Information Acquisition, Moral Hazard, and Rewarding for Bad News[*]

Hector Chade and Natalia Kovrijnykh
Department of Economics
Arizona State University

August 14, 2011

Abstract

This paper analyzes a principal-agent problem with moral hazard where a principal searches for an opportunity of uncertain return, and hires an agent to evaluate available options. The agent's effort affects the informativeness of a signal about an option's return. Based on the information provided by the agent, the principal decides whether to exercise the option at hand. We derive properties of the optimal contract in both static and dynamic versions of the problem. We show that sometimes the agent is rewarded for delivering 'bad news' about the quality of an option. We characterize distortions (relative to the first best) on the implemented effort level and optimal stopping decision. We also study dynamics induced by the optimal contract. Finally, we discuss applications that exhibit rewarding for bad news.

Keywords: Moral Hazard, Principal-Agent Model, Contracts, Information Acquisition, Search, Limited Liability.

---

[*] We are grateful to Larbi Alaoui, Ying Chen, Peter DeMarzo, Amanda Friedenberg, Marco Ottaviani, Kevin Reffett, Richard Rogerson, Edward Schlee, Jeroen Swinkels, Balázs Szentes, and William Zame for helpful comments and suggestions. We also benefited from comments by seminar participants at Arizona State University, University of British Columbia, CEA-Universidad de Chile, IESE, UCLA, Universitat Pompeu Fabra, UCSD, Toulouse School of Economoics, Econometric Society World Congress 2010, 2010 Society for Economic Dynamics meetings, 2010 Southwest Economic Theory Conference, and 2011 University of Hong Kong Summer Microeconomics Seminars.

1 Introduction

Many situations of economic interest feature a client (principal) who searches for an opportunity/option of uncertain return and hires an expert (agent) to evaluate potential options. The agent's effort, which he chooses privately, affects the quality of information he receives about the options. Based on the information the agent provides, the principal compensates him and decides whether to exercise an option or to continue the search.

Interactions of this sort abound in real-world applications. For instance, consider an uninformed investor who hires an expert to evaluate potential investment opportunities. Or take an individual seeking to purchase a house, who hires a home inspector to provide her with information about houses that become available. Another application is that of a firm that hires a headhunter to evaluate potential candidates for a job. Yet another one is that of a potential buyer hiring a real-estate agent to provide her with information about houses available for sale. Finally, take a worker who performs quality control for a company. Although other features may be present in these applications, information acquisition as described above certainly plays an important role.

The purpose of this paper is to analyze how to structure incentives in these environments, where the quality of information is endogenous and hidden, and to shed light on the main distortions that this moral-hazard problem impinges on information acquisition and on the principal's decision to exercise an option.

To focus on the main trade-offs involved in this problem, we build a tractable principal-agent model with risk-neutral parties, moral hazard, and limited liability. The agent's role is to acquire information about the unknown quality of an available opportunity, which can be high or low. Information consists of a noisy binary signal whose informativeness depends on the agent's effort. The principal conditions her decision to accept or reject an opportunity on the information acquired by the agent.

In this setting, we derive properties of the optimal contract in both the static case, with only one option to evaluate, and the dynamic case, where options arrive stochastically over time and are evaluated accordingly before the principal decides to stop the search.

The model affords a sharp characterization of the optimal compensation scheme based on the prior belief that an option is of high quality. We show that in both the static and dynamic cases, the agent is rewarded for producing a high signal — that leads the principal to exercise the option at hand — if and only if the prior belief is above a threshold level. This cut-off value is higher when the relative effect of effort on the signal distribution in the low- vs. high-quality state is higher. For instance, if effort increases the probability of observing a signal consistent with the true state equally in both states, then the threshold is equal to one half. But if a high-quality option always generates a high signal regardless of effort, while higher effort makes it more likely to observe a low signal if the option quality is low, then the threshold is equal to one.

Notice that when the prior belief is below the threshold, the agent is rewarded for delivering bad news about the quality of an option. That is, the optimal contract rewards the agent in the event in which the principal decides *not* to exercise it. We call this feature the *rewarding-for-bad-news* property.[1] We argue that it arises naturally when the agent's effort generates information about a characteristic that is payoff-relevant for the principal. In contrast, in the standard principal-agent model with moral hazard (where effort typically affects the distribution of output, i.e., a payoff-relevant characteristic), the prediction is that the agent is always rewarded for good news.

In the static setting, rewarding for bad news means that the agent gets paid when the principal chooses not to exercise an option. Interestingly, this feature disappears in the dynamic setting, because the principal *prefers to postpone* payments until she stops the search. Thus, the agent *only* gets paid at the end, that is, when the principal exercises an option. The size of that payment depends on the history observed up to that point. Hence, rewarding for bad (good) news in the dynamic case means that the agent's expected future payments increase (decrease) if he generates a low signal.

The contrast between the static and dynamic settings is also apparent when we allow the agent to *hide* the presence of an option and claim that none is available.[2] While this has no bite in the static case, it substantively affects incentive provision in the dynamic case.[3] If option arrival is contractible, to avert incentives for hiding, the contract punishes the agent with lower expected future payments if he claims that there is no option available.

---

[1] In our model, the signal provides information about both the effort chosen by the agent and the quality of the option. As a signal of effort, a low realization can be good or bad news depending on the prior. But as a signal of quality, a low realization is *always* bad news (the option's expected return is lower). It is in this sense that bad news should be interpreted. Moreover, the principal always receives a lower payoff if a low signal is realized (even after paying the agent), which reinforces the bad news interpretation.

[2] The possibility of hiding is more plausible in some applications (such as the investor example) than others (such as the home inspector example). As will become clear, the properties of the optimal contract are nontrivial even if hiding is not possible.

[3] We show in Section 6.3 that if the agent could hide the presence of information instead, distortions would occur in the static case as well. Also notice that if the agent could misreport a signal, no effort could be induced, as the agent's payoff depends only on his report. Section 6.2 reveals that this conclusion is overturned if the principal observes an additional signal about the quality of an option.

The possibility of hiding is particularly relevant when the agent is rewarded for bad news, for, as we describe below, it crucially affects equilibrium dynamics. And if option arrival is non-contractible, then hiding has no effect whatsoever on the optimal contract when the agent is rewarded for good news. In sharp contrast, it precludes any incentive provision if he is rewarded for bad news.

The implemented effort and the principal's decision to exercise an option are both distorted compared to the first best (i.e., an environment without moral hazard or limited liability). In the static environment the principal implements a positive effort level less often in the second- than in the first-best case. And when she does, the implemented effort is *lower*. This immediately implies that the decision of whether to exercise an option is distorted as well, since she makes that decision without information more often than in the first-best case. Moreover, when her decision depends upon the agent's information, the signal about option quality is less precise.

Additional distortions arise in the dynamic case. For instance, in the first-best case effort is constant over time, while in the second-best case it varies over time and can fluctuate. Also, in the first-best case the principal's decisions regarding what effort level to implement and whether to exercise an option do not involve randomization, while it is part of the optimal contract in the second-best case.

The sources of these distortions can be traced back to two important features of our environment. First, we show that the discrete nature of the problem makes the use of *lotteries* or random contracts beneficial to the principal. Sometimes the principal randomizes between implementing zero effort along with a decision to stop the search, and implementing positive effort along with a decision to exercise an option only if the signal is high. The second feature, which is responsible for *fluctuations*, comes from the possibility that the agent can hide the presence of an option. This is especially relevant when the optimal contract exhibits the rewarding-for-bad-news property. In this case, the contract punishes the agent with lower future payments if he claims that there is no option available, and rewards him with higher future payments when the signal realization about option quality is low. This introduces fluctuations over time as the value to the agent (and thus future effort levels) can increase or decrease depending on the history.

Naturally, the structure of the optimal contract has implications for the equilibrium dynamics of effort, agent's compensation, and principal's decision to exercise an option. One interesting feature that arises when the optimal contract punishes the agent for bad news is that, if the search takes long enough, the principal eventually fires the agent and decides whether to exercise an option solely based on the prior. This represents a distortion relative to the first best, where the principal, once she starts the search, only ends it when a high signal is eventually generated.

To gauge the robustness of our findings, we analyze three variations of the model. The first one studies implications of a richer signal structure with a continuum of possible realizations, the second one allows the principal to observe an additional contractible signal, and the third one varies the events that are contractible. We show that all of them exhibit the rewarding-for-bad-news property as well as the distortions mentioned above. A new feature that emerges from the first and third variations is that the principal can benefit from *committing* to a decision to exercise an option that is ex-post inefficient.

Finally, we discuss some suggestive evidence of real-world contracts in environments where effort plays an important role in acquiring information. In each case, we assess whether they exhibit some features predicted by the model, in particular, the rewarding-for-bad-news property. One example we discuss is that of software vendors who award monetary prizes to researchers who find bugs in source codes of the companies' products.

Related Literature. Our dynamic model is related to the large literature on repeated moral hazard (e.g., Rogerson, 1985b, Spear and Srivastava, 1987, Sannikov, 2008). The main features of the environment that distinguish our model are that effort affects the informativeness of a signal about option quality and not the distribution of output, and that the principal uses information provided by the agent to make a decision.

The paper is also related to a recent literature on delegated search and project completion, such as Toxvaerd (2006), Mason and Valimaki (2008), and Lewis and Ottaviani (2008). In that literature, effort affects the probability of completing a project, or the arrival of opportunities that lead the agent to optimally stop the search. That is, the principal delegates the search activity to the agent. In our model, it is the principal who conducts the search, and delegates the acquisition of information to the agent.

There is a vast literature on optimal decisions with information acquisition (see, e.g., Chade and Schlee, 2002). In the standard problem, before making a decision, an individual acquires a signal, whose cost is exogenous and is increasing in the signal's informativeness. In contrast, the cost is endogenous in our model, as it comes from the incentive problem created by the aforementioned delegation of information acquisition to an agent.

Among papers that analyze *private* information acquisition, the following are the most related ones. Szalay (2009) and Krahmer and Strausz (2010) study static procurement problems with adverse selection where the agent privately acquires information. Both papers show that information acquisition leads to distortions in effort and production, which bears some resemblance with our findings in the static case. Unlike them, we derive the rewarding-for-bad-news property and also study in detail the dynamic case. Also, the last section of Inderst and Ottaviani (2010) analyzes a static contracting problem where an agent exerts hidden effort that affects the informativeness of a signal about the suitability of a match between a firm's product and a buyer's valuation for it. Unlike in our model, it is the buyer and not the firm (the principal) who uses the information provided by the agent before purchasing, and they focus on different distortions than we do.

Another related literature is that on delegated expertise, such as Dewatripont and Tirole (1999), Gromb and Martimort (2007), and Eso and Szentes (2007) (see also the references cited therein). The first paper focuses on the creation of advocates by organizations. It shows that it can be better to have two experts, each acquiring information about a particular alternative (an advocate of that alternative), instead of having only one who acquires information about both. The second paper analyzes a similar optimal organizational design problem but allowing experts to collude among themselves or with the principal. Both papers deal exclusively with static settings, and focus on different issues than we do. However, the one-agent-one-signal case analyzed in the second paper has similarities with the second variation of our model mentioned above, except that effort is binary, and the agent can misreport his signal. Finally, the third paper studies a static problem between a client and an expert consultant. The expert provides the client with information that reduces the noise in her estimate of the value of an investment opportunity. The client has private information and the expert designs the contract. In our case, the principal is the 'client' and designs the contract. Also, she does not have private information, and moral hazard is on the expert's side instead.

Predictions similar in flavor to our rewarding-for-bad-news result appear in the following papers. Levitt and Snyder (1997) analyze a model where the agent's effort affects the distribution of a privately-observed signal that is correlated with an outcome of a project. Early access to that information is valuable to the principal because she can terminate the project. They show that an agent who predicts a bad outcome receives a higher wage than an agent who incorrectly predicts a good outcome. That is, the agent is rewarded for coming forward with bad news. Compared to our result, theirs emerges from adverse selection and relies on the benefits from early information revelation.

Similarly, Manso (2011) shows that rewards or tolerance for early failure can be part of an optimal contract that motivates innovation (see also Ederer and Manso, 2011). He analyzes a two-period model where the agent can direct his effort towards innovating activities or routine practices, and shows that compensation can be higher when failure is observed in the first period and success in the second, than when the sequence is reversed.

Our prediction that the agent is rewarded for producing a signal consistent with the prior is also reminiscent of Gentzkow and Shapiro (2006). Using a reputation model, they show that a low-type media firm (that has a noisy signal about the true state of the world) chooses to bias its report towards its customers' prior so as to pass for a high type (that has a signal that perfectly reveals the state). In our model there is no bias, as the agent truthfully reveals the signal realizations. Moreover, unlike in Mullainathan and Shleifer (1995), who assume that readers *prefer* to see their beliefs confirmed by newspapers, in our paper the client prefers to see a high report, but rewards the expert for issuing the report most consistent with him exerting effort.

The outline of the paper is as follows. Section 2 describes the model. Section 3 derives main properties of the optimal contract in the static case, while Section 4 focuses on the dynamic case. In Section 5 we discuss some evidence that suggests that the rewarding-for-bad-news property is present in real-world contracts. Section 6 presents variations of the model. Section 7 concludes. The Appendix contains the proofs of all the results.

2 The Model

The dynamic principal-agent model that we analyze, which subsumes the static version as a special case, has the following characteristics.

The Environment. Each period over a potentially infinite horizon, an option arrives with probability $0 < p \leq 1$. Arrivals over time are independent events.

The quality q of an option is unobservable, and it can be low or high, i.e., $q \in \{\ell, h\}$. If exercised, a high-quality option yields a return $y_h > 0$ to the principal, while the return of a low-quality option is $y_\ell < 0$. Once an option is exercised, the relationship ends.[4]

In any period when an option arrives, the prior belief that it is of high quality is $0 \leq \gamma \leq 1$, which can also be interpreted as the fraction of high-quality options in the pool.

The role of the agent in this relationship is that of an expert who provides the principal with information about the quality of an option at hand. More precisely, when an option arrives, the agent observes a signal $\theta \in \{\theta_\ell, \theta_h\}$ that is correlated with the quality of that option. How informative the signal is about quality depends on the level of effort $e \geq 0$ that

---

[4] For convenience, we use the terms 'exercising an option,' 'buying,' and 'purchasing' interchangeably.

the agent exerts. To capture this dependence, we focus on a class of information structures described by the following conditional probabilities:

$$\Pr\{\theta = \theta_h | q = h, e\} = \alpha + \beta_h \eta(e) \quad \text{and} \quad \Pr\{\theta = \theta_h | q = \ell, e\} = \alpha - \beta_\ell \eta(e), \tag{1}$$

where $0 \leq \alpha \leq 1$, $\beta_i \geq 0$, $i = \ell, h$, $\beta_\ell + \beta_h > 0$, and the function η satisfies $\eta(0) = 0$, $\eta'(e) > 0$, $\eta''(e) \leq 0$ for all e, and $\lim_{e\to\infty} \eta(e) = \bar{\eta} < \infty$. Also, to ensure that the expressions in (1) are between zero and one, we require $\beta_h \leq (1-\alpha)/\bar{\eta}$ and $\beta_\ell \leq \alpha/\bar{\eta}$.

Note that if effort is zero, the conditional distribution of the signal is the same regardless of an option's quality, and therefore the signal is uninformative. Conditional on the option being of a certain quality, the probability of observing a signal consistent with that quality is increasing in the agent's effort. That is, higher effort makes the signal more informative.[5]

As will become clear below, this class of information structures delivers some nice features. First, it affords a clean and intuitive characterization of the properties of the optimal contract in terms of the prior belief γ. Indeed, we show in the Appendix that this is the largest class that has this property. Second, under this class we can be confident that the agent's incentives to choose effort are fully captured by the first-order condition of his problem. More precisely, the agent's problem becomes strictly concave in effort.[6] Finally, from the standpoint of economic applications, this class subsumes many cases of interest, including the extreme cases with $\beta_h = 0$ and $\beta_\ell = 0$, and the symmetric case with $\beta_h = \beta_\ell$. When $\beta_h = 0$, the agent's effort only affects the distribution of the signal if the option's quality is low, so the agent's effort is important in recognizing faults (e.g., as in the home inspector or quality control examples). The situation is reversed if instead $\beta_\ell = 0$. And when $\beta_h = \beta_\ell$, the agent's effort increases — by the same amount in both states — the likelihood of observing a signal consistent with the unknown option quality.

