Investing in Failure: Belief Manipulation and Strategic Experimentation[1]

Peter Buisseret

Abstract

A Principal hires an Agent to experiment with a technology of unknown quality. The Principal invests funds and the Agent supplies costly, hidden effort. Higher investment by the Principal raises the probability of success when the Agent works and the project is good, but increases the Agent's dynamic information rent from shirking. I show that the Principal may over-invest relative to the first best. This commits the Principal to sufficiently pessimistic beliefs upon initial failure that she discontinues the project, denying the Agent rents from the repeated interaction. When the Principal can write a long-term contract, over-investment persists and inefficiency may increase, relative to a setting in which only short-term contracts can be written.

Keywords— strategic experimentation, learning, agency

[1] *Author affiliation and address*: Department of Economics, Warwick University, Coventry, CV4 7AL, UK. *Email*:p.buisseret@warwick.ac.uk. *Telephone*: +44 (0)24 7652 8248. *Fax*: +44(0)2476 523032.

Acknowledgements: I am indebted to Steven Callander for his generous and extremely helpful discussion of an early draft of this project. I also thank Avidit Acharya, Pablo Beker, Dan Bernhardt, Marina Halac, Niall Hughes, Navin Kartik, Motty Perry and Carlo Prato for suggestions and encouragement. Wentao Fu provided outstanding research assistance. All errors and omissions are solely my responsibility.

1. Introduction

A fundamental problem of experimentation is knowing when to quit. I explore how this decision is affected by the presence of an Agent to whom a Principal must delegate a role in the experiment. Previous work highlights *under-investment* and early abandonment of productive projects (Bergemann and Hege (2005) or Hörner and Samuelson (2013)) as a consequence of experimentation in a Principal-Agent setting. This paper provides a rationale for why *over-investment* might occur relative to the efficient benchmark, whilst still inducing early project abandonment.

In my model, the Principal invests funds in a project which is either *good*, with probability v_0, or *bad*, with probability $1 - v_0$. At each date, she chooses investment and offers the Agent a share of the one-off surplus that may be generated from the project. The role of the Agent is to exert costly, hidden effort. If the project is good and the Agent works, the probability of success is proportional to the funds invested by the Principal at that date; otherwise, it fails for sure.

This technology induces the following agency problem: suppose that the Principal observes an initial failure and believes that the Agent worked on the project. Then, her posterior belief about its quality falls to $\hat{v} < v_0$, and the difference $v_0 - \hat{v}$ strictly increases in the initial investment. If the Principal wishes to continue for another period she will offer a contract that holds the Agent-type who shares her belief to her participation constraint. But suppose that, in reality, the Agent shirked, so that initial failure was inevitable. Then, the Agent's belief about the project is in fact v_0. But the compensation package designed for the Agent-type $\hat{v}$ gives rent to an Agent-type $v_0 > \hat{v}$. This creates a moral hazard problem which stems from the private information about the project quality created by the Agent's deviation.

The Principal faces a trade-off: choosing higher initial levels of investment at the outset of the project increases the probability of success, conditional on the project being good and the Agent working hard. But higher initial levels of investment also increase the divergence between the Principal's prior and posterior beliefs after observing failure, since failure is more informative when initial investment is larger. This, in turn, raises the continuation value to the Agent from shirking, since her information rent is increasing in the divergence of v_0 and $\hat{v}$.

Relative to the first-best, I show that the Principal may choose an inefficiently high level of initial investment. Given her belief that the Agent worked, the informativeness of failure under such a large initial investment leads her to discontinue the project, eliminating the Agent's opportunity to acquire dynamic information rent. When the Principal can only write a short-term contract (i.e., within-period commitment) this choice of investment partly serves as a *belief commitment* device for the Principal to ensure that she prefers to quit after early failure. Nonetheless, over-investment also occurs when the Principal can commit to a long-term contract (i.e. between-period commitment); in fact, there may be instances in which the Principal invests more under long-term commitment than *either* the first-best contract *or* short-term commitment. Moreover, the long-term contract may increase inefficiency since it allows the Principal to credibly destroy even more of the surplus associated with the two-date investment solution in order to deny rents to the Agent.

The agency problem that I study builds on Bhaskar (2012).[2] In his model, an Agent chooses privately observed effort, and the Principal's problem is solely to provide pecuniary incentives to induce effort. Bhaskar (2012) characterizes the Principal's optimal contract when the Agent can create private information in the same way as the present model, in a setting where the project is always profitable for any belief. I show that when there are beliefs about project quality for which the Principal prefers to abandon the project, and she is able to manipulate the distribution of the public signal - and thus the evolution of her posterior belief upon observing failure - she may do so as a means to soften the Agent's incentive constraint. The interpretation of this manipulation as a decision about how much to invest connects the model to a large literature on strategic experimentation in which the Principal chooses a level of funding and delegates the funds to the Agent (Bergemann and Hege (2005) or Hörner and Samuelson (2013)). These papers also characterize investment distortions as a consequence of the Agency problem, however the funding inefficiency in these models manifests as *under-investment*, or a delay in the release of funds. Inefficient over-investment is ruled out in Bergemann and Hege (2005) because the Principal's cost of funds is linear; in Hörner and Samuelson (2013) it is ruled out because the Principal's investment decision has binary support. As a result, early quitting in existing models arises despite the fact that it would be socially optimal to continue experimenting for additional periods. In my model, by contrast, the Principal's decision not to experiment in the second period is socially optimal *given her investment decision in the first period*: the initial choice of investment effectively destroys tomorrow's surplus in order to allow the Principal to withhold rents from the Agent, today.

[2]See also Bhaskar (2014). There is a large and growing literature on strategic experimentation: in addition to work discussed explicitly, recent contributions include Halac, Kartik and Liu (2013), who study a model with adverse selection and moral hazard, and Manso (2011).

2. The Model

I consider a two-date economy, with dates 0 and 1. There is a Principal and an Agent. The timing unfolds as follows. At the start of date zero, Nature selects a project quality $\omega \in \{0, 1\}$, according to $\Pr(\omega = 1) = v_0 \in (0, 1)$ which persists across both dates. This choice is not observed by the players. Next, the Principal chooses $\lambda_t \in [0, 1]$ at cost $\tau(\lambda_t)$ which is strictly increasing, strictly convex, twice differentiable and satisfies $\tau(0) = 0$. This investment choice is observed by both players. She then offers the Agent a contract, described in greater detail, below. The Agent subsequently chooses a distribution over effort levels $e_t \in \{0, 1\}$ at cost $c(e_t)$ satisfying $c(0) = 0$ and $c(1) = c > 0$. Only the Agent observes this choice. A public signal $s_t \in \{s_t^L, s_t^H\}$ is then realized according to the distribution:

$$\Pr(s_t^H | \lambda_t, e_t, \omega) = \lambda_t e_t \omega. \tag{1}$$

The public signal s_t^H is interpreted as a 'success' and s_t^L a 'failure' at date t. If a success is realized, a one-off surplus of W is generated and the game ends at date zero.[3] Otherwise, no surplus is generated and, at date zero, the game proceeds to date one. The timing of the game at date one is the same as at date zero, beginning with the Principal's investment decision. The common discount factor is $\beta \in (0, 1)$.

Contracts and Histories

A date one *history* is $h_1 = (\lambda_0, s_0, e_0)$. A date one *public history* is $h_1^p = (\lambda_0, s_0)$. The set of all date one public histories is H_1^p, and the set of all date one histories is H_1. The date zero history is the null history. Let v_1^p denote the belief of the Principal at the start of date one that the project is good. The Agent's belief that the project is good at the start of date one is v_1. The strategy of the Agent at date t is $\sigma_t \in [0, 1]$, a distribution over effort choices, $e_t \in \{0, 1\}$.

The Principal's contract with the Agent specifies a bonus to be paid after a success at each date.[4] I initially assume that the Principal has short-term (intra-period) commitment. I assume that all transfers from the Principal must be weakly positive, and the Agent's outside option is zero

[3]The results are not affected by an alternative assumption that the project is implemented again at date one in the event of success at date zero. However, the assumption that experimentation ends after a breakthrough is more common in the literature: see, for example, Bergemann and Hege (2005) and Hörner and Samuelson (2013).

[4]This is without loss of generality.

at each date. The static surplus at date t when the Agent chooses high effort and the Principal invests λ_t is:

$$f(\lambda_t; v_t^p) \equiv v_t^p \lambda_t W - \tau(\lambda_t) - c. \tag{2}$$

I let $\lambda_t^*(v_t^p)$ denote the maximizer of $f(\lambda_t; v_t^p)$ and $S(v_1^p) = \max\{0, f(\lambda_1^*(v_1^p); v_1^p)\}$ denote the surplus in the continuation game beginning at date one, after failure at date zero, in the first-best. I impose restrictions on primitives which imply (a) $\lambda_t^*(1) < 1$ and (b) $f(\lambda_t^*(1); 1) > 0$.[5] The first restriction ensures that there is no prior belief for which the Principal's optimal one-shot investment decision would render failure a probability zero event, conditional on the Agent choosing high effort with probability one. The second restriction rules out a trivial setting in which the Principal would never wish to make positive investment at either date, in the first-best, for even the most optimistic belief about project quality.

I do not restrict either the Principal or the Agent to employ pure strategies. In fact, in the setting with unobserved effort, equilibrium requires that the Agent randomize over effort levels at date zero, at the very least off the equilibrium path, and possibly also on the path.[6] Moreover, the setting with unobserved effort may yield a multiplicity of mixed strategy equilibria in the continuation game after the Principal has announced an investment choice and bonus at date zero. So, my focus is on establishing properties of the most efficient equilibrium in the setting with unobserved effort (the "second best"). Over-investment also arises in the equilibrium which maximizes the welfare of the Principal, where these equilibria do not coincide.

The assumption that the probability of success increases in funds seems natural, in practice. So long as the project is of high quality, projects which are better-funded are more likely to succeed than those which are starved of resources. That higher initial funds make failure more informative is a consequence of this assumption, but is also natural in its own right. For example, one could interpret λ as the scale of the experiment: conducting a pilot on a larger sample or running more trials provides more opportunities to learn about its quality.[7]

[5] The conditions are (a) $\tau_\lambda(1) > W$ and (b) $W\tau_\lambda^{-1}(W) - \tau(\tau_\lambda^{-1}(W)) - c > 0$, where $\tau_\lambda(\cdot)$ is the derivative of τ with respect to λ.

[6] This issue also arises in Hörner and Samuelson (2013).

[7] Another model in which the Principal can manipulate the scale of the project is Biais et al. (2010).

