

Contracting for Research: Moral Hazard and the Incentive to Overstate Significance*

Samuel Häfner
University of Basel
samuel.haefner@unibas.ch

Curtis R. Taylor
Duke University
crtaylor@duke.edu

August 27, 2018

Abstract

A principal contracts with an agent, who is protected by limited liability, to acquire information concerning the desirability of investing in a project. To motivate the agent to perform the required research, it is necessary to offer him a schedule of contingent rewards that depend on his reported unverifiable findings and on the project's ultimate outcome. While the contingent rewards can be calibrated to solve the moral hazard problem *ex ante*, they endogenously create an adverse selection problem *ex post*. In particular, they generate an incentive for the agent to exaggerate the significance of his research findings, leading to another source of agency rents. The principal mitigates these rents by committing to ignore reports of extremely positive or extremely negative findings; i.e. extreme reports of either kind are bunched. Thus, the principal commits to under-utilize some of the agent's potential information.

JEL Classifications: C73, D81, D82, D86, L14.

Keywords: information acquisition, data dredging, p-hacking, reporting bias, reproducibility crisis, scoring rule.

*The authors thank Joe Farrell, Winand Emons, Marc Möller, Georg Nöldeke, Luis Rayo, Philipp Sadowski, Todd Sarver, Peter Norman Sorensen, and Steve Tadelis, for helpful comments and suggestions. We are also grateful to seminar participants at Berne, Bonn, Berkeley, Copenhagen, Utah (Eccles School) and Zurich. Häfner gratefully acknowledges financial support by the Swiss National Science Foundation (Grant P300P1-167648), and thanks the Economics Department at Duke for its hospitality. Ryan Westphal provided excellent research assistance. The authors are responsible for any remaining errors.

1 Introduction

The widespread use of 'statistical significance' (generally interpreted as 'p=0.05') as a license for making a claim of a scientific finding (or implied truth) leads to considerable distortion of the scientific process.

Statement by the ASA's Board of Directors ([Wasserstein et al., 2016](#), p. 129)

The authentic representation of research findings is a major concern across numerous disciplines ([Ioannidis, 2005](#)). For example, publication biases among top academic journals have been found to lead to exaggerated claims of significance for results in economics ([De Long and Lang, 1992](#)), in political science ([Gerber et al., 2008](#)), and in the cognitive sciences ([Ioannidis et al., 2014](#)). Indeed, an article in psychology ([Simmons et al., 2011](#), p. 1359) maintains that “undisclosed flexibility in data collection and analysis allows presenting *anything* as significant” (emphasis added). The problem of disingenuous reporting of results is also present in applied fields such as biomedical research, where selective reporting (i.e., the presentation of only a subset of the analysis proposed in the original study protocol) seems to be the rule rather than the exception ([Glasziou et al., 2014](#)). Also, as the analysis by [Vinkers et al. \(2015\)](#) shows, the relative frequency of positive wording over more neutral wording in scientific PubMed abstracts has increased over the last thirty years. One might argue that for applied research the problem of inflated significance is especially troubling because the particular presentation of results has a direct impact on the actions of practitioners, and in the case of biomedical research, also on those of patients, who have to rely on the accuracy of the reported research outcomes.

Moreover, detecting fraudulent practices or inflated claims in sophisticated research is notoriously difficult as illustrated by the high-profile cases of Diederik Stapel, formerly a professor of social psychology at Tilburg University, Anil Potti, formerly a cancer researcher at Duke University, and Brian Wansink, the founder and current director of the Cornell Food and Brand Lab.¹ The misconduct and/or misleading statistical practices (e.g., so-called data dredging or *p*-hacking) of each of these researchers went undetected by their universities, funding agencies, and professions for a number of years. Moreover, these three cases, while particularly egregious, do not appear to be otherwise atypical. In 2009, a systematic meta-analysis of survey data ([Fanelli, 2009](#)) reported that up to 33.7% of scientists admitted to having engaged in behaviors that “distort scientific knowledge.” The authors of the study go on to observe: “Considering that these surveys ask sensitive questions and have other limitations, it appears likely that this is a conservative estimate of the true prevalence of scientific misconduct.”

In this paper we investigate the incentives first to diligently conduct research and second to honestly report the significance of the findings in a model of optimal contracting. Specifically, we suppose that a principal engages an agent, who is protected by limited liability, to supply information about the state of a project prior to deciding how much to invest in it. For a given level of investment, a *good* project succeeds with higher marginal probability than a *bad* one. Thus, the more likely the project is good, the more the principal should invest in it.

The principal makes an upfront payment that the agent can either use to acquire information or surreptitiously divert for personal benefit. If the agent uses the payment to conduct an experiment

¹See respectively <https://www.nytimes.com/2013/04/28/magazine/diederik-stapels-audacious-academic-fraud.html>, <http://www.sciencemag.org/news/2015/11/potti-found-guilty-research-misconduct>, and <https://www.forbes.com/sites/stevensalzburg/2017/10/02/cornells-food-lab-is-cooking-up-fake-news/#1e9929e725cd>.

(i.e., collect and analyze the relevant data), then he privately observes the unverifiable realization of a two-part signal (k, α) concerning the state of the project. The first component $k \in \{0, 1\}$ is the research *finding* that indicates whether the experimental outcome is positive ($k = 1$) or negative ($k = 0$), while the second component $\alpha \in [1/2, 1]$ corresponds to the informativeness or *significance* of the research finding. Thus, if $k = 1$, then first-best investment in the project is increasing in α , and if $k = 0$, then first-best investment is decreasing in α . Because good projects sometimes fail and bad projects sometimes succeed, the actual state is never fully revealed to the principal, and the project outcome serves as an imperfect test of the authenticity of the agent's reported experimental results. Assuming that the project outcome functions can be ranked according to the monotone likelihood ratio property, the principal increases the informativeness of the project outcome by distorting investment down if the agent reports a positive finding and up if he reports a negative finding. In a sense, therefore, the principal commits to under-utilize the information supplied by the agent in order to economize on paying him rents. This is a manifestation of the celebrated Informativeness Principle of **Hölmstrom (1979)**.

We show that the two agency frictions in the model – hidden actions concerning whether the agent conducts research and hidden information regarding reporting of the results – give rise to two distinct sources of rent. Specifically, to overcome the agent's ex-ante incentive to divert the up-front payment, he must be offered a schedule of contingent rewards w that induce him to conduct research. Thus the agent receives $w_1(\alpha)$ if he reports observing $(1, \alpha)$ and the project succeeds and he receives $w_0(\alpha)$ if he reports observing $(0, \alpha)$ and the project fails. While these contingent rewards can be calibrated to solve the moral hazard problem ex ante, they create an adverse selection problem ex post. In particular, they endogenously generate an incentive for the agent to exaggerate the significance of his findings; i.e. to over-state the value of α . By inflating the significance of positive (negative) findings the agent induces the principal to invest more (less) in the project, raising the probability of success (failure) and thereby the probability of receiving a contingent payment.

To overcome the agent's ex-post incentive to inflate the significance of his findings, it is also necessary to offer him a schedule of non-contingent payments ϕ . Thus if he reports experimental results of (k, α) , then he receives a fixed payment $\phi_k(\alpha)$ and a chance at receiving a contingent reward $w_k(\alpha)$ as well. Because the incentive to exaggerate is greatest when α is smallest, the schedule of non-contingent payments $\phi_k(\alpha)$ is decreasing and hits 0 at the maximum value of α that the agent is permitted to report. To implement the first-best investment schedule x^* , the agent obviously must be allowed to report any $\alpha \in [1/2, 1]$ he observes. Under an optimal contract for the principal, however, the maximum value of α that the agent is permitted to report is capped at some value $\underline{\alpha}_k < 1$. In other words, the principal optimally commits to *bunch* all extreme significance levels $\alpha \in [\underline{\alpha}_k, 1]$. Thus, the distortions mentioned above manifest in a particular way. The optimal investment schedule associated with positive findings (negative findings) coincides with first-best investment $x_1^*(\alpha)$ ($x_0^*(\alpha)$) and is therefore increasing (decreasing) up to the critical level of significance, $\underline{\alpha}_1$ ($\underline{\alpha}_0$), after which it is flat. Such distortions are always profitable for the principal because truncating investment at the extremes generates a second-order loss in expected revenue, while shifting down the entire non-contingent payment schedule generates a first-order savings in expected agency cost.

The optimal contract turns out to have a very simple and intuitive structure, giving rise to interim utility for the agent that is linear in his posterior (increasing in the case of positive findings and decreasing in the case of negative ones). While this payoff structure resembles that of a simple scoring rule, it certainly cannot be implemented as such. Indeed, an important implication of this paper is that using a scoring rule to induce both acquisition and truthful reporting of information often requires the principal to commit not to use the information thus obtained. Of course,

information has value only to the extent it can be utilized!

The remainder of the paper is organized as follows. Related literature is discussed in the following section. The production and information acquisition technologies are presented in Section 3, where we also characterize the first-best investment schedule and show that the first-best policy is to acquire information provided that the cost of running an experiment is not too high. In Section 4, the agency environment is described and the three critical sets of constraints – limited liability, ex-ante moral hazard, and ex-post truth telling – are presented. Because the constraints for the original problem are especially unwieldy, we use a change of variables along with some familiar envelope methods to reformulate the contracting problem in an equivalent way that is more amenable to analysis. The main results appear in Section 5. In particular, we characterize both the contract that implements the first-best investment schedule at the least possible cost as well as the optimal contract for the principal. In fact, these contracts turn out to be closely related, the only substantial difference being that agency costs are reduced under the optimal contract because the principal bunches extreme signal realizations for both positive and negative findings. Interestingly, the first-best contract and the optimal contract are both only weakly ex-post implementable in the sense that the usual monotonicity restriction binds at every point. This gives rise to a simple and intuitive form for the contractual terms. Of course, it also implies that the agent only weakly prefers reporting the true significance of his results because his interim utility is independent of this report. In Section 6, we discuss the structure of the optimal contractual payments; provide additional intuition for the form of the investment distortions; briefly consider issues involving implementation; and suggest some policy implications of our findings. All proofs have been relegated to the appendix.

2 Related Literature

This paper contributes to a rich and robust literature on the collection and dissemination of information in the context of an agency relationship. The pioneering article in this area is [Demsky and Sappington \(1987\)](#), who were the first to analyze contracting with an expert agent. Other work that followed in this vein includes: [Biais and Casamatta \(1999\)](#), [Feess and Walzl \(2004\)](#), [Inderst and Ottaviani \(2009\)](#), [Malcomson \(2009\)](#), and [Szalay \(2009\)](#). Although these papers differ along a number of dimensions regarding their focus and application, they all feature models of *delegated expertise*. Specifically, an agent is hired to collect information and then make a decision on the principal's behalf. Because only outcomes are contractible, the incentives for information acquisition and decision making are necessarily confounded. By contrast, we analyze an environment where research and production are uncoupled – the agent is hired only to supply information which the principal then uses to make her own decision. We identify – what we believe to be – a novel insight in this context: providing incentives to perform research endogenously generates incentives to over-state the significance of the findings.

Contracting for advice has been studied by [Esō and Szentes \(2007\)](#) and [Krishna and Morgan \(2008\)](#).² These papers are primarily concerned with the provision of incentives for an individual to honestly report what he already knows, while our focus is on a two stage setting where an individual must first be induced to acquire information and then to truthfully disclose it. In [Esō and Szentes \(2007\)](#), a consultant has information that is useful for updating his client's belief about the profitability of undertaking a project. Their main result is that even though the consultant has an imperfect understanding of the impact of his information, he nevertheless extracts the same

²There is also a somewhat older related branch of literature in law and economics on the use of contingent fees to motivate attorneys. See, for example, [Dana and Spier \(1993\)](#) and [Rubinfeld and Scotchmer \(1993\)](#).

rents as if he fully understood it. By contrast, the agent in our model fully understands the impact of his report but he earns distinct rents from both aspects of the incentive problem: hidden actions and hidden information. The article by [Krishna and Morgan \(2008\)](#) combines a hidden information contracting setting with a sender–receiver game of the type introduced by [Crawford and Sobel \(1982\)](#). In a sense, they weaken the standard sender–receiver model by supposing that the receiver (principal) retains decision rights, although she can pay the sender (agent) for his private information. Interestingly, they find that the principal should never induce the agent to fully reveal what he knows and should never pay the agent for imprecise information. In our setting, the agent has no intrinsic preferences over the principals actions or the project outcome. Nevertheless, the endogenously created incentives for the agent to inflate the significance of his findings gives rise to a second source of agency rents which the principal attenuates by compelling the agent to give imprecise advice in the case of extreme research outcomes.