Exerting effort e is a costly activity for the agent, which entails a disutility given by $\psi(e)$. The function ψ satisfies $\psi(0) = 0$, $\psi'(e) > 0$ and $\psi''(e) > 0$ for all e.

We also assume that the ratio ψ'/η' satisfies $\psi'(0)/\eta'(0) = 0$, $(\psi'/\eta')'(0) = 0$, and ψ'/η' is convex in e. The first assumption ensures that the first-order condition of the agent's

---

[5] An increase in effort makes observing the signal a more informative experiment in Blackwell's sense. For a classical exposition, see Blackwell and Girshick (1954), chapter 12. See also recent contributions by Athey and Levin (2001), Jewitt (2007), Quah and Strulovici (2009), and Ganuza and Penalva (2010).

[6] Our main results still obtain if instead we assume the most general binary case with $\alpha + \eta_h(e)$ and $\alpha - \eta_\ell(e)$ (by a suitable change of origin, the intercept is without loss of generality). But the two advantages mentioned are lost in this case, thus complicating the derivation and presentation of the results. By focusing on (1), we avoid all these analytic battles that are not central to the paper.

effort choice problem fully captures its solution even if it is zero, and together with the second one it guarantees that the principal induces positive effort if she buys only after observing a high signal realization. Finally, the third one makes the principal's problem strictly concave in effort in the static case.

The agent cannot lie about a signal realization, but he can hide the existence of an option. More precisely, after observing a signal he can either claim that there is an option, in which case the signal is publicly observed, or that no option has arrived. Clearly, if $p = 1$ then no hiding is possible and signals are automatically observed by both parties.

Both the principal and the agent are risk neutral and each maximizes the expected sum of discounted payoffs using a discount factor $0 \leq \delta < 1$. The principal's per-period payoff is equal to the expected payoff from her buying decision (equal to zero if she decides not to buy) minus the expected wages paid to the agent. In turn, the agent's per-period payoff is equal to the expected wage received minus the disutility of effort.

If the agent does not work for the principal, he can instead enjoy an outside opportunity that yields an expected payoff equal to zero.[7] To avoid a trivial solution, we assume that the agent is protected by limited liability, and thus wages are restricted to be nonnegative.

CONTRACTS. At the beginning of the first period, both parties sign a long-term contract that covers the duration of the relationship. It specifies wages to be paid in each period (and a recommended level of effort) after every possible history until the principal decides to stop and exercise an option, thereby ending the relationship. The principal can use lotteries, i.e., randomize over levels of wages and effort. We assume that only the arrival of an option and the decision to exercise it are contractible events. Therefore, wages depend on the history of arrivals and the purchase decision. In the binary-signal case, this is equivalent to conditioning payments on arrivals and signal realizations. Finally, we assume that the principal can commit to long-term contracts. The agent, however, can terminate the relationship in any period and enjoy his outside option. Thus, the agent's expected value from the contract in each period should be no less than his reservation utility.

TIMING. In each period while the game continues, the timing is as follows. If an option does not arrive, then the agent receives the wage specified in the contract and waits for the next period (unless the contract specifies ending the relationship). If an option arrives, then the agent exerts effort, observes a signal realization, and then reports whether or not

---

[7] We make this assumption to simplify the presentation of the results. With some minor modifications, the analysis can accommodate any nonnegative reservation utility.

there is an option at hand.[8] If the agent reports that there is one, the signal becomes public. The principal decides whether or not to exercise the option, and the agent receives the wage prescribed by the contract in accordance with this decision.[9] If the agent claims that there is no option, the events unfold as in the case of no arrival.

3 The Static Case

We first analyze the one-period version of the model. For simplicity, we assume that $p = 1$, because in the static case the agent never has incentives to hide an option as long as the arrival of an option is a contractible event.

To simplify the notation, we define $\pi(e) \equiv \gamma(\alpha + \beta_h \eta(e)) + (1-\gamma)(\alpha - \beta_\ell \eta(e))$ as the unconditional probability that a signal is high when the agent's effort level is e. Its derivative, which plays a key role below, is thus given by $\pi'(e) = \eta'(e)(\gamma\beta_h - (1-\gamma)\beta_\ell)$.

Notice that the principal implements a positive level of effort only if she intends to exercise an option after observing a high signal, and not exercise it after observing a low signal. Indeed, if the principal's decision is independent of the signal realization, then information is wasted, and hence putting effort into acquiring it is not optimal.

If the principal exercises the option after observing a high signal, her expected return conditional on observing a high signal given a level of effort is given by

$$E[y|\theta_h, e] \equiv \frac{\gamma(\alpha + \beta_h \eta(e))y_h + (1-\gamma)(\alpha - \beta_\ell \eta(e))y_\ell}{\pi(e)}.$$

3.1 Observable Effort

Let us consider the first-best case in which the agent's effort is observable and there is no limited liability. Since the solution for this case is equivalent to the one in which the principal acquires information herself, we focus on this problem. The principal solves

$$\max\left\{0,\ \gamma y_h + (1-\gamma)y_\ell,\ \max_e -\psi(e) + \pi(e)E[y|\theta_h, e]\right\}.$$

---

[8] If the contract involves randomization, then the outcome of the lottery is first observed.

[9] Since conditioning payments on the purchase decision is equivalent to conditioning them on a signal realization, the principal does not benefit from committing to the decision to buy. Thus without loss of generality we can assume that she makes this decision after observing a signal realization. In Section 6 we consider variations of the model where this is not the case.

The three terms inside the maximization represent the relevant alternatives available to the principal, namely, exert no effort and do not buy (which yields a payoff of zero), exert no effort and buy (with expected payoff $\gamma y_h + (1-\gamma)y_\ell$), and exert an optimal level of effort and buy only if a signal realization is high (yielding $\max_e -\psi(e) + \pi(e)E[y|\theta_h, e]$).

The optimal level of effort under the third alternative uniquely solves $\psi'(e) = \eta'(e)(\gamma\beta_h y_h - (1-\gamma)\beta_\ell y_\ell)$. We will denote by e^* the principal's optimal effort choice in the first-best case.

Intuition suggests that the principal's optimal decision is to never buy if γ is small enough and to always buy if γ is high enough. For intermediate values of γ, she acquires information and buys only if the signal is high. The following result formalizes this intuition.

Lemma 1 (Static First Best) *There exist thresholds* $\underline{\gamma}^*$ *and* $\bar{\gamma}^*$*, satisfying* $0 \leq \underline{\gamma}^* < -y_\ell/(y_h - y_\ell) < \bar{\gamma}^* \leq 1$*, with strict inequalities if* $\alpha \in (0,1)$*, so that the principal*

(i) Exerts no effort, i.e., $e^* = 0$*, and does not exercise the option if* $\gamma \in [0, \underline{\gamma}^*]$*;*

(ii) Exerts no effort, i.e., $e^* = 0$*, and exercises the option if* $\gamma \in [\bar{\gamma}^*, 1]$*;*

(iii) Exerts a positive level of effort, i.e., $e^* > 0$*, and exercises the option only when the signal realization is high if* $\gamma \in (\underline{\gamma}^*, \bar{\gamma}^*)$*.*

3.2 Unobservable Effort

We now turn to the case with unobservable effort — i.e., moral hazard — and limited liability. As before, the principal has three relevant alternatives: not to hire the agent and not to buy, not to hire the agent and buy, and hire the agent, induce him to exert a positive level of effort and buy only if the signal realization is high.

Under the third alternative, the principal designs a compensation scheme (w_0, w_1), where w_0 is the wage paid if the principal does not exercise the option and w_1 is the corresponding wage if she does. The optimal contracting problem is

$$\begin{aligned} \max_{e, w_0, w_1} \quad & \pi(e)(E[y|\theta_h, e] - w_1) - (1-\pi(e))w_0 \\ \text{s.t.} \quad & -\psi(e) + \pi(e)w_1 + (1-\pi(e))w_0 \geq 0, && (2) \\ & \psi'(e) = \pi'(e)(w_1 - w_0), && (3) \\ & e \geq 0, \;\; w_0 \geq 0, \;\; w_1 \geq 0. && (4) \end{aligned}$$

Constraint (2) is the agent's participation constraint. Constraint (3) summarizes the incentive constraints using the first-order condition of the agent's effort choice problem. Constraints (4) reflect the nonnegativity of effort and presence of limited liability.

Let us take a closer look at the agent's optimization problem under (w_0, w_1), namely, $\max_e -\psi(e) + \pi(e)w_1 + (1 - \pi(e))w_0$, whose first-order condition (3) is

$$\psi'(e) = \eta'(e)(\gamma\beta_h - (1 - \gamma)\beta_\ell)(w_1 - w_0). \tag{5}$$

There are two noteworthy features embedded in equation (5). First, let $\hat{\gamma} \equiv \beta_\ell/(\beta_\ell + \beta_h)$. Notice that $\hat{\gamma} \in [0, 1]$, taking the extreme values when $\beta_\ell = 0$ and $\beta_h = 0$, respectively, and the value 1/2 when $\beta_h = \beta_\ell$ (i.e., the symmetric information structure). From (5), $\pi'(e) > 0$ if $\gamma > \hat{\gamma}$ and $\pi'(e) < 0$ if $\gamma < \hat{\gamma}$. Thus, the information structure of the problem neatly partitions the solution of the principal's problem into *two distinct cases*: (*i*) $\gamma > \hat{\gamma}$, in which case $w_1 > w_0$ in order to induce any positive level of effort from the agent, and (*ii*) $\gamma < \hat{\gamma}$, in which case $w_0 > w_1$. In both cases the agent's problem is strictly concave in effort and thus the first-order approach (replacing the agent's optimization problem by its first-order condition) is valid (Rogerson, 1985a). Second, (5) reveals that as γ approaches $\hat{\gamma}$, it becomes *infinitely costly* to induce the agent to choose any positive level of effort. Hence, no effort is implemented for values of γ close to $\hat{\gamma}$. Intuitively, the reason is that at this threshold, effort has no impact on the unconditional distribution of the signal. Therefore, it is impossible to statistically distinguish whether or not the agent has exerted effort.

The following result summarizes main properties of the optimal contract.

Proposition 1 (Optimal Static Contract) *The optimal contract satisfies:*

(i) If $\gamma > \hat{\gamma}$, then $w_1 \geq w_0 = 0$, and if $\gamma < \hat{\gamma}$, then $w_0 \geq w_1 = 0$. In both cases, the inequality is strict if the implemented level of effort is positive.

(ii) The optimal effort level is lower than the first-best level e^*, and strictly so if $e^* > 0$.

(iii) The optimal buying decision is as follows:

(*a*) *There is a threshold $\underline{\gamma} \geq \underline{\gamma}^*$, with strict inequality unless $\alpha = 0$, such that the principal does not hire the agent and does not exercise the option when $\gamma \in [0, \underline{\gamma}]$;*

(*b*) *There is a threshold $\bar{\gamma} \leq \bar{\gamma}^*$, with strict inequality unless $\alpha = 1$, such that the principal does not hire the agent and exercises the option when $\gamma \in [\bar{\gamma}, 1]$;*

(*c*) *Unless $\hat{\gamma}y_h + (1 - \hat{\gamma})y_\ell = 0$, there exists an interval around $\gamma = \hat{\gamma}$ such that the principal does not hire the agent, and exercises the option if and only if $\hat{\gamma}y_h + (1 - \hat{\gamma})y_\ell > 0$;*

(*d*) *If $\hat{\gamma} \neq -y_\ell/(y_h - y_\ell)$, then there exists an interval of values of γ such that the optimal effort is positive at the optimum.*

As in any moral-hazard problem, compensation is structured according to the relative likelihood of observing a high and a low signal under the implemented effort. The appropriate likelihood ratios involve the unconditional probabilities of the signal realizations. Formally, the following statement holds:

$$w_1 \gtreqless w_0 \Leftrightarrow \pi'(e)/\pi(e) \gtreqless -\pi'(e)/(1-\pi(e)) \Leftrightarrow \pi'(e) \gtreqless 0 \Leftrightarrow \gamma \gtreqless \hat{\gamma} \Leftrightarrow \beta_h \gtreqless ((1-\gamma)/\gamma)\beta_\ell.$$

Notice that when $\gamma < \hat{\gamma}$, the optimal compensation scheme is such that the agent gets paid *only* when the principal does *not* exercise the option. We call this feature of the optimal compensation scheme *rewarding for bad news.* As the statement above reveals, $\gamma < \hat{\gamma} \Leftrightarrow \beta_h < ((1-\gamma)/\gamma)\beta_\ell$. Thus, the agent is rewarded for bad news if and only if the impact of his effort on the distribution of the signal is sufficiently stronger in the low vs. high state. The intuition underlying this property is easily explained in the extreme case where $\beta_h = 0$, and thus $\pi(e) = \alpha - (1-\gamma)\beta_\ell\eta(e)$. Since the agent's effort decreases the probability of a high signal, it makes observing a low signal relatively more likely when the agent chooses the effort level that the principal wants to implement, than when he chooses a lower level of effort. Hence, his wage should be higher when the principal does not buy.

If $\gamma > \hat{\gamma}$, then the agent is rewarded when the principal exercises an option; i.e., the agent is rewarded for good news. This is a typical result in standard principal-agent models with moral hazard, where effort stochastically increases the distribution of 'output' instead of affecting the informativeness of the signal as in our case. Within our framework, this is akin to assuming that effort affects the distribution of available options (i.e., higher effort makes a draw of a high-quality option more likely). Formally, let $\gamma(e)$ be the probability that an option's quality is high given effort level e, with $\gamma'(e) > 0$ for all e. Suppose that the agent observes a signal of quality $\theta \in \{\theta_\ell, \theta_h\}$, which is informative (i.e., the probability of observing a high signal when the quality is high is bigger than when it is low), and is independent of effort. Then it is immediate that $\pi'(e) > 0$, and the optimal contract *always* exhibits $w_1 \geq w_0 = 0$, i.e., the agent gets paid only when the principal exercises the option.

Regarding the implemented effort, moral hazard and limited liability distort it downward from its first-best level. This is a standard result in moral-hazard problems in which contracts are conditioned on a binary outcome (e.g., purchase or no purchase).

More interesting is the distortion in the principal's decision to buy. As in the first-best case, she never exercises an option when γ is close to zero, and always exercises it when γ is close to one. One can show that the intervals of γ at both ends are larger than in the first-best case, due to the presence of moral hazard. Also noteworthy is what happens around $\gamma = \hat{\gamma}$. The endogenous cost of inducing positive effort diverges to infinity and precludes

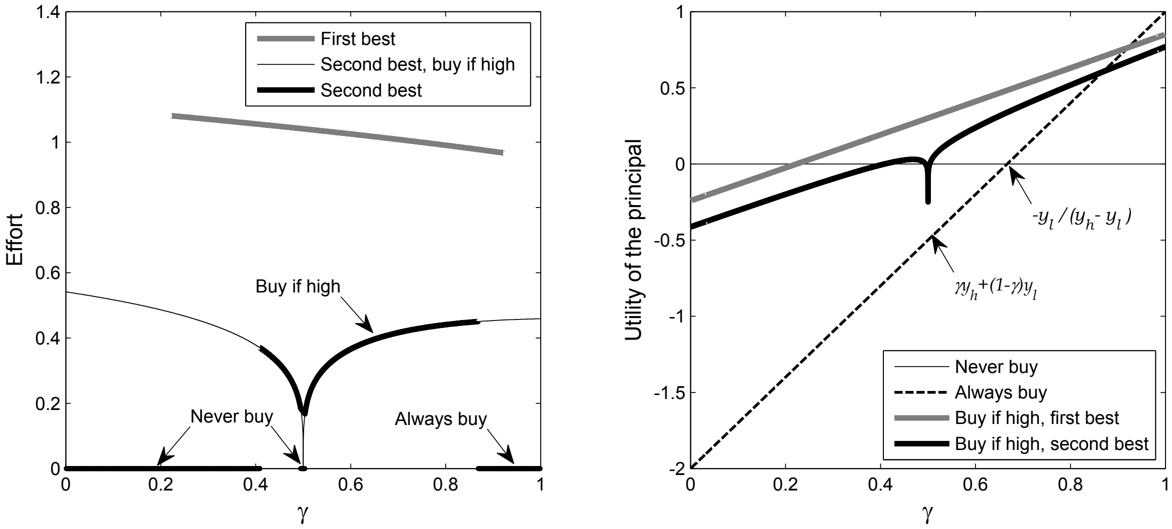

Figure 1: **Effort and Buying Decision in Symmetric Case.** The left panel depicts the first- and second-best effort choices. The right panel shows the principal's value function, i.e., the upper envelope of the expected payoff of the three relevant alternatives, in the first- and second-best cases as a function of γ. Computations were performed using $\beta_h = \beta_\ell$, power functional forms for η and ψ, and particular values for the option returns.

any incentive provision. Overall, the principal implements a positive level of effort (i.e., induces nontrivial information acquisition) less often than in the first-best case.