3. First-Best Experimentation

I begin with an efficient benchmark in which the Principal observes the Agent's effort choice at each date. In this case, the public history and the history coincide. I start by characterizing properties of the Principal's investment at date one, as a function of her date zero investment decision. Since the results of Lemma 1 will be relevant to the setting in which the Agent's effort is unobserved, I characterize properties of the first-best with reference to the public history, even though this coincides with the (true) history when effort is observed. So, let $v_1^p(\lambda_0, \sigma_0)$ denote the Principal's belief that the project is high quality, given that a failure occurred at date zero after investment choice λ_0 and the Agent is conjectured to choose high effort with probability σ_0. That is:

$$v_1^p(\lambda_0, \sigma_0) = \frac{v_0(1 - \lambda_0\sigma_0)}{1 - v_0\lambda_0\sigma_0}. \tag{3}$$

My first Lemma characterizes properties of the first-best at date one.

Lemma 1. In the first-best, the Principal's date one investment satisfies the following properties:

i. after any two date one public histories with respective date zero investment decisions $\lambda_0 < 1$ and $\bar{\lambda}_0 < 1$, and in which failure occurred at date zero, $\bar{\lambda}_0 > \lambda_0$ implies $\lambda_1^*(v_1^p(\bar{\lambda}_0, 1)) \leq \lambda_1^*(v_1^p(\lambda_0, 1))$; and,

ii. there exists a unique threshold belief $\underline{v} \in (0, 1)$ such that if $v_1^p \leq \underline{v}$, $f(\lambda_1^*(v_1^p), v_1^p) \leq 0$ and for all $v_1^p > \underline{v}$, $f(\lambda_1^*(v_1^p), v_1^p) > 0$.

If and only if $v_1^p \leq \underline{v}$, the Principal chooses $\lambda_1(v_1^p) = 0$. If $v_1^p > \underline{v}$ the Principal chooses $\lambda_1^*(v_1^p)$, and offer the Agent a bonus $b_1(v_1^p, \lambda_1) = \frac{c}{\lambda_1 v_1^p}$ in the event of success.

Point (i) states that the more the Principal invested at date zero, so long as the Agent is believed to have exerted effort, initial failure becomes increasingly bad news about project quality. And more pessimism induces a lower flow of funds at date one. Point (ii) states that if the Principal becomes sufficiently pessimistic after initial failure, she may choose to make no subsequent investment at date one, and instead abandon the project. For the rest of the paper, I restrict attention to $v_0 > \underline{v}$, so that there is a rationale for strictly positive investment in at least one of the two dates. This implies that the Agent chooses high effort in the first-best at date zero, i.e., $\sigma_0 = 1$.

In the first-best, the Principal observes the Agent's date zero effort. This has two implications. First, the Principal's date one beliefs about project quality coincide with the Agent's beliefs.

Second, the Principal's beliefs *about* the Agent's beliefs are correct. This allows the Principal to hold the Agent to her participation constraint, and extract all of the date one surplus. This will also be a feature of the second-best, so long as the Agent plays a pure strategy at date zero.

We have obtained a threshold value of the posterior ($\underline{v}$) below which no further investment occurs at date one. It is useful to characterize for each date zero prior belief v_0 the level of investment at date zero which will induce a posterior $\underline{v}$ upon initial failure. This level of investment is:

$$\hat{\lambda}_0(v_0) = \frac{v_0 - \underline{v}}{v_0(1 - \underline{v})}. \tag{4}$$

A consequence of Lemma 1 is that if the Principal invests below this threshold at date zero, she will wish to continue with positive investment at date one. If, instead, she invests above this threshold at date zero, early failure will lead her to abandon the project and make no further investment at date one.

I briefly provide some intuition for the shape of $\hat{\lambda}_0(v_0)$. First, $\hat{\lambda}_0(v_0)$ is increasing and strictly concave in v_0. Second, $\hat{\lambda}_0(\underline{v}) = 0$ and $\hat{\lambda}_0(1) = 1$. That is: at the prior belief which makes the one-off investment solution break-even ($\underline{v}$), any strictly positive investment choice at date zero will induce strictly greater pessimism, conditional on initial failure. This will induce the Principal to quit after date zero. As initial optimism about the project rises (i.e., v_0 rises), the degree of initial investment which is needed in order induce sufficient pessimism for the Principal to quit after failure increases.

Next, I give some intuition for the relationship between $\hat{\lambda}_0(v_0)$, and the one-shot investment solution, $\lambda_0^*(v_0)$. Recall $\hat{\lambda}_0(\underline{v}) = 0$ and $\hat{\lambda}_0(1) = 1$. By construction, $f(\lambda_0^*(\underline{v}); 0) = 0$, so $\lambda_0^*(\underline{v}) > 0$. Finally, I have imposed restrictions such that $\lambda_0^*(1) < 1$. Continuity implies that there is at least one prior belief, $v_0 \in (\underline{v}, 1)$, at which $\lambda_0^*(v_0)$ and $\hat{\lambda}_0(v_0)$ intersect.[8] Moreover, on the first occasion where they intersect, $\lambda_0^*(v_0)$ intersects $\hat{\lambda}_0(v_0)$ from above.

The crux of the subsequent analysis will be to understand how the Principal's choice of investment above or below the threshold $\hat{\lambda}_0(v_0)$ is distorted under unobservable effort, relative to the benchmark of complete information.

I now turn to the Principal's date zero investment decision. Since we have assumed that the Principal's prior belief is not so unfavorable as to merit no experimentation (i.e., $v_0 > \underline{v}$),

[8]For an arbitrary cost function, I cannot rule out the possibility of more than one intersection, but this is not important for any of my results.

the first-best implies a strictly positive investment at date zero.[9] The Principal's optimal date zero investment may take two forms: either she invests above the threshold $\hat{\lambda}_0(v_0)$ and makes no subsequent investment upon initial failure, or she invests below the threshold $\hat{\lambda}_0(v_0)$ and makes a positive investment at date one upon initial failure.

Lemma 2. In the first-best, the Principal's investment decision solves:

$$S(v_0) = \max_{\lambda_0 \in [0,1]} \left(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta S(v_1^p(\lambda_0, 1)) \right) \tag{5}$$

For specific functional forms of the investment cost, $\tau(\lambda)$, a stronger characterization can be obtained. I adopt the following parameterized example throughout the paper.[10]

Example 1. Set $\tau(\lambda) = \frac{1}{2}\lambda^2$, $\beta = \frac{3}{4}$, $W = 1$ and $c = \frac{1}{15}$. Then, $\underline{v} = \sqrt{\frac{2}{15}}$. In the first-best, the Principal invests at date zero only, if and only if $v_0 \leq .524$. Otherwise, she invests both at date zero and also at date one in the event of initial failure.

In this example, when the prior belief about project quality (v_0) is low, it is relatively cheap to invest at or above the rate such that failure is sufficiently informative to induce discontinuation of the project. As the Principal's initial optimism rises, so does the threshold $\hat{\lambda}_0(v_0)$. This implies that the cost of investing in sufficiently informative failure is also rising. Since the cost of investing is convex, the Principal faces a trade-off between the cost of investing more, and acquiring more information from the initial experiment. There exists a critical belief v^* at which point it is too costly to invest in a sufficiently informative experiment to induce quitting after initial failure. At that point, the Principal spreads the cost of investing over both dates.

4. Second-Best Experimentation

I now characterize properties of the most efficient equilibrium in the setting where the Agent's effort is unobserved. I begin with date one. At that date, $v_1^p(\lambda_0, \sigma_0)$ denotes the Principal's belief that the project is good after a date zero failure, as a function of the investment at date zero and the Principal's conjecture that the Agent chose high effort with probability $\sigma_0 \in [0, 1]$. The

[9] Given $v_0 > \underline{v}$, it is easy to show that in the first-best, the Principal will always choose $\lambda_0 > 0$. This is not guaranteed when effort is unobserved.

[10] In this, and all other examples, all approximations are to three decimal places. Note that a closed-form solutions cannot be obtained, even for the simplest case of quadratic costs and two dates.

Agent's belief that the project is good is $v_1(\lambda_0, e_0)$, which is a function of both the investment at date zero and the Agent's realized choice of effort, $e_0 \in \{0,1\}$. Notice that $v_1^p(\lambda_0, \sigma_0)$ identifies both the Principal's belief about project quality, as well as her distribution over the Agent's belief: she believes that the Agent's belief (i.e., type) is either $v_1(\lambda_0, 1)$ or $v_1(\lambda_0, 0)$.

Suppose that the Principal's date zero contract induces the Agent to choose high effort with probability 1 (i.e., $\sigma_0 = 1$). Then, the Principal's posterior belief about both the project is $v_1^p(\lambda_0, 1)$, and her belief distribution over the Agent's belief places probability one on the Agent sharing her own belief, i.e., she believes that the Agent's belief is $v_1 = v_1(\lambda_0, 1)$ with probability one. Since the Principal's beliefs (on the equilibrium path) must be correct, her investment and compensation strategies satisfy the properties in Lemma 1. That is, the Principal's strategy at date one replicates her strategy in the first-best.

Recall that $\hat{\lambda}_0(v_0)$ is a threshold such that if the Principal invests more than $\hat{\lambda}_0(v_0)$ at date zero *and* believes that the Agent worked with probability one, she prefers to quit at date one. Suppose, then, that the Principal initially chooses $\lambda_0 \geq \hat{\lambda}_0(v_0)$ and offers a bonus conditioned on success, $b_0 \geq 0$. In that case, if she believes that the Agent will work hard, it is indeed a best response for the Agent to work if:

$$v_0 \lambda_0 b_0 - c \geq 0, \tag{6}$$

since, if the Agent deviates, the Principal will observe only the failure and terminate the project. This implies that the Agent can be induced to work hard while being held to her participation constraint. More generally, we have the following result.

Lemma 3. If the Principal's payoff in the first-best is $f(\lambda_0^*(v_0); v_0)$, there exists an equilibrium of the continuation game after the Principal offers a contract $\lambda_0 = \lambda_0^*(v_0)$ and $b_0 = \frac{c}{v_0 \lambda_0^*(v_0)}$ which yields the Principal the same payoff, when she cannot observe the Agent's effort.

The crucial observation is that when the first-best investment calls on the Principal to quit after initial failure, she can achieve the associated payoff under the second-best, even when she does not directly observe the Agent's effort. The reason is that by investing above the threshold $\hat{\lambda}_0(v_0)$ and believing that the Agent worked hard, the Principal commits her posterior belief about project quality to be sufficiently unfavorable that she strictly prefers to quit after failure. So, if the Agent had shirked, the Principal would observe only that the project failed; believing that the Agent worked hard, the Principal simply quits the project and the Agent receives no further compensation. This implies that the Agent can be induced to choose full effort at date zero without the provision of any rent: the Principal fully appropriates the surplus from her investment.