Another strand of literature to which our paper contributes is the recent work on contracting for experimentation. Most papers in this area concentrate on multiple rounds of research, and often focus on results concerning downward distortions in the optimal stopping time. The articles closest to ours in this literature are [Gerardi and Maestri \(2012\)](#) and [Halac et al. \(2016\)](#).³ In the baseline version of the model in [Gerardi and Maestri \(2012\)](#), the agent initially possesses no private information and he must be induced to run costly experiments in each period. An experiment either provides conclusive evidence that the state is good or it does not. Thus, experimentation stops if a conclusive signal is obtained or if a long enough string of inconclusive signals is observed. In the first case, the principal takes one action and in the second case she takes a different one. The basic setup in [Halac et al. \(2016\)](#) is similar, except that the agent also possesses private information about his ability to conduct research. We abstract from dynamics, assuming that the agent performs only a single experiment, which is conclusive only in a limiting case. Thus we do not address questions of optimal stopping, but explore a richer signal space and decision set. Indeed, with a bivariate signal an agent can only misrepresent the type of the findings he observes (positive or negative), but not their significance.

The three papers most closely related to this one are [Gromb and Martimort \(2007\)](#), [Carroll \(2017\)](#), and [Häfner and Taylor \(2018\)](#). In each of these papers, as in this one, a principal contracts with an agent who must be induced to acquire information that the principal then uses to make a decision. The decision facing the principal in [Gromb and Martimort \(2007\)](#) is binary and their focus is on organizational structure, specifically whether it is cheaper to hire one agent to acquire two bivariate signals or to hire two agents to each acquire one bivariate signal. They show that the two-agent organization is cheaper in the absence of collusion. The principal in our setup faces a continuum of possible signal realizations and a continuous decision of how much to invest in her project. Moreover, our focus is on contractual – rather than organizational – methods for mitigating agency rents. [Carroll \(2017\)](#) uses the methodology developed in his earlier work, [Carroll \(2015\)](#), to analyze robust contracting in a setting similar to ours with two key exceptions. First, the principal in [Carroll \(2017\)](#) is aware of only a subset of the experiments the agent can employ. Second, the true state of nature is fully revealed ex post in Carroll’s model. Our principal fully understands the agent’s technological opportunities, but the state of nature is only imperfectly revealed in our model. Indeed, our principal commits to under-utilize the agent’s findings precisely in order to obtain better information about the underlying state and thereby reduce payment of agency rents. In a companion paper, [Häfner and Taylor \(2018\)](#), we study a setting in which an agent must be induced to acquire a bivariate signal. Because truth-telling constraints never

³Other papers in this line of research include [Hörner and Samuelson \(2013\)](#), [Klein \(2016\)](#), [Henry and Ottaviani \(2017\)](#).

bind in this environment, the focus of the analysis is on methods for solving the underlying moral hazard problem, particularly on the tradeoff between using a static contract and using a dynamic reputational mechanism. By contrast, truth-telling constraints always bind in the continuous-signal model of the current paper and we, therefore, concentrate on solving the agent's endogenous incentive to exaggerate his findings.

Finally, this paper has features in common with those encountered in the literature on scoring rules and other methods for eliciting an individual's private information or beliefs. Suggestions of how to elicit subjective beliefs go back at least to [Savage \(1971\)](#) (more recently see [Dillenberger et al., 2014](#)). The most closely related article in the scoring rule literature to ours is [Hossain and Okui \(2013\)](#), which introduces a binary payment scheme to induce incentive compatible revelation of an individual's beliefs. An overview of scoring rules presently utilized appears in [Schlag et al. \(2015\)](#). We depart from the standard scoring rule environment in two important ways. First, our agent must be induced to acquire information before it can be elicited from him. Second, our agent's information has instrumental value to our principal who wishes to use it in a subsequent decision. These elements, which are the hallmarks of applied research, unfortunately undermine the efficacy of simple scoring rules in many cases.

3 The Model and First-best Solution

There are two technologies: one for information acquisition and one for production. The outcome of the information acquisition is used to determine the level of investment in the production process. Thus, we outline the technologies in reverse chronological order and then solve for the first-best policy.

3.1 Production

A risk-neutral principal must decide how much to invest in a project without knowing the underlying state $\omega \in \{B, G\}$. If she invests $x \geq 0$ and the state is ω , then the project succeeds with probability $F_\omega(x)$ and fails with probability $1 - F_\omega(x)$. Formally, we suppose the random variable Z is distributed according to the c.d.f. F_ω , with support \mathbb{R}_+ . The project succeeds if its realization satisfies $z \leq x$ and fails if $z > x$. A successful project yields revenue of 1, while an unsuccessful one yields 0. The constant marginal cost of investment is $c > 0$. The success functions are continuous and twice differentiable. For notational convenience we denote $f_\omega \equiv F'_\omega$, for $\omega \in \{B, G\}$. The production technology also satisfies the assumptions listed below.

Assumption 1 (Regularity Conditions). *For $\omega \in \{B, G\}$:*

- (i) $F_\omega(0) = 0$, and $\lim_{x \rightarrow \infty} F_\omega(x) = 1$,
- (ii) for all $x > 0$, $0 < f_\omega(x) < \infty$,
- (iii) for all $x > 0$, $-\infty < f'_\omega(x) < 0$.

Thus, the success functions are increasing, strictly concave, and have no masspoints.

Assumption 2 (Monotone Likelihood). *The likelihood ratio $\frac{f_B(z)}{f_G(z)}$ is strictly increasing.*

This assumption says that large realizations of the random variable Z are relatively more likely to have been generated by the bad distribution than the good one. This makes sense because for any investment level $x > 0$, the project fails if $z > x$. The following three implications of Assumption

2 will be useful below and are well known (cf., e.g., Theorem 1.C.1. in [Shaked and Shanthikumar, 2007](#)):

Lemma 1 (Implications of MLRP). *For all $x > 0$:*

- (i) $F_B(x) < F_G(x)$
- (ii) $\frac{f_B(x)}{F_B(x)} > \frac{f_G(x)}{F_G(x)}$,
- (iii) $\frac{f_B(x)}{1 - F_B(x)} < \frac{f_G(x)}{1 - F_G(x)}$.

The first implication points to the difference between *good* and *bad* projects. The second and third implications say that the success function for a bad project has a larger left-tail hazard rate (part ii) and a smaller right-tail hazard rate (part iii) than the success function for a good project. This is equivalent to the statement that success is more informative when investments are small (part ii) and failure is more informative when investments are large (part iii). This interpretation plays a critical role in the distortions we identify below (see Corollary 1 and the ensuing discussion).

Furthermore, Assumptions 1 and 2 (trivially) imply $f_G(0) > f_B(0)$ and that there exists a unique $\tilde{x} \in (0, \infty)$ at which f_G and f_B cross.

Assumption 3 (Investment Cost). *The marginal cost of investment satisfies:*

- (i) $c > f_\omega(\tilde{x})$,
- (ii) $c < f_B(0)$.

Because $f_G(x) > f_B(x)$ for all $x < \tilde{x}$, part (i) of this assumption implies that it is optimal to invest more in a good project than in a bad one. Part (ii) indicates that it is (nevertheless) optimal to invest a positive amount in a bad project. We assume part (ii) purely for expositional ease.

3.2 Information Acquisition

The prior probability that the project is good ($\omega = G$) is $p \in (0, 1)$. Before deciding how much to invest in the project, the principal may, however, conduct research to update these beliefs. Specifically, if the fixed information-acquisition cost, $\kappa > 0$, is sunk, then the realization of a two-part signal $s = (k, \alpha)$ is observed. The first component, $k \in \{0, 1\}$, indicates either a negative experimental finding ($k = 0$) or a positive one ($k = 1$), and the second component, $\alpha \in [1/2, 1]$, indicates the precision or significance of the finding. Specifically $\Pr\{k = 1 | \alpha, \omega = G\} = \Pr\{k = 0 | \alpha, \omega = B\} = \alpha$. Thus, at one extreme $\alpha = 1/2$ means that the experimental outcome $k \in \{0, 1\}$ is pure noise, while at the other extreme $\alpha = 1$ means that the outcome of the experiment is perfectly informative about the state of the project. We denote the signal space by $S \equiv \{0, 1\} \times [1/2, 1]$.

Assumption 4 (The Signal Distribution). *The marginal distribution of α , denoted H , has full support on $[1/2, 1]$, and possesses continuous strictly positive density h .*

For notational convenience, the probabilities of observing positive and negative findings conditional on α respectively are written

$$\begin{aligned}\varphi_1(\alpha) &\equiv p\alpha + (1 - p)(1 - \alpha) \\ \varphi_0(\alpha) &\equiv p(1 - \alpha) + (1 - p)\alpha.\end{aligned}$$

The probability that the state is good conditional on observing signal $s = (k, \alpha)$ is then denoted

$$\Pr\{\omega = G | s = (k, \alpha)\} = \frac{p(k\alpha + (1-k)(1-\alpha))}{k\varphi_1(\alpha) + (1-k)\varphi_0(\alpha)} \equiv p_k(\alpha).$$

Finally, denote the expected revenue from the project conditional on observing $s \in S$ when investment is $x \geq 0$ by

$$G_k(x, \alpha) \equiv p_k(\alpha)F_G(x) + (1 - p_k(\alpha))F_B(x)$$

And similarly define marginal expected revenue by

$$g_k(x, \alpha) \equiv p_k(\alpha)f_G(x) + (1 - p_k(\alpha))f_B(x).$$

3.3 The First-Best Policy

Assume for the moment that the principal possesses the expertise to conduct the research herself. If she sinks κ and observes (k, α) , then her ensuing investment problem is

$$\max_{x \geq 0} G_k(x, \alpha) - cx.$$

Lemma 2 (First-Best Investment). *For any $(k, \alpha) \in S$, the first-best investment level, $x_k^*(\alpha)$, is defined implicitly by the first-order condition*

$$g_k(x_k^*(\alpha), \alpha) = c. \tag{1}$$

Moreover $x_1^*(1/2) = x_0^*(1/2) \equiv x_{1/2}^*$, while $\frac{dx_1^*}{d\alpha} > 0$ and $\frac{dx_0^*}{d\alpha} < 0$.

Hence, if the principal happens to receive a purely noisy signal $\alpha = 1/2$, then she invests the same amount $x_{1/2}^*$ as if she had not bothered to acquire information in the first place. The more informative is the signal she observes (i.e., the higher is α) the more she invests in the project in the event of good news and the less in the event of bad news.

For any $(k, \alpha) \in S$, define the principal's indirect utility (i.e., value function) by

$$V(p_k(\alpha)) \equiv p_k(\alpha)F_B(x_k^*(\alpha)) + (1 - p_k(\alpha))F_B(x_k^*(\alpha)) - cx_k^*(\alpha).$$

Note that $V(p)$ represents the principal's indirect utility when investing optimally without acquiring information (or observing $\alpha = 1/2$).

Proposition 1 (The Value of Information). *The expected value of information gross of acquisition cost is positive, $E[V(p_k(\alpha))] > V(p)$.*

This is an example of the celebrated result regarding the positive value of information stemming from the convexity of the value function in the posterior.⁴ In the present context, Proposition 1 implies that if κ is sufficiently small, then the first-best policy is to conduct research and upon observing signal (k, α) to invest $x_k^*(\alpha)$. Thus we maintain the following assumption throughout.

Assumption 5 (Moderate Cost of Research). *It is first-best optimal to acquire information,*

$$\kappa < E[V(p_k(\alpha))] - V(p).$$

⁴See, for example, Arrow (1987).

4 Contracting

Now suppose that the principal must hire an agent to conduct research on her behalf. The agent is risk-neutral but has initial wealth of zero and thus is protected by limited liability. Therefore, if the principal hires the agent, she must advance him the research cost κ . The agent may, however, surreptitiously divert this payment to his own ends, yielding him a benefit of $b \in [0, \kappa]$. (The difference $\kappa - b$ may be thought of as the cost to the agent of hiding the misappropriation and falsifying a report.) If the agent acquires information, then he privately observes the signal realization $s \in S$. Whether or not the agent acquires information, he makes an unverifiable report $\hat{s} \in S$ to the principal. As usual, we suppose that the principal makes a take-it-or-leave-it contract offer to the agent and that she has full power of commitment.

