

Collegio Carlo Alberto



A Principal-Agent Model of Sequential Testing

Dino Gerardi
Lucas Maestri

No. 115
June 2009

Carlo Alberto Notebooks

www.carloalberto.org/working_papers

© 2009 by Dino Gerardi and Lucas Maestri. Any opinions expressed here are those of the authors and not those of the Collegio Carlo Alberto.

A Principal-Agent Model of Sequential Testing*

Dino Gerardi[†] Lucas Maestri[‡]

June 2009

Abstract

This paper analyzes the optimal provision of incentives in a sequential testing context. In every period the agent can acquire costly information that is relevant to the principal's decision. Neither the agent's effort nor the realizations of his signals are observable. First, we assume that the principal and the agent are symmetrically informed at the time of contracting. We construct the optimal mechanism and show that the agent is indifferent in every period between performing the test and sending an uninformative message which continues the relationship. Furthermore, in the first period the agent is indifferent between carrying out his task and sending an uninformative message which ends the relationship immediately. We then characterize the optimal mechanisms when the agent has superior information 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 thank Dirk Bergemann, James Choy, Johannes Hörner, Rick Mansfield, Ben Polak, Larry Samuelson, and participants in several seminars for their helpful comments.

[†]Department of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06511, donato.gerardi@yale.edu.

[‡]Department of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06511, lucas.maestri@yale.edu.

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, for example, a financial institution that contemplates the possibility of investing in a pharmaceutical company. The future value of the company depends on whether it will develop a new drug to treat a certain disease. The financial institution may hire an expert (a scientist) to execute a number of costly tasks (perform independent tests, read the scientific literature, etc.) and figure out whether the company will be successful or not.

Or consider a politician who has to take a stand on a controversial issue. The politician prefers to support a certain position unless it alienates some of his constituency. Thus, the politician seeks advice from a team of consultants. The consultants conduct a number of interviews to determine how the voters feel about the issue. Then they make a recommendation.

To give a third and final example, suppose that a sports team has the possibility of hiring a young player. It is not clear whether the player will turn out to be a star or not. The team consults a professional sports scout. The expert tries to predict the player's future by analyzing his performances and comparing him to other players observed in the past.

The above examples and most of the cases in which an expert (the agent) acquires information on behalf of a decision maker (the principal) share some similarities. First, the process of information acquisition is dynamic. Experts usually do not base their recommendations on a single piece of evidence. On the contrary, they often accumulate information and refine their beliefs over several periods of time before advising the principal. Second, the effort that the agents exert to become informed and the evidence that they find are not observable to the principal. It is difficult (if not impossible) for the principal to monitor the agent and check that he carries out his task. Similarly, the principal may lack the skills to interpret the agent's findings correctly. Finally, the principal and the agent are not equally informed at the beginning of their relationship. Because of his past experience the agent may know facts that the principal ignores or he may interpret the publicly available information in a more sophisticated way. Therefore, it is reasonable to assume that the agent has superior information.

In spite of this, the principal can still motivate the agent to invest in information acquisition and share his discoveries. In fact, after the principal makes a decision some information becomes publicly available. In the above examples, it will become evident whether the pharmaceutical company successfully develops the new drug, how the position

taken by the politician affects his popularity, or whether the young player is a star. 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 our model, a principal has to make a risky decision by a certain deadline. The unknown state of the world can be either good or bad. The principal hires an agent to perform a number of costly tests. The agent can complete at most one test in every period. 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.

Our signal structure fits well in many situations, like the examples mentioned above, where it is possible to find definitive evidence in favor of a certain state. For example, a scientist who discovers that the drug has serious side effects can safely conclude that it will not be produced. The politician's uncertainty is resolved once the experts become aware of voters who strongly oppose the suggested position. In our third example, only a very talented player can offer an exceptional performance. However, even great players perform poorly in some occasions.¹

The principal has the ability to commit to a mechanism. This is a contract which specifies all the possible payments to the agent. The payments depend on the agent's messages and on the state of the world which is realized after the principal makes his decision. The payments cannot depend on the agent's effort or on the realizations of the signals since these are not observable (or they are observable but not verifiable). Both the principal and the agent are risk neutral. However, the agent is protected by limited liability and cannot make transfers to the principal. Thus, it is impossible to sell the project to the agent. The goal of the principal is 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 (clearly, information acquisition becomes useless once the agent is certain that the state is good).

Since neither the effort nor the signals are observable, the principal's 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,

¹For reasons of exposition, we refer to the states and the signals as "good" or "bad" but these labels have no intrinsic meaning. As our examples suggest, depending on the particular application that one has in mind it could be more appropriate to identify as good one state (signal) or the other.

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. Suppose that the agent Shirks and announces a message in favor of the bad state. Compared to the agent, the principal is more pessimistic that the state is good. In fact, the two players interpret the same message differently. While the principal believes that the message reflects the signal accurately, the agent is clearly aware that it carries no informational content.

To ease the exposition and to develop some intuition, we start with the simpler case in which the agent has no private information at the outset of the relationship. In other words, the principal and the agent start the game with the same prior about the state. Without loss of generality, we restrict attention to contracts in which the agent receives a payment only when the relationship ends. This can occur in any period $t = 0, \dots, T - 1$ if the agent announces a message in favor of the good state. Otherwise the relationship ends in the final period $T - 1$. Thus, a contract consists of $T + 2$ possible payments. For $t = 0, \dots, T - 1$, there is the payment that the agent receives in t if he announces, in that period, that the state is good and he is correct. Furthermore, there are two additional payments in the last period, one for each state. These are the payments that the agent receives when all his messages are in favor of the bad state.