Suppose, by contrast, that the Principal chooses an investment which—if she believed that the Agent worked hard at date one—would yield a positive continuation payoff from experimenting again at date one after an initial failure. That is, suppose that the Principal chooses $\lambda_0 < \hat{\lambda}_0(v_0)$. At date one, the Principal believes that the Agent shares her own belief about project quality—$v_1^p(\lambda_0, 1)$—with probability one. She therefore satisfies the Agent's participation constraint at date one. The Agent's date one continuation payoff can be written $V(v_1(\lambda_0, e_0), v_1^p(\lambda_0, 1))$, as a function of the Agent's true belief, $v_1(\lambda_0, e_0)$, and the Principal's belief $v_1^p(\lambda_0, \sigma_0)$, where $\sigma_0 = 1$:

$$V(v_1(\lambda_0, e_0), v_1^p(\lambda_0, 1)) = \max\left\{ c\left(\frac{v_1(\lambda_0, e_0)}{v_1^p(\lambda_0, 1)} - 1\right), 0\right\}. \tag{7}$$

Under the presumption that the Agent worked at date zero, beliefs coincide and the Agent's continuation payoff at date one is zero. Suppose, however, that the Agent initially deviated from choosing high effort with probability one to instead choosing no effort, unbeknownst to the Principal. Then, the Agent's true belief is $v_1(\lambda_0, 0) = v_0$, and:

$$V(v_0, v_1^p(\lambda_0, 1)) = c\frac{\lambda_0}{1-\lambda_0}(1 - v_0), \tag{8}$$

which is strictly increasing in λ_0: more initial investment by the Principal induces ever greater pessimism in the light of initial failure. Therefore, so long as the Principal wishes to continue experimenting at date one (i.e., $v_1^p > \underline{v}$) greater pessimism induces a larger bonus in order to satisfy the (presumed) v_1^p-type Agent's participation constraint. This generates a powerful temptation on the part of the Agent to shirk at date zero. If, on the other hand, $\lambda \geq \hat{\lambda}_0(v_0)$, the Agent's continuation payoff from shirking is zero, since the Principal quits. Thus, the Agent's value from shirking is non-monotonic in the initial investment.

I now turn to the initial date. Suppose that the Principal's choice of investment satisfies $\lambda_0 < \hat{\lambda}_0(v_0)$, so that if she believes the Agent worked and observed a failure, she would not be too pessimistic and therefore choose positive experimentation at date one. In order to induce the Agent to choose high effort with probability one, her bonus offer must satisfy:

$$v_0\lambda_0 b_0 - c \geq \beta V\left(v_0, \frac{v_0(1-\lambda_0)}{1-v_0\lambda_0}\right). \tag{9}$$

The Principal must compensate the Agent for her instantaneous effort. But she must also induce the Agent not to shirk in order to manipulate the Principal's beliefs, and thereby capture information rent at date one. Substituting from (8) yields the constraint:

$$b_0 \geq \frac{c}{\lambda_0}\left(\frac{1}{v_0} + \beta\left(\frac{1}{v_1^p(\lambda_0, 1)} - \frac{1}{v_0}\right)\right) \equiv b_H(v_0, \lambda_0). \tag{10}$$

The bonus contains a convex combination of the inverse of two beliefs: the prior (v_0) and the posterior conditioned on the public history in the event of failure (v_1^p). The convex weight is the discount factor, β: as the Agent cares more about the future, the value of holding out for the more favorable contract by shirking today increases. An alternative form in which $b_H(v_0, \lambda_0)$ can be written is:

$$b_H(v_0, \lambda_0) = c\bigg(\frac{1}{v_0\lambda_0} + \beta\frac{1 - v_1^p(\lambda_0, 1)}{v_1^p(\lambda_0, 1)}\bigg) \tag{11}$$

where the likelihood ratio $\frac{1-v_1^p}{v_1^p}$ increases in the Principal's initial investment, since on the equilibrium path higher initial investment makes failure more informative.

The key observation is that the Principal's implied commitment to continue experimentation in the future raises her cost of providing incentives to the Agent. If she chooses the contract described above, then her equilibrium value is:

$$\max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \bigg(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta f(\lambda_1^*(v_1^p(\lambda_0, 1), v_1^p(\lambda_0, 1)) - \beta V(v_0, v_1^p(\lambda_0, 1))\bigg). \tag{12}$$

Comparison of this payoff with the corresponding first best payoff shows that there is a wedge induced by the need to hand the Agent the discounted information rent she would obtain from shirking. This also induces a distortion in the Principal's investment strategy, since she claims less of the surplus.

We must also consider the possibility that, the Principal chooses a contract which induces the Agent to play a mixed strategy at date zero. In fact, these contracts are a necessity of *every* equilibrium, although they may possibly appear only off the equilibrium path. To see why, suppose that the Principal chooses an initial investment level $\lambda_0 \in (0, \hat{\lambda}_0(v_0))$ and a bonus:

$$b_0 \in \bigg(\frac{c}{v_0\lambda_0}, \frac{c}{v_0\lambda_0} + \beta c\frac{1 - v_1^p(\lambda_0, 1)}{v_1^p(\lambda_0, 1)}\bigg). \tag{13}$$

If the Principal conjectures that the Agent works at date zero with probability one, then the Agent strictly prefers to shirk. If, instead, the Principal conjectures that the Agent works at date zero with probability zero, the Agent strictly prefers to work.

After an offer (13) has been made, at least one mixed strategy equilibrium exists. In every mixed strategy equilibrium of the continuation game, the Agent randomizes over effort levels at date zero ($\sigma_0 \in (0,1)$). At date one, the Principal believes that the Agent is either an *optimistic* type who shirked at date zero and thus retains the prior belief v_0, or a *pessimistic* type who worked at date zero and holds the belief $v_1(\lambda_1, 1)$. The Principal subsequently randomizes over two kinds of contract at date one. One contract—an *inclusive* contract—satisfies the participation constraint of the pessimistic Agent and gives rent to the optimistic Agent. The other contract—an *exclusive* contract—satisfies the participation constraint of the optimistic Agent and induces the pessimistic Agent to shirk with probability one.

For a general cost of investing function $\tau(\lambda_0)$, there is no guarantee that contracts which induce the Agent to mix are solely an out-of-equilibrium phenomenon. However, all that is needed for my result is that the Principal's value from such a contract is strictly less than the value that she would attain from making the same investment when the Agent's effort is observed and she were able to hold the Agent to her participation constraint at each date. This simple observation is formally stated in the following Lemma.

Lemma 4. In the setting with unobserved effort, if the Principal chooses an investment level $\lambda_0 < \hat{\lambda}_0(v_0)$ and offers any bonus $b_0 \geq 0$, an equilibrium of the continuation game exists, and the Principal's value is strictly less than:

$$\max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \Big(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta f(\lambda_1^*(v_1^p(\lambda_0, 1); v_1^p(\lambda_0, 1)) \Big). \tag{14}$$

Mixed strategy equilibria—potentially multiple—may also exist in the continuation game after the Principal invests $\lambda_0 \geq \hat{\lambda}_0(v_0)$. The reason is that the threshold $\hat{\lambda}_0(v_0)$ is itself calculated under the presumption that the Agent chooses $\sigma_0 = 1$. If, instead, the Principal believes that the Agent is employing a non-degenerate distribution over effort choices at date zero, she may strictly prefer to invest again at date one, so long as the date zero investment choice is not too large relative to the induced mixture of the Agent. However, whenever such an equilibrium of the continuation game exists, there also exists an equilibrium of the continuation game in which the Principal holds the Agent to her participation constraint and the Agent chooses $\sigma_0 = 1$: i.e., the Agent chooses high effort with probability one at date zero. In particular, if $\lambda_0 = \lambda_0^*(v_0)$, this equilibrium generates the surplus $f(\lambda_0^*(v_0); v_0)$, which strictly dominates the surplus generated by a mixed strategy equilibrium. Since this surplus coincides with the Principal's payoff, this continuation equilibrium also yields a strictly higher payoff for the Principal. It can be shown that if β is sufficiently close to one, then in an equilibrium, the Principal's contract at date zero induces the Agent to play a pure strategy.

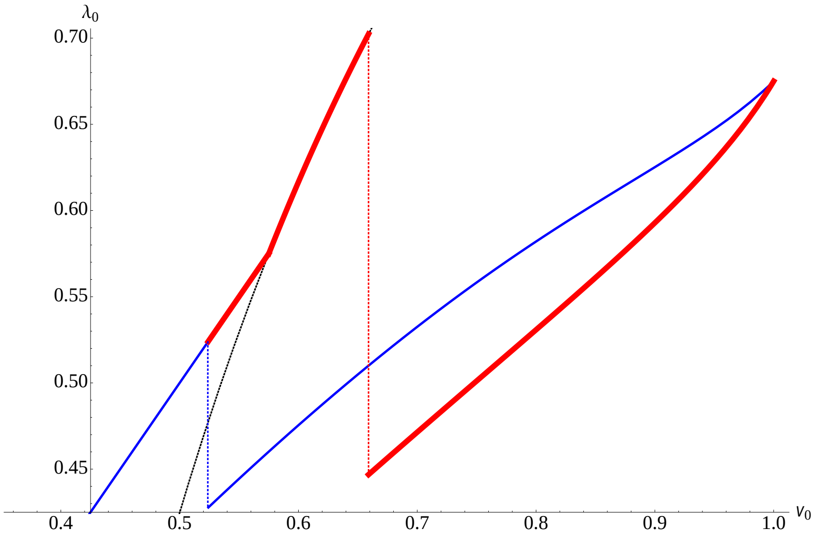

Figure 1: Illustration of the Principal's solution for λ_0 under *observed* (blue) and *unobserved* (red) effort, with $\tau(\lambda) = \frac{\lambda^2}{2}$. The dashed ray is $\hat{\lambda}_0(v_0)$. Parameter values are $\beta = \frac{3}{4}$, $W = 1$, $c = \frac{1}{15}$. Inefficient over-investment arises for $v_0 \in \Phi = (.524, .659]$. For $v_0 \leq .524$, the solutions coincide.

I now state my main result. Let $\lambda_0^{SB}(v_0)$ denote the Principal's optimal date zero investment in the second-best equilibrium.

Proposition 1. Whenever $\lambda_0^{FB}(v_0) \geq \hat{\lambda}_0(v_0)$, $\lambda_0^{SB}(v_0) = \lambda_0^{FB}(v_0)$. Moreover, there exists a non-empty set of prior beliefs about project quality, $\Phi \subset [0,1]$, such that $v_0 \in \Phi$ implies:

$$\lambda_0^{SB}(v_0) \geq \hat{\lambda}_0(v_0) > \lambda_0^{FB}(v_0). \tag{15}$$

If the Principal's first-best level of investment induces her to quit after initial failure, so will her second-best investment choice. However, there exists a set of prior beliefs about project quality, Φ, such that the Principal would initially invest to induce a positive value from future experimentation under the first-best, but under the second-best she *over-invests* so that after initial failure, no further investment occurs.

For $v_0 \in \Phi$, the Principal invests too much relative the first-best, and destroys surplus by doing so. Nonetheless, such over-investment is superior for the Principal since her subsequent decision to discontinue the experiment after initial failure denies the Agent rents from the repeated interaction. In effect, she destroys surplus at date one in order to capture a larger share at date zero.