A contract in this setting is denoted $[x, \phi, w] : S \rightarrow \mathbb{R}^3$ and is interpreted as follows. If the agent reports signal $\hat{s} = (\hat{k}, \hat{\alpha})$, then the principal commits to invest $x_{\hat{k}}(\hat{\alpha})$ in the project. The agent receives a fixed payment of $\phi_{\hat{k}}(\hat{\alpha})$ and a contingent payment of $w_{\hat{k}}(\hat{\alpha})$ if his reported research finding is supported by the outcome of the project. That is, for $\hat{k} = 1$ ($\hat{k} = 0$) he receives $w_1(\hat{\alpha})$ ($w_0(\hat{\alpha})$) if the project succeeds (fails). We show below that contingent payments must always be non-negative and hence limited liability simply requires

$$\phi_{\hat{k}}(\hat{\alpha}) \geq 0, \quad \forall \hat{s} \in S. \quad (\text{LLC})$$

Remark 1. *Implicit in the definition of a contract is the understanding that the agent is neither paid nor punished if his research finding is not supported by the outcome of the project; that is, he receives a contingent payment of zero if $\hat{k} = 1$ and the project fails or $\hat{k} = 0$ and it succeeds. In general, a contract in this setting would involve fixed payments $(\tilde{\phi}_1(\cdot), \tilde{\phi}_0(\cdot))$ and contingent payments $(\tilde{w}_{1,1}(\cdot), \tilde{w}_{1,0}(\cdot), \tilde{w}_{0,0}(\cdot), \tilde{w}_{0,1}(\cdot))$, where the agent is rewarded according to $\tilde{w}_{\hat{k},\hat{k}}(\hat{\alpha})$ if his report matches the project outcome and penalized according to $\tilde{w}_{\hat{k},1-\hat{k}}(\hat{\alpha})$ if it does not. In the presence of limited liability, however, it is without loss to suppose that fixed payments are $\phi_{\hat{k}}(\hat{\alpha}) = \tilde{\phi}_{\hat{k}}(\hat{\alpha}) + \tilde{w}_{\hat{k},1-\hat{k}}(\hat{\alpha})$, rewards are $w_{\hat{k}}(\hat{\alpha}) = \tilde{w}_{\hat{k},\hat{k}}(\hat{\alpha}) - \tilde{w}_{\hat{k},1-\hat{k}}(\hat{\alpha})$, and penalties are zero.*

As we will see, the contingent payments provide incentives at the ex ante stage for information acquisition (i.e., for hidden actions), and the fixed payment provides incentives at the ex post stage for truthful reporting of the observed signal (i.e., for hidden information). Below we restrict attention to *regular* contracts.

Definition 1 (Regular Contracts). *The contract $[x, \phi, w] : S \rightarrow \mathbb{R}^3$ is said to be regular if $[x, \phi, w]$ is continuous in $\hat{\alpha}$, and bounded.*

Given a contract $[x, \phi, w]$, define the agents expected utility when observing signal (k, α) and reporting $(\hat{k}, \hat{\alpha})$ by

$$u_{\hat{k},k}(\hat{\alpha}, \alpha) \equiv \phi_{\hat{k}}(\hat{\alpha}) + [1 - \hat{k} + (2\hat{k} - 1)G_k(x_{\hat{k}}(\hat{\alpha}), \alpha)]w_{\hat{k}}(\hat{\alpha}). \quad (2)$$

Even if the agent is induced to acquire information, he will not generally possess incentives to report it honestly. In fact, he typically will prefer to exaggerate the strength of the finding he observes in order to influence the principal's investment decision and thereby raise his probability of receiving a contingent payment. The contract, must, therefore, provide incentives for truthful reporting. Thus, suppose the agent acquires information and observes (k, α) . The condition ensuring that he reports this honestly (i.e., the *ex post incentive constraint*) is then written

$$(k, \alpha) = \arg \max_{(\hat{k}, \hat{\alpha}) \in S} u_{\hat{k},k}(\hat{\alpha}, \alpha). \quad (\text{EPIC})$$

If a contract satisfies (EPIC) and the agent observes (k, α) , then his interim (indirect) utility $u_{k,k}(\alpha, \alpha)$ can be written

$$U_k(\alpha) \equiv \phi_k(\alpha) + [1 - k + (2k - 1)G_k(x_k(\alpha), \alpha)]w_k(\alpha). \quad (3)$$

Then, noting that an agent who chooses not to acquire a signal is in the same reporting situation as one who observes signal $(k, 1/2)$, the condition ensuring information acquisition (i.e., the *ex ante incentive constraint*) can be written

$$\sum_{k \in \{0,1\}} \int_{\alpha=1/2}^1 U_k(\alpha) \varphi_k(\alpha) dH(\alpha) \geq b + \max_{(\hat{k}, \hat{\alpha}) \in S} u_{\hat{k},k}(\hat{\alpha}, 1/2). \quad (\text{EAIC})$$

The left side of (EAIC) is the expected utility to the agent under the contract from acquiring information and reporting it truthfully. The right side is the payoff from diverting the acquisition cost κ and making an optimal report when having observed no signal.

The principal's contract design problem can now be stated. In particular, if she chooses to contract with the agent, then she solves

$$\max_{[x, \phi, w]} \mathbb{E}[G_k(x_k(\alpha), \alpha) - cx_k(\alpha) - U_k(\alpha)] - \kappa$$

subject to the limited liability constraints, (LLC), ex post incentive constraints for honest reporting, (EPIC), and ex ante incentive constraints for acquiring information, (EAIC).

As written, the incentive constraints are especially unwieldy – so our next step is to reformulate them as conditions more amenable to analysis. We first note that if the agent happens to observe a purely noisy signal $\alpha = 1/2$, then he must be indifferent between reporting good news and bad news. That is, in any optimal contract it must hold

$$U_1(1/2) = U_0(1/2) \equiv U_{1/2}.$$

Next, we introduce a convenient change of variables and invoke some familiar envelope methods to reformulate the ex post truth-telling constraints.

Lemma 3 (EPIC Reformulated). *Consider the contract $[x, \phi, w]$ and define*

$$\theta_k(\alpha) \equiv w_k(\alpha)[F_G(x_k(\alpha)) - F_B(x_k(\alpha))]. \quad (4)$$

Then $[x, \phi, w]$ satisfies (EPIC) iff it satisfies

$$U_k(\alpha) = U_{1/2} + \int_{1/2}^{\alpha} (2k - 1)\theta_k(\hat{\alpha})p'_k(\hat{\alpha}) d\hat{\alpha}, \quad (5)$$

$$\alpha > \hat{\alpha} \implies \theta_k(\alpha) \geq \theta_k(\hat{\alpha}) \quad (6)$$

and

$$\theta_k(\alpha) \geq 0. \quad (7)$$

Noting that $(2k - 1)p'_k(\alpha) > 0$, condition (5) says that the agents interim expected utility under an implementable contract consists of a base level, $U_{1/2}$ plus a term that is increasing in the signal precision α . The function $\theta_k(\alpha)$ is the sensitivity of the incentive scheme to the project outcome at the signal realization (k, α) . Condition (6) is an especially clean version of the usual monotonicity condition required for truthful implementation. In the present context, it simply says that the

sensitivity of the incentive scheme must be non-decreasing – that is, for either realization of k , higher realizations of α must result in weakly better gambles for the agent. Finally, condition (7) just says that the sensitivity of the incentive scheme to the project outcome must be non-negative; i.e., the agent cannot be allocated gambles with a negative expected payoff.

Our next step in simplifying the incentive constraints is to show that (EAIC) must hold with equality under an optimal contract.

Lemma 4 (EAIC Binds). *Under an optimal contract, $[x, \phi, w]$ the ex ante incentive constraint (EAIC) binds.*

Now, consider a contract $[x, \phi, w]$ that satisfies (EPIC) and satisfies (EAIC) with equality: (EPIC) implies that the optimal report if the agent does not acquire information must be to claim that he observed a completely noisy signal. The right side of (EAIC) must, therefore, equal $b + U_{1/2}$. Let $\bar{\alpha} \equiv E[\alpha]$; then, the fact that

$$[p\bar{\alpha} + (1-p)(1-\bar{\alpha})]U_1(1/2) + [p(1-\bar{\alpha}) + (1-p)\bar{\alpha}]U_0(1/2) = U_{1/2},$$

together with equation (5), allows us to express the binding ex ante incentive compatibility constraint (EAIC) as

$$b = \sum_{k \in \{0,1\}} \int_{1/2}^1 \left[\int_{1/2}^{\alpha} (2k-1)\theta_k(\tilde{\alpha})p'_k(\tilde{\alpha}) d\tilde{\alpha} \right] \varphi_k(\alpha) dH(\alpha). \quad (\text{EAIC}')$$

Last, we turn to the limited liability constraints (LLC). Equating the right sides of (3) and (5) and using the relationship (4) between w_k , x_k and θ_k yields expressions for the non-contingent payments that must hold in any implementable contract:

$$\phi_k(\alpha; x, \theta, U_{1/2}) = U_{1/2} + \int_{1/2}^{\alpha} (2k-1)\theta_k(\tilde{\alpha})p'_k(\tilde{\alpha}) d\tilde{\alpha} - \theta_k(\alpha)[\gamma_k(x_k(\alpha)) + (2k-1)p_k(\alpha)], \quad (8)$$

where

$$\gamma_1(x) \equiv \frac{F_B(x)}{F_G(x) - F_B(x)} \text{ and } \gamma_0(x) \equiv \frac{1 - F_B(x)}{F_G(x) - F_B(x)}.$$

For future reference, we note that γ_1 is strictly increasing and γ_0 is strictly decreasing by parts (ii) and (iii) of Lemma 1 respectively; also $\gamma_k(x) + (2k-1)p > 0$ by inspection.

Therefore, the limited liability constraints simply require that $\phi_k(\alpha; x, \theta, U_{1/2})$ be non-negative for all $(k, \alpha) \in S$. For reasons that will become clear shortly, we write these conditions in the equivalent (albeit slightly more complicated) form

$$U_{1/2} - \theta_k(\alpha)(\gamma_k(x_k(\alpha)) + (2k-1)p) \geq \int_{1/2}^{\alpha} (2k-1)(\theta_k(\alpha) - \theta_k(\tilde{\alpha}))p'_k(\tilde{\alpha}) d\tilde{\alpha}, \quad \forall (k, \alpha) \in S. \quad (\text{LLC}')$$

These observations allow us to restate the principal's contract design problem as one of optimally choosing $[x, \theta, U_{1/2}]$ rather than $[x, \phi, w]$,

$$\max_{[x, \theta, U_{1/2}]} E[G_k(x_k(\alpha), \alpha) - cx_k(\alpha)] - \kappa - b - U_{1/2} \quad (\text{P})$$

subject to (EAIC'), (LLC'), (6), and (7).

A key advantage of this reformulation is that x does not appear in the constraints (EAIC'), (6), or (7). Moreover, the constraints are much easier to handle than their earlier counterparts.

Nevertheless, program (P) is still a fairly complex problem of optimal control featuring isoparametric constraints. Our strategy for solving it is to consider a relaxed problem (i.e., with weaker constraints) and then show that the solution to the relaxed problem actually satisfies the more stringent restrictions. Specifically, we consider the following relaxed version of program (P),

$$\max_{[x, \theta, U_{1/2}]} \mathbb{E}[G_k(x_k(\alpha), \alpha) - cx_k(\alpha)] - \kappa - b - U_{1/2} \quad (\text{RP})$$

subject to

$$\sum_{k \in \{0,1\}} \int_{1/2}^1 (2k-1)\theta_k(\alpha)p'_k(\alpha)\eta_k(\alpha) d\alpha = b \quad (9)$$

$$U_{1/2} - \theta_k(\alpha)[\gamma_k(x_k(\alpha)) + (2k-1)p] \geq 0 \quad (10)$$

$$\theta_k(\alpha) \text{ increasing.} \quad (11)$$

The ex-ante incentive constraint (9) is equivalent to (EAIC') and follows from it by reversing the order of integration and writing

$$\eta_k(\alpha) \equiv \int_{\alpha}^1 \varphi_k(\tilde{\alpha}) dH(\tilde{\alpha}). \quad (12)$$

Constraint (10) is necessary but not sufficient for (LLC') because $\theta_k(\alpha)$ is increasing; i.e., the right side of (LLC') is weakly positive. Constraint (11) is the same as (6) and we ignore constraint (7) altogether. In the next section, we characterize the solution to program (RP) and show that the solution satisfies the actual limited liability constraints (LLC') and the neglected non-negativity constraints (7). That is, the solution to (RP) is also the solution to (P).