We show that the optimal contract is unique and we characterize it. The contract is such that in every period the agent is indifferent between carrying out the test (and revealing its realization) and sending an uninformative message in favor of the bad state. Furthermore, in the first period the agent is indifferent between the equilibrium strategy (i.e., acquiring and revealing the signal in every period) and guessing that the state is good. However, the agent has a strict incentive to reveal his information as soon as he discovers definitive evidence in favor of the good state.

The agent obtains a positive information rent which can be divided 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 component is due to the presence of hidden information, i.e. the fact that the results of the tests are unobservable.² 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. Thus, the principal needs to make larger payments in

²We 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.

the later periods to motivate the agent to execute the test. Because of these larger payments near the deadline, the agent finds it more profitable to shirk in the initial periods. And larger incentives to deviate translate into 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. It is much safer to acquire the signal and make a very accurate prediction of the state. Consequently, the signal accuracy tends to decrease the hidden information rent.

In the second part of the paper, we analyze the case in which the agent has superior information at the time of contracting. For tractability, we assume that 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. In this case the mechanism consists of a pair of contracts, one for each type.

The contract of the low type is almost identical to the optimal contract when the prior of the low type is known to the principal. Only the payment in the bad state is distorted upwards. All the other payments coincide under the two contracts. This implies that the low type continues to be indifferent between working and shirking in every period. On the other hand, all the payments to the high type in the good state are distorted upwards. Moreover, he has a strict incentive to execute the test in every period.

We show that when the priors of the two types are sufficiently close to each other the principal prefers to offer the same contract to both types. However, when the distance between the priors is sufficiently large the principal prefers to separate the two types and offers two different contracts. In any case, the information rent of each type is strictly larger than the rent obtained when the prior is known to the principal. That is, the new source of private information generates an additional rent. 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.

Our study is closely related to Gromb and Martimort (2007) who analyze a similar problem of delegated expertise. In Gromb and Martimort (2007) there are two noisy signals that can be observed before the principal makes a decision. The principal does not have direct access to them. Instead, he must hire one or two experts to acquire and report the signals. As in our model, moral hazard, hidden information and limited liability are all features of the agency problem. The principal can either hire one expert to acquire two signals or he can hire two experts to acquire one signal each. Gromb and Martimort compare the two organizational forms under various forms of collusion. In contrast, in this paper the principal hires a single expert and tries to induce him to acquire an arbitrary number of signals. Our focus is on the dynamics of the optimal contract and on the provision of intertemporal incentives. Furthermore, we consider the case of adverse selection while in Gromb and Martimort (2007) the principal and the experts are symmetrically informed at

the time of contracting.³

This paper contributes to the extensive literature on dynamic agency. Initial contributions such as Green (1987), Holmstrom and Milgrom (1987), Spear and Srivastava (1987) and Atkenson (1991) focus on moral hazard. In every period, the agent's unobservable action affects the probability distribution of the observable outcome. The literature has also analyzed the case in which the agent receives private information over time. For example, DeMarzo and Sannikov (2006), Biais, Mariotti, Plantin, and Rochet (2007), and DeMarzo and Fishman (2007) consider dynamic contracting environments in which the agent (entrepreneur) privately observes the cash flows of a certain project. The probability distribution of the cash flows is publicly known and, thus, there is no learning. In contrast, the interaction among learning, moral hazard and adverse selection (the initial information asymmetry) is a key element in our analysis.

Among the articles that analyze the effect of learning in dynamic agency, Bergemann and Hege (1998, 2005) are particular close to our work, especially in terms of information structure. They study the provision of venture capital. The quality of the project is unknown to the entrepreneur and the venture capitalist. The successful completion of the project depends both on its quality and the volume of financing it receives. Investing without success induces the entrepreneur and the venture capitalist to lower their expectations about the quality of the project. The entrepreneur controls the allocation of the funds which is, therefore, subject to moral hazard. There are three differences between Bergemann and Hege (1998, 2005) and our study. First, the outcome of the entrepreneur's investment is verifiable. In particular, the entrepreneur cannot generate a success unless he invests in the project. In contrast, in our model the agent invests in information acquisition and the realization of the signal is private information. Therefore, the agent can guess that the state is good. Second, we allow for the possibility that the principal and the agent are asymmetrically informed at the time of contracting. In our context, a mechanism must give incentives to the agent to reveal his initial information. Such incentives are also absent in the venture capital model. The third difference concerns who has the bargaining power. In Bergemann and Hege (1998, 2005), there is a competitive market of venture capitalists and the entrepreneur (i.e., the agent) has all the bargaining power. In contrast, we assume that the principal chooses the mechanism to maximize his welfare.

DeMarzo and Sannikov (2008) study a dynamic model in which the investor and the agent (who are equally informed at the outset of the relationship) learn over time about the profitability of the firm. The two players have different beliefs if the agent diverts the resources of the firm to private ends. The goal of the paper is to design the optimal

³We also use a different information structure. Our optimal contract with two signals (and no adverse selection) differs from the optimal sequential contract with one expert of Gromb and Martimort (2007).

dynamic contract. Manso (2007) studies a principal-agent model of experimentation. In each period the agent can choose among different activities. The productivity of one of them is unknown. In the first period the agent can experiment, i.e. he can choose the innovative activity to learn about its productivity. The random outcome of each activity is verifiable. Manso constructs the optimal contract that induces experimentation.

