Dual Moral Hazard and the Innovator's Dilemma

Enrique Ide*

13th September, 2022

Abstract

I explain why current success can undermine an organization's ability to innovate. I consider a standard two-arm bandit experimentation problem with two modifications. First, a principal allocates resources. Second, the risky arm must be installed to deliver profits, which requires upfront effort from an agent. With incomplete contracts, a dual moral hazard problem emerges. Its severity increases with the payoffs of the safe arm, so the organization can end up worse off having a more profitable status quo. The model predicts which innovations are more difficult to adopt and supports the idea that independent business units can help incumbents innovate.

Keywords: Innovation, Principal-Agent, Dynamic Contracts, Experimentation.
JEL classifications: D23, D83, O30.

*IESE Business School, C/Arnús i Garí 3-7, 08034 Barcelona, Spain (email: eide@iese.edu). I am extremely grateful to my advisors, Andy Skrzypacz, Mike Ostrovsky, and Eddie Lazear. I have also benefited enormously from the suggestions of Nano Barahona, Daniel Barron, Alessandro Bonatti, Isa Chavez, Matías Covarrubias, Nicolás Figueroa, George Georgiadis, Nicolas Lambert, Ivan Marinovic, Juan-Pablo Montero, Cristóbal Otero, Sebastián Otero, Andrés Perlorth, Luis Rayo, Eduard Talamàs, Jean Tirole, Felipe Varas, Akhil Vohra, and participants at venues where this work (or previous incarnations of it) has been presented.

1 Introduction

Empirical studies have shown that new ventures and smaller firms introduce a large share of innovations.[1] Furthermore, numerous case studies suggest that successful incumbents usually struggle to innovate first and are sometimes driven to extinction by technological change.[2] The most prominent explanation for these two stylized facts is that of Arrow (1962): by successfully innovating, incumbents cannibalize their own profits. Hence, new firms have more to gain from innovating, as upon succeeding, they would replace an incumbent, not themselves.

However, increasing evidence indicates that Arrow's argument cannot account for the full picture; organizational factors also appear to play an important role. First, numerous case studies suggest that an organization's success in exploiting its current business can undermine its ability to explore new opportunities. The result is thus a tradeoff—a *dilemma*—between the organization's current profits and its long-term prospects.[3] The latter cannot be rationalized using pure profit-cannibalization arguments and has led some management scholars to argue in favor of setting independent business units to aid established firms in developing innovations (see, e.g., Christensen, 1997; Gans, 2016).

Second, systematic empirical evidence shows that successful firms struggle in particular with innovations that require new organizational knowledge (e.g., new processes, routines, or a different code/technical language). Henderson (1993), for instance, studies the photolithographic alignment equipment industry to empirically disentangle pure industrial organization theories of innovation failure (where incumbents optimally decide not to innovate) from theories stressing organizational factors. She shows that established firms were less effective than new firms at investing in necessary innovations, but only when the innovation required new organizational capabilities.

In this paper, I explain why current success can undermine an organization's ability to innovate, leading to a dilemma between an organization's current profits and its long-term prospects. My theory provides testable implications about which innovations will be more difficult for successful firms to adopt and is consistent with Henderson's (1993) findings that successful incumbents will struggle with innovations that require new organizational knowledge. Finally, my model supports the idea that setting up an independent business unit can help successful firms be more effective at innovating.

The Baseline Model.— I consider a standard experimentation problem in continuous-time where an organization must allocate resources between a safe arm S and a risky arm R. The safe arm gives a known flow payment while the risky arm might yield payoffs after exponentially distributed random times. To this setting, however, I incorporate two modifications. First, a (female) principal allocates the organization's resources. Second, the risky arm must be "installed" to have a chance of delivering profits. The latter requires costly non-contractible effort from a (male) agent that must be exerted upfront.

[1] See Scherer (1965), Gellman Research Associates (1976, 1982), Futures Group (1984), Pavitt, Robson and Townsend (1987), Acs and Audretsch (1988), and Igami (2017).

[2] See Christensen (1993) Henderson (1993), Klepper and Sleeper (2005), Garicano and Rayo (2016), Gans (2016).

[3] This idea has been equivalently dubbed the "innovator's dilemma" (Christensen, 1997), the "disruption dilemma" (Gans, 2016), or the "success syndrome" (O'Reilly and Tushman, 2016), among others.

The organization is completely described by these two players. Both are risk-neutral and discount the future at the same constant rate. Only the principal obtains direct benefits from either arm. The agent has no wealth and is protected by limited liability. Hence, to incentivize the agent to install R, the principal offers the agent a contract tied to the arm's success. I assume that players observe each other's actions and have access to the same information, so they hold common posteriors throughout the game. I also assume that R's first success provides irrefutable evidence that the arm is profitable.[4]

The safe and the risky arm thus represent, respectively, the current business and an uncertain innovation. The need for installation symbolizes that the innovation in question requires the development of new organizational capabilities to succeed (e.g., new processes, routines, or a new organizational language). The principal—having the authority to allocate the organization's resources and design the agent's compensation—embodies the organization's upper management. The agent, in turn, epitomizes the rest of the organization, i.e., middle managers and employees. Finally, the installation cost represents the effort that must be exerted upfront to develop the new organizational capabilities needed for the innovation's success.

Dual Moral Hazard and the Innovator's Dilemma.— To develop intuition, I begin by considering the case where the allocation of resources is not contractible, and the principal offers the agent a single time-independent bonus $B \geq 0$ payable upon R's first success ("a breakthrough").

The equilibrium, in this case, is always inefficient compared to the first-best outcome. First, there is a region of parameters where the agent installs R, but the organization experiments less with this arm than is efficient. Second, there is another region where installing R is efficient, but motivating the agent to execute the task is impossible.

These two inefficiencies have the same underlying explanation: the emergence of a dual moral hazard problem. Indeed, to motivate the agent to install R, the principal promises the agent a payment contingent on R's success. However, conditional on its installation, this payment induces the principal to stop experimenting with R at a higher belief—she is no longer the full residual claimant of R's profits—making it less likely that the agent receives the bonus promised in the contract. Anticipating the principal's incentives, the agent then demands a higher compensation to install the arm in the first place, which then feeds back into the problem as it induces the principal to experiment with R even less, and so on.