Figure 1 continues my earlier example; it shows optimal experimentation at date zero under the first-best (blue) and moral hazard (red). The dashed black ray is $\hat{\lambda}_0(v_0)$. In the first-best, the

Principal's optimal plan depends on her initial optimism. When v_0 is relatively low, it is sufficiently 'cheap' for her to experiment above the threshold which would induce her to quit, in the event of failure. When v_0 is sufficiently large, however, the Principal is better off spreading the cost of investment over both dates: it is too expensive to initially invest the funds required to generate a sufficiently informative failure that would lead her to quit after the initial round.

In the second-best, the Principal chooses an inefficiently high level of investment when $v_0 \in [.524, .659]$, which can be further classified into two relevant regions. When $v_0 \in [.524, .575]$, the Principal in the second-best chooses the solution $\lambda_0^*(v_0)$, which is the one-shot unconstrained optimal investment, and which satisfies $\lambda_0^*(v_0) \geq \hat{\lambda}_0(v_0)$. However, when $v_0 \geq .575$, this solution is no longer consistent with the threat to end the experiment after initial failure, since the one-off investment solution has the property $\lambda_0^*(v_0) < \hat{\lambda}_0(v_0)$, and thus the Principal would wish to experiment again at date one. In that case, the Principal invests at date zero just enough to reduce the date one surplus to zero, i.e., at the constraint $\lambda_0^{SB}(v_0) = \hat{\lambda}_0(v_0)$. This commits her to prefer discontinuing the experiment after a date zero failure. Though this destroys future surplus, it allows the Principal to retain a larger share of the date zero surplus. Finally, when investing at the constraint $\hat{\lambda}_0(v_0)$ becomes sufficiently expensive, i.e., $v_0 > .659$, the Principal reverts to experimenting over both dates, inducing high effort at date zero and date one and paying the Agent the requisite information rent.

Notice that in the second-best, when the Principal experiments at date zero in such a way as to induce a positive continuation pay-off from further experimentation, she nonetheless chooses an initial investment that is too low, relative to the first-best. In this case, she trades off the benefits of early discovery with the incentive cost that arises from giving the Agent a higher continuation pay-off from shirking - both of which are increasing in her initial investment. This inefficient under-investment is similar to Bergemann and Hege (2005).

This model is the first to propose over-investment relative to the efficient benchmark as a way to soften the Agent's incentive constraint - by contrast, existing work emphasizes under-investment as the second-best solution. The latter prescription is a necessity of a framework in which the Principal's investment costs are linear, such as Bergemann and Hege (2005). But Proposition 1 offers a second important contrast with existing work. In Bergemann and Hege (2005) and Hörner and Samuelson (2013), the Principal stops funding the experiment too soon, despite the fact that there is a positive surplus from continued investment.[11] In this model, by contrast, when the Principal's initial investment departs from the first-best, a decision not to provide additional funds is socially optimal at that date, conditional on her choice of investment at date zero. The inefficiency arises from the fact that this initial choice of investment destroys future surplus in order to avoid making transfers to the Agent at date zero.

5. Over-investing with Commitment

For $v_0 \in \Phi$, the Principal's investment induces sufficiently pessimistic beliefs upon initial failure that she can credibly commit to undertake no further experimentation. This allows her to extract all of the surplus from the one-shot investment, despite the inefficiency of such a contract. It is natural to conjecture that investing above the threshold $\hat{\lambda}_0(v_0)$ is solely a means of *belief commitment*, since it allows the Principal to credibly threaten to quit the project after initial failure. It is also natural to conjecture that if the Principal could commit to writing a long-term contract at date zero which commits her to investment and bonus decisions at both dates—and in particular, allows her to pre-commit to quitting at the end of date zero—the inefficiency would disappear.

Both of these conjecture are false: when the Principal can write long-term contracts, over-investment still arises - possibly in situations where it would not under short-term contracts - and inefficiency may *increase* relative to the setting in which only short-term contracts can be written.

I now assume that the Principal can write a two-date contract specifying both investment and bonus decisions at both dates, at the start of date zero ("long-term commitment"). I let λ_t^{CO} denote her optimal choice of investment under the contract with long-term commitment. I begin with three preliminary observations which provide the intuition for my next result.

First, if the Principal's first-best investment calls on her to quite after initial failure—i.e., $\lambda^{FB}(v_0) = \lambda_0^*(v_0) \geq \hat{\lambda}_0(v_0)$—then her optimal contract both *with* and *without* inter-period commitment replicates the first-best. The reason is simple: by offering this contract she maximizes and fully appropriates the surplus from the relationship.

Second, even with commitment, a contract which induces positive experimentation at both dates must continue to satisfy the Agent's incentive constraints at each date. For example, if she

[11] Hörner and Samuelson (2013) also study non-Markov equilibria of their model, in which the Principal front-loads the Agent's effort in early periods until some threshold belief is attained, then switches to the worst equilibrium which is either funding with delay, or stopping experimentation. However, in the initial period in which front-loading takes place, it is the same investment that would arise in the first-best.

wishes to induce high effort at both dates, her promised bonuses at each date must replicate those made under the short-term contracts. More generally, for any equilibrium of the continuation beginning after the Principal offers a contract in which $\lambda_0 < \hat{\lambda}_0(v_0)$, she will obtain a smaller fraction of the induced surplus than she would from the continuation equilibrium after offering a contract in which she invests positively only at date zero, holds the Agent to her participation constraint in that period, and commits to no further investment at date one.

Third, when the Principal has long-term commitment, she no longer needs to invest above the threshold $\hat{\lambda}_0(v_0)$ in order to commit herself to quit upon failure. Together, these observations yield the following conclusion:

Proposition 2. Whenever, $\hat{\lambda}_0^{FB}(v_0) \geq \hat{\lambda}_0(v_0)$, $\lambda_0^{CO}(v_0) = \lambda_0^{FB}(v_0)$. Moreover, there exists a non-empty set of prior beliefs about project quality $\Psi \subset [0,1]$, such that $v_0 \in \Psi$ implies:

$$\lambda_0^{CO}(v_0) > \lambda_0^{FB}(v_0). \tag{16}$$

That is: when $v_0 \in \Psi$, the Principal's investment in a long-term contract is strictly higher than the first-best, at date zero.

Note that the existence of a set of prior beliefs for which over-investment occurs does not rely on any equilibrium selection: in that sense, the result for a long-term contract is even stronger than in the case with short-term contracts.

Over-investment under long-term commitment need not imply that the Principal's value from investing again at date one be zero, in contrast with the setting with only short-term commitment. To see this, suppose that in the setting with unobserved effort, the Principal's payoff from investing the static optimum $\lambda_0^*(v_0)$ at date zero and then quitting gives her a higher payoff than any contract which induces positive investment at date one. Suppose, moreover, that the static optimum $\lambda_0^*(v_0)$ lies below the threshold $\hat{\lambda}_0(v_0)$, so that the Principal's value from experimentation at date one is strictly positive, conditional on investing $\lambda_0^*(v_0)$ and believing that the Agent worked hard at date zero.

Under short-term commitment, offering $\lambda_0^*(v_0)$ requires the Principal to compensate the Agent with the information rent $\beta V(v_0, v_1^p(\lambda_0, 1))$ in order to induce high effort, since the Principal cannot commit to quit through this choice of investment: the Agent knows that the Principal will wish to invest again at date one, and so the Principal must provide the Agent with costly upfront incentives to work hard. The only way for the Principal to make this commitment while still inducing full

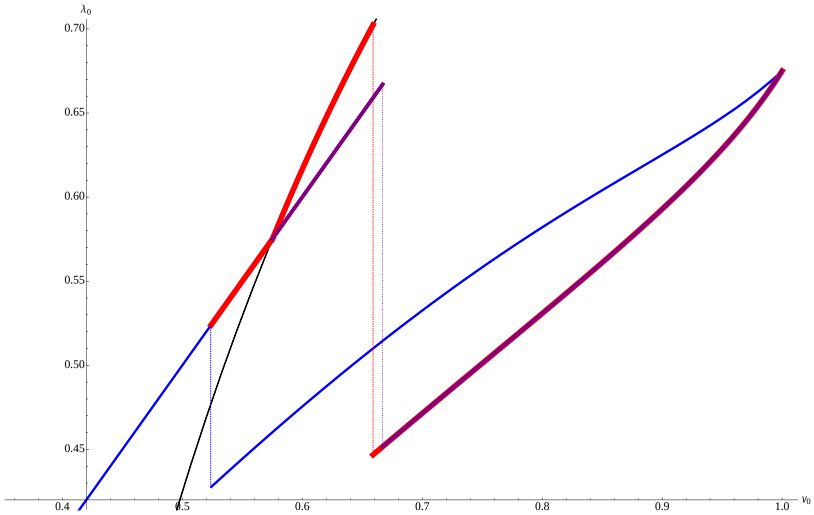

Figure 2: Illustration of the Principal's solution for λ_0 under *observed* (blue) and unobserved effort with *short-term* contracts (red), and unobserved effort with a *long-term* contract (purple) with $\tau(\lambda) = \frac{\lambda^2}{2}$. The dashed ray is $\hat{\lambda}_0(v_0)$. Parameter values are $\beta = \frac{3}{4}$, $W = 1$, $c = \frac{1}{15}$.

effort from the Agent at date zero is to invest precisely at the threshold $\hat{\lambda}_0(v_0)$, in order to reduce her own continuation payoff at date one to zero. If this strategy becomes too costly, the Principal reverts to experimenting over both dates and pay the Agent the requisite incentives at date zero.

Under long-term commitment, however, the Principal can pre-commit to investing at $\lambda_0^*(v_0) < \hat{\lambda}_0(v_0)$ and quitting upon failure. This allows her to induce high effort solely by satisfying the Agent's date zero participation constraint, and appropriating all of the rent. Since the Principal's value from this contract is strictly greater than that associated with investing $\hat{\lambda}_0(v_0)$, there may be prior beliefs for which she is prepared to make the one-off investment in the long-term commitment setting, when she would instead revert to experimenting over both dates in the absence of commitment. Figure 2 continues my earlier example, in which I have added the Principal's solution when she has long-term commitment.

Under long-term commitment, when $v_0 \in [.524, .575]$, the Principal's one-off investment solution is the same as in the setting where only short-term contracts are available. Likewise, when $v_0 >$ .667, the two-date investment solution coincides across both short-term and long-term contracts.

However, when the prior belief takes on an intermediate value, i.e., $v_0 \in (.575, .667]$, the Principal faces different trade-offs across each commitment setting. With only short-term commitment, the Principal cannot invest at the static optimum $\lambda_0^*(v_0)$ and commit to quit. She must either choose a contract which induces the Agent to work at both dates but requires her to hand her the associated information rent at date zero, or instead invest at just the amount necessary to reduce her continuation from further experimentation to zero (i.e., invest $\hat{\lambda}_0(v_0)$). When $v_0 \in (.575, .659]$, the Principal opts for the latter. When $v_0 > .659$, the cost of investing at an amount sufficient to make quitting credible is too large, and she instead reverts to a contract which induces full effort and positive investment at both dates.