In our model, the principal's final decision depends entirely on the messages announced by the agent. In other words, the principal delegates his decision to the (potentially) better informed agent. Laffont and Tirole (1986) assume that the agent has superior information at the time of contracting and study delegation in a static model. Lewis and Sappington (1997), and Cremer, Khalil, and Rochet (1998) consider the case of costly information acquisition. In a recent paper, Lewis and Ottaviani (2008) extend the analysis to a dynamic setting. Lewis and Ottaviani analyze a model in which the principal offers short-term contracts to motivate the agent to search for innovations. They allow for the possibility that the agent's information is partially verifiable, so that the agent has the opportunity to conceal discoveries.

The paper is organized as follows. In Section 2, we describe the model when the prior is commonly known and characterize the optimal mechanism. In Section 3, we analyze the case in which the agent's prior is private information. Section 4 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 of each action depends on the state of the world which we assume to be binary: $\omega = B, G$. The principal's preferred action in state $\omega = B, G$ is $A = \omega$. The prior probability that the state is $\omega = G$ is denoted by $p_0 \in (0, 1)$.

The principal can make his decision (i.e., choose between B and G) in any period $t = 0, \dots, T - 1$, $T \geq 2$.⁴ In period 0 the principal can hire an agent to perform a number of tests. In this section, we assume that the agent does not have any private information at the outset of the game. Thus, in period 0 the principal and the agent share the same prior p_0 .

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

⁴We consider the case $T = 1$ at the end of Subsection 2.2.

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, the signal G provides definitive evidence in favor of the state G . Conditional on the state, the signals are independent across periods.

For every $t = 0, \dots, T$, we let

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

denote the agent's belief that the state is $\omega = G$ if he observes t signals equal to B . To make the problem interesting we assume that the agent's signals are beneficial to the principal. That is, given the belief p_T the principal prefers to choose action 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 goal of the principal is to induce the agent to acquire the signal and to reveal it in every period $t \leq T - 1$, until the agent finds definitive evidence $s = G$ in favor of the state G . 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).

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 $\delta \in (0, 1]$.

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 observes the signal G . In this case we say that the contract is incentive compatible.

The length of the contract T is given and is not a part of the mechanism. There are two different reasons why it is interesting to adopt this approach. First, suppose that there is a certain deadline $T - 1$ by which the principal has to choose between the two risky actions B and G . If the marginal benefit of a signal in period $t = 0, \dots, T - 1$ is much larger than the cost of the signal c , then it is indeed optimal for the principal to motivate the agent to work in every period. Second, if the length of the contract is endogenous, the optimal mechanism can be computed in two steps. In the first step we determine the

optimal contract for any given length. In the second step we compute the optimal length by comparing the principal's expected payoffs under contracts of different lengths. The analysis below shows how to solve the first and more challenging step. We discuss the second step at the end of Subsection 2.2.

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 the signal G). It is therefore without loss of generality to assume that the set of messages available to the agent in every period is $\{B, G\}$.

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 equal to B . This event can occur only if the agent deviates and lies about the realization of his signal. Thus, it is optimal to let the corresponding payment be equal to zero (the payment cannot be negative since the agent is protected by limited liability).⁵

In what follows we restrict attention to evidence-based contracts (under which the agent receives a positive payment in $t < T - 1$ only if the state is G). To simplify the exposition, we refer to them simply as contracts (or mechanisms). Formally, a contract w is the collection

⁵Of course, if w is an incentive compatible mechanism and all the constraints in which the payment appears are satisfied with strict inequalities, then we can increase the value of the payment by a small amount and the new mechanism is still incentive compatible. In fact, it is easy to construct optimal mechanisms w in which the principal makes a positive payment if the agent lies about his signal. In our opinion, these mechanisms are not particularly interesting since their distinctive feature is to make the agent's deviations more valuable.

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 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 consider 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 the signal B or the 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.

To simplify the exposition, we have defined a strategy as a contingent plan of actions. However, the off-path behavior of a strategy (what the agent does after he deviates) does not play any role in the analysis. This is because the principal commits to a mechanism in period 0 and cannot react to the agent's choices. In particular, all the strategies that

induce the same on-path behavior (but different off-path behavior) are equivalent in the sense that they all generate the same outcome.

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 the strategy σ .⁶ 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

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 sufficient to guarantee incentive compatibility. In particular, we can safely ignore multiple deviations and restrict attention to one-period deviations.⁷

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

⁶In the next section we consider the case in which the agent's prior is private information. For this reason it is convenient to make explicit the dependence of the agent's utility on the prior p_0 .

⁷We should perhaps point out that this result is not an application of the one-shot deviation principle. Clearly, the objective of the principal is not to find the strategy in Σ^* that is sequentially rational. In fact, as it will become evident from Lemma 1, we only consider one-period deviations from the on-path behavior.

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 $t' \geq t$:

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

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

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

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

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

It is certainly more interesting to see why the constraints (4)-(6) provide sufficient conditions for incentive compatibility. It follows from constraint (5) for $t = T-1$ that the agent prefers to tell the truth if he discovers in period $T-1$ that the true state is G . This and the fact that he discounted sequence of payments $\{\delta^t w(t)\}_{t=0}^{T-1}$ is (weakly) decreasing imply that it is optimal for the agent to announce the message G as soon as he observes the signal G . In other words, our constraints guarantee that agent does not lie after the signal G . Can the agent lie after the signal B ? The answer is no. Acquiring the signal and lying after B means that the agent ignores his signal (and sends the message G). But

then it would be more profitable to shirk, save the cost c and send the message G . We can therefore restrict attention to the strategies under which the agent reveals truthfully the realization of all the signals he acquires.