This begs the question of whether there is an interval of values of γ where the principal implements positive effort at all. A natural candidate is a neighborhood of $\gamma = -y_\ell/(y_h - y_\ell)$, since at that point the principal is indifferent between buying and not buying without information, and thus information is valuable. Clearly, the principal will not implement positive effort if $-y_\ell/(y_h - y_\ell) = \hat{\gamma}$. Once this is ruled out, positive effort is optimal if the cost of inducing a small amount of effort is less than its marginal value, which is implied by $\psi'(0)/\eta'(0) = 0$ and $(\psi'/\eta')'(0) = 0$.

Figure 1 illustrates the optimal effort and decision to buy as functions of the prior belief γ in the first- and second-best cases. Notice that the effort level in the latter is lower than in the former, and it is positive for intermediate values of γ outside an interval around $\hat{\gamma} = 1/2$. Also, the principal exercises the option after the high signal realization on a smaller set of values of γ in the second best.

3.2.1 Rewarding for Bad News and the Standard Principal-Agent Model

An insight that percolates throughout the paper is that the optimal contract sometimes rewards the agent when the principal does not buy, which, in the binary case, occurs when the signal realization is low. A simple explanation is that, as in any principal-agent model with moral hazard, the signal plays a dual role (see Grossman and Hart, 1983, p. 9). First, it affects the principal's decision to buy, and in this role she prefers a higher signal realization. Second, it provides information about the agent's effort, and in this role she might prefer a higher or a lower signal depending on which one is more informative — i.e., has a higher likelihood ratio — about effort. Thus, when the principal's preferences regarding these roles of the signal are not aligned, rewarding for bad news ensues. We claim that the paper provides an interesting class of principal-agent problems with moral hazard where a conflict between these two roles emerges *endogenously* in a natural and robust way.

To support this claim, we illustrate the pitfalls of obtaining a similar result in a standard principal-agent problem with moral hazard and two outcomes, success and failure. As primitives of the problem, suppose that success yields revenue $q > 0$ while failure, for simplicity, yields 0. Also, let $\pi(e)$ be the probability of success. The rest is as in our model.

If one *exogenously* assumes that the probability of success is *decreasing* in effort, i.e., $\pi'(e) < 0$, then clearly the optimal contract that implements any positive level of effort pays a higher wage if failure is observed. That is, if the principal decides to implement a positive level of effort, then the optimal contract will exhibit the rewarding-for-bad-news feature. But notice that in this case the optimal level of effort that the principal implements is trivially zero, since expected revenue $\pi(e)q$ decreases with effort.

To be sure, one could enrich the model by assuming that success yields a random output whose mean $\bar{q}(e)$ is increasing in e, and that contracts can only be conditioned on success or failure but not on the level of output observed. But since $\pi(e)$ decreases with e, one needs to make sure — via another exogenous assumption — that $\pi(e)\bar{q}(e)$ does *not* decrease in effort, for otherwise the optimal choice for the principal would again be $e = 0$.

In short, obtaining a nontrivial rewarding-for-bad-news result in a standard principal-agent problem with moral hazard requires a host of ad-hoc assumptions that have little economic content. By contrast, our information acquisition setup satisfies all of these conditions endogenously and in an economically meaningful way.

4 The Dynamic Case

We now turn to the dynamic version of the model, and assume that options arrive stochastically over time, with per-period arrival probability $p \in (0, 1]$. As we will see, the time dimension makes the problem much richer than its static counterpart.

4.1 Observable Effort

As in the static case, we start with the first-best case in which the principal acquires information herself. The problem becomes a simple optimal stopping exercise. Since the environment is stationary, so is the principal's strategy. Thus, there are three possible scenarios: (a) she does not exert effort and never buys, which yields zero profits; (b) she does not exert effort and buys whenever an option is available, which yields $[\gamma y_h + (1-\gamma)y_l]p/[1-(1-p)\delta]$; and ($c$) she always exerts a constant positive level of effort whenever an option is available, and buys the first time a signal realization is high.

Let S^* be the expected discounted profits for the principal under strategy (c). Then S^* satisfies $S^* = \max_e p(-\psi(e) + \pi(e)E[y|\theta_h, e]) + [p(1-\pi(e)) + 1 - p]\delta S^*$. The first-order condition to this problem is $\psi'(e) = \eta'(e)(\gamma\beta_h y_h - (1-\gamma)\beta_\ell y_\ell - (\gamma\beta_h - (1-\gamma)\beta_\ell)\delta S^*)$. It is easy to verify that the level of effort that solves this equation is positive.

Let e^{**} be the optimal level of effort in the dynamic first-best case. The following result, describes main properties of the solution to the principal's problem.

Lemma 2 (Dynamic First Best) *There exist threshold* $\underline{\gamma}^{**}$ *and* $\bar{\gamma}^{**}$*, satisfying* $0 \leq \underline{\gamma}^{**} < -y_\ell/(y_h - y_\ell) < \bar{\gamma}^{**} \leq 1$*, with strict inequalities if* $\alpha \in (0, 1)$*, so that the principal*

(i) *Exerts no effort, i.e.,* $e^{**} = 0$*, and does not exercise the option if* $\gamma \in [0, \underline{\gamma}^{**}]$*;*

(ii) *Exerts no effort, i.e.,* $e^{**} = 0$*, and exercises the option if* $\gamma \in [\bar{\gamma}^{**}, 1]$*;*

(iii) *Exerts a positive level of effort, i.e.,* $e^{**} > 0$*, and exercises the option if and only if the signal realization is high, if* $\gamma \in (\underline{\gamma}^{**}, \bar{\gamma}^{**})$*;*

(iv) *If* $\gamma \gtreqless \hat{\gamma}$ *then* $e^{**} \lesseqgtr e^*$ *with strict inequality if* $e^* > 0$ *and* $e^{**} > 0$*),* $\underline{\gamma}^{**} \gtreqless \underline{\gamma}^*$ *(with strict inequality unless* $\alpha = 0$*) and* $\bar{\gamma}^{**} \lesseqgtr \bar{\gamma}^*$ *(with strict inequality unless* $\alpha = 1$*).*

We stress the simple nature of the problem's solution in the first-best case, derived from an extremely basic search problem that the principal solves. This simplicity is all the more striking when compared to the richness of the dynamics and distortions that we obtain below, which are due to the moral-hazard problem that delegating the information acquisition task to the agent introduces.

4.2 Unobservable Effort

We now proceed to analyze the principal's optimal contracting problem in the second-best case, i.e., when effort is unobservable and there is limited liability.

4.2.1 Recursive Formulation of the Contracting Problem

Following a standard argument, we can formulate the optimal contracting problem recursively using the *promised value* to the agent as a state variable. To this end, let $U(V)$ be the principal's value function when the promised value to the agent is V. Also, denote by $U_{10}(V)$ the value to the principal if she chooses to exercise an option in the current period only after observing a high signal. Similarly, let $U_{11}(V)$ and $U_{00}(V)$ denote, respectively, the value to the principal if she buys (given that an option is available) and does not buy in the current period independently of the signal realization, and continues to search in the next period if no purchase is made. The principal can also choose to stop the search without buying, in which case she simply pays V to the agent. Then the principal's value function can be written as

$$U(\cdot) = \operatorname{cav}\max\{U_{00}(\cdot), U_{11}(\cdot), U_{10}(\cdot), -(\cdot)\}, \tag{6}$$

where the operator cav denotes the concavification of a continuous function; i.e., cav$f$ is the least concave function g such that $g(x) \geq f(x)$ for all x in the domain of f. Thus, U is the least concave function such that $U(V) \geq \max\{U_{00}(V), U_{11}(V), U_{10}(V), -V\}$ for all V. The economic interpretation of this is as follows. Although the four functions on the right-hand side are concave, the maximum of them need not be. Thus, the principal may benefit from using lotteries over contracts and the decision to buy. We will show that she indeed uses them in a nontrivial way.

Let us turn now to a detailed description of the functions inside the max operator. Notice that as in the static case, if the principal's decision to buy today does not depend on the signal realization, then it is optimal for her to implement zero effort. Also, since both the principal and the agent are risk neutral, all payments can be postponed until the end of the relationship. In particular, if the principal does not buy today, then she simply rolls over the agent's promised value to the next period. These considerations lead to the following expressions for U_{00} and U_{11}:

$$\begin{aligned} U_{00}(V) &= \delta U(V/\delta), \\ U_{11}(V) &= p[\gamma y_h + (1-\gamma)y_\ell - V] + (1-p)\delta U(V/\delta). \end{aligned}$$

Since in both cases the expected value to the agent is the same whether an option is available or not (equal to V), he does not have incentives to hide the presence of an option.

The description of U_{10} is much more complex, as the principal uses the information revealed by the agent in her optimal stopping decision. This implies a nontrivial incentive provision to induce the agent to exert effort and to reveal the presence of an option in the current period. In this case the agent receives compensation w_1 if the principal buys (i.e., if the signal realization is high). If a low signal is observed and hence the principal does not buy, then the agent receives a wage w_0 and his promised continuation value becomes V_0. If no option is available, then he receives w_n and continuation value V_n. The maximization problem of the principal when she buys after a high signal is as follows:[10]

$$\begin{aligned} U_{10}(V) = \max_{e,w_1,w_0,w_n,V_0,V_n} & \; p[\pi(e)(E[y|\theta_h,e] - w_1) + (1-\pi(e))(-w_0 + \delta U(V_0)] \\ & + (1-p)(-w_n + \delta U(V_n)) \\ \text{s.t.} \quad & p\left[-\psi(e) + \pi(e)w_1 + (1-\pi(e))(w_0 + \delta V_0)\right] + (1-p)(w_n + \delta V_n) = V, && (7) \\ & \psi'(e) = \eta'(e)(\gamma\beta_h - (1-\gamma)\beta_\ell)(w_1 - w_0 - \delta V_0), && (8) \\ & w_1 \geq w_n + \delta V_n, && (9) \\ & w_0 + \delta V_0 \geq w_n + \delta V_n, && (10) \\ & e \geq 0, \;\; w_1 \geq 0, \;\; w_0 \geq 0, \;\; w_n \geq 0, \;\; V_0 \geq 0, \;\; V_n \geq 0. && (11) \end{aligned}$$

Constraint (7) is a *promise-keeping* constraint; i.e., the principal has to deliver the promised value V with her choice of $(e, w_1, w_0, w_n, V_0, V_n)$.[11] The second constraint is the incentive constraint, which ensures that the agent optimally chooses the effort implemented by the

---

[10] To simplify the exposition, we formulate this problem as if the principal can only use deterministic contracts. But since we want to ensure that U_{10} is concave, in the Appendix we set up the problem assuming that the principal can use lotteries. That is, her choice variables are, in general, probability mixtures over deterministic contracts $(e, w_1, w_0, w_n, V_0, V_n)$. The outcome of the lottery is observed before the agent exerts effort. Our numerical computations suggest that lotteries are not used in this problem, i.e., U_{10} formulated above is in fact concave.

[11] Note that, as is standard in models with moral hazard, the promise-keeping constraint has to be imposed with equality in all periods except the first period, when it is imposed with weak inequality.

principal. As in the static case, it is summarized by the first-order condition of the agent with respect to effort, obtained by differentiating the left-hand side of (7) with respect to e. Constraints (9)–(10) are the *no-hiding* constraints: they guarantee that the agent will not have incentives to claim that there is no option available instead of revealing the current signal realization. Finally, constraints (11) ensure that wages are nonnegative (limited liability) and agent promised values are higher than the (zero) reservation utility.

A technical hurdle in solving this problem is that, although U is concave by construction, it need not be differentiable everywhere. Its presence in the objective function makes the problem nonsmooth, and we tackle it in the Appendix using superdifferential calculus.

As in the static case, the cost of implementing any positive effort level goes to infinity as γ approaches $\hat{\gamma}$. Also, it is easy to verify that $U(V) = U_{00}(V) = -V$ if γ is close to zero, and $U(V) = U_{11}(V) = (\gamma y_h + (1-\gamma)y_\ell)p/(1-(1-p)\delta) - V$ if γ is close to one.

4.2.2 Optimal Contract when $U(V) = U_{10}(V)$

Suppose the promised value to the agent V is such that $U(V) = U_{10}(V)$. That is, at V the principal finds it optimal to exercise an option so long as the signal realization is high.

Before we derive the main properties of the optimal contract in this case, let us go over a couple of useful results that are proved in the Appendix.

First, our assumptions ensure that the implemented effort level must be positive. It then follows from the incentive constraint that the solution again partitions into two cases: (i) $\gamma > \hat{\gamma}$, which implies that $w_1 > w_0 + \delta V_0$, and (ii) $\gamma < \hat{\gamma}$, which implies that $w_1 < w_0 + \delta V_0$.

Second, the value function U has slope greater than minus one. For a somewhat loose intuition at this point, when V is large enough the first best is achieved and $U(V)$ coincides with the first-best frontier, which has slope equal to minus one. Concavity then implies that the value function has slope greater than minus one for all V, with strict inequality unless the implemented effort is first best.

A straightforward implication of this result is that the principal weakly prefers to set $w_0 = w_n = 0$, as the marginal cost of compensating the agent with a current wage is one, while the marginal cost of compensating him using continuation values (given by the absolute value of the slope of $U(V_i)$, $i \in \{0, n\}$) is less than one. Moreover, as long as the inequality is strict, setting $w_i > 0$, $i \in \{0, n\}$, is not optimal, and the principal *strictly prefers* to postpone payments.

This observation reveals an important difference between the static and dynamic cases. In the static case, the agent is compensated in the event of purchase if and only if $\gamma > \hat{\gamma}$.

By contrast, in the dynamic case, the agent gets paid at the *end* of the relationship for *all* values of γ. As long as the principal implements positive effort, this event *coincides with the event of purchase.* To provide the agent with incentives to exert effort, the level of the payment depends on the history, summarized by the promised value V.

Clearly, when the slope of U is minus one, the principal is indifferent between making payments now or in the future. We will see that the first best obtains in this case.

Define $V^* = \inf\{V | U'(V) = -1\}$. The next result characterizes the optimal contract.

Proposition 2 (Optimal Contract when Buying after a High Signal) *Let* $p \in (0,1)$. *If at* V *the principal finds it optimal to buy only after a high signal (i.e.,* $U(V) = U_{10}(V)$*), then the optimal contract has the following properties:*

(*i*) *Suppose that* $V \in (0, V^*)$.

  (*a*) *If* $\gamma > \hat{\gamma}$, *then* $w_1 > \delta V_n = \delta V_0$ *and* $V_n = V_0 < V$.

  (*b*) *If* $\gamma < \hat{\gamma}$, *then* $w_1 = \delta V_n < \delta V_0$ *and* $V_n < V < V_0$.

(*ii*) *If* $V = 0$ *then* $e = w_1 = V_0 = V_n = 0$.

(*iii*) *If* $V^* > 0$, *then for all* $V \geq V^*$ *effort level* e^{**} *is implemented in every period.*[12]

As in the static case, when $\gamma < \hat{\gamma}$, the agent is *rewarded for bad news* (part (*i*)–(*b*)). An interesting feature of the dynamic case is that he is not compensated immediately when revealing a low signal, but via changes in his continuation value (i.e., future payments). Rewarding for bad news makes the principal vulnerable to a false claim of no option availability when the signal is high. To prevent this, she sets $w_1 = \delta V_n$. The agent's continuation value increases after a low signal and drifts downwards if he claims that no option is available.

Similarly, when $\gamma > \hat{\gamma}$, the agent is rewarded for good news (part (*i*)–(*a*)) and has incentives to hide low signals. Thus, in this case the principal makes the agent indifferent between reporting a low signal and claiming that no option is available, $V_0 = V_n$. The agent is punished in both cases by a decrease in his promised continuation value.

It is instructive to compare this result to the case in which the agent cannot hide an option because its arrival is observed by both parties (e.g., as in the home inspector example). In this case the agent's value remains constant when an option is not available, i.e., $V_n = V$. In contrast, Proposition 2 reveals that if hiding is possible, then the agent is always punished by a lower continuation value when he claims that no option has arrived.