Under long-term commitment, by contrast, the Principal can freely commit to quitting after a single date of positive investment; so, she invests $\lambda_0^*(v_0)$ and commits to making no further investment.

So, on the interval $v_0 \in (.575, .659]$, the Principal over-invests under both short- and long-term contracts, although the over-investment is more severe in the former setting. However, when $v_0 \in (.659, .667]$, the Principal *over-invests* under the long-term contract and *under-invests* with short-term contracts—despite generating a positive continuation payoff from further experimentation at date one in both cases!

Commitment Power Can Increase Inefficiency

The question arises as to the welfare consequences of moving from short- to long-term contracts. Clearly, the Principal must be better off under long- than short-term commitment. Moreover, whenever $\lambda_0^*(v_0) < \hat{\lambda}_0(v_0)$ but the Principal chooses a corner solution $\lambda_0^{SB}(v_0) = \hat{\lambda}_0(v_0)$, the surplus must also increase under the long-term contract, since the difference in surplus between the long- and short-term contracts is $f(\lambda^*(v_0); v_0) - f(\hat{\lambda}_0(v_0); v_0) > 0$. This is the case in my running example when $v_0 \in (.575, .659]$.

Suppose, however, that under short-term contracts, at some prior belief, the Principal's date zero investment generates a positive continuation payoff at date one, and that she chooses a contract which induces the Agent to work hard with probability one at date zero. Then, her investment decision is:

$$\lambda_0^{SB}(v_0) = \arg \max_{\lambda_0 \leq \hat{\lambda}_0(v_0)} (f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta S(v_1^p(\lambda_0, 1)) - \beta V(v_0, v_1^p(\lambda_0, 1))), \tag{17}$$

and the surplus is:

$$f(\lambda_0^{SB}(v_0); v_0) + (1 - v_0\lambda_0^{SB})\beta S(v_1^p(\lambda_0^{SB}, 1)), \tag{18}$$

since $\beta V(v_0, v_1^p(\lambda_0^{SB}, 1))$ is a transfer from the Principal to the Agent. If, under long-term commitment, the Principal chooses to invest $\lambda_0^{CO} = \lambda_0^*(v_0) < \hat{\lambda}_0(v_0)$ and $\lambda_1^{CO} = 0$, the surplus is instead:

$$f(\lambda_0^*(v_0); v_0). \tag{19}$$

It is possible that (18) exceeds (19), but that the difference is strictly less than $\beta V(v_0, v_1^p(\lambda_0^{SB}, 1))$, the Principal's transfer to the Agent. In that case, the Principal makes a one-off investment in the setting with commitment, but invests over both periods without commitment; so, commitment power raises the inefficiency induced by the Principal's desire to retain as much of the date one surplus as possible.

Example 1 (continued). Maintaining the running parameterization, consider $v_0 = \frac{33}{50}$; in the second-best with short-term commitment the Principal invests $\lambda_0^{SB}(v_0) = .511$, and the surplus is:

$$f(\lambda_0^{SB}(v_0); v_0) + (1 - v_0\lambda_0^{SB})\beta S(v_1^p(\lambda_0^{SB}, 1)) = .164. \tag{20}$$

However, in order to induce effort at date zero, the Principal must hand the Agent the rent $\beta V(v_0, v_1^p(\lambda_0^{SB}, 1)) = .014$ and so appropriates only the value .150, for herself. With long-term commitment, the Principal invests $\lambda_0^*(v_0) = \frac{33}{50}$, holds the Agent to her participation constraint, and generates the surplus:

$$f(\lambda_0^*(v_0); v_0) = .151, \tag{21}$$

which the Principal keeps for herself. So, giving the Principal long-term commitment power *worsens* the inefficiency arising from the agency problem. This efficiency ordering holds for all $v_0 \in (.659, .667]$.

Intuitively: the Principal is able to appropriate the full share of instantaneous surplus created by her one-off investment decision under long-term commitment—that is: she fully appropriates (19). But this solution now leaves a positive continuation payoff on the table which is claimed by neither player. If, instead, the Principal invests over both dates, she hands a portion of the date zero surplus to the Agent, but appropriates all of the date one surplus. If the difference of (18) and (19) is positive, but less than $\beta V(v_0, v_1^p(\lambda_0^{SB}, 1))$, long-term commitment power for the Principal worsens the inefficiency associated with the agency problem. This is because formal commitment power allows the Principal to credibly destroy even more of the surplus associated with the two-date investment solution in order to deny rents to the Agent at date zero, without forcing her to partly internalize this by holding her initial investment to the corner solution, $\hat{\lambda}_0(v_0)$.

6. Discussion

Incentives to over-invest arise from two features of the model. First, there exist beliefs about project quality for which the Principal would prefer to abandon the project. This gives her the ability to make initial investments which imply no further experimentation upon early failure. Second, the Agent's rents are inter-temporal: even when effort is unobserved, the Principal can hold the Agent to her participation constraint when she invests at date zero and quits after initial failure. This gives the Principal a motive to distort investments in order to capture the surplus.[12]

The results have application to project evaluation: observing that a project was initially given a large budget but that it was discontinued at an early stage without success may constitute evidence that the project was poorly conceived or badly managed. However, this analysis suggests another interpretation: the early release of large volume of funds followed by early abandonment may represent the solution of an incentive problem arising from the Agent's desire to prolong the life of the contract, in order to capture information rents.

Although two periods are sufficient to illustrate the main argument, I argue that the addition of more periods aggravates the Agency cost and is likely to increase incentives to over-invest at the outset. Suppose that there are three dates; 0, 1 and 2, and that the first-best calls on the Principal to invest positively at each of the three dates, and that the Principal wishes to induce full effort by the Agent at all three dates. On the equilibrium path, the Principal's belief coincides with the Agent's at each date. At date two, the Principal holds the Agent with belief $v_2^p = \frac{v_1^p(1-\lambda_1)}{1-v_1^p\lambda_1}$ to her participation constraint. At date one, the Principal offers the Agent her reservation payoff, in

[12]Over-investment can also arise when the Principal's investment is not observed by the Agent, so that there is scope for the former to manipulate the latter's beliefs. Here is a partial intuition. At date one, suppose that the Agent believes that the Principal invests $\lambda_1 = 0$, and subsequently randomizes with positive probability over all bonus offers on $[0, W]$. Then, for any realized bonus announcement, the Agent believes the Principal invested zero, and optimally exerts zero effort at date one. If the Principal deviates to positive investment, the Agent will not observe this, and the costly investment is wasted. So, zero investment by the Principal at date one is optimal, and the Principal's is indeed indifferent over all bonus offers since she never has to pay out. This implies that both players receive a payoff of zero in the continuation equilibrium at date one. Then, we can construct strategies and beliefs at date zero which support a one-off investment by the Principal, and the Agent is held to her participation constraint at that date. However, the initial investment may not coincide with the case where the Principal's investment is observed.

addition to her discounted benefit from manipulating the Principal's beliefs by secretly shirking:

$$b_1 = \frac{c}{v_1^p \lambda_1} + \beta c \frac{1 - v_2^p}{v_2^p}. \tag{22}$$

Now consider the bonus that she offers the Agent at date zero. If the Agent secretly shirks, she creates a wedge between her beliefs and the Principal's beginning at date one. The wedge in these beliefs also persists at date two, and so the value of the deviation at date zero rises since the Agent now obtains a stream of information rents, since each subsequent bonus is too generous relative to the Agent's true beliefs. Using a similar procedure as in the two-date example, we obtain the requisite bonus at date one:

$$b_0 = \frac{c}{v_0 \lambda_0} + \beta c \frac{1 - v_1^p}{v_1^p} + \beta^2 c \frac{1 - v_2^p}{v_2^p}. \tag{23}$$

If, instead, the Principal invests only over the first two dates, the date zero bonus is only the first term in this expression, and the date one bonus need only satisfy the Agent's participation constraint. Thus, adding more periods of experimentation generates a discrete increase in Agency costs *at each date in which experimentation takes place prior to the final date.* This suggests that incentives to over-invest early on in order to trade-off the inefficiency of the short horizon with the associated agency costs of experimenting over the longer horizon may appear even more strongly.

7. Appendix

Proof of Lemma 1

To prove (i), I show $\lambda_1^*(v_1^p)$ increases in v_1^p and v_1^p decreases in λ_0. Fix $v_1^p = v \in (0,1)$ and $\lambda_1^*(v)$, which I shorthand λ. Take $v' > v$ and associated solution $\lambda_1^*(v') = \lambda'$. Since λ' is optimal for v' and λ is optimal for v:

$$v'(\lambda' - \lambda)W + \tau(\lambda) - \tau(\lambda') \geq 0 \tag{24}$$

$$v(\lambda - \lambda')W + \tau(\lambda') - \tau(\lambda) \geq 0 \tag{25}$$

and combining yields:

$$(v - v')(\lambda - \lambda') \geq 0$$

so $v < v'$ implies $\lambda \leq \lambda'$. Next, $\lambda_0 > 0$ if and only if $e_0 = 1$, in the first-best, so $v_1^p(\lambda_0, 1) = \frac{v_0(1-\lambda_0)}{1-v_0\lambda_0}$ which strictly decreases in λ_0. This completes the proof of (i).

I now prove (ii). Recall $S(v_1^p) = \max\{0, f(\lambda_1^*(v_1^p); v_1^p)\}$. It is easy to show $f(\lambda_1^*(v_1^p); v_1^p)$ is strictly increasing in v_1^p. Thus, $S(v_1^p) \geq 0$ implies $S(\hat{v}_1^p) > S(v_1^p)$ for $\hat{v}_1^p > v_1^p$. Next, $f(\lambda_1^*(0); 0) < 0$ and $f(\lambda_1^*(1); 1) > 0$. So, there is a unique $v_1^p \in (0,1)$ such that (a) $S(v_1^p) = f(\lambda_1^*(v_1^p); v_1^p) = 0$ and (b) for any $(v_1^p)' > v_1^p$, $S((v_1^p)') > 0$, which I call $\underline{v}$.

Proof of Lemma 3

I first show that if the Principal announces an investment $\lambda_0 \geq \hat{\lambda}_0(v_0)$, then for any date zero bonus that she offers, an equilibrium of the continuation game beginning with the Agent's date zero effort choice exists. Recall that σ_t denotes the Agent's distribution over effort levels at date t.