The next step is to show that another class of strategies can be ignored. Consider a strategy $\hat{\sigma}$ under which the agent guesses the state G in some period $\hat{t} > 0$. This strategy is weakly dominated by the strategy σ' under which the agent guesses the state G in the first period. Both strategies are such that the agent receives a positive payment if and only if the state is G . However, under the strategy σ' the agent receives the payment $w(0)$ in the first period and does not exert any effort. Under the strategy $\hat{\sigma}$ the agent may receive the payment after the first period and he may have to exert effort. Since effort is costly and the discounted sequence of payments $\{\delta^t w(t)\}_{t=0}^{T-1}$ is decreasing, the agent weakly prefers σ' to $\hat{\sigma}$.

Once we rule out the “guessing” strategies similar to $\hat{\sigma}$ above, we are 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 the message B . The constraints (5) tell us that the strategies under which the agent shirks in a single period do not constitute profitable deviations. It turns out that those constraints have stronger implications and are enough to prevent deviations under which the agent shirks in 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 the signal G). Consider two different scenarios. In the first scenario, the agent has acquired the signal in every period $t' < t$. The 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 (evidence-based) 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*

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

Under the optimal contract the agent is indifferent between acquiring and revealing the signal in every period and guessing the state G in the first period. Moreover, suppose that the agent has acquired and revealed the signal in every period $t' < t$. Suppose also that the agent plans to adopt the same behavior in every period $t' > t$. Then in period t the agent is indifferent between shirking (and announcing B) and acquiring and revealing the signal. In other words, the constraints (5) are all binding under the optimal mechanism. Finally, the principal sets the payment $w^*(G)$ equal to zero. Therefore, the agent receives a payment only when his message matches the state of the world.

In the proof of Proposition 1 we solve a relaxed problem in which we ignore the non-negativity constraints and the constraints (6). We show that the remaining constraints (4) and (5) must be binding.

The agent receives $w(B)$ only if the state is B . Thus, the probability of receiving this payment is the same for all the strategies under which the agent is truthful when he acquires the signal and announces the message B when he shirks. In other words, $w(B)$ does not appear in the constraints (5). The only role of $w(B)$ is to prevent the agent from guessing the state G in the first period. Clearly, the principal will choose it to make the agent indifferent. That is, the constraint (4) will bind.

Suppose that in a certain period $t > 0$ the agent has a strict incentive to acquire the signal. Suppose now that the principal lowers the the payment $w(t)$ by a small amount ε . Clearly this will not affect the incentive in every period $t' > t$. Also, if ε is sufficiently small then the agent will continue to acquire the signal in period t . Now, let us consider a period $t' < t$. The probability (computed in period t') that the agent receives the payment $w(t)$ is smaller if the agent acquires the signal in t' than if he shirks (if the agent acquires the signal in t' he could observe the signal G and the game would end). Thus, by reducing

the payment $w(t)$ the principal gives more incentive to the agent to exert effort in period $t' < t$. Clearly, a lower value $w(t)$ reduces the expected utility that the agent obtains when he acquires and reveals the signal in every period. This could induce the agent to guess the state G in the first period. However, if the principal also lowers the payment $w(0)$, the “guessing” strategy will become less appealing. In fact, it turns out that it is possible to lower both $w(t)$ and $w(0)$ in such a way that all the constraints are still satisfied. Thus the solution to the relaxed problem satisfies all the constraints (5) with equality.⁸ In a very similar way, we show that it is optimal for the principal to set the payment $w(G)$ equal to zero.

Thus, 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 $t = 0, \dots, T - 2$.⁹ Therefore, the contract w^* satisfies the constraints (6) and is optimal.

So far we have shown that w^* is the unique optimal contract in the class of evidence-based mechanisms. We now show that the uniqueness result holds for the whole class of mechanisms. To see this, suppose that there exists an optimal contract which specifies an intermediate payment $x > 0$ in period $t < T - 1$. 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 this contradicts the fact that w^* is the unique optimal evidence-based mechanism (recall that $w^*(G) = 0$). We conclude that the optimal contract is unique and is evidence-based.

2.2 The Agent’s Information Rent

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

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

⁸If the agent has a strict incentive to acquire the signal in $t = 0$, then it is enough to lower (by a small amount) the payment $w(0)$. This will still give an incentive to the agent to acquire the signal in the first period and it will make the “guessing” strategy less profitable.

⁹An additional feature of the optimal contract is that the sequence of payments $\{w^*(t)\}_{t=0}^{T-1}$ is increasing over time if $\delta < 1$ (the sequence is constant if $\delta = 1$). For brevity, we omit the proof of this result.

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 the state G or he does not. It is therefore natural to assume that if the agent shirks then he and the principal observe the 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 truth-telling 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.¹⁰ 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.¹¹

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.

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

¹⁰When the realization of the signal is verifiable, a contract is incentive compatible if and only if it satisfies the constraints (5).

¹¹In 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.

p_t (defined in equation (1)) denote the common belief. In this case it is enough to pay $c/\alpha p_t$ upon observing the 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, the higher α , the higher the incentives of the agent to shirk in t in order to get the larger payment in $t+1$. Because of this, 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 the 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 the 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 the 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.