The *dilemma* between the profits of the current business and the organization's long-term prospects is then explained because a more profitable status quo increases the principal's temptation to return to the safe arm. The latter increases the severity of the dual moral hazard problem and explains why the principal's equilibrium expected payoffs at time zero (which also turns out to be the organization's welfare in equilibrium) are generically non-monotonic in S's profits. Hence, from an ex-ante perspective, the principal can end up worse off having a more profitable safe arm, even though the latter increase the fallback payoff if the innovation turns out not to be profitable.

[4]That the agent's installation decision is observable is not crucial but does simplify the analysis. The same goes for the assumption that any positive news on R's profits provides irrefutable evidence of its profitability. See Section 3.4 for details.

I then explore the role that uncertainty plays in creating this dual moral hazard problem. Interestingly, I find that ex-ante uncertainty on R's profitability is necessary but not sufficient; uncertainty must resolve slowly over time for the incentive conflict to arise. More precisely, I show that the dual moral hazard problem disappears—even with the very simple time-independent bonus I am considering—if the organization can learn infinitely quickly whether R is profitable or not.

This last result is important because it shows that not all innovations will be equally difficult for successful firms to adopt. In particular, current success will not lead to organizational rigidities if it is possible to grasp the value of the innovation in question quickly. This would be the case, for instance, of easily scalable innovations where preliminary experiments, prototypes, and minimum viable products provide a good sense of the value of the overall opportunity. Conversely, it shows the difficulty that successful firms will face when trying to develop innovations requiring extensive time and testing.

Finally, I discuss that a similar dual moral hazard problem arises if the organization has too many attractive opportunities in the future instead of an extremely profitable status quo. As I briefly describe below, this case fits particularly well with the experience that some small laboratories have had when working with big pharmaceutical companies to develop new drugs.

The Optimal Contract: The Allocation of Resources is Contractible.— I then move on to characterize the optimal contract under different contractual assumptions. I first show that when the allocation of resources is contractible,[5] the principal can achieve her first-best outcome. This result follows because, with exponential bandits and conclusive news, the first-best resource allocation policy is a stopping-time policy with a known and fixed stopping time T^{fb}.

As a result, the principal can implement the efficient outcome by committing to a lockup period of duration T^{fb} during which all of the organization's resources are directed towards R regardless of the outcome. Once the lockup period ends, the principal regains the flexibility to allocate resources as she wishes. The agent can then be incentivized to install R with the help of a single time-independent bonus $B \geq 0$ contingent R's first success. Alternatively, the principal can create an independent business unit to develop the innovation. That is, it transfers the agent the decision rights on the allocation of resources for an overall budget of T^{fb} units of the resource (while offering the same time-independent bonus B contingent on a breakthrough).[6]

Of course, these results rely critically on the fact that the first-best experimentation policy is a stopping-time policy with a known and fixed stopping time T^{fb}. That is why, as a robustness check, I also consider an extension where players receive additional information on R's profitability once the arm is installed. The latter implies that the efficient stopping time is now unknown at the time

[5]Whether the allocation of resources is contractible or not will depend on the interpretation given to the resource. If the resource in question is monetary in nature, this is more likely to be contractible. The opposite would be the case if, instead, the resource represents the principal's priorities or attention. In reality, of course, one would expect a mixture of both interpretations.

[6]This interpretation of the optimal contract appears to be very close to what Clayton Christensen had in mind when he advocated in favor of independent business units: "How separate does the effort need to be? The primary requirement is that the project cannot be forced to compete with projects in the mainstream organization for resources... Whether the independent organization is physically separate is less important than is its independence from the normal resource allocation process" (Christensen, 1997, p.139).

of contracting.

I then characterize, drawing heavily on the tools developed by Amador, Werning and Angeletos (2006) and Guo (2016), under what conditions the first-best is implementable, and show that, when that is not the case, the principal faces a tradeoff between *commitment* and *flexibility*.[7] Nevertheless, the optimal contract can still be interpreted as setting an independent business unit with an appropriately calibrated budget.

The Optimal Contract: The Allocation of Resources is Not Contractible.— Finally, I characterize the optimal contract when the allocation of resources is not contractible. My first result states that the principal achieves, again, her first-best outcome in this case. Intuitively, the principal indirectly contracts upon the allocation of resources by committing to pay the agent damages if the safe arm delivers profits before the efficient stopping time T^{fb}.

However, this result relies critically on the fact that S's profits are deterministic (instead of known but stochastic as in, e.g., Bolton and Harris, 1999) and that players can perfectly contract on S's profits. This last assumption seems particularly optimistic. For instance, it might be difficult to contractually specify what constitutes the "current business" exactly. Or, even if that is not the case, there could be ways to allocate resources to S to manipulate or backload its profits (think of R&D investments to achieve incremental innovations in S).

For all these reasons, I also characterize the optimal contract when contracting on S's profits is unfeasible. In this case, it is without loss of generality to restrict attention to a single time-contingent bonus $B_t \geq 0$ paid upon R's first success.[8] I show that the principal's optimal time-contingent contract now implements the first outcome only when S's profits are low enough; otherwise, the dual moral hazard problem described previously claws its way back into the scene and the innovator's dilemma, unraveling included, emerges again.

Is there evidence of this dual moral hazard problem?— In my theory, an organization's success in exploiting its current business undermines its ability to explore new opportunities because it impairs upper management's ability to commit to new and unproven courses of action. Although there is no systematic empirical evidence on whether this particular mechanism is important, there is some anecdotal one.

Consider, for instance, the case of German firm BioNTech, creators of the BioNTech-Pfizer COVID-19 Vaccine (along with pharmaceutical giant Pfizer). In a recent academic symposium, Nobel Prize recipient Oliver Hart asked BioNTech founder, Dr. Ugur Sahin, why he decided to create BioNTech as a separate enterprise instead of directly working for Pfizer at an outset. His answer

[7] On the one hand, the principal can mitigate her time inconsistency problem and reduce the agent's compensation by committing to experiment with R for a minimum pre-specified amount of time. On the other hand, the latter limits her ability to react to the new information that arrives once R is installed.