Claim 1. If the Principal chooses $\lambda_0 \geq \hat{\lambda}_0(v_0)$ and $b_0 < \frac{c}{v_0\lambda_0}$, there exists an equilibrium of the continuation game in which $\sigma_0 = 0$, $\lambda_1 = \lambda_1^*(v_1(\lambda_0, 0))$, $b_1 = \frac{c}{v_1^p(\lambda_0,0)\lambda_1}$, $\sigma_1 = 1$ if $v_1^p(\lambda_0, e_0)\lambda_1 b_1 \geq c$ and $\sigma_1 = 0$ otherwise.

Proof. The Agent's date one strategy is trivially optimal. The Principal's belief at date one about project quality is $v_1^p(\lambda_0, 0) = v_0$, and she believes that the Agent's belief is $v_1(\lambda_0, 0) = v_0$ with probability one, so her investment strategy and bonus offer at that date is optimal, given the Agent's strategy. Since the Agent's date one continuation payoff is zero regardless of her action at date zero, she weakly prefers to exert no effort at date zero if $v_0\lambda_0 b_0 - c \leq 0$, which is satisfied. □

Claim 2. If the Principal chooses $\lambda_0 \geq \hat{\lambda}_0(v_0)$ and $b_0 \geq \frac{c}{v_0\lambda_0}$, there exists an equilibrium of the continuation game in which $\sigma_0 = 1$, $\lambda_1 = 0$, $b_1 = 0$, $\sigma_1 = 1$ if $v_1^p(\lambda_0, e_0)\lambda_1 b_1 \geq c$ and $\sigma_1 = 0$ otherwise.

Proof. The Agent's date one strategy is trivially optimal. At date one, the Principal's belief is $v_1^p(\lambda_0, 1) \leq \underline{v}$, so her investment strategy and bonus offer is optimal. At date one, the Agent's rent is zero regardless of her action at date zero, so at date zero she weakly prefers to exert full effort if $v_0\lambda_0 b_0 - c \geq 0$, which is satisfied. □

This establishes existence of an equilibrium of the continuation game after any offer by the Principal $\lambda_0 \geq \hat{\lambda}_0(v_0)$ and $b_0 \geq 0$. Moreover, if the Principal invests $\lambda_0^*(v_0)$ and offers $b_0 = \frac{c}{v_0\lambda_0^*(v_0)}$ at date zero, an equilibrium of the continuation game exists which generates the surplus $f(\lambda_0^*(v_0); v_0)$, which is fully appropriated by the Principal.

Proof of Lemma 4

The proof proceeds by a sequence of claims, in which I show that an equilibrium of the continuation game exists after any contract offered by the Principal satisfying $0 \le \lambda_0 < \hat{\lambda}_0(v_0)$ and $b_0 \ge 0$.

Claim 3. There exists an equilibrium of the continuation game beginning with the Agent's date zero effort decision, after the Principal announces any date zero investment $\lambda_0 \in (0, \hat{\lambda}_0(v_0)]$ and date zero bonus $b_0 \le \frac{c}{v_0 \lambda_0}$.

Proof. I claim that an equilibrium exists in which the Agent chooses $\sigma_0 = 0$ at date zero and at date one the Principal invests $\lambda_1 = \lambda_1^*(v_0)$, offers $b_1 = \frac{c}{v_1^p(\lambda_0, 0)\lambda_1}$, and the Agent chooses $\sigma_1 = 1$ if $v_1(\lambda_0, e_0)\lambda_1 b_1 - c \ge 0$, and $\sigma_1 = 0$ otherwise. The proof is similar to the previous Lemma. $\square$

Claim 4. There exists an equilibrium of the continuation game beginning with the Agent's date zero effort decision, after the Principal announces any date zero investment $\lambda_0 \in (0, \hat{\lambda}_0(v_0)]$ and date zero bonus $b_0 \ge \frac{c}{\lambda_0}\left(\frac{1}{v_0} + \beta\left(\frac{1}{v_1^p(\lambda_0, 1)} - \frac{1}{v_0}\right)\right)$.

Proof. By arguments in the text, an equilibrium exists in which $\sigma_0 = 1$, $b_1 = \frac{c}{v_1(\lambda_0, 1)\lambda_1}$, $\lambda_1 = \lambda_1^*(v_1^p(\lambda_0, 1))$ and $\sigma_1 = 1$ if $v_1(\lambda_0, e_0)\lambda_1 b_1 - c \ge 0$, and $\sigma_1 = 0$ otherwise. $\square$

Claim 5. After the Principal chooses $\lambda_0 \in (0, \hat{\lambda}_0(v_0))$ and

$$b_0 \in \left(\frac{c}{v_0 \lambda_0}, \frac{c}{\lambda_0}\left(\frac{1}{v_0} + \beta\left(\frac{1}{v_1^p(\lambda_0, 1)} - \frac{1}{v_0}\right)\right)\right), \tag{26}$$

there exists no equilibrium of the continuation game in which the Agent chooses $\sigma_0 \in \{0, 1\}$.

Proof. Suppose an equilibrium exists in which $\sigma_0 = 1$. At date one, the Principal's belief is $v_1^p(\lambda_0, 1)$, and she believes that the Agent's belief is $v_1(\lambda_0, 1)$ with probability one. Then, there is unique equilibrium of the continuation game beginning at date one, in which the Principal offers $b_1 = \frac{c}{v_1^p(\lambda_0, 1)\lambda_1}$ and invests $\lambda_1 = \lambda_1^*(v_1(\lambda_0, 1))$, and the Agent chooses $\sigma_1 = 1$ if $v_1(\lambda_0, e_0)\lambda_1 b_1 - c \ge 0$, and $\sigma_1 = 0$ otherwise. This implies that the Agent prefers to work at date zero only if:

$$v_0 \lambda_0 b_0 - c \ge \beta c \left(\frac{v_1(\lambda_0, 0)}{v_1^p(\lambda_1, 1)} - 1\right), \tag{27}$$

which is violated by (26), where I have substituted $v_1(\lambda_0, 0) = v_0$.

Next, suppose that an equilibrium exists in which $\sigma_0 = 0$. At date one, the Principal's belief is $v_1^p(\lambda_0, 0)$, and she believes that the Agent's belief is $v_1(\lambda_0, 0)$ with probability one. Then, there is a unique equilibrium of the continuation game beginning at date on, in which the Principal offers $b_1 = \frac{c}{v_1(\lambda_0,0)\lambda_1}$ and $\lambda_1 = \lambda_1^*(v_1(\lambda_0,0))$, and the Agent chooses $\sigma_1 = 1$ if $v_1(\lambda_0,e_0)\lambda_1 b_1 - c \geq 0$, and $\sigma_1 = 0$ otherwise. The Agent therefore prefers to exert zero effort at date zero only if:

$$v_0\lambda_0 b_0 - c \leq 0, \tag{28}$$

which is violated by (26). □

Claim 6. After the Principal chooses $\lambda_0 \in (0, \hat{\lambda}_0(v_0)]$ and

$$b_0 \in \left(\frac{c}{v_0\lambda_0}, \frac{c}{\lambda_0}\left(\frac{1}{v_0} + \beta\left(\frac{1}{v_1^p(\lambda_0,1)} - \frac{1}{v_0}\right)\right)\right), \tag{29}$$

there exists an equilibrium of the continuation game in which the Agent chooses $\sigma_0 \in (0,1)$.

Proof. At date one, the Principal's belief about project quality is $v_1^p(\lambda_0,\sigma_0)$. The Agent's belief is $v_1(\lambda_0,e_0)$. I call the Agent with belief $v_1(\lambda_0,1)$ a *pessimistic* Agent-type and the Agent with belief $v_1(\lambda_0,0)$ an *optimistic* Agent-type.

If the Principal chooses $\lambda_1 > 0$, I specify that the Principal randomizes over the following two contracts:

(i) with probability $\eta \in [0,1]$, she chooses an *inclusive* contract, which satisfies the participation constraint of the pessimistic Agent, and gives positive rent to the optimistic Agent: $b_1^I(\lambda_1) \equiv \frac{c}{v_1(\lambda_0,1)\lambda_1}$. The Principal's investment solves:

$$\max_{\lambda_1\in[0,1]} \left(\frac{v_0(1-\sigma_0\lambda_0)}{1-v_0\lambda_0\sigma_0}\lambda_1(W - b_1^I(\lambda_1)) - \tau(\lambda_1)\right), \tag{30}$$

and I specify that both Agent-types select high effort with probability one.

(ii) with probability $1-\eta$, she chooses an *exclusive* contract, which satisfies the participation constraint of the optimistic Agent: $b_1^E(\lambda_1) \equiv \frac{c}{v_1(\lambda_0,0)\lambda_1}$. This gives the pessimistic Agent an negative expected payoff from exerting high effort, so that she strictly prefers low effort at date one. The Principal's investment solves:

$$\max_{\lambda_1\in[0,1]} \left(\frac{v_0(1-\sigma_0)}{1-v_0\lambda_0\sigma_0}\lambda_1(W - b_1^E(\lambda_1)) - \tau(\lambda_1)\right), \tag{31}$$

and I specify that the optimistic Agent selects high effort with probability one, while the pessimistic Agent selects low effort with probability one.

The strategy of each Agent-type after the Principal's realized date one contract is trivially optimal. Let η denote the probability that the Principal at date one offers an inclusive contract. Then, the Agent at date zero is indifferent between high effort and low effort only if:

$$v_0\lambda_0 b_0 - c = \beta\eta c\left(\frac{v_0}{v_1(\lambda_0, 1)} - 1\right) \tag{32}$$

If $b_0 = b_0^L(v_0, \lambda_0) \equiv \frac{c}{v_0\lambda_0}$, we must therefore have $\eta = 0$. If $b_0 = \frac{c}{\lambda_0}\left(\frac{1}{v_0} + \beta\left(\frac{1}{v_1^p(\lambda_0,1)} - \frac{1}{v_0}\right)\right) \equiv b_0^H(v_0, \lambda_0)$, we must have $\eta = 1$. It follows that for any interior bonus, there exists a unique mixture (η) which satisfies (32).

Next, I show that there exists at least one mixture, σ_0, which makes the Principal indifferent between each of the *inclusive* and *exclusive* contract at date one. Suppose $\sigma_0 = 0$. Then, we have:

$$\max_{\lambda_1\in[0,1]}\left(v_0\lambda_1 W - c\frac{1-\lambda_0 v_0}{1-\lambda_0} - \tau(\lambda_1)\right) < \max_{\lambda_1\in[0,1]}\left(v_0\lambda_1 W - c - \tau(\lambda_1)\right) \tag{33}$$

where the LHS is the Principal's value from an inclusive contract, and the RHS is the Principal's value from an exclusive contract, since the Agent is always optimistic at date one. Suppose $\sigma_0 = 1$. Then, we have:

$$\max_{\lambda_1}\left(\frac{v_0(1-\lambda_0)}{1-v_0\lambda_0}\lambda_1 W - c - \tau(\lambda_1)\right) \geq 0, \tag{34}$$

where the LHS is the Principal's value from an inclusive contract, and the RHS is the Principal's value from an exclusive contract, since when $\sigma_1 = 0$, the Agent is always pessimistic at date one. The LHS is weakly positive by supposition of $\lambda_0 \leq \hat{\lambda}_0(v_0)$. So, there exists at least one $\sigma_0 \in (0, 1]$ which creates indifference, and thus a mixed strategy equilibrium of the continuation game exists. □

Claim 7. After the Principal chooses $\lambda_0 = 0$ and any bonus offer $b_0 \geq 0$, an equilibrium of the continuation game exists.