5 The First-Best and Optimal Contracts

The structure of the constraint set (9)–(11) allows us to derive some crucial insights regarding the characteristics of the solution to (RP), and ultimately to (P). We record these in the following result.

Lemma 5 (Partial Characterization). *Let $[x, \theta, U_{1/2}]$ be a solution to the relaxed problem (RP). Then, there are thresholds $\underline{\alpha}_k \in [1/2, 1]$ such that, if $\underline{\alpha}_k < 1$, then, for all $\alpha \geq \underline{\alpha}_k$,*

- (i) *The relaxed limited liability constraint (10) for k holds with equality.*
- (ii) *The investment $x_k(\alpha)$ is non-increasing for $k = 1$ and non-decreasing for $k = 0$.*

Furthermore, if $\underline{\alpha}_k > 1/2$, then, for all $\alpha < \underline{\alpha}_k$,

- (iii) *The relaxed limited liability constraint (10) for k is slack.*
- (iv) *The investment x_k satisfies $x_k(\alpha) = x_k^*(\alpha)$.*

This result provides a helpful partial characterization of the solution to (RP). Specifically, investments are first-best up to some critical value for the signal precision $\underline{\alpha}_0$ and $\underline{\alpha}_1$ if $\underline{\alpha}_k > 1/2$. Moreover, if $\underline{\alpha}_k < 1$, then the limited liability constraints bind at $\underline{\alpha}_k$ and at all higher values of α .

Before stating our main results, it is helpful to introduce one further piece of notation and to define a particular class of contracts. Thus, let $\underline{x} = \inf x_0(\alpha)$ and $\bar{x} = \sup x_1(\alpha)$ and define the function

$$\tilde{U}_{1/2}(\underline{x}, \bar{x}) \equiv \frac{b(\gamma_0(\underline{x}) - p)(\gamma_1(\bar{x}) + p)}{p(1-p)(2\bar{\alpha} - 1)(\gamma_0(\underline{x}) + \gamma_1(\bar{x}))}. \quad (13)$$

As we shall see, this function represents the expected rents accruing to the agent associated with the private knowledge of the research outcome when the principal commits to invest at least \underline{x} if he reports a negative finding and at most \bar{x} if he reports a positive one.

Definition 2 (WEPI Contracts). *A contract $[x, \theta, U_{1/2}]$ is called weakly ex post implementable (WEPI) if $\theta_k(\alpha) = C_k$ for positive constants C_1 and C_0 .*

Under a WEPI contract, the monotonicity constraint (11) associated with ex post incentive compatibility is satisfied with equality (i.e., weakly) for all $\alpha \in [1/2, 1]$. Moreover, a WEPI contract satisfying the relaxed limited liability constraints (10) also satisfies the actual ones (LLC') because the right side of (LLC') is 0. Also, a WEPI contract trivially satisfies the neglected non-negativity constraint (7). Thus, if the optimal contract turns out to be WEPI, then the solution to problem (RP) is also the solution to the actual contract design problem (P).

We now establish that the contract implementing first-best investments at the least possible cost is, indeed, WEPI. Using Lemma 5, we solve the relaxed problem (RP) keeping investments fixed at their first-best level x^* . Then, by the observations just made, this is also the solution under the (actual) more stringent constraints.

Proposition 2 (Implementing the First-Best). *The contract that implements the first-best investment schedules x^* at the least expected cost is WEPI and is given by*

$$C_k^* \equiv \frac{\tilde{U}_{1/2}(x_0^*(1), x_1^*(1))}{\gamma_k(x_k^*(1)) + (2k - 1)p}, \quad k \in \{0, 1\},$$

This yields agency cost of $b + \tilde{U}_{1/2}(x_0^(1), x_1^*(1))$.*

Implementing the first-best investment schedule x^* at the least possible cost boils down to finding the lowest value of $U_{1/2}$ that can be achieved under the incentive constraints. Constraint (EAIC') can be thought of as a budget balance condition that requires the expected rents associated with the contingent payments to equal the agent's benefit b from diverting the information acquisition cost κ . To reduce the agency rents associated with eliciting truthful reporting, $U_{1/2}$, the principal generally wishes to *push down* the values of θ_k associated with high values of α . Budget balance then requires her to *pull up* θ_k at low values of α . Given that $\theta_k(\alpha)$ must be non-decreasing, the cost-minimizing schedule, therefore, turns out to be completely flat; i.e., the contract is WEPI.

As is usually the case in monopolistic screening models, the principal in this setting can raise her expected payoff by distorting the first-best allocation; i.e., her investments. This is stated in the next result, which is a corollary to Proposition 2.

Corollary 1 (Distorted Investment). *Under an optimal contract $\underline{\alpha}_1 < 1$ and $\underline{\alpha}_0 < 1$.*

To understand this result, recall the interpretation of Lemma 1 that says project success is more informative when investments are small (part ii) and project failure is more informative when investments are large (part iii). By distorting the investment schedule down in the case of good news and up in the case of bad news the outcome of the project is, therefore, more informative concerning the agent's reported findings and provides a better check regarding whether or not he conducted




Figure 1: The investment schedule x_k in the optimal contract follows the first-best investment x_k^* up to a threshold $\underline{\alpha}_k$ and then remains flat. The efficiency loss from bunching signals $\alpha \geq \underline{\alpha}_k$ is shaded in gray.

research. (Of course, in equilibrium the agent always actually acquires a signal.) In accordance with the Informativeness Principle (Hölmstrom, 1979), the principal better monitors the agent's compliance by distorting investments as prescribed and thereby reduces payment of expected rents. In particular, expected rents fall from $\tilde{U}_{1/2}(x_0^*(1), x_1^*(1))$ to some lower level $\tilde{U}_{1/2}(x_0^*(\underline{\alpha}_0), x_1^*(\underline{\alpha}_1))$.⁵

We are finally in a position to characterize the optimal contract for the principal subject to one caveat discussed below. Because we show that this caveat turns out not to matter, the contract specified in the following result is, as a matter of fact, the globally optimal one.

Proposition 3 (The Optimal Contract). *Suppose the optimal contract involves $\underline{\alpha}_k^\dagger > 1/2$ for $k \in \{0, 1\}$ and let $x_k^\dagger \equiv x_k^*(\underline{\alpha}_k^\dagger)$. Then the optimal contract for the principal is WEPI and is given by*

$$x_k(\alpha) = \begin{cases} x_k^*(\alpha), & \text{if } \alpha \in [1/2, \underline{\alpha}_k^\dagger] \\ x_k^\dagger, & \text{if } \alpha \in [\underline{\alpha}_k^\dagger, 1] \end{cases},$$

and

$$C_k^\dagger \equiv \frac{\tilde{U}_{1/2}(x_0^\dagger, x_1^\dagger)}{\gamma_k(x_k^\dagger) + (2k - 1)p}, \quad k \in \{0, 1\},$$

where

$$\mathbb{E}[g_{\hat{k}}(x_k^\dagger, \alpha) - c | \hat{k} = k, \alpha \geq \underline{\alpha}_k] (1 - H(\underline{\alpha}_k)) = \frac{\partial \tilde{U}_{1/2}}{\partial x_k^\dagger}, \quad k \in \{0, 1\}. \quad (14)$$

This yields agency cost of $b + \tilde{U}_{1/2}(x_0^\dagger, x_1^\dagger)$. Further, it holds for $k \in \{0, 1\}$ that $\underline{\alpha}_k^\dagger \rightarrow 1$ as $b \rightarrow 0$; i.e., the solution converges to first-best as agency costs vanish.

The optimal values $\underline{\alpha}_0^\dagger$ and $\underline{\alpha}_1^\dagger$ are implicitly defined by (14), which has a natural interpretation. The left side is the expected marginal cost of increasing the distortion in investment by tightening the cutoff for reporting significance $\underline{\alpha}_k$. Specifically, it is the average distortion away from first-best investment over the range $[\underline{\alpha}_k^\dagger, 1]$ times the probability of drawing a significance level in this

⁵Corollary 1 is an example of the distortions encountered in multi-dimensional monopolistic screening models of the kind studied by Armstrong (1996).




Figure 2: The figure depicts the fixed payments ϕ_k^* in the first-best contract together with fixed payments ϕ_k in the optimal contract. The optimal payments are shifted downwards and remain zero for signals $\alpha \geq \underline{\alpha}_k$. The agency cost gains from such shifts are shaded in gray. Comparing the gains to the losses in Figure 1, we see that the gains are first-order while the losses are second-order.

range. (In accordance with Corollary 1, this value approaches 0 as $\underline{\alpha}_k^\dagger \rightarrow 1$ both by the Envelope Theorem, $g_k(x_k^*(1), 1) - c = 0$, and because $\alpha = 1$ is a negligible event, $1 - H(1) = 0$.) The right side of (14) is the marginal benefit from tightening the cutoff $\underline{\alpha}_k$: namely, the savings in agency rent associated with reducing the schedule of non-contingent payments $\phi_k(\alpha)$. We return to this point in the following section where other aspects of Proposition 3 are also discussed.

Proposition 3 does not *quite* establish that the specified contract is globally optimal because of the qualification that $\underline{\alpha}_k^\dagger > 1/2$ for $k \in \{0, 1\}$. This stipulation is necessary because the principal has access to another feasible type of contract that is not nested in the class of contracts for which part (iv) of Lemma 5 applies. Specifically, she can set $\underline{\alpha}_k = 1/2$ for $k \in \{0, 1\}$.⁶ This corresponds to simply asking the agent for a *thumbs-up-or-down* report. That is, the principal requires the agent to report k but not α . This type of contract is WEPI by construction and involves only two investment levels \bar{x}_1 and \bar{x}_0 .⁷

Given Propositions 2 and 3, a reasonable conjecture is that the set of thumbs-up-or-down contracts corresponds to the unique one achieved in the limit as $\underline{\alpha}_k \rightarrow 1/2$ for $k \in \{0, 1\}$. While this limiting contract clearly belongs to the set of implementable thumbs-up-or-down contracts, it is definitely not the optimal one! As $\underline{\alpha}_k \rightarrow 1/2$, the contract specified in Proposition 3 converges to one in which $x_k^\dagger = x_{1/2}^*$ for $k \in \{0, 1\}$; i.e., the principal does not use any of the information observed by the agent whatsoever. Of course, there would be no point in hiring the agent in this case.

While it is straightforward – though somewhat tedious – to characterize the optimal thumbs-up-or-down contract, Proposition 4 below obviates the need for doing so. Before formally stating the claim, we need only to point out two basic observations. First, within the class of thumbs-up-or-down contracts, one involving $\bar{x}_k = x_k^*(1/2)$ is never optimal because – as mentioned above – it uses none of the agent’s information. Second, a contract involving $\bar{x}_k > x_k^*(1)$ is clearly dominated by one with $\bar{x}_k = x_k^*(1)$.

⁶The mixed cases of $\underline{\alpha}_k = 1/2$ and $\underline{\alpha}_{1-k} > 1/2$ are handled analogously.

⁷Formally, in order to invoke the Revelation Principal we must allow a third investment level $\bar{x}_{1/2}$ corresponding to a report by the agent that he did not acquire information, but this will not occur on-path.




Figure 3: The figure depicts investments \bar{x}_k under a feasible thumbs-up-or-down contract. The proof to Proposition 4 shows that any such contract is dominated by an alternative, feasible WEPI contract involving the same agency rents yet first-best investments x_k^* up to the thresholds $\underline{\alpha}_k'$. The gains from using the alternative contract over the thumbs-up-or-down contract are shaded in gray.