[8] When neither the allocation of resources nor S's profits are contractible, the agent's compensation can be made contingent on time, the number of successes with R, and the payments the agent has received in the past. That focusing on single time-contingent bonus $B_t \geq 0$ paid upon a breakthrough is without loss follows because players are risk-neutral and R's first success provides definitive evidence that the arm is profitable.

was revealing:

> [Before BioNTech] we tried to do collaborations with the pharmaceutical industry... and the collaborations were always fine in the first two, three years. But after a few years the management... decided to stop the collaboration because the company was going to a different direction.
>
> And it was very frustrating to see that you invest time into a collaboration, the collaboration is running fine and it is still stopped. We understood that if we wanted to develop medicines we had to start our own company.
>
> Academic Symposium on Relational Contracts (2021).[9]

Ughur Sahin's experience resonates well with the theory I develop. This is especially so, considering how drug discovery and development is a complex and time-intensive process involving not only a high degree of uncertainty but also extensive testing.[10]

1.1 Related Literature

First, my paper contributes to the literature stressing the importance of upper management commitment to employee motivation (Prendergast, 1993; Rotemberg and Saloner, 1993, 1994; Almazan and Suarez, 2003; Van den Steen, 2005; Bolton, Brunnermeier and Veldkamp, 2012). This literature has used the emergence of a dual moral hazard problem—similar in broad strokes to the one developed here—to explain a variety of phenomena, from the role of promotions in inducing specific human capital acquisition to the benefits of CEO entrenchment for optimal corporate governance.

Among all these papers, the ones by Rotemberg and Saloner (1993, 1994; from hereafter RS93/94) are the closest to mine. RS93/94 study a static principal-agent model with incomplete contracts in which an agent can generate projects of uncertain value but the principal retains control over whether to implement the projects after their value is revealed.

The authors then show that if the project's value is drawn from a distribution with continuous support, then a dual moral hazard problem emerges. However, if the agent's effort leads to only one good outcome (he obtains a "good" idea with some probability and a "bad," worthless idea otherwise), then the incentive friction disappears. They then argue that organizations will therefore benefit from following "narrow business strategies," since pursuing a limited number of activities is more likely to generate "home-run" innovations in which the agent's effort generates either a very profitable or a worthless idea.

---

[9] Video of the exchange available at: https://www.youtube.com/watch?v=2Njd5TiW8Do&ab_channel=SaidBusinessSchool%2CUniversityofOxford — min 37.15 (last accessed: September 2nd, 2022).

[10] For instance, according to the FDA, of those drugs that make it to clinical trials, 94% fail, while according to Brown et al. (forthcoming), of those that succeed, the average clinical development time is 9.1 years. The entire process costs drugmakers $2.6 billion in 2013 dollars (DiMasi, Grabowski and Hansen, 2016). For the FDA data, see https://www.fda.gov/patients/drug-development-process/step-3-clinical-research (last accessed: September 12, 2022).

My paper also contributes to the literature studying how organizational factors could handicap innovation efforts by established firms.[11] For instance, Christensen (1993, 1997) asserts that established firms fail to react to new-entrant innovations because they come to the incorrect conclusion that their current customers would not want the new offering. In contrast, Henderson and Clark (1990) posit that established firms will find it hard to recognize "architectural innovations" (i.e., changes in how the different components of a product are linked together) since the current product architecture is embedded in the organization's structure and information-processing procedures.

Finally, my paper is related to the literature on experimentation in principal-agent settings (e.g., Bergemann and Hege, 1998, 2005; Manso, 2011; Horner and Samuelson, 2013; Guo, 2016; Halac, Kartik and Liu, 2017; Moroni, 2020). This literature, however, is usually concerned with how to design contracts to keep the agent motivated through time. My paper, in contrast, deals with how to design contracts so as to overcome the principal's time inconsistency problem that arises due to limited commitment.

The remainder of the paper is structured as follows. Section 2 formally lays out the baseline model and the first-best outcome. Section 3 characterizes the equilibrium under simple contracts, showing the emergence of the dual moral hazard problem and the dilemma between the organization's current profits and its long-term prospects. Section 4 characterizes the optimal contract under different contractual assumptions. Finally, Section 5 concludes.

2 The Baseline Model

The Environment.— Time $t \in [0, \infty)$ is continuous,[12] and the discount rate is $r > 0$. There are two players, a (female) principal, and a (male) agent. They are collectively known as "the organization."

The principal is endowed with one unit of a perfectly divisible resource per unit of time. She continually splits the resource between a safe arm S and a risky arm R. I begin by assuming that the allocation of resources is non-contractible and consider the contractible case in Section 4.

Only the principal obtains direct benefits from either arm; the agent benefits solely from the compensation delivered by the principal. The safe arm is ready to deliver profits, generating a known flow payoff that depends on the resources allocated to it. The risky arm, in contrast, must first be installed to have a chance of producing profits. Once installed, it generates positive lump-sum payoffs stochastically over time with an arrival rate that depends on an unobservable and persistent binary state $\theta \in \{0, 1\}$ and the amount of the resource allocated to it.

Formally, if R is installed and the principal allocates a fraction k_t of the resource to it over the interval $[t, t + dt)$ (and, consequently, she allocates $1 - k_t$ to S), she then obtains a lump-sum amount $\Pi > 0$ from R with instantaneous probability $\lambda\theta k_t$ plus an instant flow payment equal to $(1 - k_t)s$

[11] Awaya and Krishna (forthcoming) provide an alternative explanation based on neither cannibalization of profits nor organizational factors. Their theory, instead, relies on established firms having better information about the innovation's prospects, which dulls their incentives to learn from rivals.