Proof. It is easy to show that the following is an equilibrium of the continuation game: the Agent chooses $\sigma_0 = 0$, the Principal offers $\lambda_1 = \lambda_1^*(v_1^p(0,0))$ and $b_1 = \frac{c}{v_1^p(0,0)\lambda_1^*(v_1^p(0,0))}$, and the Agent chooses $\sigma_1 = 1$ if $v_1(0, e_0)\lambda_1 b_1 - c \geq 0$, and $\sigma_1 = 0$, otherwise. □

We conclude that for any contract satisfying $\lambda_0 < \hat{\lambda}_0(v_0))$ and $b_0 \geq 0$, an equilibrium of the continuation game exists.

For a general cost function $\tau(\cdot)$, I cannot rule out the possibility that there is more than one σ_0 which generates indifference between each of the two contracts, in the continuation game beginning after an offer $\lambda_0 < \hat{\lambda}_0(v_0)$ and a bonus which is drawn from the interval (26). However, it is sufficient to note the surplus generated in any such continuation equilibrium is strictly less than the surplus generated when the Principal can directly observe the Agent's effort, holds the Agent to her participation constraint at each date, and the Agent chooses high effort at both dates:

$$\max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \left(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta f(\lambda_1^*(v_1^p(\lambda_0, 1); v_1^p(\lambda_0, 1)) \right). \tag{35}$$

Proof of Proposition 1

When the Agent's effort choice is not observed by the Principal, for a given prior belief v_0, let $\underline{S}^{SB}(v_0)$ denote the supremum of the surplus that can be attained in a continuation equilibrium when the prior is v_0, beginning with the Agent's date zero effort choice where the supremum is taken over all pairs (λ_0, b_0) satisfying $\lambda_0 \in [0, \hat{\lambda}_0(v_0))$ and $b_0 \geq 0$. We have:

$$\underline{S}^{SB}(v_0) < \max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \left(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta f(\lambda_1^*(v_1^p(\lambda_0, 1); v_1^p(\lambda_0, 1)) \right). \tag{36}$$

Similarly, let $\overline{S}^{SB}(v_0)$ denote the supremum of the surplus that can be achieved in any continuation equilibrium beginning with the Agent's date zero effort choice, where the supremum is taken over all pairs (λ_0, b_0) satisfying $\lambda_0 \in [\hat{\lambda}_0(v_0), 1]$ and $b_0 \geq 0$.

Claim 8. If $\lambda_0^*(v_0) \geq \hat{\lambda}_0(v_0)$:

$$\overline{S}^{SB}(v_0) = \max_{\lambda_0 \geq \hat{\lambda}_0(v_0)} (v_0\lambda_0 W - \tau(\lambda_0) - c) = f(\lambda_0^*(v_0); v_0). \tag{37}$$

Moreover, $\overline{S}^{SB}(v_0)$ is generated in an equilibrium of the continuation game after the Principal chooses λ_0 and b_0 only if $\lambda_0 = \lambda_0^*(v_0)$ and $b_0 = \frac{c}{v_0\lambda_0^*(v_0)}$.

Proof. That the surplus $f(\lambda_0^*(v_0); v_0)$ can be generated in a continuation equilibrium after $\lambda_0 = \lambda_0^*(v_0)$ and $b_0 = \frac{c}{v_0\lambda_0^*(v_0)}$ is established in Lemma 3. That it can only be generated after this investment and bonus offer is trivial. Finally, since $\lambda_0^*(v_0) \geq \hat{\lambda}_0(v_0)$, we have:

$$f(\lambda_0^*(v_0); v_0) = \max_{\lambda_0 \in [\hat{\lambda}_0(v_0), 1], \sigma_0 \in [0,1]} \left(v_0\lambda_0\sigma_0 W - \tau(\lambda_0) - c + \beta(1 - v_0\lambda_0\sigma_0) S(v_1^p(\lambda_0, \sigma_0)\right), \tag{38}$$

so that $f(\lambda_0^*(v_0); v_0) = \overline{S}^{SB}(v_0)$. □

Recall that $\lambda_0^{SB}(v_0)$ denotes the Principal's investment at date zero in the equilibrium with unobserved effort which maximizes the surplus.

Claim 9. $\Theta = \{v_0 : \lambda_0^{FB}(v_0) \geq \hat{\lambda}_0(v_0)\}$ is non-empty, and $v_0 \in \Theta$ implies $\lambda^{SB}(v_0) = \lambda_0^{FB}(v_0) = \lambda_0^*(v_0)$.

Proof. $f(\lambda_0^*(v_0); v_0)$ is continuous and strictly increasing in v_0. By construction, $f(\lambda_0^*(\underline{v}); \underline{v}) = 0$, $\hat{\lambda}_0(\underline{v}) = 0$ and for all $\lambda_0 \geq 0$:

$$f(\lambda_0; \underline{v}) + \beta(1 - \underline{v}\lambda_0) f(\lambda_1^*(v_1^p); v_1^p) < 0. \tag{39}$$

Finally, $\lambda_0^*(\underline{v}) > 0 = \hat{\lambda}_0(\underline{v})$. Continuity and strict monotonicity of $\lambda_0^*(v_0) > \hat{\lambda}_0(v_0)$ implies that there exists $\delta > 0$ and $B_+^\delta(v_0) \equiv \{v_0' : 0 < v_0' - v_0 < \delta\}$ such that $v_0' \in B_+^\delta(v_0)$ implies $f(\lambda_0^*(v_0'); v_0') > 0$, $\lambda_0^*(v_0') > \hat{\lambda}_0(v_0')$ and $\max_{\lambda_0 \in [0,1]} \left(f(\lambda_0; v_0') + \beta(1 - v_0'\lambda_0) f(\lambda_1^*(v_1^p); v_1^p)\right) < f(\lambda_0^*(v_0'); v_0')$. Thus, Θ is non-empty. Next, $v_0 \in \Theta$ implies:

$$f(\lambda_0^*(v_0); v_0) \geq \max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \left(f(\lambda_0; v_0^A) + (1 - v_0^A \lambda_0)\beta f(\lambda_1^*(v_1^p); v_1^p) \right) > \underline{S}^{SB}(v_0). \tag{40}$$

By Lemma 3, $v_0 \in \Theta$ implies $\lambda_0^*(v_0) \geq \hat{\lambda}_0(v_0)$; Claim 7 implies $f(\lambda_0^*(v_0); v_0) = \overline{S}^{SB}(v_0)$. Since the Principal fully appropriates the surplus $\overline{S}^{SB}(v_0)$ and $\overline{S}^{SB}(v_0) > \underline{S}^{SB}(v_0)$, we conclude that $\lambda_0^{SB}(v_0) = \lambda_0^{FB}(v_0) = \lambda_0^*(v_0)$. □

Claim 10. $\Phi \equiv \{v_0 : \lambda_0^{SB}(v_0) \geq \hat{\lambda}_0(v_0) > \lambda_0^{FB}(v_0)\}$ is non-empty.

Proof. Recall $\hat{\lambda}_0(v_0)$ is strictly concave, $\hat{\lambda}_0(\underline{v}) = 0$ and $\hat{\lambda}_0(1) = 1$; moreover, $\lambda_0^*(v_0)$ is continuous, strictly increasing, $\lambda_0(\underline{v}) > 0$ and $\lambda_0^*(1) < 1$. Thus, there exists a largest v_0 such that $\lambda_0^*(v_0)$ implies $f(\lambda_1^*(v_1^p); v_1^p) \leq 0$. This implies there exists a largest v_0 such that $\lambda_0^*(v_0) \geq \hat{\lambda}_0(v_0)$ and:

$$f(\lambda_0^*(v_0); v_0) = \max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \left(f(\lambda_0; v_0^A) + (1 - v_0^A \lambda_0)\beta f(\lambda_1^*(v_1^p); v_1^p) \right), \tag{41}$$

which I denote v_0^A. For all $v_0 > v_0^A$, $\lambda_0^{FB}(v_0) < \hat{\lambda}_0(v_0)$. Define:

$$\epsilon_1(v_0) = \max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0^A)]} \left(f(\lambda_0; v_0^A) + (1 - v_0^A \lambda_0)\beta f(\lambda_1^*(v_1^p); v_1^p) \right) - \underline{S}^{SB}(v_0), \tag{42}$$

where $\epsilon_1(v_0) > 0$ for any v_0. Thus, $f(\lambda_0^*(v_0^A); v_0^A) > \underline{S}^{SB}(v_0^A)$.

1. Suppose, first, $\lambda_0^*(v_0^A) > \hat{\lambda}_0(v_0^A)$. Define:

$$\epsilon_2(v_0) = \max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0^A)]} \left(f(\lambda_0; v_0^A) + (1 - v_0^A \lambda_0)\beta f(\lambda_1^*(v_1^p); v_1^p) \right) - f(\lambda_0^*(v_0); v_0), \tag{43}$$

and note $\epsilon_2(v_0^A) = 0$, $\epsilon_2(v_0) > 0$ if $v_0 > v_0^A$, and $\epsilon_2(v_0)$ is continuous. Moreover, $f(\lambda_0^*(v_0^A); v_0^A)$ is continuous and strictly increasing, $\lambda_0^*(v_0)$ is continuous and strictly increasing, and $\epsilon_1(v_0) > 0$ for all v_0. So, there exists $\delta > 0$ such that $v_0' \in B_+^\delta(v_0^A)$ implies $\lambda_0^*(v_0') > \hat{\lambda}_0(v_0)$, and $\epsilon_2(v_0') < \epsilon_1(v_0')$. By Claim 7:

$$f(\lambda_0^*(v_0'); v_0') - \underline{S}^{SB}(v_0') = \overline{S}^{SB}(v_0') - \underline{S}^{SB}(v_0') = -\epsilon_2(v_0)' + \epsilon_1(v_0') > 0. \tag{44}$$

This implies $\lambda_0^{SB}(v_0') = \lambda_0^*(v_0') > \hat{\lambda}_0(v_0')$ and $\lambda_0^{FB}(v_0') < \hat{\lambda}_0(v_0')$.

