . Otherwise we set $w^h(t) = \frac{w^*(0; p_0^\ell)}{\delta^t}$ and move to step $t + 1$. Notice that the fact that the algorithm reaches step t implies that the solution $w^h(t)$ to the equation (32) must be greater than $w^*(t; p_0^\ell)$.

Finally, in step T , we set $w^h(t) = \frac{w^*(0; p_0^\ell)}{\delta^t}$ for every $t = 1, \dots, T - 1$, and choose $w^h(G)$ to solve the equation (32). It is easy to check that if the algorithm reaches step T then the solution $w^h(G)$ is positive and weakly smaller than

$$w^h(T - 1) - \frac{c}{p_{T-1}^h \alpha} = \frac{w^*(0; p_0^\ell)}{\delta^{T-1}} - \frac{c}{p_{T-1}^h \alpha}.$$

We now show that the mechanism (w^h, w^ℓ) is incentive compatible and

$$\max_{\sigma \in \Sigma'} u(\sigma, p_0^\ell; w^h) = \max\{v_0, \dots, v_T\}. \quad (33)$$

Given inequality (31), this is clearly enough to conclude that the mechanism (w^h, w^ℓ) is optimal.

It is immediate to see that the contracts w^h and w^ℓ are suitable for p_0^h and p_0^ℓ , respectively.¹⁷ And since $w^h(0) = w^\ell(0)$, it is not profitable for the type p_0^k , $k = h, \ell$, to choose $\hat{w}^{k'}$, $k' \neq k$, and guess in the first period that the state is G . Also, by construction,

$$u(0, p_0^h; w^h) = u(0, p_0^h; \hat{w}^h) \geq u(0, p_0^h; \hat{w}^\ell) = u(0, p_0^h; w^\ell).$$

It remains to check equality (33). Suppose that the algorithm used to construct w^h stops at step $\hat{t} = 1, \dots, T$. After \hat{t} , the payments of w^h in state G coincide with the payments of $w^*(p_0^\ell)$ and $w^h(\hat{t}) \geq w^*(\hat{t}; p_0^\ell)$. Thus, the low type does not have an incentive to shirk

¹⁶The solution $w^h(1)$ must be strictly greater than $w^*(1; p_0^\ell)$ if Δ is equal to zero.

¹⁷Notice that if for some t , $u(t, p_t^h; w^h) \geq \delta u(t, p_t^h; w^h)$ and $w^h(t) = \delta w^h(t + 1)$, then it must be the case that $u(t + 1, p_{t+1}^h; w^h) \geq \delta u(t + 2, p_{t+1}^h; w^h)$.

from period \hat{t} on. Formally, if $\sigma \in \Sigma'$ is a strategy under which the low type shirks in period $t \geq \hat{t}$, then there exists another strategy $\sigma' \in \Sigma'$ with $u(\sigma', p_0^\ell; w^h) \geq u(\sigma, p_0^\ell; w^h)$.

Let us now restrict attention to the strategies in Σ' under which the low type works in every period $t \geq \hat{t}$. If σ is a strategy under which the agent works in period t and shirks in some period t' , $t < t' < \hat{t}$, then σ is a weakly dominated strategy. To see this, consider the strategy σ . We must be able to find a period $\tilde{t} < \hat{t} - 1$, such that the low type works in \tilde{t} and shirks in $\tilde{t} + 1$. Let σ' be a strategy which is identical to σ except that the low type shirks in \tilde{t} and works in $\tilde{t} + 1$. At the beginning of period \tilde{t} , the low type has the same belief, say p , both when he uses σ and when he uses σ' . Thus, the difference between the agent's continuation payoffs of the strategies σ and σ' , computed at the beginning of period \tilde{t} , is equal to

$$-c + p\alpha w^h(\tilde{t}) + \delta c - p\alpha\delta w^h(\tilde{t} + 1) = -c + \delta c \leq 0,$$

where the equality follows from $w^h(\tilde{t} + 1) = \frac{w^h(\tilde{t})}{\delta}$. Thus, $u(\sigma', p_0^\ell; w^h) \geq u(\sigma, p_0^\ell; w^h)$.