[12] It is possible to obtain similar results in a three-period model. However, characterizing the equilibrium for all possible parameters requires going over numerous cases, so the analysis loses transparency. The continuous-time formulation helps avoid this issue.

from S. In contrast, if R is not installed and the principal makes the same allocation of resources, she obtains zero from R and $(1-k_t)s$ from S.[13]

Only the agent can install the risky arm.[14] The latter requires him to incur a one-time non-contractible installation cost $f > 0$. This cost symbolizes the effort needed to develop the necessary organizational capital for the innovation to have a chance of success. For instance, it could represent the effort exerted to develop new processes and routines (e.g., Nelson and Winter, 1982; Chassang, 2010) or to establish a new code/technical language to communicate within the organization (e.g., Arrow, 1974; Henderson and Clark, 1990; Crémer, Garicano and Prat, 2007).[15]

Information and Beliefs.— The parameters λ, s, Π, and f are all publicly known, but the persistent binary state $\theta \in \{0,1\}$ is unobservable, as mentioned previously. The prior probability that $\theta = 1$ is $\bar{q} \in (0,1)$. I assume that all players observe each other's actions, so they hold common posterior beliefs throughout the game. This implies that the agent's installation decision is observable (but not contractible, as stated above) and that the first lump-sum arrival (or "breakthrough") reveals to both players that $\theta = 1$.[16,17]

Denote then by q_t the common belief at t that $\theta = 1$. If at time t the risky arm has not yet been installed, then no new information arrives, so no updating takes place: $q_t = \bar{q}$. In contrast, if R was installed at some $t' \leq t$, then by Bayes' rule, q_t evolves smoothly as long as there is no breakthrough, drifting downwards at a rate proportional to k_t, i.e., "no news is bad news":

$$dq_t = -\lambda k_t q_t (1-q_t) dt \tag{1}$$

Of course, once there is a breakthrough, the posterior jumps up from q_{t^-} to 1 and stays there forever.

Regarding the model's parameters, I make the following assumptions: (i) $f < \lambda \bar{q} \Pi / r$, and (ii) $0 < s < \lambda \Pi$. Condition (i) implies that it is first-best optimal for an organization without a safe arm to install R, while condition (ii) states that conditional on installation, the risky arm is preferred to the safe arm by the organization as a whole, but only when $\theta = 1$.

Timing.— Time $t = 0$ is divided into two logically sequential dates denoted by $t = 0^-$ and $t = 0^+$. At time $t = 0^-$, the principal offers the agent an incentive scheme to induce him to install R (the nature of the scheme is described below). At $t = 0^+$, and after observing the compensation offered by the principal, the agent decides whether to install the risky arm ($a_0 = 1$), incurring the cost $f > 0$, or not ($a_0 = 0$). Finally, at every $t > 0$, the principal decides how to allocate the unit of available

[13] To focus on the frictions that arise in trying to motivate the agent to install the risky arm, I omit the costs of providing incentives to the agent to run the everyday business (i.e., the costs of incentivizing the agent to produce output). Hence, the arms' profits should be interpreted as the organization's profits when the incentive costs of producing output have been taken into account.

[14] This could be, for example, because middle managers and employees could have better organizational knowledge of the firm's processes. Alternatively, they could be the owners of some inalienable asset, such as skilled human capital.

[15] The results do not change if R's installation also requires the principal to incur in a positive fixed cost (which could symbolize, for instance, the cost of acquiring new physical capital).

[16] That R's installation is observable is not crucial but does simplify the analysis (see Section 3.4 for details).

[17] The baseline model only considers the case of "conclusive" news. As discussed in Section 3.4, results do not significantly change if we also incorporate inconclusive news into the model.

resources among the two arms; i.e., she chooses $k_t \in [0,1]$.[18]

The Incentive Scheme.— To develop intuition, I begin by considering a simple yet natural incentive scheme: the principal offers the agent a time-independent bonus $B \geq 0$ payable upon a breakthrough (i.e., R's first success).[19] These *single-bonus contracts* are common in practice, as extensively documented by Murphy (1999), Prendergast (1999), and Holmstrom (2016). In Section 4, I characterize and discuss the optimal contract under different contractual assumptions.

Equilibrium Concept, Strategies, and Payoffs.— I focus on subgame perfect equilibria. A strategy for the agent corresponds to an installation choice at $t = 0^+$, $a_0(B) \in \{0,1\}$, as a function of the bonus offered by the principal.

A strategy for the principal, in turn, consists of (i) a bonus offer at $t = 0^-$, $B \geq 0$, and (ii) a resource allocation decision for every $t > 0$: $k_t(q_t, a_0, B) \in [0,1]$. This last mapping stipulates the allocation of resources at time $t > 0$ as a function of the following:

1. The current belief that $\theta = 1$,
2. Whether the agent installed R or not at $t = 0^+$,
3. The bonus B that the principal promised to the agent.

Let N_t be the number of successes with R observed up to time t, including t, and define τ as the time at which the first success occurs, i.e., $\tau \equiv \inf\{t \geq 0 : N_t = 1\}$. Given players' actions, the principal's and the agent's expected payoffs at time 0, v_0 and u_0, expressed in per-period units, are given by:

$$\begin{aligned} v_0 =& \mathbb{E}_0\left[\int_0^\infty re^{-rt}\big[(a_0 k_t \lambda\theta)\Pi + s(1-k_t)\big]dt - re^{-r\tau}B\right] \\ u_0 =& \mathbb{E}_0\left[re^{-r\tau}B\right] - ra_0 f \end{aligned}$$

2.1 The First-Best Benchmark

Before moving on to the main analysis, it is convenient to describe the first-best outcome. Lemma 1 characterizes the optimal resource allocation decision at any $t > 0$. Proposition 1 then uses the lemma to describe the optimal installation decision at $t = 0$.

Lemma 1. *Let $g \equiv \lambda\Pi$, $\mu \equiv r/\lambda$, and define $w(q_t, a_0; s)$ as the organization's value as a function of (q_t, a_0) when S gives s per unit of allocated resource. The optimal resource allocation at any given*

[18] Restricting the principal to offering an incentive scheme only at $t = 0^-$ and the agent to disbursing f only at $t = 0^+$ is without loss of generality in this setting. This is because no useful information arrives if R remains uninstalled. Of course, this would no longer hold if the principal had access to additional information that did not depend on the agent's action. In such a model, the principal would also have to decide *when* to install R. As I discuss in Section 3.4, the results do not significantly change in this alternative model.