2. Suppose, second, $\lambda_0^*(v_0^A) = \hat{\lambda}_0(v_0^A)$. Define $\epsilon_3(v_0) \equiv f(\hat{\lambda}_0(v_0); v_0) - \beta f(\lambda_0^*(v_0); v_0)$ and $\epsilon_4(v_0) \equiv f(\lambda_0^*(v_0); v_0) - f(\hat{\lambda}_0(v_0); v_0)$ and note that $\epsilon_3(v_0^A) > 0$ for $\beta < 1$, $\epsilon_4(v_0^A) = 0$, $\epsilon_2(v_0^A) = 0$ and $\epsilon_3(\cdot)$ and $\epsilon_4(\cdot)$ are continuous at v_0^A. This implies that there exists $\delta > 0$ such that $v_0' \in B_+^\delta(v_0^A) \equiv \{v_0' : 0 < v_0' - v_0 < \delta\}$ implies $\epsilon_3(v_0') > 0$ and $\epsilon_1(v_0') > \epsilon_2(v_0') + \epsilon_4(v_0')$. This implies:

$$f(\hat{\lambda}_0(v_0')) - \underline{S}^{SB}(v_0') = f(\hat{\lambda}_0(v_0')) - f(\lambda_0^*(v_0')) + f(\lambda_0^*(v_0')) - \underline{S}^{SB}(v_0') \tag{45}$$

$$= -\epsilon_4(v_0') + \epsilon_1(v_0') - \epsilon_2(v_0') \tag{46}$$

$$> 0 \tag{47}$$

and since $\epsilon_3(v_0') > 0$, $f(\hat{\lambda}_0(v_0'); v_0') - \beta f(\lambda_0^*(v_0'); v_0') > 0$. This implies $\lambda_0^{SB}(v_0') \geq \hat{\lambda}_0(v_0') > \lambda_0^{FB}(v_0')$.

We conclude that $\Phi \equiv \{v_0 : \lambda_0^{SB}(v_0) \geq \hat{\lambda}_0(v_0) > \lambda_0^{FB}(v_0)\}$ is non-empty. □

Proof of Proposition 2

Claim 11. If, at date zero, the Principal offers the long-term contract $\lambda_0 = \lambda_0^*(v_0)$, $b_0 = \frac{c}{v_0\lambda_0}$, $\lambda_1 = 0$ and $b_1 \geq 0$, there exists an equilibrium of the continuation game in which the Agent chooses $\sigma_0 = 1$ and $\sigma_1 = 0$.

The proof is similar to the proof for short-term contracts, and omitted.

Claim 12. In the continuation game after the Principal offers any contract $(\lambda_0, \lambda_1, b_0, b_1)$ satisfying $\lambda_0 > 0$, $\lambda_1 > 0$, $b_0 \geq 0$ and $b_1 \geq 0$, an equilibrium exists.

Proof. I specify the following strategy for the Agent: at date one, choose $e_1 = 1$ if $v_1(\lambda_0, e_0)\lambda_1 b_1 - c \geq 0$ and $e_1 = 0$, otherwise; at date zero, choose $e_0 = 1$ if:

$$v_0\lambda_0 b_0 - c + (1 - v_0\lambda_0)\beta \max\{0, v_1(\lambda_0, 1)\lambda_1 b_1 - c\} \geq \beta \max\{0, v_1(\lambda_0, 0)\lambda_1 b_1 - c\} \tag{48}$$

and $e_0 = 0$, otherwise. □

Claim 13. In any equilibrium of the continuation game after the Principal offers a long-term contract satisfying $\lambda_0 \le \hat{\lambda}_0(v_0)$ and $\lambda_1 > 0$, the Principal's payoff is strictly less than:

$$\max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \left(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta f(\lambda_1^*(v_1^p(\lambda_0, 1); v_1^p(\lambda_0, 1)) \right). \tag{49}$$

Proof. Notice that (49) is the surplus that is generated from the Principal's optimal investment which induces a positive value to experimentation at date two, when she holds the Agent to her participation constraint at both dates. Suppose, to the contrary, that the Principal attains this payoff. Then, her contract is: $\lambda_1 = \lambda_1^*(v_1^p(\lambda_0, 1))$, $b_1 = \frac{c}{v_1^p(\lambda_1, 1)\lambda_1}$, $b_0 = \frac{c}{v_0\lambda_0}$, and:

$$\lambda_0 \in \arg \max_{\lambda_0 \in [0, \hat{\lambda}_0(v_0)]} \left(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta f(\lambda_1^*(v_1^p(\lambda_0, 1); v_1^p(\lambda_0, 1)) \right). \tag{50}$$

and, the Agent must choose high effort at both dates in some equilibrium of the continuation game. However, in any continuation equilibrium after this contract is offered, the Agent weakly prefers to exert high effort at date zero if and only if:

$$v_0\lambda_0 b_0 - c + (1 - v_0\lambda_0)\beta \max\{0, v_1(\lambda_0, 1)\lambda_1 b_1 - c\} \ge \beta \max\{0, v_1(\lambda_0, 0)\lambda_1 b_1 - c\}. \tag{51}$$

or:

$$b_0 \ge \frac{c}{\lambda_0} \left(\frac{1}{v_0} + \left(\frac{1}{v_1^p(\lambda_0, 1)} + \frac{1}{v_0} \right) \right) > \frac{c}{v_0\lambda_0}. \tag{52}$$

Thus, there does not exist a continuation equilibrium after the Principal offers this contract in which the Agent exerts high effort at both dates with probability one. Thus, the Principal cannot attain (49). □

Let $S^{CO}(v_0)$ denote the supremum of the Principal's payoffs over all continuation equilibria for any long-term contract $(\lambda_0, \lambda_1, b_0, b_1)$ satisfying $\lambda_0 \le \hat{\lambda}_0(v_0)$ and $\lambda_1 > 0$.

Claim 14. If:

$$f(\lambda_0^*(v_0); v_0) > S^{CO}(v_0), \tag{53}$$

then there exists a unique equilibrium (up to the specification of $b_1 \ge 0$) in which the Principal offers the Agent a contract satisfying $\lambda_0 = \lambda_0^*(v_0)$, $\lambda_1 = 0$, $b_0 = \frac{c}{v_0\lambda_0}$ and $b_1 \ge 0$.

Proof. Consider a contract $\lambda_0 = \lambda_0^*(v_0)$, $b_0 = \frac{c}{v_0\lambda_0} + \Delta(b_0)$, where $\Delta(b_0)$ is chosen to satisfy:

$$0 < v_0\lambda_0^*(v_0)\Delta(b_0) < f(\lambda_0^*(v_0); v_0) - S^{CO}(v_0), \tag{54}$$

and $\lambda_1 = 0$ and $b_1 = 0$. There exists a unique equilibrium of the continuation game after this contract is offered in which the Agent choose $\sigma_0 = 1$ and $\sigma_1 = 0$, and the Principal's payoff is $f(\lambda_0^*(v_0); v_0) - v_0\lambda_0^*(v_0)\Delta(b_0) > S^{CO}(v_0)$. So, any contract satisfying $\lambda_0 \leq \hat{\lambda}_0(v_0)$ and $\lambda_1 > 0$ is strictly dominated.

Next, consider a contract satisfying $\lambda_0 > \hat{\lambda}_0(v_0)$ and $\lambda_1 > 0$. This contract gives the Principal a payoff which is strictly less than:

$$\max_{\lambda_0 \geq \hat{\lambda}_0(v_0)} \left(f(\lambda_0; v_0) + (1 - v_0\lambda_0)\beta f(\lambda_1^*(v_1^p); v_1^p) \right) \leq \max_{\lambda_0 \geq \hat{\lambda}_0(v_0)} f(\lambda_0; v_0) \leq f(\lambda_0^*(v_0); v_0). \tag{55}$$

By a similar argument to the previous paragraph, the Principal can therefore construct another contract $\lambda_0^*(v_0)$, $\lambda_1 = 0$, $b_0 = \frac{c}{v_0\lambda_0} + \Delta(b_0)$ for $\Delta(b_0)$ arbitrarily small which yields a unique continuation equilibrium payoff for the Principal that is arbitrarily close to $f(\lambda_0^*(v_0); v_0)$, and strictly larger than her payoff from a continuation equilibrium after $\lambda_0 > \hat{\lambda}_0(v_0)$ and $\lambda_1 > 0$. Finally, it is straightforward to show that a contract satisfying $\lambda_0 = 0$ is never optimal for the Principal.

Thus, $f(\lambda_0^*(v_0); v_0) > S^{CO}(v_0)$ implies the Principal's optimal contract has the property $\lambda_0 > 0$, $\lambda_1 = 0$. I finally show that the unique equilibrium (up to the specification of $b_1 \geq 0$) is as stated in the Claim. If the Principal offers a contract $\lambda_0 = \lambda_0^*(v_0)$, $b_0 > \frac{c}{v_0\lambda_0}$ and $\lambda_1 = 0$, $b_1 \geq 0$, the Agent strictly prefers $\sigma_0 = 1$ and $\sigma_1 = 0$. Call $\Delta(b_1) = b_1 - \frac{c}{v_0\lambda_0}$. Suppose that, after the Principal chooses a contract satisfying $\lambda_0 = \lambda_0^*(v_0)$, $\Delta(b_1) = 0$, $\lambda_1 = 0$ and $b_1 \geq 0$, the Agent chooses $\sigma_0 < 1$. Then, the Principal's payoff is:

$$\sigma_0 f(\lambda_0^*(v_0); v_0) < f(\lambda_0^*(v_0); v_0) - v_0\lambda_0^*(v_0)\Delta(b_0') \tag{56}$$

where the RHS is the Principal's payoff if she replicates the contract but instead chooses some $b_0' > \frac{c}{v_0\lambda_0}$ and where $\Delta(b_0')$ is sufficiently small to satisfy the above inequality. But this contract is itself strictly dominated by another contract which replicates it but instead satisfies $\Delta(b_0'') = \frac{1}{2}\Delta(b_0')$. Thus, there is a unique equilibrium (up to the specification of $b_1 \geq 0$) in which $\lambda_0 = \lambda_0^*(v_0)$, $\lambda_1 = 0$, $b_0 = \frac{c}{v_0\lambda_0}$ and $b_1 \geq 0$, and after which the Agent chooses high effort with probability one at date zero, and low effort with probability one at date one. $\square$

The rest of the result follows a similar structure to the proof of Proposition 1.

References

Bergemann, Dirk and Ulrich Hege. 2005. "The financing of innovation: learning and stopping." *RAND Journal of Economics* pp. 719–752.

Bhaskar, V. 2012. "Dynamic Moral Hazard, Learning and Belief Manipulation." Working Paper.

Biais, Bruno, Thomas Mariotti, Jean-Charles Rochet and Stéphane Villeneuve. 2010. "Large risks, limited liability, and dynamic moral hazard." *Econometrica* 78(1):73–118.

Halac, Marina, Navin Kartik and Qingmin Liu. 2013. "Optimal Contracts for Experimentation." Working Paper.

Hörner, Johannes and Larry Samuelson. 2013. "Incentives for experimenting agents." *The RAND Journal of Economics* 44(4):632–663.

Manso, Gustavo. 2011. "Motivating innovation." *The Journal of Finance* 66(5):1823–1860.