Part (*ii*) reveals that $V = 0$ is an absorbing state, since all continuation values are also equal to zero. The agent exerts no effort and does not get paid. Notice that in this case it

---

[12] The threshold V^* is positive if the principal implements $e > 0$ at the optimum for some V.

can only be optimal to buy only after a high signal if the principal is indifferent between buying and not buying without information, i.e., when $\gamma = -y_\ell/(y_h - y_\ell)$.

Finally, part (*iii*) shows that when the agent's promised value is sufficiently high, the principal implements the first-best level of effort from that period onwards. The basic reason is that in this case limited liability is no longer binding.

A case not covered by Proposition 2 is $p = 1$, where hiding has no bite and thus the choice of V_n is irrelevant. All the other properties stated in this section extend to this case. Moreover, when $\gamma < \hat{\gamma}$ and $V \leq V^*$, one can show that $w_1 = 0$, and hence the agent only gets paid when the first best is achieved. That is, all payments are postponed until sufficiently many low signals are generated and V increases above V^*. One way to implement the first-best effort in this case is with a payment independent of the signal, i.e., $w_1 = w_0$ (e.g., a fixed payment per inspection in the home inspector example).

In the static case we showed that optimal effort was always below its first-best level. This property partially extends to the dynamic case as well.

Claim 1 (Optimal Effort) *If* $\gamma < \hat{\gamma}$, *then* $e(V) \leq e^{**}$ *for all* V.

To understand the intuition behind this result and why it is only stated for $\gamma < \hat{\gamma}$, consider the marginal cost and marginal benefit of effort in the first- and second-best cases. Regardless of γ, the marginal cost of effort is higher in the second-best case due to moral hazard. Also, when $\gamma < \hat{\gamma}$, higher effort makes continuing the search more likely. Since future social surplus in the second best is lower than in the first best, the social benefit of continuing the search is also lower, and hence the marginal benefit of effort is lower. And since the marginal cost is higher and the marginal benefit is lower, lower effort obtains. However, when $\gamma > \hat{\gamma}$, both the marginal cost and marginal benefit are lower, making the comparison of effort levels ambiguous.

4.2.3 Optimal Buying Decision and Use of Lotteries

Thus far we have derived several properties of the optimal contract under the assumption that the agent's promised value is such that it is optimal for the principal in the current period to exercise an option only when the realization of the signal is high. To complete the description of the optimal contract, we need to understand when it would be optimal for the principal to induce zero effort and simply buy or not buy. Furthermore, since the buying decision is discrete, we cannot ignore the possibility that the optimal contract involves the use of lotteries, as the value function in equation (6) suggests.

The following result fills this gap and characterizes the optimal buying decision as a function of the agent's promised value. In words, it states that the use of lotteries occurs for sufficiently low values of the agent's promised value, where the principal randomizes between inducing no effort and buying (or not buying) and inducing positive effort and buying only after the high signal realization. The value of γ determines the buying decision at the lowest levels of the agent's promised value.

Proposition 3 (Use of Lotteries in the Optimal Contract) *(i) If* $\gamma > -y_\ell/(y_h - y_\ell)$*, then there exists a threshold* $\tilde{V}_1$ *such that*

(a) $U(V) = U_{10}(V)$ *for* $V \geq \tilde{V}_1$*, and thus the principal buys only if the signal is high;*

(b) $U(0) = U_{11}(0) > 0$*, and thus the principal always buys at* $V = 0$*;*

(c) The principal randomizes between the two alternatives for values $V \in (0, \tilde{V}_1)$*.*

(ii) If $\gamma < -y_\ell/(y_h - y_\ell)$*, then there exists a threshold* $\tilde{V}_0$ *such that*

(a) $U(V) = U_{10}(V)$ *for* $V \geq \tilde{V}_0$*, and thus the principal buys only if the signal is high;*

(b) $U(0) = U_{00}(0) = 0$*, and the principal stops the search without buying at* $V = 0$*;*

(c) For $V \in (0, \tilde{V}_0)$*, randomizing between the two alternatives is optimal. Moreover,* $U_{00}(V) = U(V)$ *on* $[0, \delta\tilde{V}_0]$*, and hence the optimal strategy on* $(0, \tilde{V}_0)$ *is indeterminate. In particular, not buying and rolling over the agent's value to the next period is optimal.*

(iii) If $\gamma = -y_\ell/(y_h - y_\ell)$*, then buying after the high signal is an optimal strategy for all* V*.*

Figure 2 provides a graphical illustration of the above result. Although it is intuitive that the principal might benefit from the use of lotteries due to the discrete features of the problem, it is somewhat surprising that their use in (6) can be pinned down *exactly*, as stated in the proposition.[13] Indeed, a priori it is unclear when they are going to be used, since the problem involves the intersection of four concave functions, which can cross multiple times. As the proof reveals, however, the structure of the principal's problem is such that they all cross at most once, thereby making the analysis manageable.

Finally, notice that incentive provision requires that, as a punishment, the principal sometimes lowers the agent's continuation value. But when the promised value is low enough, it is no longer optimal to acquire information with probability one, and sometimes the principal makes the purchase decision solely based on the prior. In other words, in order to optimally provide incentives in the *current* period, the principal chooses to sacrifice *future*

---

[13] To be precise, Proposition 3 characterizes the use of lotteries in problem (6), where randomization is over both contracts and the decision to buy. As mentioned, the principal could also use lotteries over contracts in the U_{10} problem, where the decision to buy is fixed.

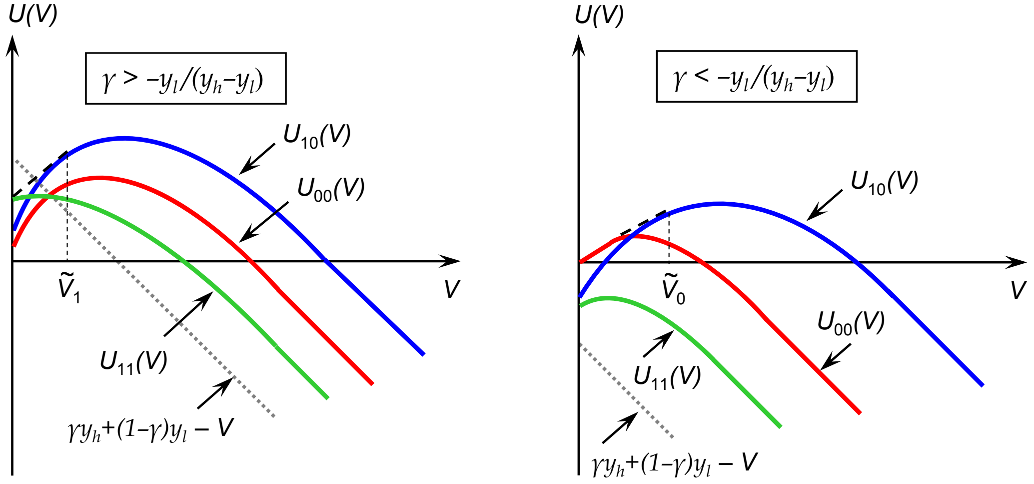

Figure 2: **Use of Lotteries.** The right panel illustrates the use of lotteries due to the intersection between U_{00} and U_{10}, which are the only relevant ones when $\gamma < -y_\ell/(y_h - y_\ell)$. The dashed line shows how the use of lotteries concavifies the principal's problem. The left panel provides a similar illustration when $\gamma > -y_\ell/(y_h - y_\ell)$.

information gains. Compared to the first-best case, this illustrates the distortion that moral hazard has on information acquisition, as well as on the principal's purchase decision.[14]

4.2.4 Dynamics Induced by the Optimal Contract

The properties of the optimal long-term contract described in Propositions 2 and 3 have important implications for the dynamics of variables such as the agent's effort, his promised value, and the principal's decision to buy.

Consider first the case of $\gamma > \hat{\gamma}$. So long as $U = U_{10}$, it follows from Proposition 2 that the agent's promised value decreases over time. This is in line with predictions that emerge in other dynamic principal-agent models with moral hazard. In our setup, however, the agent's promised value will eventually reach the lotteries region $[0, \tilde{V}_i]$. Notice that once the agent's value is in that region, it never leaves it. If $\gamma > -y_\ell/(y_h - y_\ell)$, then for a random number of periods (possibly zero), effort is constant at $e(\tilde{V}_1)$. This spell ends either when the principal buys after observing a high signal, or when the agent's value and effort drop to zero, and the principal stops and exercises the first available option. (Since $w_1 = 0$ when

---

[14] Note in passing that the principal's search problem ends in finite time almost surely. The only path along which it can continue forever is when an infinite number of consecutive low signal realizations is observed, which is clearly a zero-probability event.

$V = 0$, the ‘fired’ agent receives no payment.) That is, if the agent does not generate a high signal for many periods, then the principal eventually gives up and buys based solely on the prior. Once again, notice the sharp contrast with the first-best case, where effort is constant over time, and the principal continues until she observes a high signal.

If $\gamma < -y_\ell/(y_h - y_\ell)$, then the dynamics are similar, except that when the agent’s value drops to zero (an absorbing state), the principal stops the search without buying. Since her strategy is indeterminate for $V \in (0, \delta\tilde{V_0}]$, other patterns of dynamics can also occur. In particular, the agent’s effort can fluctuate over time between zero and $e(\tilde{V_0})$. That is, the principal can temporarily ‘suspend’ the agent (for a random number of periods) and then resume the search. Since the principal is indifferent between lotteries and the pure strategy of not buying, these two alternatives give her the same ex-ante utility.

Next, consider the case of $\gamma < \hat{\gamma}$, which is the novel case that springs from our model. Notice that now fluctuations in V and e over time can occur not only when lotteries are active, but also while $U = U_{10}$. To see this, recall that Proposition 2 reveals that the agent’s promised value increases after a low signal realization and decreases when no option is available. Therefore, the agent’s value, and consequently his effort, fluctuate over time. Whether and when the lotteries region becomes active depends on the number of periods with low signal realizations vs. no options. Moreover, the lotteries region is no longer absorbing, as the agent’s continuation value after a low signal will be above $\tilde{V_i}$ once the lottery outcome is to implement $e(\tilde{V_i})$ and buy after a high signal.

Thus, the interaction between hiding and rewarding for bad news causes the agent’s value and effort to fluctuate over time. Moreover, even in the absence of either of these features, fluctuations can occur in the region of indeterminacy when $\gamma < -y_\ell/(y_h - y_\ell)$.

The dynamics for $\gamma < \hat{\gamma}$ are different when the agent cannot hide presence of an option (e.g., as when $p = 1$). So long as in the initial period the principal implements positive effort with probability one (i.e., $U = U_{10}$), the agent’s value never decreases, and the search only ends after a high signal is eventually generated. The agent never gets fired, and the longer the search continues, the better off he is. This is in contrast to the results in the $\gamma > \hat{\gamma}$ case, where a search that generated bad news for a long enough time always ends in the principal firing the agent and making the purchase decision solely based on the prior. Intuitively, this difference in dynamics is precisely due to the fact that the agent is rewarded/punished for bad news when γ is below/above $\hat{\gamma}$.

5 Suggestive Evidence

We have analyzed a parsimonious model so as to zero in on the main trade-offs involved in contracting settings with delegated information acquisition. Among the results obtained, we have highlighted the rewarding-for-bad-news property as well as distortions in both effort provision and buying decision. In this section we review some applications of the model and discuss whether our theoretical predictions accord with features of contracts observed in reality. Since it is difficult to assess the empirical content of the aforementioned distortions absent detailed evidence on real-world contracts, we will focus on suggestive evidence of the presence of rewarding for bad news.

Recall that when the belief about the quality of an option is sufficiently high (low), the principal rewards the agent for good (bad) news. Alternatively, rewarding for good or bad news emerges depending on whether the agent is sufficiently more apt at detecting high- or low-quality options by exerting effort. Viewed in this way, it is intuitive and natural to envision situations where effort affects the informativeness of a signal about option quality and compensation is structured as mentioned. And if in a particular environment the agent's effort is instrumental in recognizing faults or errors (i.e., when β_ℓ is high relative to β_h, and hence $\hat{\gamma}$ is also high), then rewarding the agent for catching an error *should* be an essential part of his compensation scheme.

One example that fits this setting is a department of quality control. For an illustration, consider an engineer testing the reliability of particular auto parts of a car manufacturer. If an engineer testing Toyota's brake system had detected their infamous mechanical problem in a timely manner, saving the company millions of dollars in losses due to recalls and lawsuits, then he would have likely received a generous bonus from the company. Similarly, one can also surmise that structural engineers who test the foundations of buildings should also be rewarded for detecting faults which can help avert a disaster.

Although we do not observe how compensations of quality control department employees are structured in particular companies, there is direct evidence in the software industry of rewarding outside researchers for providing 'bad news.'[15] Indeed, some software vendors establish so called "bug bounty" or vulnerability reward programs that promise monetary prizes to people who find vulnerabilities (bugs) in source codes of the companies' products.

---

[15] Some sources, e.g., http://trailofbits.com/2009/03/22/no-more-free-bugs/, do in fact suggest that "software vendors pay their own employees and consultants to find [bugs] and help them fix them in their products during development". But notice that whether a researcher is employed or not by the company is not crucial for making our point that these companies reward for bad news.

The two most well-known case are Google, that rewards for bugs found in the Chrome web browser, Gmail, YouTube, Blogger and other products, and Mozilla, that has a similar program for the Firefox web browser. The standard reward is \$500 for an eligible bug, although particularly severe or clever bugs earn rewards of \$3,000 and more. Also, social networking site Facebook recently announced a bug bounty program with a base reward of \$500.[16] Notice that these programs are consistent with an optimal contract in the case where the agent's effort only affects the distribution of the signal in the low-quality state (i.e., $\beta_h = 0$), which is a very plausible assumption in this case.

One of the examples mentioned in the Introduction is that of a home inspector hired by a prospective home buyer. Arguably, in this case the home inspector's effort is more important for finding out bad rather than good news about houses.[17] In reality, home inspectors are paid a given amount for each inspection, independently of the report. Notice that this class of contracts *does* generate incentives to deliver bad news in our environment.[18] This is so because after bad news the search continues, which allows the inspector to remain employed by the prospective buyer.[19]

Another application mentioned earlier is a real-estate agent. In reality, real-estate agents are usually paid a fraction of the price of a sold house. The price of a house can be viewed as a noisy signal of the house quality (see also Section 6.2). For instance, if the posted price is particularly low, the potential buyer revises her prior about the quality of the house downwards. In the context of our model, consider a simple example where the price can be high or low, and the posterior after observing the high/low price falls above/below the threshold $\hat{\gamma}$. Consider the performance of a contract where the agent's payment is proportional to the price of a house. If the difference in prices and the probability of observing a high price are large enough, the agent will prefer an opportunity of selling a high-price house in the future rather than selling a low-price house today. Then the described compensation scheme provides the agent with incentives to deliver bad/good news about a house precisely when the beliefs about the house quality is below/above the threshold (i.e., when the price is low/high.) Thus, even though such a contract may be

---

[16] Examples of online sources that document this are http://krebsonsecurity.com/2010/11/google-extends-security-bug-bounty-to-gmail-youtube-blogger/ and http://www.facebook.com/whitehat/bounty/.

[17] Notice that several tests (as well as their results) that home inspectors perform on a house are verifiable and are difficult to misreport, so it is not unreasonable to assume that the signal is publicly observable.

[18] In fact, with an appropriate choice of payment such a contract is an optimal way of implementing the first-best level of effort.

[19] An implicit assumption here is that the home inspector is independently hired by the potential buyer rather than suggested by the real-estate agent.

suboptimal, its properties appear to be consistent with some of our model predictions.

It is important to keep in mind that our model focuses solely on the agent's role in acquiring information. In some environments, however, such as the case of a headhunter, the agent's effort can potentially affect *both* the informativeness of a signal about job candidates as well as improve the pool of available candidates. For example, a headhunter can carefully select a subset of candidates to review from a large set of job applicants, and then evaluate each candidate individually in more detail. That is, in this case the effort improves both *selection* (the quality of the draw) and *evaluation* (the informativeness of the signal about the draw). It is easy to derive implications of adding this feature to our model. Recall our discussion in Section 3 that when effort affects γ instead of information quality, the agent is always rewarded for good news. Therefore, our theory predicts that when the agent's effort affects both margins, incentives related to the two margins are aligned only when the quality of the pool is sufficiently improved, i.e., when γ exceeds $\hat{\gamma}$.