[19] Since players are risk-neutral, this is equivalent to the principal offering the agent a constant fraction $\alpha \in [0,1]$ of R's proceeds (i.e., a "linear" or "equity" contract). Furthermore, given that single lump-sum arrival provides definitive news that $\theta = 1$, these contracts are also equivalent to a compensation scheme that depends only on N_t, the number of successes with R observed up to time t.

$t > 0$ *is given by:*

$$k_t^{fb}(q_t, a_0 = 0; s) = 0 \text{ for all } q_t \in [0,1]$$
$$k_t^{fb}(q_t, a_0 = 1; s) = \begin{cases} 1 & \text{if } q_t > q^{fb} \\ 0 & \text{otherwise} \end{cases} \quad \text{where } \; q^{fb} \equiv \frac{\mu s}{\mu s + (1+\mu)(g-s)}$$

Consequently, if R is not installed, then $w(q_t, a_0 = 0; s) = s$, while if R is installed, it is given by:

$$w(q_t, a_0 = 1; s) = \begin{cases} gq_t + (s - gq^{fb}) \left(\frac{1-q_t}{1-q^{fb}}\right) \left(\frac{\Omega(q_t)}{\Omega(q^{fb})}\right)^{\mu} & \text{if } q_t > q^{fb} \\ s & \text{otherwise} \end{cases}$$

where $\Omega(q) \equiv (1-q)/q$ denotes the odds ratio at the belief q.

Proof. See Appendix for all proofs. □

If R is not installed at time $t = 0$, then the optimal policy is to allocate all of the resources to S, as R never delivers any profits. In contrast, if R is installed, the optimal policy is characterized by a cutoff belief q^{fb} such that $k_t^* = 1$ if $q_t \geq q^{fb}$ and $k_t^* = 0$ otherwise. Given that under this policy beliefs evolve deterministically in time in the absence of a breakthrough (see equation (1)), the optimal policy can be equivalently understood as a stopping-time policy with a fixed stopping time equal to:

$$k_t^*(N_t = 0) = \begin{cases} 1 & \text{if } t \leq T^{fb}(\bar{q}) = (\ln \Omega(q^{fb}) - \ln \Omega(\bar{q}))/\lambda \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad k_t^*(N_t) = 1 \text{ for all } N_t \geq 1$$

That is, if a breakthrough occurs before $T^{fb}(\bar{q})$, the organization allocates all resources to R forever; otherwise, the organization allocates all resources to R until the stopping time and switches to S from thereafter.

With this at hand, I can now characterize the efficient installation decision for the organization as a whole:

Proposition 1. *In the first-best, it is (strictly) optimal to install R whenever $s < s^{fb}$, where $s^{fb} \in (0, g)$ is given by the (unique) solution of:*

$$w(\bar{q}, a_0 = 1; s^{fb}) - rf = s^{fb}$$

Consequently, the first-best expected payoff at $t = 0$, w_0^{fb}, is:

$$w_0^{fb} = \begin{cases} w(\bar{q}, a_0 = 1; s) - rf & \text{if } s \leq s^{fb} \\ s & \text{otherwise} \end{cases}$$

As expected, installing R is optimal when s is not too large. This is because the constraint on resources makes S and R substitutes, so installing R effectively replaces S. Crucially, though, w_0^{fb} is

always strictly increasing in s, as it increases the fallback payoff in case the risky arm turns out not to be profitable. Hence, even though a more attractive status quo can explain why the organization might decide not to adopt an innovation, it does not lead to a *dilemma*: There is no tradeoff between a more profitable business today, i.e., a higher s, and the "long-term" prospects of the organization, i.e., w_0^{fb}.

3 The Equilibrium

In this section, I study the equilibrium of the game and show the emergence of a dilemma. In particular, I show that v_0^*, the principal's equilibrium expected payoffs at time 0 (which also turns out to be the organization's welfare in equilibrium), is non-monotonic on s. This is because a more profitable safe arm hinders the organization's *ability* to install the risky arm.

Thus, from hereafter I assume that s is such that it is strictly optimal for the organization to install R in the first-best outcome. Otherwise, it is never an equilibrium to install the risky arm and the equilibrium of the game is trivial.

Assumption 1. *$s < s^{fb}$, where $s^{fb} \in (0, g)$ is given by Proposition 1.*

3.1 Characterization

To characterize the equilibrium, I follow the same structure as in the first-best. First, I obtain the principal's equilibrium resource allocation at any $t > 0$. Then, I use this as an input to characterize the players' behavior at $t = 0$.

Lemma 2. *Define $v(q_t, a_0, B; s)$ and $u(q_t, a_0, B; s)$ as the principal's and the agent's value as a function of (q_t, a_0, B), respectively. The principal's equilibrium resource allocation at any given $t > 0$ is given by:*

$$k_t^*(q_t, a_0 = 0, B; s) = 0 \text{ for all } q_t \in [0, 1]$$

$$k_t^*(q_t, a_0 = 1, B; s) = \begin{cases} 1 & \text{if } q_t > q^* \\ 0 & \text{otherwise} \end{cases}, \text{ where } q^* \equiv \frac{\mu s}{\mu s + (1+\mu)(g-s) - \lambda\mu B}$$

Consequently, if the agent does not install R, then $v(q_t, a_0 = 0, B; s) = s$ and $u(q_t, a_0 = 0, B; s) = 0$, respectively. In contrast, if the agent installs R, then:

$$v(q_t, a_0 = 1, B; s) = \begin{cases} gq_t + (s - gq^*)\left(\frac{1-q_t}{1-q^*}\right)\left(\frac{\Omega(q_t)}{\Omega(q^*)}\right)^{\mu} - u(q_t, a_0 = 1, B; s) & \text{if } q_t > q^* \\ s & \text{otherwise} \end{cases}$$

$$u(q_t, a_0 = 1, B; s) = \begin{cases} \lambda B q_t \left(\frac{\mu}{1+\mu}\right)\left[1 - \left(\frac{\Omega(q_t)}{\Omega(q^*)}\right)^{1+\mu}\right] & \text{if } q_t > q^* \\ 0 & \text{otherwise} \end{cases}$$