Proposition 4 (No Thumbs). *For any implementable thumbs-up-or-down contract satisfying $x_0^*(1) \leq \bar{x}_0 < x_{1/2}^* < \bar{x}_1 \leq x_1^*(1)$, there exists an implementable contract of the kind specified in Proposition 3 that the principal strictly prefers.*

The superior contract follows the first-best investment schedules $x_0^*(\alpha)$ and $x_1^*(\alpha)$ from $x_{1/2}^*$ until the points where they respectively hit the investment levels of the original contract, \bar{x}_0 and \bar{x}_1 and coincide with these constant levels of investment thereafter. This contract satisfies all of the constraints (because the original contract satisfies them), generates the same level of agency cost, $\tilde{U}_{1/2}(\bar{x}_0, \bar{x}_1)$, and involves better investments for a non-negligible set of signals (see Figure 3). Thus, the qualification that $\underline{\alpha}_k^\dagger > 1/2$ is without loss, and the unique globally optimal contract is the one specified in Proposition 3.

Finally, we note that the principal prefers the contract specified in Proposition 3 to autarky if the agent's private benefit b from committing moral hazard is small enough. As b approaches 0, agency costs vanish and the investment schedules converge to the first-best. Thus, for b sufficiently small, the principal prefers to contract with the agent by Assumption 5.

6 Discussion

6.1 Payments

Having derived the optimal contract $[x, \theta, U_{1/2}]$, we need to convert its terms back into payments. In particular, we can invert (4) to get the contingent rewards

$$w_k^\dagger(\alpha) = \begin{cases} \frac{C_k^\dagger}{F_G(x_k^*(\alpha)) - F_B(x_k^*(\alpha))}, & \text{for } \alpha \in [1/2, \underline{\alpha}_k^\dagger] \\ \frac{C_k^\dagger}{F_G(x_k^\dagger) - F_B(x_k^\dagger)}, & \text{for } \alpha \in [\underline{\alpha}_k^\dagger, 1]. \end{cases}$$

Suppose momentarily that the principal attempted to incentivise the agent purely through these contingent payments. If the agent acquired information and observed (k, α) , then his expected

payoff from reporting $(k, \hat{\alpha})$ would be

$$(1 - k + (2k - 1)G_k(x_k(\hat{\alpha}), \alpha))w_k^\dagger(\hat{\alpha}) = \begin{cases} C_k^\dagger(\gamma_k(x_k^*(\hat{\alpha})) + (2k - 1)p_k(\alpha)), & \text{for } \hat{\alpha} \in [1/2, \underline{\alpha}_k^\dagger) \\ C_k^\dagger(\gamma_k(x_k^\dagger) + (2k - 1)p_k(\alpha)), & \text{for } \alpha \in [\underline{\alpha}_k^\dagger, 1] \end{cases}$$

This is increasing in $\hat{\alpha}$ up to $\hat{\alpha} = \underline{\alpha}_k^\dagger$. In other words, the agent would have incentives to exaggerate the significance of his research findings in order to induce the principal to invest more in the case of positive findings and less in the case of negative findings. Indeed, this is true of any contract that is sensitive to reports of α that does not involve fixed payments. Moreover, contracts that are not sensitive to reports of α are dominated by ones that are sensitive (except at the extremes) by Proposition 4.

To resolve the agent's incentive to exaggerate, the principal must make a non-contingent payment according to (8) of

$$\phi_k^\dagger(\hat{\alpha}) = \begin{cases} C_k^\dagger(\gamma_k(x_k^\dagger) - \gamma_k(x_k^*(\hat{\alpha}))), & \text{for } \hat{\alpha} \in [1/2, \underline{\alpha}_k^\dagger) \\ 0, & \text{for } \alpha \in [\underline{\alpha}_k^\dagger, 1]. \end{cases} \quad (15)$$

Adding together the two preceding expressions results in expected utility of

$$u_{k,k}(\hat{\alpha}, \alpha) = \tilde{U}_{1/2}(x_0^\dagger, x_1^\dagger) + C_k^\dagger(2k - 1)(p_k(\alpha) - p), \forall \alpha.$$

Thus, under an optimal contract, the agent's interim expected utility is independent of the reported precision of his information, $\hat{\alpha}$. This holds because the implementability restriction (6) binds at every point; i.e., the contract is only weakly implementable in terms of providing incentives for truthful reporting of α . It is easy to check that the difference in expected utility from misreporting k (i.e., $\hat{k} = 1 - k$) is $\frac{b(2k-1)(p-p_k(\alpha))}{p(1-p)(2\hat{\alpha}-1)}$, which is strictly negative whenever $\alpha > 1/2$. Thus, unless the agent observes a purely noisy signal, he possesses a strict incentive to report truthfully the type of finding (positive or negative) he observes.

6.2 Distortions

As just noted, the reason the principal must use non-contingent payments under an optimal contract is to deter the agent from overstating the significance of his findings. The expression for ϕ_k given in (15) is decreasing precisely because the weaker are his findings, the more the agent must be paid not to exaggerate. Of course, the most tempting exaggeration is to claim the highest feasible level of significance, $\alpha = 1$ if the contract allows this. As is demonstrated in Corollary 1, however, an optimally-designed contract truncates the highest reportable level of significance at $\underline{\alpha}_k^\dagger < 1$. This permits the principal to shift down the entire payment schedules $\phi_k(\alpha)$, resulting in a first-order reduction in expected agency cost from $\tilde{U}_{1/2}(x_0^*(1), x_1^*(1))$ to $\tilde{U}_{1/2}(x_0^*(\underline{\alpha}_0^\dagger), x_1^*(\underline{\alpha}_1^\dagger))$. Thus the principal sacrifices extreme reports and the concomitant investments in order to economize on expected agency costs (review Figures 1 and 2).

The agent's expected utility (expected agency cost from the principal's perspective) under an optimal contract is scaled by the factor $\frac{b}{p(1-p)(2\hat{\alpha}-1)}$. Comparative statics concerning the components of this term are intuitive. Clearly, if b is small, then the agent does not benefit much from diverting κ , and the rent required to induce him to acquire information and report it truthfully is correspondingly modest. Indeed, as Proposition 3 demonstrates, the principal achieves her first-best payoff in the limit as b vanishes.

Next, note that expected agency cost is high when p is either close to 1 or close to 0. In either case, it is expensive to provide incentives for the agent to acquire information because he is

tempted to simply pocket κ and go with the prior, reporting $\hat{k} = 1$ if p is close to 1 or $\hat{k} = 0$ if p is close to 0. Similarly, expected rent is high when $\bar{\alpha}$ is close to 1/2. In this case, it is expensive to incentivise the agent because the average informativeness of the experiment is not very high and he can, therefore, do nearly as well by remaining ignorant and making a false report. A downward shift in H over an open set of $\alpha \in [1/2, 1]$ corresponds to first-order stochastic dominance and leads to a rise in the average quality of information, $\bar{\alpha}$. Thus, when the agent has access to better information on average, the principal can reduce his expected compensation. (Of course, shifting the distribution H down might well correspond to a more expensive experiment requiring more data or better analytics; i.e., raising $\bar{\alpha}$ might well necessitate a rise in κ and b .)

6.3 Implementation

Although we have – as is usually done – solved the principal’s problem for an optimal direct truthful mechanism, we could implement the same incentive scheme in a number of equivalent ways. For instance, the principal could ask the agent to report his posterior beliefs after conducting the research, $\hat{\pi} \in [p_0(\underline{\alpha}_0^\dagger), p_1(\underline{\alpha}_1^\dagger)]$ or to recommend a level of investment $\hat{x} \in [x_0^\dagger, x_1^\dagger]$. Suppose, for example, that the principal asked the agent to report his posterior $\hat{\pi}$. Then the agent’s interim indirect utility under an optimal contract would be

$$b + U(\pi) = \begin{cases} b + C_0^\dagger(\gamma_0(x_0^\dagger) - \pi), & \text{for } \pi \in [p_0(\underline{\alpha}_0^\dagger), p] \\ b + C_1^\dagger(\gamma_1(x_1^\dagger) + \pi), & \text{for } \pi \in [p, p_1(\underline{\alpha}_1^\dagger)] \end{cases}$$

This is a piecewise linear function that attains a minimum at $b + U_{1/2}^\dagger$ when $\pi = p$. Thus, the agent’s interim expected payoff is convex in the posterior π . In other words, the optimal contract induces him to have risk-loving preferences or, what amounts to the same thing, demand for information. Indeed, this is why he (weakly) prefers to conduct the experiment rather than divert the advanced payment κ . Nevertheless, the principal cannot gain from adding more randomness to the mechanism, as this would be tantamount to garbling the signal (k, α) , necessitating an increase in his expected reward.

Finally, while this payoff function resembles that of a simple scoring rule, it essentially cannot be implemented as such. In particular, offering the agent a base payment of $b + U_{1/2}^\dagger$ and contingent payments that are linear in his reported posterior would almost surely not induce him to acquire information, and if – by chance – it did, then it would certainly not induce him to honestly report his findings.

6.4 Policy Prescriptions and Practical Considerations

While the model presented in this paper is admittedly stylized, we believe it nevertheless suggests some robust policy implications. First, researchers should be rewarded not only for strong positive findings, but compelling negative ones as well (e.g., studies failing to replicate prior published findings should garner serious consideration for publication themselves). Additionally, a report of weak findings should also be compensated, albeit primarily through non-contingent means. For example, a scientist whose current research does not result in findings that are promising enough to justify applying for a major grant should still receive baseline compensation, perhaps from the institution with which he is affiliated. Failure to provide such compensation generates incentives for researchers to dredge their data in an effort to discover (spurious) significant results.

Finally, we believe that our results shed some light on the current debate on how best to deal with the reproducibility crisis in science (Baker, 2016). First, there is a school of thought maintaining

that researchers should be discouraged or even banned from reporting p -values altogether. Indeed, in 2015 the editors of the *Journal of Basic and Applied Social Psychology* did just this, when they announced that “From now on, *BASP* is banning the NHSTP [Null Hypothesis Significance Testing Procedure]” (Trafimow and Marks, 2015). Our findings suggest that this policy is overkill. While banning the NHSTP eliminates the incentive to inflate significance performe, Proposition 4 indicates that using some threshold level of significance greater than 0 is generally desirable.

A second school of thought advocates the opposite policy, recommending that the standards for reporting statistical significance be raised, say from 5% to 0.5% (Benjamin et al., 2018). This could potentially help the current situation by increasing the time and effort required to generate spurious findings through data dredging and similar practices (it would be interesting to augment our baseline model by including a cost of exaggeration $e(\hat{\alpha} - \alpha)$). Even so, our results suggest a caveat. Raising the standards for statistical significance will also raise the incentive for researchers to inflate their findings and the concomitant expense associated with deterring them from doing so.

Finally, a third camp of scholars recommends using a continuous criterion such as the Bayes factor rather than a fixed threshold for evaluating the validity of experimental findings (Etz and Vandekerckhove, 2016). Our analysis is – as are the vast majority of game-theoretic investigations – couched in a Bayesian framework. Thus, while we admittedly sidestep the vexing issues surrounding adoption of an appropriate prior, the decision maker in our model adjusts her optimal policy continuously over a range of possible reported experimental outcomes. Nevertheless, cutoffs that closely resemble the fixed thresholds of classical hypothesis testing do emerge as part of an optimal mechanism in our Bayesian environment.