As anticipated in Section 2, we now investigate the case in which the length of the contract is endogenous. The principal chooses 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 $T \geq 1$

$$C(T) = \tilde{C}(T) + \frac{c}{\alpha} + c(1 - p_0) \sum_{t=1}^{T-1} \left(\frac{\delta}{1-\alpha}\right)^t$$

$$\tilde{C}(T) = c + c \sum_{t=1}^{T-1} \delta^t [p_0(1-\alpha)^t + 1 - p_0]$$

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 $V(T) - C(T)$. As one would expect, T^* is (weakly) smaller than the efficient length, i.e. the length that maximizes $V(T) - \tilde{C}(T)$. Because of the presence of agency problems, testing is stopped too early and the gains of additional signals are not realized.

3 Asymmetric Initial Information

In Section 2, we assume that the principal and the agent share the same prior p_0 at the beginning of the game. 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 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. With probability $\rho \in (0, 1)$ the agent is a high type and believes that the state is G with probability $p_0^h \in (0, 1)$. With probability $1 - \rho$ the agent is a low type and believes that the state is G with probability $p_0^\ell \in (0, p_0^h)$. We denote the two types by their beliefs.

We assume that the number of available signals is $T \geq 2$.¹² 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 the signal G). 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 mechanisms and, therefore, $w^k = \left((w^k(t))_{t=0}^{T-1}, w^k(G), w^k(B) \right)$, for $k = h, \ell$.¹³

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 the message G (or in period $T - 1$ if he reports the message B in every period). The agent receives the payment specified by the contract that he chose.

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 the 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 the state

¹²We 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 $T = 1$.

¹³We return to this point at the end of the section.

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 the 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 the 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 the 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 he acquires and reveals the signal in every period $t' \geq t$. Also, $u(\sigma, p_0; w)$ 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:*

$$\begin{aligned}u(0, p_0^k; w^k) &\geq p_0^k w^{k'}(0), \quad k = h, \ell, \quad k' \neq k \\ u(0, p_0^h; w^h) &\geq u(0, p_0^h; w^\ell) \\ u(0, p_0^\ell; w^\ell) &\geq u(\sigma, p_0^\ell; w^h), \quad \sigma \in \Sigma'\end{aligned}$$

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

3.1 The Optimal Mechanism with Asymmetric Initial Information

The principal's problem is given by

$$\begin{aligned} \min_{w^h \geq 0, w^\ell \geq 0} & \rho u(0, p_0^h; w^h) + (1 - \rho) u(0, p_0^\ell; w^\ell) \\ \text{s.t.} & (w^h, w^\ell) \text{ is incentive compatible} \end{aligned} \tag{8}$$

The problem admits a solution. 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 low type when the prior p_0^ℓ is common knowledge, except that $\bar{w}(B)$ is such that the high type is indifferent between acquiring and revealing the signal in every period and guessing the 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 shall 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

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

Among all the contracts that are suitable for p_0^ℓ , $w^*(p_0^\ell)$ specifies the lowest payments in state G . 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 prior 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*

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

If a contract w is suitable for some $p \in (0, 1)$, then $w(B)$ must be at least $w^*(B; p)$. This follows from the optimality of the contract $w^*(p)$ when the principal and the agent share the same prior p . Notice that $w^*(B; p)$ is such that an agent who has a belief equal to zero and who acquires and reveals the signal in every period obtains a strictly positive rent. We therefore have $u(0, 0; w) > 0$.

Suppose that the mechanism (w^h, w^ℓ) is incentive compatible. Clearly, $u(0, p_0^h; w^h) \geq p_0^h w^\ell(0)$, i.e. the high type (weakly) prefers to acquire and reveal the signal with the contract w^h rather than guess the state G in the first period with the contract w^ℓ . Notice that an agent with a belief equal to zero strictly prefers the first alternative to the second one. In fact, the agent obtains a positive rent if he works and is honest in every period (w^h is suitable for p_0^h), and a payoff equal to zero if he guesses the state G . Finally, the agent's incentives are linear functions of his beliefs. Therefore, we must have $u(0, p; w^h) > p w^\ell(0)$ for every $p < p_0^h$.

Given an incentive compatible mechanism (w^h, w^ℓ) , the deviation under which the low type chooses the contract w^ℓ and guesses the state G in the first period is not the most profitable one (since $u(0, p_0^\ell; w^h) > p_0^\ell w^\ell(0)$). 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 case in which the prior p_0^ℓ is common knowledge, 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 the 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 an intuition for the second result. Suppose that $u(0, p_0^h; w^h) > p_0^h w^h(0)$. If $u(0, p_0^\ell; w^\ell) = p_0^\ell w^h(0)$ then the high type strictly prefers the contract w^h to the contract w^ℓ (formally, $u(0, p_0^h; w^h) > p_0^h w^\ell(0)$ and $u(0, p_0^h; w^h) > u(0, p_0^h; w^\ell)$). But then the principal can lower the payment $w^h(B)$ by a small amount and all the incentive constraints are still satisfied.