The principal's equilibrium behavior at $t > 0$ is very similar to the optimal resource allocation in the first best. In particular, when R is installed, the principal also follows a cutoff strategy, i.e., $k_t^* = 1$ if $q_t \geq q^*$ and $k_t^* = 0$ otherwise. The equilibrium cutoff q^*, however, is strictly higher than the efficient cutoff (at least when $B > 0$) and, furthermore, is strictly increasing in B. This is because the principal does not internalize the totality of R's profits due to the compensation promised to the agent. Thus, the stopping-time associated with this strategy is now lower:

$$T^*(\bar{q}; B) = \frac{1}{\lambda} \ln \left(\frac{\Omega(q^*)}{\Omega(\bar{q})} \right) < T^{fb}(\bar{q})$$

With these results in hand, I can now obtain the players' equilibrium behavior at time $t = 0$. Consider first the agent's installation decision at $t = 0^+$ given a $B \geq 0$. If the agent does not install R, then he obtains a continuation value of $u(\bar{q}, a_0 = 0, B; s) = 0$, and the principal obtains $v(\bar{q}, a_0 = 0, B; s) = s$. In contrast, if the agent installs R, he obtains $u(\bar{q}, a_0 = 1, B; s) - rf$, and the principal obtains $v(\bar{q}, a_0 = 1, B; s)$. Clearly, then:

$$a_0^*(B; f, s) = \begin{cases} 1 & \text{if } rf \leq u(\bar{q}, a_0 = 1, B; s) \\ 0 & \text{otherwise} \end{cases}$$

and so at $t = 0^-$, the principal solves:

$$\max_{B \geq 0} \left\{ s + a_0^*(B; f, s) \left[v(\bar{q}, a_0 = 1, B; s) - s \right] \right\}$$

Proposition 2. *The equilibrium bonus offered by the principal is:*

$$B^* = \begin{cases} \underline{B}(f, s) & \text{if } s \leq s^* \\ [0, \infty) & \text{otherwise} \end{cases}$$

where $\underline{B}(f, s) \equiv \inf\{B \geq 0 : u(\bar{q}, a_0 = 1, B; s) = rf\}$ *and* $s^* \in (0, s^{fb})$ *is given by the (unique) solution to:*

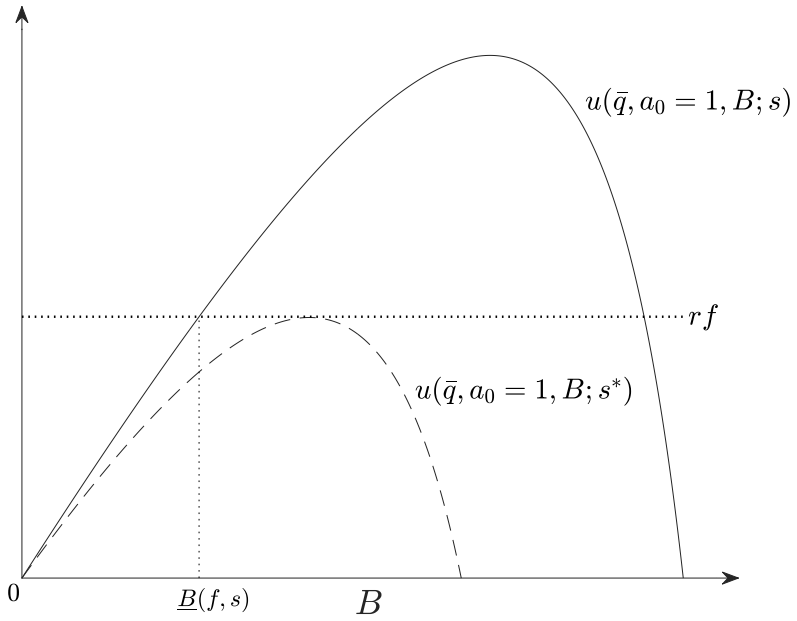
$$\max_{B \geq 0} u(\bar{q}, a_0 = 1, B; s^*) = rf$$

In equilibrium, the agent installs R whenever $s \leq s^*$, and players' equilibrium expected payoffs at $t = 0$ are:

$$u_0^* = 0 \qquad \text{and} \qquad v_0^* = \begin{cases} v(\bar{q}, a_0 = 1, \underline{B}(f, s); s) > s & \text{if } s \leq s^* \\ s & \text{otherwise} \end{cases}$$

The key to understanding Proposition 2 is that the agent's continuation value of installing R, $u(\bar{q}, a_0 = 1, B)$, is non-monotonic in B (see Figure 1). This implies that offering the agent a higher bonus does not necessarily translate into stronger incentives to install R. Intuitively, although a higher B increases the amount the agent receives in the event of a breakthrough, it also decreases the principal's incentives to experiment with R, making it less likely that the bonus will actually be paid.

Figure 1: The Equilibrium Contract when $s < s^*$



The fact that $u(\bar{q}, a_0 = 1, B; s)$ is non-monotonic in B, implies that when s is sufficiently high (i.e., $s > s^*$), then it is *impossible* to incentivize the agent to install R. The principal then has no option but to play the safe arm so her equilibrium expected payoff is $v_0^* = s$. In contrast, when $s \leq s^*$, incentivizing the agent is feasible. The optimal compensation scheme then entails offering the lowest value of B that leaves the agent indifferent between installing R or not.