A less standard application of our model is that of news media shows or publications, where viewers or readers play the role of the principal, and the show or publication that of the agent. It is not far-fetched to posit that a significant part of viewers or readers would prefer the economy or political process to be in a good shape and thus would welcome good news — i.e., similar to our model, where the principal prefers a high signal. But if the perception is that things are not going well (pessimistic beliefs), then they will tend to *view* or *read* more (i.e., reward) those shows or publications that deliver bad news. Notice that no preference bias is needed to obtain such a prediction, as readers in this example do not prefer to see their pessimistic beliefs being confirmed by these media outlets.

6 Variations of the Model

We now explore some variations of the basic model. The purpose is twofold: first, to assess the robustness of our main results, and second, to understand if additional insights emerge.

6.1 Continuous-Signal Case

Our model assumes a binary-signal structure. Allowing for a continuous-signal structure introduces an additional margin in the contract design problem: the principal's continuous choice of the set of signal realizations under which she exercises an option. To illustrate main implications of this extension, we analyze a simple information structure with a continuum of signal realizations. For simplicity, we will restrict attention to the static case.

Let $\theta \in [\underline{\theta}, \bar{\theta}]$, and assume conditional densities $f(\theta|h,e) = (\bar{\theta}-\underline{\theta})^{-1} + \beta_h g(e)(\theta-\mu)$ and $f(\theta|\ell,e) = (\bar{\theta}-\underline{\theta})^{-1} - \beta_\ell g(e)(\theta-\mu)$, where $\beta_i \geq 0$, $i = \ell, h$, $\beta_\ell + \beta_h > 0$, and $g(e)$ is strictly increasing and differentiable, with $g(0) = 0$, and $\mu = \int_{\underline{\theta}}^{\bar{\theta}} \theta d\theta / (\bar{\theta}-\underline{\theta}) = (\bar{\theta}+\underline{\theta})/2$. Notice that the signal is uninformative if $e = 0$ (both densities collapse to the uniform distribution on $[\underline{\theta}, \bar{\theta}]$), and it is more informative the higher the effort level (signal realizations become more discriminating about the true quality of the option).

In this setting, when the principal hires the agent, she exercises the option if the signal realization is above a threshold, denoted by θ^*. Thus, the probability that she buys is

$$\pi(e,\theta^*) = \int_{\theta^*}^{\bar{\theta}} [\gamma f(\theta|h,e) + (1-\gamma) f(\theta|\ell,e)] d\theta = \frac{\bar{\theta}-\theta^*}{\bar{\theta}-\underline{\theta}} + (\gamma\beta_h - (1-\gamma)\beta_\ell) g(e) \frac{(\bar{\theta}-\theta^*)(\theta^*-\underline{\theta})}{2}.$$

The principal chooses (w_0, w_1, e, θ^*) to maximize expected profits subject to the usual constraints. The incentive constraint, $\pi_e(e,\theta^*)(w_1 - w_0) = \psi'(e)$, depends on θ^*. Thus, the principal might want to manipulate θ^* to ease the provision of incentives for effort.

In the first best, the principal sets this threshold at the signal realization that makes her indifferent between buying and not buying. In the second best, however, assuming that she can *commit* to a purchase decision, she *distorts* θ^* relative to the ex-post optimal level (i.e., the signal that makes her indifferent between buying and not buying).[20]

Claim 2 (Continuous Signal) *The optimal contract exhibits the following properties:*

(*i*) *If $\gamma > \hat{\gamma}$, then $w_1 > w_0 = 0$ and θ^* is distorted upwards.*

(*ii*) *If $\gamma < \hat{\gamma}$, then $w_0 > w_1 = 0$ and θ^* is distorted downwards.*

As before, the optimal contract exhibits the rewarding-for-bad-news property when $\gamma < \hat{\gamma}$. Also, the principal finds it optimal to forgo some options with positive expected payoff when $\gamma > \hat{\gamma}$, and accept some options with negative expected payoff when $\gamma < \hat{\gamma}$.

To understand the intuition behind the distortion in θ^*, assume $\gamma > \hat{\gamma}$ and let the signal at which the principal is just indifferent between buying and not buying be below the mean μ. We now argue that if she uses this ex-post optimal threshold to implement a given effort level, then she can do better by marginally increasing that threshold. To see this, notice that $\pi_{e\theta^*} = (\gamma\beta_h - (1-\gamma)\beta_\ell) g'(e)(\mu - \theta^*)$, so that when $\gamma > \hat{\gamma}$, $\pi_e(e,\theta^*)$ is increasing in θ^* for θ^* below μ. Thus, if the principal marginally increases the threshold from its ex-post

---

[20] The signals at which the principal is ex-post indifferent between buying and not buying in the first- and second-best cases are different, as the corresponding levels of effort are different.

optimal level (which has a negligible effect on her expected revenue from buying since it is at its maximum) and at the same time lowers w_1, then she can still induce the same level of effort at a cheaper cost, as she makes a lower payment with a lower probability. Hence, forgoing some options with positive expected payoff is profitable for the principal.

6.2 The Principal Observes an Additional Signal

In our model, the agent's signal is the only source of information available to the principal. In certain environments, she may observe an additional (contractible) signal about an option upon exercising it. We now analyze the effects of this extra source of information.

Suppose that if the principal exercises an option, then she observes an additional signal $\sigma \in \{\sigma_\ell, \sigma_h\}$ whose informativeness is exogenous, given by $\Pr\{\sigma_h | y = i\} = \rho_i$, $i = \ell, h$, with $0 \leq \rho_\ell \leq \rho_h \leq 1$.[21] In the investor example, the quality of an option can be the project type, with a project of type i succeeding with probability ρ_i. After investing in the project, the investor observes whether it succeeded or failed, and in addition to the purchase decision, she can condition the payment to the expert on this event.

For simplicity, we restrict attention to the static case. Let w_{1i} be the agent's payment if the principal buys and the additional signal realization is σ_i. As before, w_0 is the payment if the principal does not buy. The next result shows how incentives are structured.

Claim 3 (Additional Signal) *Suppose the principal buys only after observing a high signal. Then there is a threshold $\tilde{\gamma}$ such that*

(i) If $\gamma > \tilde{\gamma}$, then $w_{1h} > w_0 = w_{1\ell} = 0$ at the optimal contract.

(ii) If $\gamma < \tilde{\gamma}$, then $w_0 > w_{1h} = w_{1\ell} = 0$ at the optimal contract.

The Appendix derives $\tilde{\gamma}$, which depends on α, β_h/β_ℓ, and ρ_h/ρ_ℓ. To convey the intuition of the claim in the simplest way, let $\beta_h = \beta_\ell$ and $\alpha = 1/2$. Then $\tilde{\gamma} = 1/(1 + \sqrt{\rho_h/\rho_\ell})$.

Notice that the principal does not compensate the agent when she exercises the option but her signal turns out to be low. Whether the agent is compensated for buying followed by success or for not buying depends on the prior. As long as $\rho_\ell > 0$, there is an interval of γ where $w_0 > w_{1h} = 0$ and thus the agent is rewarded for bad news. For this to be a feature of the optimal contract, it must be the case that the principal finds it optimal to implement a positive level of effort for those γ. Recall that she implements zero effort near both $\gamma = 0$ and $\gamma = 1$. But if ρ_ℓ is sufficiently high — i.e., if the principal's signal generates a false

---

[21] A special case is observing option quality without noise, i.e., when $\rho_h = 1$ and $\rho_\ell = 0$.

positive with a sufficiently high probability —, then $1/(1+\sqrt{\rho_h/\rho_\ell})$ exceeds the threshold of γ below which $e = 0$, and hence rewarding for bad news emerges in equilibrium.

When the principal can observe an additional signal, rewarding for bad news is less likely to obtain. To see this, notice that $\tilde{\gamma} \leq \hat{\gamma} = 1/2$, with strict inequality so long as the additional signal is informative, i.e., $\rho_\ell < \rho_h$. The more informative the principal's signal is, the lower $\tilde{\gamma}$ is, and it becomes zero if the signal is perfectly informative.

In our model, no effort can be induced if the agent can misreport his signal. The reason is that the agent's payoff depends only on his report. This is not the case in this extension, as the agent's payoff also depends on the additional (contractible) signal, which in turn is correlated with the option quality. It is easy to show that if the agent intends to misreport, he will choose zero effort. Thus, there are two relevant truth-telling constraints: the agent must prefer to exert effort and report truthfully to exerting no effort and always reporting a high (low) signal. Intuitively, one should expect the first (second) constraint to bind when $\gamma > \tilde{\gamma}$ ($\gamma < \tilde{\gamma}$). Moreover, w_{1h} and w_0 must be positive in both cases.

As mentioned, Gromb and Martimort (2007) analyze a model which is similar to this extension. Compared to their setup, continuous effort allows us to clearly separate effects of moral hazard (unobservable effort) and adverse selection (possibility of misreporting a signal). The technical reason is that we have the incentive constraint which is separate from the truth-telling constraints. In their model, both truth-telling constraints bind in equilibrium, and the prior plays no essential role. In our model, the prior is crucial; the principal rewards the agent with w_{1h} (w_0) when $\gamma > \tilde{\gamma}$ ($\gamma < \tilde{\gamma}$) in order to provide him with incentives to exert effort, and pays w_0 (w_{1h}) in order to eliminate incentives to misreport.

Instead of observing a signal after exercising an option, one can envision economic situations where an option carries with it a contractible characteristic, say, its price, that could also be informative about its unknown quality: e.g., the price of a house in the real-estate agent example. Without going into tedious analytical details, it is easy to assess how the principal's contracting problem changes with this additional signal. Consider for simplicity the static case.[22] First, the price times the probability of buying is subtracted from the principal's objective. Second, the parameter γ is replaced with the posterior belief about the option quality given the observed price. Hence the price affects the optimal contract both directly as well as through the updating of the prior. The rest is as summarized in Proposition 1, except that the thresholds $\underline{\gamma}$ and $\bar{\gamma}$ will be functions of the price observed.

---

[22] We have already discussed in Section 5 a dynamic implication of adding prices.

6.3 Non-Contractible Option or Information Arrival

We have assumed that whether or not an option arrives is a contractible event. Suppose instead that the buying decision is contractible but option arrival is not. Since the principal cannot exercise an option if one is not available, the compensation scheme does not distinguish between no option arrival and no purchase. That is, the promised values under no option arrival and no purchase are equal, i.e., $V_n = V_0$.

Interestingly, this has *no effect* on the optimal contract when $\gamma > \hat{\gamma}$: as Proposition 2 shows, the principal sets $V_n = V_0$ even when option arrival is contractible. The situation, however, is drastically different when $\gamma < \hat{\gamma}$. Indeed, setting $V_n = V_0$ destroys *all* incentive provision. To see this, notice that $V_n = V_0$ implies that the agent prefers to hide the presence of a high signal realization and claim that no option arrived. Then $\gamma < \hat{\gamma}$ and the no-hiding constraint $w_1 \geq \delta V_n$ imply that the right-hand side of the incentive constraint, $\psi'(e) = \eta'(e)(\gamma\beta_h - (1-\gamma)\beta_\ell)(w_1 - \delta V_n)$, is nonpositive. Therefore, only $e = 0$ is implementable. We have thus proved the following claim:

Claim 4 (Non-Contractible Arrival) *Non-contractible option arrival has no effect on the optimal contract when $\gamma > \hat{\gamma}$ but precludes any incentive provision when $\gamma < \hat{\gamma}$.*

A key feature underlying this result is that if no option arrives then the principal's buying decision is trivial. It is interesting to contrast this case with a variation of our model in which p is instead interpreted as the probability that the agent observes an informative signal after exerting effort (with probability $1-p$ his effort does not generate information). That is, an option arrives in every period but, unlike in our model, effort only generates information with probability p. Suppose that whether a signal is generated is known only to the agent, and that he can hide its presence. Also, assume that, while the purchase decision is contractible, information arrival is not.

Notice that, unlike in the case with non-contractible option arrival, now the principal has a nontrivial purchase decision to make even if the agent claims that there is no information available. This difference overturns the aforementioned no-incentive-provision result when $\gamma < \hat{\gamma}$, but it requires the principal to *commit* to a buying decision with no information that might be ex-post inefficient. To see this, suppose there is a value of $\gamma \in (0, \hat{\gamma})$ at which the principal prefers not to buy without information. Then the agent has incentives to hide a high signal realization and claim that no information arrived, since he gets rewarded in the event of no purchase. But this incentive is mitigated if the principal commits to buy without information when $\gamma < \hat{\gamma}$. A similar analysis holds for $\gamma > \hat{\gamma}$: in this case the principal needs to commit not to buy when she has no information.

In short, if commitment of this sort is possible, then the effects of non-contractibility of information arrival are not as severe as those of option arrival.

7 Concluding Remarks

This paper studies a contracting problem where a principal searches for an opportunity of uncertain return. She hires an agent to acquire information about potential options that arise over time. How precise the acquired information is depends on the agent's effort, which is unobservable. Based on the information provided by the agent, the principal decides whether to stop and exercise the option, or to continue the search. We further assume that the agent can hide the presence of an option.

We provide a complete analysis of the optimal contract both in the static and dynamic cases. Among the main features that emerge in this setting, we have highlighted two that are robust across all the variations explored in the paper. First, effort provision and the optimal buying decision are *distorted* compared to the first best. This is due to the cost of providing incentives for effort. Second, there is a region of parameters where the optimal contract calls for *rewarding for bad news*, i.e., the agent receives a higher compensation in the event in which the information revealed induces the principal to pass on the option at hand. This property can be traced back to the role of the agent's effort in affecting the *informativeness* of the signal about option quality, and has no counterpart in standard moral-hazard models where an agent's effort affects the distribution of output.

We built a simple model that allowed us to analyze what we consider to be the main trade-offs present in environments with delegated information acquisition. Although we explored several variations and extensions of the model in Section 6, there are still many interesting avenues for future research. An important extension worth exploring is to relax the assumption that the fraction of high-quality options is known. Then both the principal and the agent update their beliefs over time upon observing the history of signal realizations. Notice that the principal might now implement positive effort solely with the purpose of learning, even if she does not intend to buy in the current period. A delicate issue arises, however: if the agent deviates and chooses a different effort level than the one the principal wants to implement, then the beliefs of the principal and the agent will diverge, as the agent knows the level of effort he chose while the principal does not. This is a well-known problem that substantially complicates the analysis of contracting situations (see Bergemann and Helge, 2005, Horner and Samuelson, 2010, and DeMarzo and Sannikov, 2011 for contracting problems with learning).

Another interesting extension is to allow the principal to recall options that she passed on. If an option is available in every period, then recall will never have any bite. To see this, notice that the principal passes on an option if the signal realization is low, which lowers her belief about the quality of the option. Thus, all options that are passed on have a lower probability of being of high quality than a new untried option. The situation is much different if an option arrives with probability less than one, for then she might recall passed options in periods in which no option arrives. Now, each passed-on option has attached to it the posterior belief that it is of high quality given a low signal and the level of effort implemented in that period (which could be zero). Whenever recall is used, the principal picks the option with the highest posterior. But this means that the state space needed to analyze this problem recursively becomes much larger, which complicates the analysis. More importantly, a feature similar to the one described in the learning extension arises: when choosing effort, the agent takes into account that an option discarded today can be recalled tomorrow based on its posterior belief. Thus, if he chooses a different effort level than the one the principal wants to implement, the problem of diverging beliefs emerges.

Since both learning and recall are potentially important in some of the economic applications mentioned, they constitute interesting open problems to tackle in future research.

A Appendix: Omitted Proofs

Define $u(e) \equiv \pi(e)E[y|\theta_h, e] = \gamma(\alpha + \beta_h\eta(e))y_h + (1-\gamma)(\alpha - \beta_\ell\eta(e))y_\ell$.

A.1 Remarks on the Class of Information Structures Used

As mentioned, the assumed class of information structures affords a nice characterization of the properties of the optimal contract in terms of the value of the prior belief γ. Indeed, it is the largest class that allows for such a parametrization. To see this, assume that $\Pr\{\theta = \theta_h | q = h, e\} = f_h(e)$ and $\Pr\{\theta = \theta_\ell | q = \ell, e\} = f_\ell(e)$, where $f_h(e)$ and $f_\ell(e)$ are between 0 and 1 for all e, C^2, increasing, and concave with $f_h(0) = 1 - f_\ell(0)$ and $f_h'(e) + f_\ell'(e) > 0$ for all e. Then $\pi(e) = \gamma f_h(e) + (1-\gamma)(1 - f_\ell(e))$, and its derivative is $\pi'(e) = \gamma f_h'(e) - (1-\gamma)f_\ell'(e) = (f_h'(e) + f_\ell'(e))[\gamma - f_\ell'(e)/(f_h'(e) + f_\ell'(e))]$. Define $\hat{\gamma}(e) \equiv$

$f_\ell'(e)/(f_h'(e) + f_\ell'(e))$. Notice that the threshold $\hat{\gamma}(e)$ is the same for all e if and only if $f_h'(e) = cf_\ell'(e)$, where $c \geq 0$. That is, f_h and f_ℓ are affine transformations of each other, and without loss of generality we can write $f_h(e) = \alpha + \beta_h\eta(e)$ and $f_\ell(e) = \xi + \beta_\ell\eta(e)$, with α, ξ, β_h, and β_ℓ nonnegative, and η increasing and concave with $\eta(0) = 0$. Since $f_h(0) = 1 - f_\ell(0)$, $\xi = 1 - \alpha$, which delivers the class used in the paper.