A Proofs

A.1 Proofs for Section 3

Proof of Lemma 2. Consider the expected marginal benefit function

$$g_k(x, \alpha) = p_k(\alpha)f_G(x) + (1 - p_k(\alpha))f_B(x).$$

This is positive and decreasing by Assumption 1 and crosses c once by part (ii) of Assumption 3. Next observe that $p_1(1/2) = p_0(1/2) = p$ implies $x_1^*(1/2) = x_0^*(1/2)$. Finally, implicit differentiation of $g_k(x_k^*, \alpha) - c = 0$ yields

$$\frac{dx_1^*}{d\alpha} = -\frac{p(1-p)(f_G(x_1^*(\alpha)) - f_B(x_1^*(\alpha)))}{\varphi_1(\alpha)(p\alpha f'_G(x_1^*(\alpha)) + (1-p)(1-\alpha)f'_B(x_1^*(\alpha)))}.$$

This is strictly positive because $f_G(x_1^*(\alpha)) > f_B(x_1^*(\alpha))$ by part (i) of Assumption 3 and $f'_\omega(\cdot) < 0$ by part (xiii) of Assumption 1. Likewise

$$\frac{dx_0^*}{d\alpha} = \frac{p(1-p)(f_G(x_0^*(\alpha)) - f_B(x_0^*(\alpha)))}{\varphi_0(\alpha)(p(1-\alpha)f'_G(x_0^*(\alpha)) + (1-p)\alpha f'_B(x_0^*(\alpha)))}$$

is strictly negative. □

Proof of Proposition 1. If $E[p_k(\alpha)] = p$ and $V(\cdot)$ is strictly convex, then the claim will follow from Jensen’s inequality. Observe that

$$E[p_k(\alpha)] = \int_{\alpha=1/2}^1 \sum_{k \in \{0,1\}} \varphi_k(\alpha)p_k(\alpha) dH(\alpha).$$

Direct substitution reveals

$$\varphi_1(\alpha)p_1(\alpha) + \varphi_0(\alpha)p_0(\alpha) = p,$$

from which it follows that $E[p_k(\alpha)] = p$. Next, by the Envelope Theorem

$$\frac{dV}{d\tilde{p}} = F_G(x_k^*(\alpha)) - F_B(x_k^*(\alpha)).$$

Differentiating again gives

$$\frac{d^2V}{d\tilde{p}^2} = (f_G(x_k^*(\alpha)) - f_B(x_k^*(\alpha))) \frac{\partial x_k^*}{\partial \tilde{p}}.$$

Implicitly differentiating (1) gives

$$\frac{\partial x_k^*}{\partial \tilde{p}} = -\frac{f_G(x_k^*) - f_B(x_k^*)}{\tilde{p}f'_G(x_k^*) + (1 - \tilde{p})f'_B(x_k^*)}.$$

Therefore

$$\frac{d^2V}{d\tilde{p}^2} = -\frac{(f_G(x_k^*) - f_B(x_k^*))^2}{\tilde{p}f'_G(x_k^*) + (1 - \tilde{p})f'_B(x_k^*)}.$$

This is strictly positive by part (iii) of Assumption 1. \square

A.2 Proofs for Section 4

Proof of Lemma 3. We start by showing that, if $[x, \phi, w]$ satisfies (EPIC) then (5) – (7) hold. First, observe that $u_{k,k}(\hat{\alpha}, \alpha)$ is differentiable in α for all $\hat{\alpha}$, and hence absolutely continuous. Further, because both the payments $[\phi, w]$ as well as $|p'_k(\alpha)|$ are uniformly bounded, there is a finite number $d > 0$ such that $|\partial u_{k,k}/\partial \alpha| < d$ for all $\hat{\alpha}$. Last, because $[x, \phi, w]$ is continuous in $\hat{\alpha}$, so is $u_{k,k}(\hat{\alpha}, \alpha)$, and hence it attains a maximum on $[1/2, 1]$ for any α . Consequently, Theorem 2 in [Milgrom and Segal \(2002\)](#) applies, yielding (5). Next, (EPIC) requires for the good signal case (the bad signal case is analogous) that

$$\begin{aligned} & u_{1,1}(\alpha, \alpha) - u_{1,1}(\hat{\alpha}, \alpha) + u_{1,1}(\hat{\alpha}, \hat{\alpha}) - u_{1,1}(\alpha, \hat{\alpha}) \geq 0 \\ \iff & [u_{1,1}(\alpha, \alpha) - u_{1,1}(\hat{\alpha}, \alpha)] - [u_{1,1}(\alpha, \hat{\alpha}) - u_{1,1}(\hat{\alpha}, \hat{\alpha})] \geq 0 \\ \iff & [p_1(\alpha) - p_1(\hat{\alpha})][\theta_1(\alpha) - \theta_1(\hat{\alpha})] \geq 0, \end{aligned}$$

from which (6) follows because $p_1(\alpha)$ is strictly increasing.

To show that (EPIC) implies (7), we prove the contrapositive. Thus, assume that $\theta_1(\alpha') < 0$ for some $\alpha' \in [1/2, 1]$ (the argument for $k = 0$ is analogous). Then, because $\theta_1(\cdot)$ is continuous and non-decreasing, there exists $(\underline{\alpha}, \bar{\alpha})$ such that $\theta_1(\alpha) < 0$ for all $\alpha \in (\underline{\alpha}, \bar{\alpha})$. Moreover, it follows from (5) that $U_1(\cdot)$ is strictly decreasing on $(\underline{\alpha}, \bar{\alpha})$, or

$$\alpha < \hat{\alpha} \in (\underline{\alpha}, \bar{\alpha}) \implies \phi_1(\alpha) + G_1(x_1(\alpha), \alpha)w_1(\alpha) > \phi_1(\hat{\alpha}) + G_1(x_1(\hat{\alpha}), \hat{\alpha})w_1(\hat{\alpha}).$$

Observing that $\hat{\alpha} > \alpha$ implies

$$\phi_1(\alpha) + G_1(x_1(\alpha), \hat{\alpha})w_1(\alpha) > \phi_1(\alpha) + G_1(x_1(\alpha), \alpha)w_1(\alpha),$$

we obtain $u_{1,1}(\alpha, \hat{\alpha}) > u_{1,1}(\hat{\alpha}, \hat{\alpha})$ which contradicts (EPIC).

Next, we show that (EPIC) follows from (5)–(7).⁸ First, observe that (5) gives us

$$U_k(\alpha) = U_k(\hat{\alpha}) + \int_{\hat{\alpha}}^{\alpha} (2k-1)\theta_k(\tilde{\alpha})p'_k(\tilde{\alpha})d\tilde{\alpha}, \quad (\text{A.1})$$

for any $\alpha, \hat{\alpha} \in [1/2, 1]$. Now, if we assume that $\alpha > \hat{\alpha}$, then (A.1) and (6) together give us that

$$U_k(\alpha) \leq U_k(\hat{\alpha}) + (2k-1)\theta_k(\alpha)[p_k(\alpha) - p_k(\hat{\alpha})].$$

From the fact that

$$u_{k,k}(\alpha, \alpha) = u_{k,k}(\alpha, \hat{\alpha}) + (2k-1)\theta(\alpha)[p(\alpha) - p(\hat{\alpha})], \quad (\text{A.2})$$

we obtain $u_{k,k}(\alpha, \hat{\alpha}) \leq u_{k,k}(\hat{\alpha}, \hat{\alpha})$. Next, if we assume $\alpha < \hat{\alpha}$, then (A.1) and (6) together give us

$$U_k(\alpha) \leq U_k(\hat{\alpha}) - (2k-1)\theta_k(\alpha)[p_k(\hat{\alpha}) - p_k(\alpha)].$$

From (A.2), we again obtain $u_{k,k}(\alpha, \hat{\alpha}) \leq u_{k,k}(\hat{\alpha}, \hat{\alpha})$. Together, we thus have

$$u_{k,k}(\alpha, \alpha) \geq u_{k,k}(\hat{\alpha}, \alpha), \forall \alpha, \hat{\alpha}. \quad (\text{A.3})$$

It remains to show that $u_{1-k,k}(\hat{\alpha}, \alpha) \leq u_{k,k}(\alpha, \alpha)$ holds for all $\hat{\alpha}$. To see this, observe that

$$\begin{aligned} u_{1-k,k}(\hat{\alpha}, \alpha) &= u_{1-k,1-k}(\hat{\alpha}, 1/2) - (2k-1)\theta_{1-k}(\hat{\alpha})[p_k(\alpha) - p] \\ &\leq U_{1/2}, \end{aligned}$$

where the inequality follows from (7) and (A.3). But (5) together with (7) gives $u_{k,k}(\alpha, \alpha) \geq U_{1/2}$, thus giving us the claim, and completing the proof. \square

Proof of Lemma 4. Suppose to the contrary that there is an optimal contract $[x, \phi, w]$ for which (EAIC) is slack; i.e.,

$$\begin{aligned} &\int_{\alpha=1/2}^1 \left[\varphi_1(\alpha) (\phi_1(\alpha) + G_1(x_1(\alpha), \alpha)w_1(\alpha)) \right. \\ &\quad \left. + \varphi_0(\alpha) (\phi_0(\alpha) + (1 - G_0(x_0(\alpha), \alpha))w_0(\alpha)) \right] dH(\alpha) > b \\ &\quad + \max_{\hat{\alpha} \in [1/2, 1]} \left\{ \max\{\phi_1(\hat{\alpha}) + G_1(x_1(\hat{\alpha}), 1/2)w_1(\hat{\alpha}), \phi_0(\hat{\alpha}) + (1 - G_0(x_0(\hat{\alpha}), 1/2)w_0(\hat{\alpha}))\} \right\}. \end{aligned}$$

Because $\{\phi_k(\cdot), w_k(\cdot)\}$ enter linearly on either side of the above inequality, there exists $\eta \in (0, 1)$ such that we can multiply $\phi_k(\cdot)$ and $w_k(\cdot)$ on both sides of the inequality by $(1 - \eta)$ and it will continue to hold. Moreover, the fact that the original contract $[x, \phi, w]$ satisfies (LLC) and (EPIC) evidently implies that the rescaled contract $[x, (1 - \eta)\phi, (1 - \eta)w]$ satisfies them. Thus, there is an incentive compatible contract respecting limited liability that induces the same level of investment as the optimal contract that is cheaper for the principal, a contradiction. \square

⁸The first part of the argument is borrowed from Myerson (1981).

A.3 Proofs for Section 5

Proof of Lemma 5. To begin, observe that (i) implies (ii); i.e., whenever (10) binds at (k, α) , then x_k must be weakly decreasing (increasing) for $k = 1$ ($k = 0$). To see this, recall that $\gamma_k(\alpha)$ is strictly increasing (decreasing) for $k = 1$ ($k = 0$) by parts (ii) and (iii) of Lemma 1. The claim then follows from the monotonicity constraint (11) on θ_k .

Next, we show that (iii) implies (iv); i.e., whenever (10) does not bind for some (k, α) , then the investment $x_k(\alpha)$ in the optimal contract corresponds to $x_k^*(\alpha)$. To see this, fix $k = 1$ (the argument for $k = 0$ is analogous), and suppose (10) does not bind for some α' . Then, by continuity, (10) does not bind in a neighborhood of this α' . Now, suppose $x_k(\alpha')$ does not correspond to the first-best investment $x_k^*(\alpha')$, and consider an alternative investment \tilde{x}_k corresponding to x_k on all α except for those in a neighborhood of α' , where $|x_1^*(\alpha) - \tilde{x}_1(\alpha)| < |x_1^*(\alpha) - x_1(\alpha)|$ holds for all α in this neighborhood and \tilde{x}_1 is chosen to lie sufficiently close to x_1 such that (10) is still satisfied. Then, the alternative investment \tilde{x}_k is also feasible but, by concavity of $G_1(z, \alpha) - cz$ in z , yields a higher payoff to the principal. Moreover, the alternative contract still satisfies (9) because θ is unchanged.

Now, before we show (i) and (iii), we make two observations. First, we observe that if (10) binds at (k, α) , then $\theta_k(\alpha) > 0$: Fix $k = 1$ (the argument for $k = 0$ is analogous), and by way of contradiction suppose that (10) binds at some $(1, \alpha)$ for which $\theta_1(\alpha) = 0$. Evaluating (10) at $(1, \alpha)$ then reveals $U_{1/2} = 0$. Consequently, there cannot exist $\hat{\alpha} \neq \alpha$ such that $\theta_1(\hat{\alpha}) > 0$, because this would violate (10) and thus contradict feasibility of the contract. Yet, if no $\hat{\alpha}$ for which $\theta_1(\hat{\alpha}) > 0$ exists, then there must exist $\hat{\alpha}$ such that $\theta_0(\hat{\alpha}) > 0$ by ex-ante incentive compatibility, (9). But then, the limited liability constraint (10) for $k = 0$ is violated, contradicting again feasibility of the contract.

Second, we observe that whenever (10) does not bind for k on some interval $[\alpha, \hat{\alpha}]$ for which $\theta_k(\alpha) > 0$, then the left side of the respective (10) is strictly decreasing: For $k = 1$,

$$\theta_1(\hat{\alpha}) [\gamma_1(x_1(\hat{\alpha})) + p] - \theta_1(\alpha) [\gamma_1(x_1(\alpha)) + p] > 0,$$

where the strict inequality follows from our assumption that $\theta_1(\alpha) > 0$ and the facts that θ_1 is weakly increasing, γ_1 is strictly increasing (by part (ii) of Lemma 1), and x_1 is strictly increasing because $x_1 = x_1^*$ (by part (i) above) and x_1^* is strictly increasing (by Lemma 2). An analogous argument can be made for $k = 0$.