3.2 Dual Moral Hazard and The Innovator's Dilemma

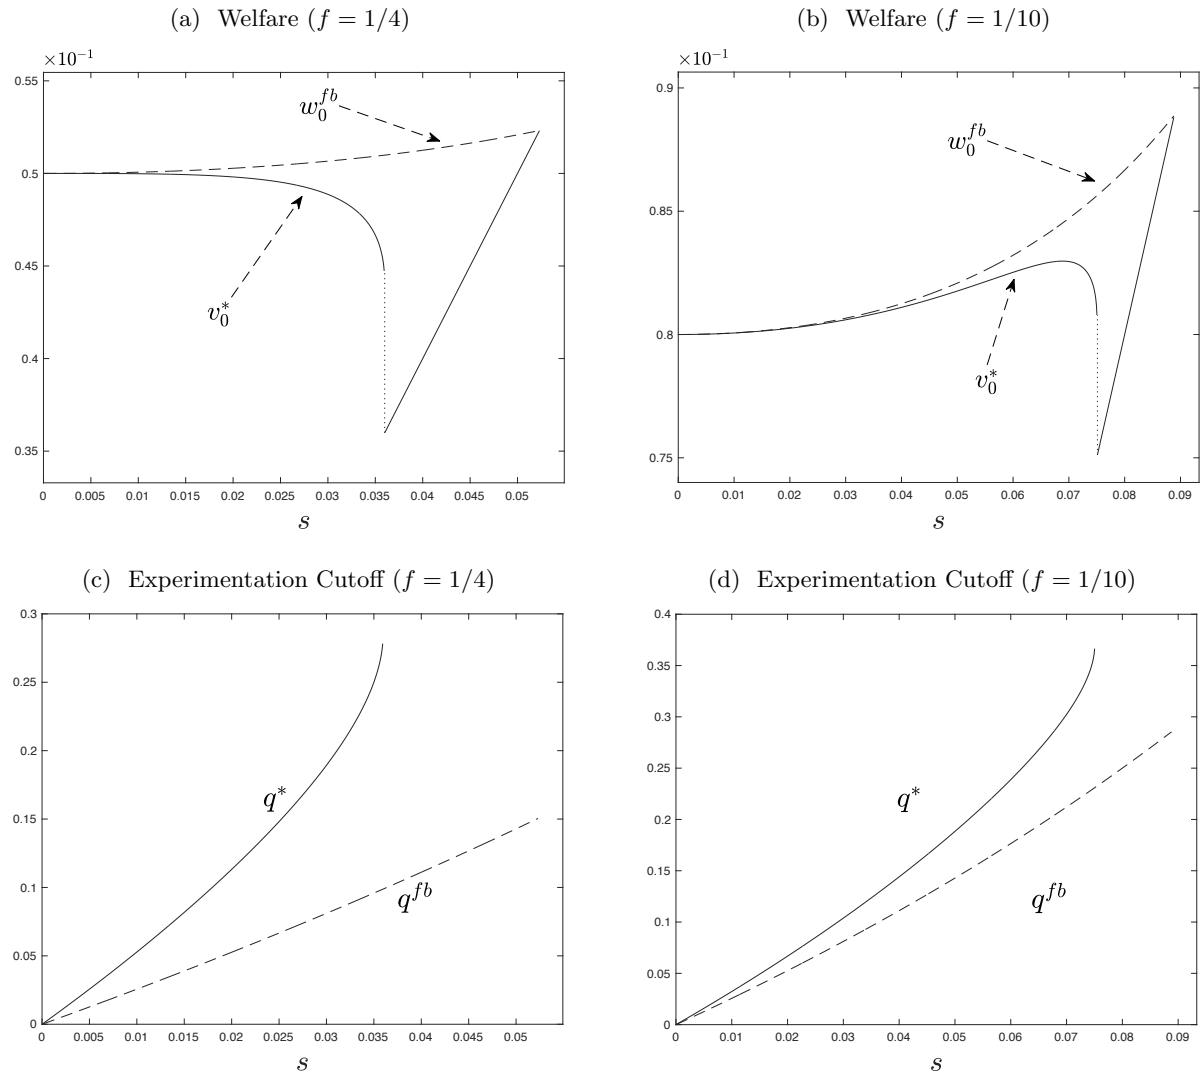
Compared to the first-best outcome, the equilibrium is inefficient, i.e., $w_0^{fb} > v_0^* + u_0^* = v_0^*$ (see Figure 2).[20] First, when $s \in (0, s^*]$, the agent installs R, but the organization experiments less with this arm than what is efficient; i.e., $q^* > q^{fb}$. This is an inefficiency on the "intensive" margin. Second, when $s \in (s^*, s^{fb})$, installing R is efficient, but motivating the agent to execute the task is impossible. This is an inefficiency on the "extensive" margin.

These two different inefficiencies have the same underlying explanation: the emergence of a dual moral hazard problem. The intuition is as follows. To motivate the agent to install R, the principal promises the agent a bonus contingent on R's success. However, conditional on R being installed, this promise alters the attractiveness of allocating resources to the risky arm relative to the status quo, as the principal is no longer the full residual claimant of R's profits. The latter incentivizes the principal to stop allocating resources to R at a higher belief, making it less likely that the agent receives the bonus offered. The result is a disagreement ex-post (after R is installed) between the principal and the agent over the appropriate allocation of resources.

The problem is then compounded by the fact that anticipating this disagreement, the agent

[20] Check `https://www.dropbox.com/s/km4fquqgwhdg6sx/2022_07_16_MATLAB.zip?dl=0` for the different `MATLAB` codes used to create the figures presented throughout the paper.

Figure 2: The Equilibrium vs. The First-Best



Notes. Remaining parameters: $\Pi = 1$, $r = 1/5$, $\lambda = 1/5$, $\bar{q} = 1/2$.

requires a higher bonus to install the arm in the first place. This then feeds back into the problem, as it induces the principal to experiment with R even less, and so on. Because in equilibrium the agent always breaks even, all the inefficiency ends up falling back on the principal through the higher bonus needed to incentivize the agent. The principal is thus the victim of a time-inconsistency problem.

In a sense, the incentive conflict resembles the one that arises in the endogenous default model in finance (e.g., Leland, 1994). Indeed, because the innovation requires the agent to front-load effort, and given that effort is not contractible, the agent is effectively making a loan to the principal when installing the risky arm: he is incurring f today in expectation of future benefits. However, once the loan is made (i.e., the risky arm is installed) and given that the principal's resource allocation policy is not contractible either, these committed benefits induce the principal to "default" (return to the status quo) sooner than is efficient when she grows pessimistic about the innovation's prospects. This default risk is then priced ex-ante, making it harder and more expensive to motivate the agent

to install R in the first place.