A.2 Static First Best: Proof of Lemma 1

Assume first that $\alpha \in (0, 1)$. The principal solves $\max\{0, \gamma y_h + (1-\gamma)y_\ell, \max_e u(e) - \psi(e)\}$. Notice that $\gamma y_h + (1 - \gamma)y_\ell \geq 0$ if and only if $\gamma \geq -y_\ell/(y_h - y_\ell)$. At $\gamma = -y_\ell/(y_h - y_\ell)$ the principal is indifferent between exerting no effort and not exercising the option, and exerting no effort and exercising it. But at that value of γ she strictly prefers to exert positive effort and exercise the option if the signal realization is high. To see this, notice that at $\gamma = -y_\ell/(y_h - y_\ell)$ the first-order condition becomes $\psi'(e) = -\eta'(e)(\beta_h + \beta_\ell)y_\ell y_h/(y_h - y_\ell)$, which has a unique solution $e > 0$. Since the principal can always obtain zero profits by choosing $e = 0$, and her problem is strictly concave in effort, it must be the case that $u(e) - \psi(e) > 0$ at the optimal level of effort when $\gamma = -y_\ell/(y_h - y_\ell)$.

At $\gamma = 1$ ($\gamma = 0$) the principal's payoff if she exerts no effort and buys (does not buy) is y_h (zero), which is strictly higher than the other two alternatives. Hence, at $\gamma = 0$ the principal does not exert effort and does not buy; at $\gamma = -y_\ell/(y_h - y_\ell)$ she exerts a positive level of effort and buys if the signal is high; and at $\gamma = 1$ she exerts no effort and buys.

We now show that $\max_e u(e) - \psi(e)$ is strictly increasing and convex in γ. By the Envelope Theorem, $\partial(u(e) - \psi(e))/\partial\gamma = (\alpha + \beta_h\eta(e))y_h - (\alpha - \beta_\ell\eta(e))y_\ell > 0$. Thus, $u(e) - \psi(e)$ is strictly increasing in γ. Differentiating again yields $\partial^2(u(e) - \psi(e))/\partial\gamma^2 = \eta'(e)(\beta_h y_h + \beta_\ell y_\ell)\partial e/\partial\gamma$. Differentiating the first-order condition $\eta'(e)(\gamma\beta_h y_h - (1-\gamma)\beta_\ell y_\ell) = \psi'(e)$ with respect to γ yields $\partial e/\partial\gamma = -\eta'(e)(\beta_h y_h + \beta_\ell y_\ell)/[\eta''(e)\psi'(e)/\eta'(e) - \psi''(e)]$. Hence, $\partial^2(u(e) - \psi(e))/\partial\gamma^2 = -\eta'(e)^2(\beta_h y_h + \beta_\ell y_\ell)^2/[\eta''(e)\psi'(e)/\eta'(e) - \psi''(e)] \geq 0$, and $u(e) - \psi(e)$ is convex in γ. Thus, it must single-cross zero at $\underline{\gamma}^* \in (0, -y_\ell/(y_h - y_\ell))$ and $\gamma y_h + (1-\gamma)y_\ell$ at $\bar{\gamma}^* \in (-y_\ell/(y_h - y_\ell), 1)$, which proves the interval structure and effort/buying decisions stated in parts (*i*)–(*iii*) of the lemma when $\alpha \in (0, 1)$.

If $\alpha = 0$, then $\beta_\ell = 0$ and hence $u(e) - \psi(e) = \gamma\beta_h\eta(e) - \psi(e)$. It is easy to verify that $u(0) - \psi(0) = 0$ and $\max_e u(e) - \psi(e) \geq 0$, with equality only at $\gamma = 0$. Hence, for all $\gamma > 0$, exerting positive effort and exercising the option after a high signal dominates exerting zero effort and never exercising it. Thus, $\underline{\gamma}^* = 0$ and $\bar{\gamma}^* < 1$ in this case. A similar argument reveals that if $\alpha = 1$ (and thus $\beta_h = 0$), then $\underline{\gamma}^* > 0$ and $\bar{\gamma}^* = 1$. $\square$

A.3 Optimal Static Contract: Proof of Proposition 1

Let λ and μ be the Lagrange multipliers associated with constraints (2) and (3). The first-order conditions with respect to w_1, w_0, and e are

$$\lambda - 1 + \frac{\pi'(e)}{\pi(e)}\mu \leq 0, \qquad w_1 \geq 0, \tag{12}$$

$$\lambda - 1 - \frac{\pi'(e)}{1-\pi(e)}\mu \leq 0, \qquad w_0 \geq 0, \tag{13}$$

$$u'(e) - \psi'(e) + \mu\left(\pi''(e)\frac{\psi'(e)}{\pi'(e)} - \psi''(e)\right) \leq 0, \qquad e \geq 0, \tag{14}$$

all with complementary slackness, where we substituted constraint (3) into (14).

The following results will be used in the proof of the proposition:

Lemma 3 *At the optimum, $w_0 = 0$ if $\gamma > \hat{\gamma}$, and $w_1 = 0$ if $\gamma < \hat{\gamma}$.*

Proof. We only prove the first case (the other case is analogous). If $\gamma > \hat{\gamma}$, then $w_1 \geq w_0$, and $w_0 \geq 0$ is the only relevant limited-liability constraint. Suppose $w_0 > 0$. Then $w_1 > 0$, and (12)–(13) are binding, implying $\mu = 0$ and $\lambda = 1$. Thus (2)–(3) yield $w_0 = \psi(e) - \pi(e)\psi'(e)/\pi'(e)$, where the right-hand side equals zero at $e = 0$ and is strictly decreasing in e. Hence $w_0 \leq 0$, a contradiction. $\square$

Lemma 4 *At the optimum, constraint (2) does not bind unless $e = 0$.*

Proof. Let $\gamma > \hat{\gamma}$ (the other case is analogous), and suppose that (2) binds and $e > 0$. Then (3) and Lemma 3 yield $w_1 = \psi'(e)/\pi'(e)$, and (2) becomes $-\psi(e) + \pi(e)\psi'(e)/\pi'(e) = 0$. But the left-hand side is positive if $e > 0$, a contradiction. $\square$

Lemma 5 *If $\alpha = 0$ ($\alpha = 1$), then the principal prefers to hire the agent and buy after a high signal to not hiring him and not buying (buying), and strictly so if $\gamma > 0$ ($\gamma < 1$).*

Proof. We only prove the case of $\alpha = 0$ (the other case is analogous), which implies that $\beta_\ell = 0$ and thus $\hat{\gamma} = 0$. If the principal hires the agent, then Lemma 3 yields $w_1 = \psi'(e)/\pi'(e)$ and $w_0 = 0$, and the problem becomes $\max_e \gamma\beta_h\eta(e) - \pi(e)\psi'(e)/\pi'(e)$.

The objective function equals zero at $e = 0$, and it is easy to verify that the (unique) solution to this problem is positive if and only if $\gamma > 0$. Thus, positive effort and buying after a high signal yields higher profits than not hiring the agent and not buying. $\square$

Lemma 6 *As γ approaches $\hat{\gamma}$, the effort implemented by the principal when she buys only after the high signal realization goes to zero.*

Proof. If $\alpha \in \{0, 1\}$, then the result is an easy implication of Lemma 5. Assume $\alpha \in (0, 1)$, and let $\gamma < \hat{\gamma}$ (the other case is analogous). Using the first-order conditions, we obtain $\mu = -(1 - \pi(e))/\pi'(e)$, which strictly increases in γ. Rewrite (14) as follows: $\gamma\beta_h y_h - (1-\gamma)\beta_\ell y_\ell \leq \psi'(e)/\eta'(e) - [(1-\pi(e))/\pi'(e)]d(\psi'(e)/\eta'(e))/de$. The right-hand side is continuous, increasing in γ, and goes to infinity as γ goes to $\hat{\gamma}$, while the left-hand side approaches $(y_h - y_\ell)\beta_\ell\beta_h/(\beta_\ell + \beta_h)$. Hence, effort goes to zero as γ goes to $\hat{\gamma}$. $\square$

We are now ready to prove the proposition:

Proof of Proposition 1. *(i)* Let $\gamma > \hat{\gamma}$ (the other case is analogous). Since $\pi'(e) > 0$, constraint (3) and Lemma 3 yield $w_1 \geq w_0 = 0$, with strict inequality if $e > 0$.

(ii) If $e = 0$, then it is trivially lower than e^* and strictly so if $e^* > 0$. Suppose $e > 0$ and let $\gamma > \hat{\gamma}$ (the other case is analogous). From Lemma 4, $\lambda = 0$ and hence $\mu = \pi(e)/\pi'(e) > 0$. Then (14) becomes $u'(e) - \psi'(e) = -(\pi(e)/\pi'(e))(\pi''(e)\psi'(e)/\pi'(e) - \psi''(e)) > 0$. Since e^* solves $u'(e^*) - \psi'(e^*) = 0$, it follows that $e < e^*$.

(iii)–*(a)* If $\alpha = 0$, then Lemma 5 implies that $\underline{\gamma} = 0$. Assume $\alpha > 0$. At $\gamma = 0$ it is optimal for the principal not to hire the agent and not to exercise the option (the other two alternatives yield a negative payoff as $y_\ell < 0$ and $\alpha > 0$). By continuity, this holds in a neighborhood of $\gamma = 0$. Let $\underline{\gamma}$ be the largest value such that the principal does not hire the agent and does not exercise the option for all $\gamma \in [0, \underline{\gamma}]$. To show that $\underline{\gamma} > \underline{\gamma}^*$, notice that at $\underline{\gamma}^*$, $u(e^*) - \psi(e^*) = 0$. Since $0 < e < e^*$, $u(e) - \psi(e) < 0$, thus proving that $\underline{\gamma} > \underline{\gamma}^*$.

(iii)–*(b)* If $\alpha = 1$, then Lemma 5 implies that $\bar{\gamma} = 1$. Assume $\alpha < 1$. Then at $\gamma = 1$ it is optimal for the principal not to hire the agent and to exercise the option (it yields y_h, strictly higher than the payoff of the other two alternatives as $\alpha < 1$). By an analogous argument as in *(iii)*–*(a)* it follows that $\bar{\gamma} < \bar{\gamma}^*$.

(iii)–*(c)* Suppose $\gamma = \hat{\gamma}$. Since the cost of implementing any positive level of effort is infinite, $e = 0$. Hence, if the principal exercises the option only if the signal realization is high, then her payoff is $\alpha(\hat{\gamma} y_h + (1 - \hat{\gamma}) y_\ell)$. Notice, however, that in this case the principal

prefers (strictly unless $\alpha \in \{0,1\}$ or $\hat{\gamma} = -y_\ell/(y_h - y_\ell)$) not to hire the agent, and to exercise the option if and only if $\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell \geq 0$.

We now show that the same holds in a neighborhood of $\hat{\gamma}$. If $\alpha \in \{0,1\}$, then $\hat{\gamma} \in \{0,1\}$. It follows from Lemma 5 that there is a right or left neighborhood of $\hat{\gamma}$ such that the principal strictly prefers not to hire the agent. Assume $\alpha \in (0,1)$, and let $\gamma < \hat{\gamma}$ (the other case is analogous). The principal's problem in this case can be written as $v(\gamma) \equiv \max_{e\in[0,e^*]} u(e) - (1-\pi(e))\psi'(e)/[\eta'(e)(\gamma\beta_h - (1-\gamma)\beta_\ell)]$, where by part (*ii*) we restrict attention to $e \leq e^*$. Take $\varepsilon > 0$ small, and consider $\gamma \in [0, \hat{\gamma} - \varepsilon]$. The Theorem of the Maximum implies that v is continuous on $[0, \hat{\gamma}-\varepsilon]$. We will use this fact to show that there is a left neighborhood of $\hat{\gamma}$ such that the principal chooses not to hire the agent.

There are two cases to consider, $\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell < 0$ and $\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell > 0$. If $\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell < 0$, then the principal's profits at $\gamma = \hat{\gamma}$ if she only buys after a high signal are $\alpha(\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell) < 0$. From Lemma 6, the optimal effort level under this policy goes to zero as γ goes to $\hat{\gamma}$. Thus, if ε is small enough, then $u(e)$ evaluated at the optimal effort level will be negative, and hence $v(\gamma) < 0$. Thus, not hiring the agent and not buying dominates hiring him and buying only if the signal is high. Hence, there is a left neighborhood of $\hat{\gamma}$ such that the principal prefers not to hire the agent. If $\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell > 0$, then the principal's profits at $\gamma = \hat{\gamma}$ if she buys after a high signal are $\alpha(\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell) > 0$. But as $\alpha(\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell) < \hat{\gamma}y_h + (1-\hat{\gamma})y_\ell$, she strictly prefers to buy without hiring the agent. Proceeding as when $\hat{\gamma}y_h + (1-\hat{\gamma})y_\ell < 0$, one can show that if ε is small enough then $v(\gamma) < \gamma y_h + (1-\gamma)y_\ell$. Thus, there is a left neighborhood of $\hat{\gamma}$ where the principal strictly prefers to buy without hiring the agent.

(*iii*)–(*d*) Suppose $\hat{\gamma} \neq -y_\ell/(y_h - y_\ell)$. At $\gamma = -y_\ell/(y_h - y_\ell)$ the principal can obtain zero expected profits if she buys only after the high signal realization and implements $e = 0$, for it yields $\alpha(\gamma y_h + (1-\gamma)y_\ell) = 0$. Also, the proof of Lemma 6 shows that if $d(\psi'/\eta')/de = 0$ at $e = 0$, then the principal implements a positive level of effort for all $\gamma \neq \hat{\gamma}$ when she buys only after the high signal realization. By revealed preference, she achieves a positive level of expected profits by doing so. Thus, the optimal effort level at $\gamma = -y_\ell/(y_h - y_\ell)$ is positive. By continuity, it is positive for values of γ in a neighborhood of $-y_\ell/(y_h - y_\ell)$. $\square$

A.4 Dynamic First Best: Proof of Lemma 2

We only sketch the proof since it is very similar to its static counterpart. The proofs of (*i*)–(*iii*) are analogous to those of Lemma 1. The first implication in (*iv*) follows from the first-order condition in the text, as $(\gamma\beta_h - (1-\gamma)\beta_\ell)\delta S^*$ is positive, negative, or zero, respectively, if γ is greater, less, or equal to $\hat{\gamma}$. Finally, the proof of the threshold comparison in *(iv)* is analogous to those of parts *(iii)*–*(a)* and *(iii)*–*(b)* of Proposition 1. ☐

A.5 Optimal Dynamic Contract: Proof of Proposition 2

Assuming that the principal can use lotteries, her problem when she buys only after a high signal can be rewritten as

$$\begin{aligned}
U_{10}(V) = \max_{e_i, w_{1i}, w_{0i}, w_{ni}, V_{0i}, V_{ni}, s_i \in [0,1], i=1,2, s_1+s_2=1} & \sum_{i=1}^{2} s_i \{ p[\pi(e_i)(E[y|\theta_h, e_i] - w_{1i}) \\
& + (1-\pi(e_i))(-w_{0i} + \delta U(V_{0i}))] + (1-p)(-w_{ni} + \delta U(V_{ni}))\} \\
\text{s.t.} \quad & \sum_{i=1}^{2} s_i \{p[-\psi(e_i) + \pi(e_i) w_{1i} + (1-\pi(e_i))(w_{0i} + \delta V_{0i})] + (1-p)(w_{ni} + \delta V_{ni})\} = V, \\
& \psi'(e_i) = \eta'(e_i)(\gamma\beta_h - (1-\gamma)\beta_\ell)(w_{1i} - w_{0i} - \delta V_{0i}), \ \ i = 1,2, \\
& w_{1i} \geq w_{ni} + \delta V_{ni}, \ \ w_{0i} + \delta V_{0i} \geq w_{ni} + \delta V_{ni}, \ \ i = 1,2, \\
& e_i \geq 0, \ \ w_{1i} \geq 0, \ \ w_{0i} \geq 0, \ \ w_{ni} \geq 0, \ \ V_{0i} \geq 0, \ \ V_{ni} \geq 0, \ \ i = 1,2.
\end{aligned}$$

Due to the use of lotteries, the value function U_{10} is concave. Also, if V is such that $U(V) = U_{10}(V)$, then the following 'envelope condition' holds:[23]

$$-\lambda \in \partial U(V), \tag{15}$$

where λ is the Lagrange multiplier on the promise-keeping constraint in the above problem.