To sum up, given the contract w^h , it is optimal for the low type to use one of the following strategies: $\sigma^0, \dots, \sigma^{\hat{t}}$. For $t = 0, \dots, \hat{t}$, σ^t denotes the strategy under which the low type starts to acquire and reveal the signal in period t (the agent shirks and sends message B before t). Recall that $w^h(t) = \frac{w^*(0; p_0^\ell)}{\delta^t}$ for every $t < \hat{t}$ since we are considering the case in which the algorithm stops at step \hat{t} . Thus, for every $t = 0, \dots, \hat{t}$, we have

$$u(\sigma^t, p_0^\ell; w^h) = \delta^t u(t, p_0^\ell; w^h) = v_t.$$

Therefore, we conclude that the mechanism (w^h, w^ℓ) is optimal.

Finally, we show that $w^h(1) > w^*(1; p_0^\ell) = w^\ell(1)$. This implies $w^h(B) < w^\ell(B)$. In fact, if $w^h(B) \geq w^\ell(B)$ then the mechanism (w^h, w^ℓ) is not incentive compatible (recall that $w^h(t) \geq w^\ell(t)$ for every t , with a strict inequality at $t = 1$, and $w^h(G) \geq w^\ell(G)$).

We need to distinguish between two cases. First, suppose that $p_0^\ell < p_1^h$. If the contract w^h constructed using the above algorithm is such that $w^h(1) = w^*(1; p_0^\ell)$, then it is optimal to offer the same contract to both types. But this contradicts Proposition 3.

Second, suppose that $p_0^\ell = p_1^h$. It is possible to show that if $p_0^\ell = \frac{p_0^h(1-\alpha)^t}{p_0^h(1-\alpha)^t + 1 - p_0^h}$ for some $t = 1, 2, \dots$, then there exists an optimal mechanism $(\hat{w}^h, \hat{w}^\ell)$ such that

$$u(0, p_0^h; \hat{w}^h) - \delta u(1, p_0^h; \hat{w}^h) = 0. \quad (34)$$

For brevity, we omit the proof of this claim. We then start from $(\hat{w}^h, \hat{w}^\ell)$ and construct an optimal mechanism using the above algorithm. Equality (34) guarantees that the payment $w^h(1)$ identified in step 1 is strictly greater than $w^*(1; p_0^\ell)$ (see footnote 16).

References

- [1] Baron, D. P., and R. B. Myerson (1982): “Regulating a Monopolist with Unknown Costs,” *Econometrica*, 50, 911-930.
- [2] Bergemann, D., and U. Hege (1998): “Venture Capital Financing, Moral Hazard, and Learning,” *Journal of Banking and Finance*, 22, 703-735.
- [3] Bergemann, D., and U. Hege (2005): “The Financing of Innovation: Learning and Stopping,” *Rand Journal of Economics*, 36, 719-752.
- [4] Cremer, J., F. Khalil, and J.-C. Rochet (1998): “Contracts and Productive Information Gathering,” *Games and Economic Behavior*, 25, 174-193.
- [5] Gromb, D., and D. Martimort (2007): “Collusion and the Organization of Delegated Expertise,” *Journal of Economic Theory*, 137, 271-299.
- [6] Holmstrom, B., and P. Milgrom (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica*, 55, 303-328.
- [7] Hörner, J. and L. Samuelson (2010): “Incentives for Experimenting Agents,” mimeo.
- [8] Klein, N. and T. Mylovannov (2010): “Expert Experimentation,” mimeo.
- [9] Laffont, J.-J., and J. Tirole (1986): “Using Cost Observation to Regulate Firms,” *Journal of Political Economy*, 94, 614-641.
- [10] Lewis, T. R., and M. Ottaviani (2008): “Search Agency,” mimeo.
- [11] Lewis, T. R., and D. E. M. Sappington (1997): “Information Management in Incentive Problems,” *Journal of Political Economy*, 105, 796-821.
- [12] Manso, G. (2010): “Motivating Innovation,” *Journal of Finance*, forthcoming.
- [13] Mussa, M., and S. Rosen (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18, 301-317.
- [14] Olszewski, W., and M. Pęski (2010): “The principal-agent approach to testing experts,” *American Economic Journals: Microeconomics*, forthcoming.