Thus, let us assume that $u(0, p_0^\ell; w^\ell) > p_0^\ell w^h(0)$, i.e. the constraint in which the low type chooses w^h and guesses the state G immediately is not binding. In this case 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 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 the 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 obtains when the prior p_0^h is common knowledge. 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 priors 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.*

Let us investigate when the principal can improve upon the mechanism (\bar{w}, \bar{w}) . Because of Lemma 4 we can restrict attention to mechanisms (w^h, w^ℓ) such that the contract w^h gives to the high type the same rent as the contract \bar{w} (and $w^h(0) = \bar{w}(0)$). The mechanism (w^h, w^ℓ) can do strictly better than (\bar{w}, \bar{w}) only if $w^h(B) < \bar{w}(B)$. In fact, if $w^h(B) \geq \bar{w}(B)$, then an agent who has a belief smaller than p_0^h and plans to acquire and reveal the signal in every period prefers w^h to \bar{w} . In particular, $u(0, p_0^\ell; w^h) \geq u(0, p_0^\ell; \bar{w})$. But then w^ℓ must give to the low type at least the same rent as \bar{w} , otherwise he has an incentive to choose the contract w^h .

Consider \bar{w} and notice that given this contract it is optimal for the low type to follow the strategy σ^1 under which he exerts effort in every period $t \geq 1$ (of course, it is also

optimal to start to exert effort in period zero). Suppose that the principal decreases $\bar{w}(B)$ and increases $\bar{w}(1)$ keeping constant the high type's rent. Let w^h denote the new contract. The change will increase (decrease) the payoff of the strategy σ^1 when the low type's prior p_0^ℓ is larger (smaller) than p_1^h . Intuitively, let us compare the low type when he follows the strategy σ^1 and the high type when he starts to acquire the signal in the first period. The low type is relatively more likely to receive the larger payment $w^h(1)$ and less likely to receive the smaller payment $w(B)$ when his prior p_0^ℓ is larger than p_1^h . The opposite result holds when p_0^ℓ is smaller than p_1^h .

This argument is sufficient to show that (\bar{w}, \bar{w}) is not optimal when $p_0^\ell < p_1^h$.¹⁴ To prove the optimality of (\bar{w}, \bar{w}) when $p_0^\ell \geq p_1^h$, we also need to consider changes of \bar{w} which involve the payments $\bar{w}(2), \dots, \bar{w}(T-1)$ and $\bar{w}(G)$. However, the same logic applies. Any attempt to decrease $\bar{w}(B)$ and adjust the remaining payments (to keep constant the high type's rent) will necessarily make the low type better off when he uses the strategy σ^1 . More precisely, the low type is weakly better off if $p_0^\ell = p_1^h$ and strictly better off if $p_0^\ell > p_1^h$. Thus, in the latter case (\bar{w}, \bar{w}) is the unique optimal mechanism.

However, when $p_0^\ell \leq p_1^h$ there are multiple optimal mechanisms. We describe one in the next proposition. Then we briefly 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)

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

¹⁴The principal can lower the payment $\bar{w}(B)$ to the low type by a small amount and he will not find it profitable to choose the contract w^h .

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 (28)).

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

$$u(0, p_0^h; w^h) - \delta u(1, p_0^h; w^h) = \Delta \tag{9}$$

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 (9).

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.¹⁵ 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

¹⁵Recall 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.

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 when the prior p_0^ℓ is known) 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 (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 contracts that are evidence-based. Allowing for intermediate payments has no consequences for the contract of the low type. Recall 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 even if we allow for mechanisms that are not-evidence based. 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.

4 Conclusions

This paper analyzes the optimal provision of incentives in a sequential testing context. 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 principal and the agent are symmetrically informed at the outset of their relationship. Under the optimal contract, the agent is indifferent in every period between performing the test and shirking. Furthermore, in the first period the agent is indifferent between carrying out his task and sending an uninformative message which ends the relationship immediately. 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 agent with a low prior is minimally distorted. The principal prefers to offers different contracts if and only if the agent's types are sufficiently diverse. Finally, all the types benefit from their initial private information.

We view our analysis as a first step in exploring sequential testing with agency problems. Thus, there are various ways to extend our results. 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. Of course, the literature on sequential testing (see Wald (1947), Arrow, Blackwell, and Girshick (1949), and, more recently, Moscarini and Smith (2001)) has analyzed much more general environments. An obvious extension is to assume that all the realizations of the signal contain some noise. In this case the agent's beliefs may move in either direction over time without becoming degenerate. If the horizon is finite, the final payments could depend on the entire history of messages, which makes the problem less tractable. It is probably more convenient to assume an infinite horizon and use recursive techniques. With an infinite horizon the relationship between the principal and the agent ends when the belief becomes sufficiently small or sufficiently large.

Another direction for future research is to analyze the case in which the agent has more than two initial types. While we do not have a complete characterization of the optimal mechanism in this case, it is easy to see that some of our results generalize. First, the contract offered to the type with the lowest belief about the good state is minimally distorted. All the payments that this type receives in the good state coincide with the payments that he receives when his prior is commonly known. Second, all agent types receive an additional rent because of their initial private information.

Finally, in the paper we assume that the marginal value of each signal is so large that the principal prefers to motivate every type of the agent to exert effort in every period. That is, the maximal length of the relationship is the same for all the types. In general, the principal may prefer to stop testing and make a decision before the deadline if the belief falls below a certain threshold. Of course, the number of signals (in favor of the bad

state) needed to reach the threshold depends on the agent's initial belief. It is, therefore, interesting to investigate the case in which the principal offers contracts of various lengths to the different types of the agent.

Appendix

Proof of Lemma 1.

We have only to prove that if a mechanism w satisfies the constraints (4)-(6) then it is incentive compatible.