We can proceed with the analysis by taking the first-order conditions with respect to $e_i, w_{1i}, w_{0i}, w_{ni}, V_{0i}, V_{ni}, s_i$, $i = 1,2$. To avoid clutter, we will treat the problem as if the principal uses a deterministic contract, as formulated in subsection 4.2.1: i.e., with $s \in \{0,1\}$. It is straightforward to verify that most statements hold for each realization of a random contract if the principal finds it optimal to use one.[24]

Let μ, ξ_1, ξ_0, and ν_i be the Lagrange multipliers on constraints (8), (9), (10), and $V_i \geq 0$, $i = 0, n$, respectively. Denote by $U'_+(V)$, $U'_-(V)$, and $\partial U(V)$ the right derivative, left derivative, and superdifferential, respectively, of the concave function U at V.

---

[23] See, for instance, Aubin (1979), Chapter 5, pp. 130-140.

[24] The only exception is that $V_0 > V$ when $\gamma < \hat{\gamma}$ might hold only in expectation, i.e., $s_1 V_{01} + s_2 V_{02} > V$.

The first-order conditions with respect to w_1, w_0, w_n, and e are:

$$p\pi(e)(\lambda-1)+\mu\pi'(e)+\xi_1 \le 0, \qquad w_1 \ge 0, \tag{16}$$
$$p(1-\pi(e))(\lambda-1)-\mu\pi'(e)+\xi_0 \le 0, \qquad w_0 \ge 0, \tag{17}$$
$$(1-p)(\lambda-1)-\xi_1-\xi_0 \le 0, \qquad w_n \ge 0, \tag{18}$$
$$p(u'(e)-\pi'(e)(w_1-w_0+\delta U(V_0)))+\mu\left(\pi''(e)\frac{\psi'(e)}{\pi'(e)}-\psi''(e)\right) \le 0, \qquad e \ge 0, \tag{19}$$

all with complementary slackness. Using standard rules of superdifferential calculus,[25] we obtain the following first-order condition with respect to V_0:

$$0 \in p(1-\pi(e))(\partial U(V_0)+\lambda)-\mu\pi'(e)+\xi_0+\nu_0. \tag{20}$$

Similarly, the first-order condition with respect to V_n is

$$0 \in (1-p)(\partial U(V_n)+\lambda)-\xi_1-\xi_0+\nu_n. \tag{21}$$

We invoke the following results below.

Lemma 7 *For all $V \ge 0$ and all $x \in \partial U(V)$, $x \ge -1$.*[26]

Proof. From (16) and (17), $1-\lambda \ge (\xi_1+\xi_0)/p \ge 0$. If U is differentiable at V so that $\partial U(V)=\{U'(V)\}$, then from (15), $U'(V)=-\lambda \ge -1$, with strict inequality unless both (9) and (10) are slack. Since U is concave, it is differentiable everywhere but on a countable set. It follows that $U'_+(V) \ge -1$ for all values of V (which in turns implies the statement of the lemma since $\partial U(V)=[U'_+(V),U'_-(V)]$ for any $V>0$, and $\partial U(0)=\{U'_+(0)\}$ if $U'_+(0)$ is finite). To see this, suppose that $U'_+(V')<-1$ for some V'. Take $V''>V'$ such that $U'(V'')$ exists. By the previous argument, $U'_+(V'')=U'(V'')\ge -1$, which contradicts concavity of U, since a concave function has a decreasing right derivative. □

Lemma 8 *For all $V>0$, any effort level that satisfies (19) is positive.*

Proof. We will show that $e=0$ cannot satisfy the first-order condition (19). Adding and subtracting δV_0 from $w_1-w_0+\delta U(V_0)$, using the incentive constraint (8), and noticing that the last term equals $-\mu\eta'(e)(\psi'/\eta')'(e)$, allow us to write (19) evaluated at $e=0$ as

---

[25] See, for instance, Borwein and Lewis (2006), Chapter 6, p. 134.

[26] That is, the slope of the value function is always greater than minus one.

$p[u'(0)-\psi'(0)-\pi'(0)\delta(V_0+U(V_0)))]-\mu\eta'(0)(\psi'/\eta')'(0) \le 0$. Since $(\psi'/\eta')(0) = (\psi'/\eta')'(0) = 0$, after some algebra this inequality becomes

$$\gamma\beta_h y_h - (1-\gamma)\beta_\ell y_\ell - (\gamma\beta_h - (1-\gamma)\beta_\ell)\delta(V_0 + U(V_0)) \le 0. \tag{22}$$

Notice that $0 \le V_0 + U(V_0) \le S^*$, where the first inequality follows from (6) (i.e., $U(V) \ge -V$ for all V) and the second from the fact that the second-best surplus cannot exceed the first-best surplus. It is then immediate that (22) is violated if $\gamma < \hat{\gamma}$, as both terms are positive and the first one strictly so. And if $\gamma > \hat{\gamma}$, (22) is violated as $\gamma\beta_h y_h - (1-\gamma)\beta_\ell y_\ell > (\gamma\beta_h - (1-\gamma)\beta_\ell)S^* \ge (\gamma\beta_h - (1-\gamma)\beta_\ell)\delta(V_0 + U(V_0))$, where the first inequality follows from the first-order condition in the first best, where effort is strictly positive when the principal buys only after a high signal. Thus, $e = 0$ cannot be a solution. □

We now turn to the proof of the proposition:

Proof of Proposition 2. *(i)*–*(a)* The incentive constraint (8) and Lemma 8 imply that $w_1 > \delta V_0$. To show that $V_0 = V_n$, suppose to the contrary that $V_0 > V_n$. Then the principal can decrease w_1 and δV_0 by the same amount so that the implemented effort does not change, and at the same time increase δV_n to keep (7) satisfied, that is, $0 > dw_1 = \delta dV_0 = -\delta dV_n(1-p)/p$. Such a variation increases the principal's payoff, violating optimality of the original contract. To see this, take any $x \in \partial U(V_0)$ and $y \in \partial U(V_n)$. The principal's payoff changes by $dS = -p\pi dw_1 + p(1-\pi)x\delta dV_0 + (1-p)y\delta dV_n = pdw_1[\pi(-1-x)+x-y]$. By Lemma 7, $x \ge -1$. Also, $V_0 > V_n$ and concavity of U imply $x \le y$. Hence $dS \ge 0$, with strict inequality unless $\partial U(V_0) = \partial U(V_n) = \{-1\}$. And if $\partial U(V_n) = \{-1\}$, (21)–(20) and $\xi_0 = \xi_1 = 0$ imply $\partial U(V) = \{-1\}$. Thus $dS > 0$ as long as $V < V^*$.

Notice that $\mu \ge 0$. Indeed, $\xi_1 = 0$ and (16) imply $\mu\pi' = p\pi(1-\lambda)$. From (15) and Lemma 7, $\lambda \le 1$. Since $\pi' > 0$, we have $\mu \ge 0$. Moreover, $\mu > 0$ if $V < V^*$ so that $\lambda < 1$.

Next, we prove that $V_0 < V$. Since $V > 0$, if $V_0 = 0$ we are done, so assume $V_0 > 0$ so that $\nu_0 = 0$. Since $V < V^*$, $\mu > 0$. Suppose that $V_0 \ge V$ so that $U'_-(V_0) \le U'_+(V)$. Using $\xi_1 = 0$, from (15) and (20) we have $-\lambda \in [\partial U(V_0) - \mu\pi'/(p(1-\pi))] \cap \partial U(V)$, which is an empty set, a contradiction. Thus $V_0 < V$. Since $V_0 = V_n$, $V_n < V$ follows.

(i)–*(b)* The proof of $w_1 = \delta V_n$ is analogous to the proof of $V_0 = V_n$ in part *(i)*. Moreover, the argument of the proof implies $\xi_1 > 0$. To implement $e > 0$, $\delta V_0 > w_1 = \delta V_n$, and hence (10) holds with strict inequality so that that $\xi_0 = 0$. We now prove that $V_n < V$. Proceeding in a similar way as in part (i), assuming $V_n \ge V$ yields $\nu_n = 0$ and $-\lambda \in [\partial U(V_n) - \xi_1/(1-p)] \cap \partial U(V) = \emptyset$, a contradiction. Thus $V_n < V$.

Next we show that $V_0 > V$. Suppose that $V_0 \leq V$. Since $w_1 < \delta V_0 < V$ and $\delta V_n \leq \delta V_0 < V$, the value delivered to the agent is $p(-\psi(e) + \pi(e)w_1 + (1 - \pi(e))\delta V_0) + (1 - p)\delta V_n < V$, violating (7), a contradiction. Thus $V_0 > V$.

Notice that $\mu \geq 0$. Indeed, $V_0 > V$ implies $V_0 > 0$ so that $\nu_0 = 0$, and also $U'_-(V_0) \leq U'_+(V)$. Suppose that $\mu < 0$. Then from (15) and (20), $-\lambda \in [\partial U(V_0) - \mu\pi'(e)/(p(1 - \pi(e))] \cap \partial U(V) = \emptyset$ since $\pi'(e) < 0$, a contradiction.

(ii) Suppose that $\gamma > \hat{\gamma}$ (the other case is analogous). From part *(i)*–*(a)*, $V_0 = V_n$. Substituting this into (7) and using (8) yield $w_1 = V + p\psi(e) + [p(1 - \pi(e)) + 1 - p]\psi'(e)/\pi'(e)$ and $\delta V_0 = V + p[\psi(e) - \pi(e)\psi'(e)/\pi'(e)]$. We can then substitute these equations into the objective function and solve the principal's problem for the optimal e, subject to $V_0 \geq 0$, i.e., $V + p[\psi(e) - \pi(e)\psi'(e)/\pi'(e)] \geq 0$. If $V = 0$, the only value of e that satisfies this constraint is $e = 0$. To see this, notice that at $e = 0$ the term in the square brackets vanishes, and its derivative with respect to e is negative (so this term is negative for all $e > 0$). Hence, if $V = 0$ then $e = 0$ and thus $w_1 = V_0 = V_n = 0$.

(iii) Let $V \geq V^* > 0$. The second inequality implies, by Lemma 8, that optimal effort is positive if the principal buys only after a high signal realization. Assume first that $V > V^*$. Then $U'(V)$ exists and equals -1. From (15), $\lambda = 1$ and hence the first-order condition (16) becomes $\mu\pi'(e) + \xi_1 \leq 0$. Let $\gamma > \hat{\gamma}$. As we have shown in the proof of Proposition 2, in this case $\mu \geq 0$ and $\xi_1 = 0$. Thus $\pi' > 0$ implies $\mu = \xi_1 = 0$. From (20), $0 = x + 1 + (\xi_0 + \nu_0)/(p(1 - \pi(e)))$ for some $x \in \partial U(V_0)$. But since $\xi_0 \geq 0$, $\nu_0 \geq 0$, and $x \geq -1$ for all $x \in \partial U(V_0)$ by Lemma 7, the above equality can only hold if $x = -1$ and $\xi_0 = \nu_0 = 0$. Now suppose that $\gamma < \hat{\gamma}$. As we have shown in the proof of Proposition 2, in this case $\mu \geq 0$ and $\xi_0 = 0$. From (20), $0 = x + 1 + (-\mu\pi'(e) + \nu_0)/(p(1 - \pi(e)))$ for some $x \in \partial U(V_0)$. If $\mu > 0$ or $\nu_0 > 0$, then from $-\mu\pi'(e) + \nu_0 > 0$ and Lemma 7 it follows that $x + 1 + (-\mu\pi'(e) + \nu_0)/(p(1 - \pi(e))) > 0$ for all $x \in \partial U(V_0)$, a contradiction. Hence $\mu = \nu_0 = 0$. Then from $\mu\pi'(e) + \xi_1 \leq 0$, $\xi_1 = 0$. Thus in both cases $\mu = \xi_1 = \xi_0 = \nu_0 = 0$. Also, (16)–(18) holding at equalities and (21) imply that the limited-liability constraints do not bind, and $\nu_n = 0$. Therefore Lagrange multipliers on constraints (8)–(11) are all zero, and the problem of the principal is effectively unconstrained (it is only subject to the promise-keeping constraint). Moreover, from (20) and (21) it follows that $U'(V_0) = U'(V_n) = -1$ as well, and hence the principal solves the unconstrained problem in all subsequent periods as well. Thus the first best is achieved.

Assume now that $V = V^*$. At that point $U(V^*)$ need not exist and λ is not necessarily equal to -1. But the argument above shows that there exists a solution to the first-order conditions that achieves the first best. Since for any V profits are higher in the first-best case, the constructed solution is optimal. $\square$

A.6 Optimal Dynamic Contract: Proof of Claim 1

Recall that the first-order condition with respect to effort in the first-best case is $f^* \equiv p(u'(e) - \psi'(e) + \pi'(e)\delta S^*) = 0$. The corresponding first-order condition in the second-best case, (19), can be written as $f \equiv pu'(e) - p\pi'(e)[w_1 + \delta U(V_0)] + \mu(\eta''(e)\psi'(e)/\eta'(e) - \psi''(e)) = 0$. Using (8), the second term can be rewritten as $p\pi'(e)[w_1 - \delta V_0 + \delta V_0 + \delta U(V_0)] = p\psi'(e) + p\delta S$, where $S \equiv V_0 + U(V_0)$. Then $f = f^* + \pi'(e)\delta(S - S^*) + \mu(\eta''(e)\psi'(e)/\eta'(e) - \psi''(e))$. Since $\eta''\psi'/\eta' - \psi'' < 0$, $\mu \geq 0$, $S^* \geq S$, and $\pi' < 0$ when $\gamma < \hat{\gamma}$, we have $f \leq f^*$, and hence $e \leq e^{**}$, with strict inequalities unless $\mu = 0$. $\square$

A.7 Optimal Dynamic Contract: Proof of Proposition 3

The proof relies on a fairly intuitive graphical argument. It involves several steps, and for expositional clarity we prove each step separately.

Optimal Buying Decision at $V = 0$. We first analyze the principal's optimal buying decision as a function of γ when the promised value to the agent is $V = 0$. Using part (*ii*) of Proposition 2, $U_{10}(0) = p\alpha[\gamma y_h + (1-\gamma)y_\ell] + (1-p\alpha)\delta U(0)$, $U_{11}(0) = p[\gamma y_h + (1-\gamma)y_\ell] + (1-p)\delta U(0)$, and $U_{00}(0) = \delta U(0)$. If $\gamma > -y_\ell/(y_h - y_\ell)$ so that $\gamma y_h + (1-\gamma)y_\ell > 0$, then it follows that $U(0) = U_{11}(0) > U_{10}(0) > U_{00}(0) = \delta U(0) > 0$. If $\gamma < -y_\ell/(y_h - y_\ell)$, then $U_{11}(0) < U_{10}(0) < U_{00}(0) = U(0) = \delta U(0) = 0$. Finally, if $\gamma = -y_\ell/(y_h - y_\ell)$, then $U(0) = U_{00}(0) = U_{10}(0) = U_{11}(0) = 0$.

Intersection of $U_{00}(V)$ and $U_{10}(V)$. Next, we show that U_{00} can intersect U_{10} at most once. It is enough to show that U_{00} cannot intersect U_{10} from below, for this will rule out multiple crossings. Towards a contradiction, suppose that at V', U_{00} crosses U_{10} from below. By (6), such a crossing would require that at V' the slope of U_{00} be bigger than the slope of the upper envelope, U. (And if the functions are not differentiable at V', the same holds for their right derivatives.) That is, $U'_{00+}(V') > U'_+(V')$. However, this cannot occur since by the definition of U_{00}, $U'_{00+}(V) = U'_+(V/\delta) \leq U'_+(V)$ for all V, where the inequality follows from concavity of U.