We can use these observations to argue that if (10) binds for some α_k then it must bind for all $\alpha \geq \alpha_k$: Suppose to the contrary, for $k = 1$, that there is $\alpha' \in [1/2, 1)$ such that (10) is slack for all $\alpha > \alpha'$ in a neighborhood of α' but binds at α' . Because $\theta_k(\alpha) > 0$ for all $\alpha \geq \alpha'$ by the first observation above, the second observation just established gives us that $U_{1/2} - \theta_1(\alpha) [\gamma_1(x_1(\alpha)) + p] < 0$ for all $\alpha > \alpha'$ in a neighborhood of α' , contradicting feasibility. But this implies (i) and (iii), thus completing the proof. \square

Proof of Proposition 2. We first solve the relaxed problem (RP). Consider a first-best contract in which $x_k(\alpha) = x_k^*(\alpha)$ for all $(k, \alpha) \in S$. Because investments x are kept fixed, the problem reduces to choosing θ to minimize $U_{1/2}$ subject to (9), (10), and (11). Because investments are strictly monotone, the limited liability constraints (10) bind, if at all, at $\alpha = 1$ and are thus written

$$\theta_k(1) [\gamma_k(x_k^*(1)) + 2k - 1] - U_{1/2} \leq 0.$$

Given the linearity of the constraints in θ , the solution will involve extreme values. We therefore

consider a relaxed program in which there are cutoffs a_k such that

$$\theta_k(\alpha) = \begin{cases} \theta_k(1/2), & \text{if } \alpha < a_k \\ \theta_k(1) & \text{if } \alpha \geq a_k \end{cases}$$

We verify below that the solution to the relaxed program satisfies the neglected constraints. Thus, consider minimizing the Lagrangian

$$\begin{aligned} & U_{1/2} + \mu_1[\theta_1(1)[\gamma_1(x_1^*(1)) + 1] - U_{1/2}] + \mu_0[\theta_0(1)[\gamma_0(x_0^*(1)) - 1] - U_{1/2}] \\ & + \lambda \left[b - \theta_1(1/2) \int_{1/2}^{a_1} p'_1(\alpha)\eta_1(\alpha) d\alpha - \theta_1(1) \int_{a_1}^1 p'_1(\alpha)\eta_1(\alpha) d\alpha \right. \\ & \left. + \theta_0(1/2) \int_{1/2}^{a_0} p'_0(\alpha)\eta_0(\alpha) d\alpha + \theta_0(1) \int_{a_0}^1 p'_0(\alpha)\eta_0(\alpha) d\alpha \right]. \end{aligned}$$

The first-order condition for $U_{1/2}$ yields

$$\mu_1 + \mu_0 = 1. \quad (\text{A.4})$$

Thus, one or both limited liability constraints bind. The first-order conditions for $\theta_1(1)$ and $\theta_0(1)$ reveal

$$\mu_1[\gamma_1(x_1^*(1)) + 1] = \lambda \int_{a_1}^1 p'_1(\alpha)\eta_1(\alpha) d\alpha \quad (\text{A.5})$$

$$\mu_0[\gamma_0(x_0^*(1)) - 1] = -\lambda \int_{a_0}^1 p'_0(\alpha)\eta_0(\alpha) d\alpha. \quad (\text{A.6})$$

From (A.4), the left side of one or both of the preceding two equations must be positive. This, however, implies that $\lambda > 0$ (recall that $p'_0(\alpha) < 0$). The first-order conditions for $\theta_1(1/2)$ and $\theta_0(1/2)$ are

$$\begin{aligned} 0 &= \lambda \int_{1/2}^{a_1} p'_1(\alpha)\eta_1(\alpha) d\alpha \\ 0 &= -\lambda \int_{1/2}^{a_0} p'_0(\alpha)\eta_0(\alpha) d\alpha, \end{aligned}$$

implying $a_k = 1/2$. Moreover, $a_k = 1/2$ implies that the right sides of both (A.5) and (A.6) are positive, which then implies μ_1 and μ_0 are both strictly positive; i.e., both limited liability constraints bind. Finally, the first-order conditions for a_1 and a_0 are

$$\begin{aligned} 0 &= -\lambda(\theta_1(1) - \theta_1(1/2))\eta_1(a_1) \\ 0 &= -\lambda(\theta_0(1) - \theta_0(1/2))\eta_0(a_0), \end{aligned}$$

implying $\theta_k(1) = \theta_k(1/2)$. Thus, the cost-minimizing first-best contract is WEPI. Letting $\theta_k(\alpha) = C_k$, the binding limited-liability constraints (10) become

$$U_{1/2} = C_1(\gamma_1(x_1^*(1)) + p)$$

$$U_{1/2} = C_0(\gamma_0(x_0^*(1)) - p).$$

Also, the ex-ante incentive constraint (9) becomes

$$C_1 + C_0 = \frac{b}{p(1-p)(2\bar{\alpha}-1)}.$$

These three equations can be solved for C_1 , C_0 and $U_{1/2}$. To complete the proof, we observe that the solution obtained also satisfies the constraints of the more restrictive, original problem (P). In particular, under a WEPI contract the right side of (LLC') is 0 and the neglected constraint (7) is satisfied because $\gamma_k(x_k^*(1)) + (2k-1)p > 0$. \square

Proof of Corollary 1. Fix any investment x that is strictly monotone; i.e., x_1 is strictly increasing and x_0 is strictly decreasing. (This is true for the first-best investment schedule x^* by Lemma 2.) The proof to Proposition 2 then gives us that the agency cost for implementing x only depends on $(x_1(1), x_0(1))$ and is given by $b + \tilde{U}_{1/2}(x_0(1), x_1(1))$. Taking partial derivatives gives

$$\frac{\partial \tilde{U}_{1/2}}{\partial x_k(1)} = \frac{b\gamma'_k(x_k(1))(\gamma_{1-k}(x_{1-k}(1)) + (2k-1)p)^2}{p(1-p)(2\bar{\alpha}-1)(\gamma_1(x_1(1)) + \gamma_0(x_0(1)))^2},$$

which is strictly positive for $k=1$ and strictly negative for $k=0$ by parts (ii) and (iii) of Lemma 1. Because a marginal change of the strictly monotone investment schedule x in any sufficiently small neighborhood of $\alpha=1$ has a negligible impact on the expected return from investing x , it is always beneficial for the principal to marginally decrease (increase) any strictly increasing (decreasing) investment x_1 (x_0) in a neighborhood around $\alpha=1$. A fortiori, this also holds for the first-best investment x^* , leading to the conclusion that investments in the optimal contract cannot be strictly monotone at least in a neighborhood around $\alpha=1$; i.e., $\underline{\alpha}_k < 1$ for $k \in \{0, 1\}$. \square

Proof of Proposition 3. Consider program (RP). Let λ be the multiplier for the equality constraint (9) and let $\mu_k(\alpha)$ be the multiplier functions for the constraints (10). Then, in accordance with Lemma 5, consider maximizing the following objective function

$$\begin{aligned} \sum_{k \in \{0,1\}} \int_{1/2}^1 & \left[(G_k(x_k(\alpha), \alpha) - cx_k(\alpha) - \kappa - b - U_{1/2}) \varphi_k(\alpha) dH(\alpha) \right. \\ & \left. + \lambda \left[\sum_{k \in \{0,1\}} \int_{1/2}^1 \theta_k(\alpha)(2k-1)p'_k(\alpha)\eta_k(\alpha) d\alpha - b \right] \right. \\ & \left. + \sum_{k \in \{0,1\}} \int_{\underline{\alpha}_k}^1 \mu_k(\alpha) [U_{1/2} - \theta_k(\alpha)[\gamma_k(x_k(\alpha)) + (2k-1)p] d\alpha \right] \end{aligned}$$

First, we show that the solution must be WEPI. To see this, note that because the objective is linear in $\theta_k(\alpha)$, the solution will involve extreme values. Then because $\theta_k(\alpha)$ must be increasing, to honor (11), we augment the above objective by including choice variables a_k such that

$$\theta_k(\alpha) = \begin{cases} \theta_k(1/2), & \text{if } \alpha < a_k \\ \theta_k(1) & \text{if } \alpha \geq a_k \end{cases}$$

Case 1: First suppose $a_k \leq \underline{\alpha}_k$ for $k \in \{0, 1\}$. The first-order condition for $U_{1/2}$ yields

$$1 = \mu_0 + \mu_1, \tag{A.7}$$

where

$$\mu_k \equiv \int_{\underline{\alpha}_k}^1 \mu_k(\alpha) d\alpha.$$

The first-order condition for $\theta_k(1)$ is

$$\lambda \int_{a_k}^1 (2k-1)p'_k(\alpha)\eta_k(\alpha) d\alpha = \int_{\underline{\alpha}_k}^1 \mu_k(\alpha)[\gamma_k(x_k(\alpha)) + (2k-1)p] d\alpha.$$

Together the preceding conditions imply $\mu_k > 0$ and $\lambda > 0$. The first-order condition for a_k is

$$[\theta_k(1) - \theta_k(1/2)]\lambda(2k-1)p'_k(a_k)\eta_k(a_k) = 0.$$

Because $\lambda > 0$, this yields $\theta_k(1) = \theta_k(1/2)$, i.e., the contract is WEPI.

Case 2: Next, suppose $a_k > \underline{\alpha}_k$ for $k \in \{0, 1\}$. By part (iv) of Lemma 5, the relaxed liability constraints bind for all $\alpha \geq \underline{\alpha}_k$. Therefore define $x_k^{1/2}$ and x_k^1 implicitly by

$$\gamma_k(x_k^i) \equiv \frac{U_{1/2}}{\theta_k(i)} - (2k-1)p, \quad i \in \{1/2, 1\}, k \in \{0, 1\}.$$

Now, the first-order condition for $U_{1/2}$ again yields (A.7). The first-order condition for $\theta_k(1)$ is

$$\lambda \int_{a_k}^1 (2k-1)p'_k(\alpha)\eta_k(\alpha) d\alpha = \int_{a_k}^1 \mu_k(\alpha)[\gamma_k(x_k^1) + (2k-1)p] d\alpha \quad (\text{A.8})$$

and the first-order condition for $\theta_k(1/2)$ is

$$\lambda \int_{1/2}^{a_k} (2k-1)p'_k(\alpha)\eta_k(\alpha) d\alpha = \int_{\underline{\alpha}_k}^{a_k} \mu_k(\alpha)[\gamma_k(x_k^{1/2}) + (2k-1)p] d\alpha.$$

Together the preceding expressions, along with (A.7), imply $\mu_k > 0$ and $\lambda > 0$. The first-order condition for a_k can be written

$$\begin{aligned} & [\theta_k(1) - \theta_k(1/2)]\lambda(2k-1)p'_k(a_k)\eta_k(a_k) \\ &= \mu_k(a_k) \left([U_{1/2} - \theta_k(1/2)(\gamma_k(x_k^{1/2}) + (2k-1)p)] - [U_{1/2} - \theta_k(1)(\gamma_k(x_k^1) + (2k-1)p)] \right). \end{aligned}$$

Both terms in square brackets on the right side of this expression are 0 by the above definition of x_k^i . Therefore, $\lambda > 0$ again implies $\theta_k(1) = \theta_k(1/2)$.

Case 3: The optimal contract is shown to be WEPI for the mixed cases $a_k \leq \underline{\alpha}_k$ and $a_{1-k} > \underline{\alpha}_{1-k}$ by analogous arguments.

Knowing that the contract is WEPI, (9) can be written

$$C_0 + C_1 = \frac{b}{p(1-p)(2\bar{\alpha}-1)}.$$

and (10) can be written

$$U_{1/2} = C_k(\gamma_k(x_k^\dagger) + (2k-1)p), \quad k \in \{0, 1\}.$$

These three equations can be solved for $U_{1/2}(x_0^\dagger, x_1^\dagger)$ and $C_k(x_0^\dagger, x_1^\dagger)$. Recalling that $x_k^\dagger = x_k^*(\underline{\alpha}_k)$, the only things left to pin down are $\underline{\alpha}_0$ and $\underline{\alpha}_1$.