The constraints (4)-(6) guarantee that the agent will never acquire the signal and lie about its realization. To see this, we use the constraint (5) at $T - 1$

$$\begin{aligned} 0 \leq u(T - 1, p_{T-1}) - \delta u(T, p_{T-1}) = \\ -c + p_{T-1}\alpha [w(T - 1) - w(G)] \end{aligned} \quad (10)$$

and obtain

$$w(T - 1) > w(G) \quad (11)$$

This and the constraints (6) imply that it is optimal for the agent to tell the truth when he observes the signal G . Clearly, any strategy in which the agent is honest after observing the signal G and dishonest after observing the signal B is strictly dominated. The agent can save the cost c and send the uninformed message G . In what follows, we therefore restrict attention to strategies under which the agent reveals truthfully all the signals that he acquires.

Let σ' be a strategy that prescribes guessing the state G (i.e., shirking and announcing the message G) in the first period. Let $\hat{\sigma}$ be any strategy such that the agent guesses the 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 $\hat{\sigma}$.¹⁶ We have

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

¹⁶Recall 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.

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

It remains to consider strategies under which the agent tells the truth when he acquires the signal and sends the message B when he shirks. The last step of the proof is to show that the constraints (5) imply

$$u(t, p) > \delta u(t+1, p) \quad (12)$$

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. In fact, 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 the constraints (5). Thus, the mechanism w is incentive compatible.

We now prove inequality (12). One can immediately see from inequality (10) that inequality (12) 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

$$\begin{aligned} a(t; w) &= -c + \delta a(t+1; w) \\ b(t; w) &= \alpha w(t) - \delta \alpha a(t+1; w) + \delta(1-\alpha)b(t+1; w) \end{aligned} \quad (13)$$

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 $p \in [0, 1]$

$$u(t, p) = a(t) + b(t)p \quad (14)$$

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

Thus, we have

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

The constraints (5) imply 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

$$w(t) - \delta a(t+1) - \delta b(t+1) > 0 \tag{15}$$

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 the constraints (6), the inequality (11), and, of course, $c > 0$.

We now proceed by induction. We assume that inequality (15) 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 the constraints (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 the constraints (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:

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

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

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

where $\tilde{\psi}, \psi_0, \dots, \psi_{T-1}$ are T constants. $\tilde{\psi}$ is the agent's expected cost of testing when he acquires the signal in every period. For every t , ψ_t represents the difference between the 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

$$\begin{cases} \tilde{\varphi}(\hat{w}) = 0 \\ \varphi(t, \hat{w}) = 0, \quad t = 0, \dots, T - 1 \\ \hat{w}(G) = 0 \end{cases} \quad (16)$$

Notice that $w(B)$ appears only in $\tilde{\varphi}(w)$ with a positive coefficient. If $\tilde{\varphi}(\hat{w}) > 0$ then we can lower $\hat{w}(B)$ 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 $\hat{w}(0)$ 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

$$\begin{aligned} w'(\hat{t}) &= \hat{w}(\hat{t}) - \varepsilon \\ w'(0) &= \hat{w}(0) - \alpha (1 - \alpha)^{\hat{t}-1} \delta^{\hat{t}} \varepsilon \end{aligned}$$

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 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 $t > 0$. Thus, \hat{w} cannot be a solution to the relaxed problem.

Consider now the system of linear equations (16). 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 (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 $u(t + 1, p_{t+1}; w^*) = \delta u(t + 2, p_{t+1}; w^*)$, the constraint (5) at $t + 1$. The constraints (6) are 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

$$u(0, p_0^h; w^\ell) > u(0, p_0^h; \hat{w}^\ell) \quad (17)$$

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

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

Notice also that

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

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,

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

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 the 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 (18)-(20) 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

$$u(t, p_t^h; w^h) \geq \delta u(t+1, p_{t+1}^h; w^h) \quad (21)$$

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 (18) and (19) 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 the constraint (21) is binding in every period. In this case we must have $w^h(G) > 0$ (if $w^h(G) = 0$ and all the constraints (21) 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 (13) 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 (14) 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 a prior $p_0 \in (0, 1)$ 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 the 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

$$u(0, p_0^h; w^h) = p_0^h w^h(0) \tag{22}$$

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

$$u(0, p_0^\ell; w^\ell) = p_0^\ell w^\ell(0) \tag{23}$$

We know from Lemma 3 that $u(0, p_0^\ell; w^\ell) > p_0^\ell w^\ell(0)$. This and equality (23) imply $w^h(0) > w^\ell(0)$. Furthermore, it follows from $a(0; w^\ell) > 0$ (see Claim 1) and equality (23) 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

$$u(0, p_0^\ell; w^\ell) > p_0^\ell w^\ell(0) \tag{24}$$

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 (24) 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 (17), 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 (22) 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 the 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 $\Lambda(p_0^\ell, w^h) > 0$, 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

$$\begin{aligned} 0 \leq p_0^\ell \left[\sum_{t=1}^{T-1} (1-\alpha)^{t-1} \alpha \delta^t (w^h(t) - \bar{w}(t)) + (1-\alpha)^{T-1} \delta^{T-1} (w^h(G) - \bar{w}(G)) \right] + \\ (1-p_0^\ell) \delta^{T-1} (w^h(B) - \bar{w}(B)) = u(\sigma^1, p_0^\ell; w^h) - u(\sigma^1, p_0^\ell; \bar{w}) \end{aligned} \quad (25)$$