Intersection of $U_{11}(V)$ and $U_{10}(V)$. Since $U_{11}(V) = p[\gamma y_h + (1-\gamma)y_\ell - V] + (1-p)U_{00}(V)$, we have that $U'_{11+}(V) = -p + (1-p)U'_{00+}(V) \leq U'_+(V)$ for all V, where the inequality follows from $U'_{00+}(V) \leq U'_+(V)$ shown above and Lemma 7. Hence by an argument analogous to the one above, U_{11} can intersect U_{10} at most once.

INTERSECTION OF $U_{00}(V)$ AND $-V$. We show that $U_{00}(V) \geq -V$ for all V. Note that U_{00} is concave with slope ≥ -1. This follows from $U'_{00+}(V) = U'_+(V/\delta)$ for all V, concavity of U, and Lemma 7. Moreover, at $V = 0$, $U_{00}(0) \geq 0$. Thus $U_{00}(V) \geq -V$ for all V.

COMPLETION OF PART (i). We must show that if $\gamma y_h + (1-\gamma)y_\ell > 0$, then not buying independently of the signal realization is never optimal, and there is a threshold above which the principal buys only after a high signal, and uses lotteries below it.

We proved above that in this case $U_{00}(0) < U_{10}(0) < U_{11}(0) = U(0)$. Since U_{00} cannot intersect (or be tangent to) U_{10} from below, it follows that $U_{00}(V) < U_{10}(V)$ for all V. Thus not buying independently of the signal realization is never optimal.

If γ is close enough to one, then, as in the static case, acquiring information is never optimal, and $U(V) = U_{11}(V) = (\gamma y_h + (1-\gamma)y_\ell)p/(1-(1-p)\delta) - V > U_{10}(V)$ for all V. That is, buying independently of the signal realization is optimal for all V. In this case, $\tilde{V}_1 = +\infty$. Otherwise, U_{11} and U_{10} intersect exactly once, in which case there is a threshold $\tilde{V}_1 < +\infty$ such that $U(V) = U_{10}(V)$ for $V \geq \tilde{V}_1$.

To complete the proof, we show that the principal uses lotteries for $V \in (0, \tilde{V}_1)$. To see this, note that $U_{11}(V) < U(V)$ for all $V > 0$. This follows from $U'_{11+}(V) = -p + (1-p)U'_+(V/\delta) \leq -p + (1-p)U'_+(V) \leq U'_+(V)$, with strict inequality if $V < V^*$. And if $V^* = 0$, then $U = U_{11}$ is linear with slope -1, which is the case with $\tilde{V}_1 = +\infty$ described above.

COMPLETION OF PART (ii). We must to show that if $\gamma y_h + (1-\gamma)y_\ell < 0$, then buying independently of the signal realization is never optimal, and there is a threshold above which the principal buys only after a high signal, and it is optimal to use lotteries below it.

We proved above that in this case $U(0) = U_{00}(0) > U_{10}(0) > U_{11}(0)$. Hence, similar to the argument above, U_{10} and U_{11} do not intersect.

If γ is sufficiently close to zero, then acquiring information is never optimal, and $U(V) = U_{00}(V) = -V > U_{10}(V)$ for all V. That is, not buying independently of the signal is optimal for all V. In this case, $\tilde{V}_0 = +\infty$. Otherwise, U_{10} and U_{00} intersect exactly once, in which case there is a threshold $\tilde{V}_0 < +\infty$ such that $U(V) = U_{10}(V)$ for $V \geq \tilde{V}_0$, and lotteries are used just to the left of $\tilde{V}_0$.

To complete the proof, we show that it is optimal for the principal to use lotteries for $V \in (0, \tilde{V}_0)$. This is the case if and only if U is linear on $(0, \tilde{V}_0)$. Suppose not, and let $V' = \sup\{V | V < \tilde{V}_0, U'_+(V) > U'_-(\tilde{V}_0)\}$. (Since lotteries are used just to the left of $\tilde{V}_0$, $V' < \tilde{V}_0$.) Notice that $U(V) = U_{00}(V)$ for $V \leq V'$. Next, take $V'' = V' - \varepsilon$, $\varepsilon \in [0, V'(1-\delta))$ arbitrarily small such that $U'_+(V'') > U'_-(\tilde{V}_0)$. Then $U'_+(V'') = U'_{00+}(V'') = U'_+(V''/\delta) < U'_+(V'')$, because $V''/\delta > V'$ and for all $V > V'$, $U'_+(V) = U'_-(\tilde{V}_0) < U'_+(V'')$, a contradiction. Since $U(V)$ is linear on $(0, \tilde{V}_0)$, $U'_{00}(V) = U'(V/\delta) = U'(V)$ for $V \in (0, \delta\tilde{V}_0)$. Hence for $V \in [0, \delta\tilde{V}_0]$, $U(V) = U_{00}(V)$, and thus implementing $e = 0$ with probability one, not buying, and rolling over the agent's value is also an optimal strategy.

COMPLETION OF PART (*iii*). If $\gamma y_h + (1-\gamma)y_\ell = 0$, then $U(0) = U_{11}(0) = U_{00}(0) = U_{10}(0) = 0$. We claim that in this case $U(V) = U_{10}(V)$ for all V. By an argument similar to the one in part (*ii*), there exists a threshold $\tilde{V}$ such that $U_{10}(V) = U(V)$ for $V \geq \tilde{V}$, and the principal uses lotteries on $(0, \tilde{V})$, so that U is linear on that interval. Suppose that $U_{10}(V') < U(V')$ for some $V' < \tilde{V}$. Then concavity of U_{10} implies that $U_{10}(V) < U(V)$ for all $V < V'$, which contradicts $U_{10}(0) = U(0)$. Thus $\tilde{V} = 0$. □

A.8 Extensions: Proof of Claim 3

Define $u(e, \theta^*) = \int_{\theta^*}^{\bar{\theta}} [\gamma f(\theta|h, e)y_h + (1-\gamma)f(\theta|\ell, e)y_\ell]\, d\theta$. The principal's problem is:

$$\begin{aligned} \max_{e,w_0,w_1,\theta^*} \quad & u(e,\theta^*) - \pi(e,\theta^*)w_1 - (1-\pi(e,\theta^*))w_0 \\ \text{s.t.} \quad & -\psi(e) + \pi(e,\theta^*)w_1 + (1-\pi(e,\theta^*))w_0 \geq 0, \\ & \psi'(e) = \pi_e(e,\theta^*)(w_1 - w_0), \\ & e \geq 0, \;\; w_0 \geq 0, \;\; w_1 \geq 0, \;\; \bar{\theta} \geq \theta^* \geq \underline{\theta}. \end{aligned}$$

It is easy to show that the participation constraint is slack at the optimum.

Recall that $\pi(e, \theta^*) = (\bar{\theta} - \theta^*)/(\bar{\theta} - \underline{\theta}) + (1/2)[\gamma\beta_h - (1-\gamma)\beta_\ell]g(e)(\bar{\theta} - \theta^*)(\theta^* - \underline{\theta})$.

(*i*) If $\gamma > \hat{\gamma}$, then $\pi_e(e, \theta^*) > 0$ implying that $w_1 > w_0 = 0$ at the optimum. Using $w_1 = \psi'(e)/\pi_e(e, \theta^*)$, the first-order condition with respect to θ^* is $u_{\theta^*} = \psi'\left(\pi_e\pi_{\theta^*} - \pi\pi_{e\theta^*}\right)/\pi_e^2$. Straightforward algebra yields $\pi_e\pi_{\theta^*} - \pi\pi_{e\theta^*} = -(1/2)[\gamma\beta_h - (1-\gamma)\beta_\ell]g'(e)(\bar{\theta} - \theta^*)^2/(\bar{\theta} - \underline{\theta}) < 0$. Hence, θ^* is greater than the ex-post optimal threshold which solves $u_{\theta^*} = 0$.

(*ii*) If $\gamma < \hat{\gamma}$, then $\pi_e(e, \theta^*) < 0$ implying that $w_0 > w_1 = 0$ at the optimum. Using $w_0 = -\psi'(e)/\pi_e(e, \theta^*)$, the first-order condition with respect to θ^* becomes $u_{\theta^*} = \psi'\left(\pi_e\pi_{\theta^*} + (1-\pi)\pi_{e\theta^*}\right)/\pi_e^2$. Straightforward algebra yields $\pi_e\pi_{\theta^*} + (1-\pi)\pi_{e\theta^*} = (1/2)[(1-\gamma)\beta_\ell - \gamma\beta_h]g'(e)(\theta^* - \underline{\theta})^2/(\bar{\theta} - \underline{\theta}) > 0$. Hence, θ^* is distorted downwards. □

A.9 Extensions: Proof of Claim 4

Define $\pi_h(e) \equiv \Pr\{\theta = \theta_h, \sigma = \sigma_h | e\} = \gamma(\alpha + \beta_h \eta(e))\rho_h + (1-\gamma)(1-\alpha-\beta_\ell \eta(e))\rho_\ell$. The principal's problem can be written as

$$\begin{aligned}
\max_{e,w_0,w_{1h},w_{1\ell}} \quad & \pi(e)E[y|\theta_h, e] - \pi_h(e)w_{1h} - (\pi(e) - \pi_h(e))w_{1\ell} - (1-\pi(e))w_0 \\
\text{s.t.} \quad & -\psi(e) + \pi_h(e)w_{1h} + (\pi(e) - \pi_h(e))w_{1\ell} + (1-\pi(e))w_0 \geq 0, \\
& \psi'(e) = \pi_h'(e)(w_{1h} - w_{1\ell}) - \pi'(e)(w_0 - w_{1\ell}), \\
& e \geq 0, \;\; w_0 \geq 0, \;\; w_{1h} \geq 0, \;\; w_{1\ell} \geq 0.
\end{aligned}$$

Let μ denote the Lagrange multiplier on the incentive constraint (the participation constraint is slack). The first-order conditions with respect to w_{1h}, $w_{1\ell}$, and w_0 are

$$\begin{aligned}
-\pi_h(e) + \mu\pi_h'(e) \leq 0, & \qquad w_{1h} \geq 0, \\
-(\pi(e) - \pi_h(e)) + \mu(\pi'(e) - \pi_h'(e)) \leq 0, & \qquad w_{1\ell} \geq 0, \\
-(1-\pi(e)) - \mu\pi'(e) \leq 0, & \qquad w_0 \geq 0,
\end{aligned}$$

all with complementary slackness. Compare $\pi_h'/\pi_h = \eta'[\gamma\beta_h\rho_h - (1-\gamma)\beta_\ell\rho_\ell]/[\gamma(\alpha+\beta_h\eta)\rho_h + (1-\gamma)(\alpha-\beta_\ell\eta)\rho_\ell]$, $(\pi' - \pi_h')/(\pi - \pi_h) = \eta'[\gamma\beta_h(1-\rho_h) - (1-\gamma)\beta_\ell(1-\rho_\ell)]/[\gamma(\alpha+\beta_h\eta)(1-\rho_h) + (1-\gamma)(\alpha-\beta_\ell\eta)(1-\rho_\ell)]$, and $-\pi'/(1-\pi) = \eta'[(1-\gamma)\beta_\ell - \gamma\beta_h]/[\gamma(1-\alpha-\beta_h\eta) + (1-\gamma)(1-\alpha+\beta_\ell\eta)]$. It is easy to show that $\pi_h'/\pi_h > (\pi' - \pi_h')/(\pi - \pi_h)$ as $\rho_h > \rho_\ell$. Moreover, tedious algebra reveals that $\pi_h'/\pi_h \gtrless -\pi'/(1-\pi)$ if and only if $\gamma \gtrless \tilde{\gamma}$, where $\tilde{\gamma}$ is the positive solution of a quadratic equation, and it is given by

$$\tilde{\gamma} = \frac{-\left[\frac{\left(1+\frac{\beta_h\rho_h}{\beta_\ell\rho_\ell}\right)}{\alpha\left(1+\frac{\beta_h}{\beta_\ell}\right)\left(\frac{\rho_h}{\rho_\ell}-1\right)} - 1\right] + \sqrt{\left[\frac{\left(1+\frac{\beta_h\rho_h}{\beta_\ell\rho_\ell}\right)}{\alpha\left(1+\frac{\beta_h}{\beta_\ell}\right)\left(\frac{\rho_h}{\rho_\ell}-1\right)} - 1\right]^2 + \frac{4}{\alpha\left(1+\frac{\beta_h}{\beta_\ell}\right)\left(\frac{\rho_h}{\rho_\ell}-1\right)}}}{2}.$$

The results then follow as the principal puts all the weight on the wage with the highest likelihood ratio.[27] □

---

[27] It is easy to verify that $\tilde{\gamma} \leq \hat{\gamma}$, with strict inequality if $\rho_h > \rho_\ell$ and $\alpha < 1$.

References

Athey, S., and J. Levin (2001): "The Value of Information in Monotone Decision Problems," Stanford University Working Paper.

Aubin, J. (1979): *Mathematical Methods of Game and Economic Theory.* North-Holland, New York.

Bergemann, D., and U. Helge (2005): "The Financing of Innovation: Learning and Stopping," *Rand Journal of Economics*, 36, 719–752.

Blackwell, D., and M. Girshick (1954): *Theory of Games and Statistical Decisions.* Dover Publications Inc., New York.

Borwein, J., and A. Lewis (2006): *Convex Analysis and Nonlinear Optimization.* Springer, New York.

Chade, H., and E. Schlee (2002): "Another Look at the Radner-Stiglitz Nonconcavity in the Value of Information," *Journal of Economic Theory*, 107, 421–452.

DeMarzo, P., and Y. Sannikov (2011): "Learning in Dynamic Incentive Contracts," Stanford University Working Paper.

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

Ederer, F., and G. Manso (2011): "Incentives for Innovation: Bankruptcy, Corporate Governance, and Compensation Systems," in *Handbook of Law, Innovation, and Growth.* Edward Elgar Publishing.

Eso, P., and B. Szentes (2007): "The Price of Advice," *RAND Journal of Economics*, 38, 863–880.

Ganuza, J.-J., and J. Penalva (2010): "Signal Ordering Based on Dispersion and the Supply of Private Information in Auctions," *Econometrica*, 78, 1007–1030.

Gentzkow, M., and J. Shapiro (2006): "Media Bias and Reputation," *Journal of Political Economy*, 114, 280–316.

Gromb, D., and D. Martimort (2007): "Collusion and the Organization of Delegated Expertise," *Journal of Economic Theory*, 137, 271–299.

Grossman, S., and O. Hart (1983): "An Analysis of the Principal-Agent Problem," *Econometrica*, 51, 7–45.

Horner, J., and L. Samuelson (2010): "Incentives for Experimenting Agents," Yale University Working Paper.

Inderst, R., and M. Ottaviani (2010): "Intermediary Commissions and Kickbacks," Northwestern University Working Paper.

Jewitt, I. (2007): "Information Order in Decision and Agency Problems," Nuffield College Working Paper.

Krahmer, D., and R. Strausz (2010): "Optimal Procurement Contracts with Pre-Project Planning," University of Bonn Working Paper.

Levitt, S., and C. Snyder (1997): "Is no News Bad News? Information Transmission and the Role of "Early Warning" in the Principal-Agent Model," *RAND Journal of Economics*, 28, 641–661.

Lewis, T., and M. Ottaviani (2008): "Search Agency," Northwestern University Working Paper.

Manso, G. (2011): "Motivating Innovation," *Journal of Finance*, forthcoming.

Mason, R., and J. Valimaki (2008): "On Dynamic Principal-Agent Problems in Continuous Time," University of Southampton Working Paper.

Mullainathan, S., and A. Shleifer (1995): "The Market for News," *American Economic Review*, 4, 1031–1053.

Quah, J., and B. Strulovici (2009): "Comparative Statics, Informativeness, and the Interval-Dominance Ordering," *Econometrica*, 77, 1949–1942.

Rogerson, W. (1985a): "The First-Order Approach to Principal-Agents Problems," *Econometrica*, 53, 1357–1367.

——— (1985b): "Repeated Moral Hazard," *Econometrica*, 53, 69–76.

Sannikov, Y. (2008): "A Continuous-Time Version of the Principal-Agent Problem," *Review of Economic Studies*, 75, 957–984.

Spear, S., and S. Srivastava (1987): "On Repeated Moral Hazard with Discounting," *Review of Economic Studies*, 54, 599–617.

Szalay, D. (2009): "Contracts with Endogenous Information," *Games and Economic Behavior*, 65, 586–625.

Toxvaerd, F. (2006): "Time of the Essence," *Journal of Economic Theory*, 129, 252–272.