For $\alpha < \underline{\alpha}_k$, the first-order conditions for x_k yield

$$g_k(x_k(\alpha), \alpha) = c.$$

So investments are first-best in accordance with part (i) of Lemma 5. For any $\alpha \geq \underline{\alpha}_k$ the first-order condition for $x_k(\alpha)$ is

$$(g_k(x_k^\dagger, \alpha) - c)\varphi_k(\alpha)h(\alpha) = \mu_k(\alpha)C_k(x_0^\dagger, x_1^\dagger)\gamma_k'(x_k^\dagger)$$

or

$$\mu_k(\alpha) = \frac{(g_k(x_k^\dagger, \alpha) - c)\varphi_k(\alpha)h(\alpha)}{C_k(x_0^\dagger, x_1^\dagger)\gamma_k'(x_k^\dagger)}.$$

Integrating this from $\underline{\alpha}_k$ to 1 yields

$$\mathbb{E}[g_{\hat{k}}(x_{\hat{k}}^\dagger, \alpha) - c | \hat{k} = k, \alpha \geq \underline{\alpha}_k](1 - H(\underline{\alpha}_k)) = \mu_k C_k(x_0^\dagger, x_1^\dagger)\gamma_k'(x_k^\dagger). \quad (\text{A.9})$$

The first-order condition for C_k gives

$$\mu_k(\gamma_k(x_k^\dagger) + (2k - 1)p) = \lambda \int_{1/2}^1 (2k - 1)p_k'(\alpha)\eta_k(\alpha) d\alpha, \quad k \in \{0, 1\}$$

Recalling the definition of $\eta_k(\alpha)$ from (12) and reversing the order of integration on the right side of the above expression gives

$$\mu_k(\gamma_k(x_k^\dagger) + (2k - 1)p) = \lambda p(1 - p)(2\bar{\alpha} - 1), \quad k \in \{0, 1\}.$$

Therefore

$$\mu_0(\gamma_0(x_0^\dagger) - p) = \mu_1(\gamma_1(x_1^\dagger) + p).$$

Combining this with (A.7) gives

$$\frac{b\mu_k}{p(1 - p)(2\bar{\alpha} - 1)} = C_k(x_0^\dagger, x_1^\dagger).$$

Combining this expression with (A.9) finally yields

$$\begin{aligned} \mathbb{E}[g_{\hat{k}}(x_{\hat{k}}^\dagger, \alpha) - c | \hat{k} = k, \alpha \geq \underline{\alpha}_k](1 - H(\underline{\alpha}_k)) \\ = \frac{b\gamma_k'(x_k^\dagger)}{p(1 - p)(2\bar{\alpha} - 1)} \left(\frac{\gamma_{1-k}(x_{1-k}^\dagger) + (2k - 1)p}{\gamma_0(x_0^\dagger) + \gamma_1(x_1^\dagger)} \right)^2, \quad k \in \{0, 1\}. \end{aligned}$$

This expression implicitly defines $\underline{\alpha}_1^\dagger$ and $\underline{\alpha}_0^\dagger$. Moreover, direct calculation shows that the right side is equal to $\partial\tilde{U}_{1/2}/\partial x_k^\dagger$. Also, note that as b vanishes, the right side becomes 0, but the left side equals 0 iff $\underline{\alpha}_k^\dagger = 1$.

To complete the proof, we observe that the solution obtained also satisfies the constraints of the more restrictive, original problem (P). In particular, under a WEPI contract the right side of (LLC') is 0 and the neglected constraint (7) is satisfied because $\gamma_k(x_k^\dagger) + (2k - 1)p > 0$. \square

Proof of Proposition 4. Fix an implementable thumbs-up-or-down contract

$$[\bar{x}_0, \bar{x}_1, \bar{C}_0, \bar{C}_1, \tilde{U}_{1/2}(\bar{x}_0, \bar{x}_1)]$$

with $x_0^*(1) \leq \bar{x}_0 < x_{1/2}^* < \bar{x}_1 \leq x_1^*(1)$. Define $\underline{\alpha}'_k$ implicitly by $x_k^*(\underline{\alpha}'_k) = \bar{x}_k$ for $k \in \{0, 1\}$ and consider the alternative contract in which investments are first-best for $\alpha \in [1/2, \underline{\alpha}'_k)$ and equal \bar{x}_k for $\alpha \in [\underline{\alpha}'_k, 1]$. This contract generates the same agency cost $b + \tilde{U}_{1/2}(\bar{x}_0, \bar{x}_1)$ and is WEPI with the same constants \bar{C}_k . Also, the alternative contract satisfies (10) with equality for $\alpha \geq \underline{\alpha}'_k$ because this is true for the thumbs-up-or-down contract. Moreover, for $\alpha \in [1/2, \underline{\alpha}'_k)$ we have

$$\tilde{U}_{1/2}(\bar{x}_0, \bar{x}_1) - \bar{C}_k[\gamma_k(x_k^*(\alpha)) + (2k - 1)p] > 0$$

because $\gamma'_k(x_k^*(\alpha)) \frac{dx_k^*}{d\alpha} > 0$ by Lemmas 1 and 2. Thus, this contract satisfies all the constraints, generates the same agency cost, and provides a strictly higher payoff to the principal through the improved investment schedules. \square

References

- ARMSTRONG, M. (1996): “Multiproduct Nonlinear Pricing,” *Econometrica: Journal of the Econometric Society*, 51–75.
- ARROW, K. J. (1987): “The Value of and Demand for Information,” in *Decision and organization: A volume in honor of Jacob Marschak*, University of Minnesota Press, 131–139, 2nd ed.
- BAKER, M. (2016): “1,500 Scientists Lift the Lid on Reproducibility,” *Nature News*, 533, 452.
- BENJAMIN, D. J., J. O. BERGER, M. JOHANNESSON, B. A. NOSEK, E.-J. WAGENMAKERS, R. BERK, K. A. BOLLEN, B. BREMBS, L. BROWN, C. CAMERER, ET AL. (2018): “Redefine Statistical Significance,” *Nature Human Behaviour*, 2, 6.
- BIAIS, B. AND C. CASAMATTA (1999): “Optimal Leverage and Aggregate Investment,” *The Journal of Finance*, 54, 1291–1323.
- CARROLL, G. (2015): “Robustness and Linear Contracts,” *American Economic Review*, 105, 536–63.
- (2017): “Robust Incentives for Information Acquisition,” Working Paper.
- CRAWFORD, V. AND J. SOBEL (1982): “Strategic Information Transmission,” *Econometrica*, 50, 1431–1445.
- DANA, J. D. AND K. E. SPIER (1993): “Expertise and Contingent Fees: The Role of Asymmetric Information in Attorney Compensation,” *Journal of Law, Economics, & Organization*, 9, 349–367.
- DE LONG, J. B. AND K. LANG (1992): “Are all Economic Hypotheses False?” *Journal of Political Economy*, 100, 1257–1272.
- DEMSKY, J. S. AND D. E. M. SAPPINGTON (1987): “Delegated Expertise,” *The Journal of Accounting Research*, 25, 68–89.
- DILLENBERGER, D., J. S. LLERAS, P. SADOWSKI, AND N. TAKEOKA (2014): “A Theory of Subjective Learning,” *Journal of Economic Theory*, 153, 287–312.

- ESŐ, P. AND B. SZENTES (2007): “The Price of Advice,” *The RAND Journal of Economics*, 38, 863–880.
- ETZ, A. AND J. VANDEKERCKHOVE (2016): “A Bayesian Perspective on the Reproducibility Project: Psychology,” *PloS one*, 11, e0149794.
- FANELLI, D. (2009): “How many Scientists Fabricate and Falsify Research? A Systematic Review and Meta-Analysis of Survey Data,” *PloS one*, 4, e5738.
- FEESS, E. AND M. WALZL (2004): “Delegated Expertise – When are Good Projects Bad News?” *Economics Letters*, 82, 77–82.
- GERARDI, D. AND L. MAESTRI (2012): “A Principal-Agent Model of Sequential Testing,” *Theoretical Economics*, 7, 425–463.
- GERBER, A., N. MALHOTRA, ET AL. (2008): “Do Statistical Reporting Standards Affect what is Published? Publication Bias in Two Leading Political Science Journals,” *Quarterly Journal of Political Science*, 3, 313–326.
- GLASZIOU, P., D. G. ALTMAN, P. BOSSUYT, I. BOUTRON, M. CLARKE, S. JULIOUS, S. MICHIE, D. MOHER, AND E. WAGER (2014): “Reducing Waste from Incomplete or Unusable Reports of Biomedical Research,” *The Lancet*, 383, 267–276.
- GROMB, D. AND D. MARTIMORT (2007): “Collusion and the Organization of Delegated Expertise,” *Journal of Economic Theory*, 137, 271–299.
- HÄFNER, S. AND C. R. TAYLOR (2018): “Advice Is Cheap: Information Is Not!” Working Paper, Available at SSRN: <https://ssrn.com/abstract=3052147>.
- HALAC, M., N. KARTIK, AND Q. LIU (2016): “Optimal Contracts for Experimentation,” *The Review of Economic Studies*, 83, 1040–1091.
- HENRY, E. AND M. OTTAVIANI (2017): “Research and the Approval Process: the Organization of Persuasion,” CEPR Discussion Paper No. DP11939.
- HÖLMSTROM, B. (1979): “Moral Hazard and Observability,” *The Bell Journal of Economics*, 10, 74–91.
- HÖRNER, J. AND L. SAMUELSON (2013): “Incentives for Experimenting Agents,” *The RAND Journal of Economics*, 632–663.
- HOSSAIN, T. AND R. OKUI (2013): “The Binarized Scoring Rule,” *Review of Economic Studies*, 80, 984–1001.
- INDERST, R. AND M. OTTAVIANI (2009): “Misselling Through Agents,” *American Economic Review*, 99, 883–908.
- IOANNIDIS, J. P. (2005): “Why Most Published Research Findings are False,” *PLoS medicine*, 2, e124.
- IOANNIDIS, J. P., M. R. MUNAFO, P. FUSAR-POLI, B. A. NOSEK, AND S. P. DAVID (2014): “Publication and other Reporting Biases in Cognitive Sciences: Detection, Prevalence, and Prevention,” *Trends in Cognitive Sciences*, 18, 235–241.

- KLEIN, N. (2016): “The Importance of Being Honest,” *Theoretical Economics*, 11, 773–811.
- KRISHNA, V. AND J. MORGAN (2008): “Contracting for Information under Imperfect Commitment,” *The RAND Journal of Economics*, 39, 905–925.
- MALCOMSON, J. M. (2009): “Principal and Expert Agent,” *The BE Journal of Theoretical Economics*, 9.
- MILGROM, P. AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- RUBINFELD, D. L. AND S. SCOTCHMER (1993): “Contingent Fees for Attorneys: An Economic Analysis,” *The RAND Journal of Economics*, 24, 343–356.
- SAVAGE, L. J. (1971): “Elicitation of Personal Probabilities and Expectations,” *Journal of the American Statistical Association*, 66, 783–801.
- SCHLAG, K. H., J. TREMEWAN, AND J. J. VAN DER WEELE (2015): “A Penny for your Thoughts: A Survey of Methods for Eliciting Beliefs,” *Experimental Economics*, 18, 457–490.
- SHAKED, M. AND J. G. SHANTHIKUMAR (2007): *Stochastic Orders*, Springer Science & Business Media.
- SIMMONS, J. P., L. D. NELSON, AND U. SIMONSOHN (2011): “False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant,” *Psychological Science*, 22, 1359–1366.
- SZALAY, D. (2009): “Contracts with Endogenous Information,” *Games and Economic Behavior*, 65, 586–625.
- TRAFIMOW, D. AND M. MARKS (2015): “Editorial,” *Basic and Applied Social Psychology*, 37, 1–2.
- VINKERS, C. H., J. K. TIJDINK, AND W. M. OTTE (2015): “Use of Positive and Negative Words in Scientific PubMed aAbstracts between 1974 and 2014: Retrospective Analysis,” *Bmj*, 351, h6467.
- WASSERSTEIN, R. L., N. A. LAZAR, ET AL. (2016): “The ASA’s Statement on p-values: Context, Process, and Purpose,” *The American Statistician*, 70, 129–133.