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 the 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 (25) 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

$$u(1, 1, \hat{w}^h) \geq u(1, 1, \hat{w}^\ell) \quad (26)$$

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 (26) 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

$$u(0, p_0^h; \hat{w}^h) = p_0^h \hat{w}^h(0) = p_0^h w^*(0; p_0^\ell) \quad (27)$$

have the following implications:

$$\hat{w}^h(B) = w^*(B; p_0^h) + \frac{(1-\alpha)\Delta}{(1-p_0^h)\alpha\delta^{T-1}} \quad (28)$$

and

$$u(1, 1; \hat{w}^h) = \frac{w^*(0; p_0^\ell)}{\delta} - \frac{c}{p_0^h \alpha \delta} - \frac{\Delta}{p_0^h \alpha \delta} \quad (29)$$

From equations (27)-(29) 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 the 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 (28) 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 the 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

$$u(0, p_0^\ell; \hat{w}^\ell) \geq \max\{v_0, \dots, v_T\} \quad (30)$$

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

$$u(1, 1; w^h) = \frac{w^*(0; p_0^\ell)}{\delta} - \frac{c}{p_0^h \alpha \delta} - \frac{\Delta}{p_0^h \alpha \delta} = u(1, 1; \hat{w}^h) \quad (31)$$

It follows from inequality (26) and the way the mechanism (w^h, w^ℓ) is designed in step 1 that the solution $w^h(1)$ to the equation (31) must be weakly greater than $w^*(1; p_0^\ell)$.¹⁷ 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 (31). 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 (31) 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 (31). 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

$$\max_{\sigma \in \Sigma'} u(\sigma, p_0^\ell; w^h) = \max\{v_0, \dots, v_T\} \quad (32)$$

Given inequality (30), 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.¹⁸ 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 (32). 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

¹⁷The solution $w^h(1)$ must be strictly greater than $w^*(1; p_0^\ell)$ if Δ is equal to zero.

¹⁸Notice 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. More 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 the 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

$$u(0, p_0^h; \hat{w}^h) - \delta u(1, p_0^h; \hat{w}^h) = 0 \tag{33}$$

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 (33) guarantees that the payment $w^h(1)$ identified in step 1 is strictly greater than $w^*(1; p_0^\ell)$ (see footnote 17).

References

- [1] Arrow, K. J., D. Blackwell, and M. A. Girshick (1949): “Bayes and Minimax Solutions of Sequential Decision Problems,” *Econometrica*, 17, 213-244.
- [2] Atkenson, A. (1991): “International Lending with Moral Hazard and Risk of Repudiation,” *Econometrica*, 59, 1069-1089.
- [3] Baron, D. P., and R. B. Myerson (1982): “Regulating a Monopolist with Unknown Costs,” *Econometrica*, 50, 911-930.
- [4] Bergemann, D., and U. Hege (1998): “Venture Capital Financing, Moral Hazard, and Learning,” *Journal of Banking and Finance*, 22, 703-735.
- [5] Bergemann, D., and U. Hege (2005): “The Financing of Innovation: Learning and Stopping,” *Rand Journal of Economics*, 36, 719-752.
- [6] Biais, B., T. Mariotti, G. Plantin and J.-C. Rochet (2007). “Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications,” *Review of Economic Studies*, 74, 345-390.
- [7] Cremer, J., F. Khalil, and J.-C. Rochet (1998): “Contracts and Productive Information Gathering,” *Games and Economic Behavior*, 25, 174-193.
- [8] DeMarzo, P. M., and M. J. Fishman (2007): “Optimal Long-Term Financial Contracting,” *Review of Financial Studies*, 20, 2079-2128.
- [9] DeMarzo, P. M., and Y. Sannikov (2006): “Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model,” *Journal of Finance*, 61, 2681-2724.
- [10] DeMarzo, P. M., and Y. Sannikov (2008): “Learning in Dynamic Incentives Contracts,” mimeo.
- [11] Green, E. (1987): “Lending and the Smoothing Of Uninsurable Income,” in E. Prescott and N. Wallace, Eds., *Contractual Arrangements for Intertemporal Trade*, Minneapolis: University of Minnesota Press.
- [12] Gromb, D., and D. Martimort (2007): “Collusion and the Organization of Delegated Expertise,” *Journal of Economic Theory*, 137, 271-299.

- [13] Holmstrom, B., and P. Milgrom (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica*, 55, 303-328.
- [14] Laffont, J.-J., and J. Tirole (1986): “Using Cost Observation to Regulate Firms,” *Journal of Political Economy*, 94, 614-641.
- [15] Lewis, T. R., and M. Ottaviani (2008): “Search Agency,” mimeo.
- [16] Lewis, T. R., and D. E. M. Sappington (1997): “Information Management in Incentive Problems,” *Journal of Political Economy*, 105, 796-821.
- [17] Manso, G. (2007): “Motivating Innovation,” mimeo.
- [18] Moscarini, G., and L. Smith (2001): “The Optimal Level of Experimentation,” *Econometrica*, 69, 1629-1644.
- [19] Mussa, M., and S. Rosen (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18, 301-317.
- [20] Spear, S. E., and S. Srivastava (1987): “On Repeated Moral Hazard with Discounting,” *Review of Economic Studies*, 54, 599-617.
- [21] Wald, A. (1947): “Foundations of a General Theory of Sequential Decision Functions,” *Econometrica*, 15, 279-313.