As Figure 2 also illustrates, in equilibrium the principal's ex-ante payoff/the organization's welfare is generically non-monotonic in S's profits. For s low enough, the incentive friction is of second-order importance since the principal has little temptation to return to the status quo. Hence, v_0^* is strictly increasing in s as in the first-best. However, as s continues to increase, so does the principal's ex-post temptation to switch back to the safe arm. The latter compounds the dual moral hazard problem, leading to a *dilemma* between the current business' profits and the organization's long-term prospects: v_0^* eventually becomes decreasing in s. Finally, for s high enough, there is no bonus capable of inducing the agent to install R since the incentive conflict unravels. This explains the discontinuous downward jump of v_0^* at $s = s^*$, and the fact that for $s \in (s^*, s^{fb})$, v_0^* becomes increasing in s again, as the organization is "stuck" with S despite the presence a potentially superior technology R.[21]

The dual moral hazard problem characterized in Proposition 2 provides support to Ugur Sahin's experience regarding the difficulties of developing novel medicines and vaccines as part of established pharmaceutical companies (see his exchange with Oliver Hart in the Introduction). Furthermore, the proposition can also explain two phenomena often documented in the management literature.

The first is the existence of a dilemma between current success and the organization's long-term prospects, which are given by its ability to adopt and adapt to innovations (Christensen, 1997; O'Reilly and Tushman, 2016). The second is the importance of *organization rigidities* (as opposed to pure profit-cannibalization) in explaining the way organizations react to innovations, especially those requiring new organizational knowledge (Henderson, 1993).

3.3 The Role of Uncertainty

I now explore the role of uncertainty in creating this dual moral hazard problem. In particular, I show that ex-ante uncertainty on the profitability of R is necessary but not sufficient; uncertainty must resolve slowly over time for the incentive conflict to arise.

Corollary 1.

(i) $v_0^* \to w_0^{fb}$ *as* $\bar{q} \to 1$.

(ii) For any $\bar{q} \in (0,1)$*, then* $v_0^* \to w_0^{fb}$ *as* $\lambda \to \infty$ *and* $\Pi \to 0$ *with* $g \equiv \lambda\Pi$ *constant.*

The first part of Corollary 1 shows the necessity of ex-ante uncertainty, stating that the dual moral hazard problem disappears as players become ex-ante certain that R is superior to S. Intuitively, when the principal becomes ex-ante more optimistic about R's prospects, there is a decrease in the ex-ante risk of the principal ex-post switching to S inefficiently early.

Interestingly, this result could explain some anecdotal evidence suggesting that, even though they usually struggle to innovate first, incumbent firms are relatively successful second-mover adopters,

[21] Naturally, the welfare loss in equilibrium compared to the first-best (as measured by $w_0^{fb} - v_0^*$) disappears when either $s \to 0$ or $s \to s^{fb}$ (and is, of course, non-monotonic in s). When $s \to 0$, there is no inefficiency because the principal has no temptation to return to the status quo, as mentioned above. In contrast, when $s \to s^{fb}$ the inefficiency vanishes because there is no value in installing R. Thus, "being caught in the middle" (that is, with intermediate values of s) is particularly problematic for the organization. This would be the case when R provides a significant, but not extreme profit improvement over the status quo.

see, e.g., Markides and Geroski (2005) and Gans (2016, pp. 123-124). This is because once a new entrant develops an innovation, the uncertainty surrounding the innovation decreases, ameliorating the organizational rigidities deterring successful incumbents from adopting the innovation in the first place.

The second part of the corollary shows, however, that ex-ante uncertainty is necessary but not sufficient. In particular, it states that keeping fixed at g the expected per-period profits of R conditional on $\theta = 1$, the incentive conflict also disappears as the arrival rate of R's payments, λ, goes to infinity. That is, there is no dual moral hazard problem if the organization can learn infinitely quickly whether the risky arm is profitable or not.

Intuitively, the principal's incentives to experiment with R increase with the rate at which she can learn its profitability. Hence, a higher λ induces the principal to wait for more evidence of the unprofitability of R before switching back to S. As a result, there is less scope for an ex-post disagreement between the principal and the agent, since both players agree that it is optimal to allocate resources to R only if the latter is profitable.[22]

This last result is important because it shows that not all innovations will be equally difficult for successful firms to adopt. In particular, current success will not lead to organizational rigidities if it is possible to grasp the value of the innovation in question quickly. This would be the case, for instance, of easily scalable innovations where preliminary experiments, prototypes, and minimum viable products provide a good sense of the value of the overall opportunity. Conversely, it reinforces the difficulty of developing innovations that require extensive time and testing (as is the case of novel medical treatments, for example) within successful firms.

3.4 Discussion

Before analyzing more complex contracts (and the role played by different contractual assumptions), I briefly discuss some of the modeling assumptions adopted in the baseline setting.

Observability of Effort.— The assumption that the agent's effort is observable (though not verifiable) is not crucial but it does simplify the analysis. In particular, it is easy to see that the outcome described in Proposition 2 (which is the unique equilibrium with observable effort) is also *an equilibrium* of the game in which the agent's installation is not observable to the principal.

However, with unobservable effort, there are additional Pareto-dominated equilibria (dominated, in particular, by the equilibrium described in Proposition 2). The reason is that the agent's installation decision and the principal's resource allocation are complements, so there is scope for miscoordination when the agent's effort is not observable.[23] Hence, the outcome of Proposition 2 can be understood as the equilibrium of a game in which effort is not observable, and Pareto-dominated equilibria are discarded.

Inconclusive Evidence.— The baseline model only considers the case of conclusive evidence. However,

[22] In the limit $\lambda \to \infty$ (keeping $g \equiv \lambda\Pi$ fixed), there is no disagreement since the principal only switches to S once she is certain that $\theta = 0$ (i.e., $q^* \to 0$) and the agent receives no compensation when that is the case.

[23] For instance, there is an equilibrium in which the agent never installs R and the principal never allocates resources to that arm, irrespective of the share $\alpha \in [0, 1]$ offered by the principal at $t = 0^-$.

this assumption is not essential. For instance, in a previous version of this paper, I showed that results do not change if players also observe the following "news" process that imperfectly reveals R's type when the latter is installed and the principal allocates resources to it:[24]

$$dX_t = \theta a_0 k_t dt + k_t^{1/2} \sigma dZ_t, \text{ where } \{Z_t\}_{t \geq 0} \text{ is a standard Brownian motion} \tag{2}$$

Timing.— As noted in Section 2, the organization's problem in the present model is to decide whether to install R, not when to install it. This is without loss since no useful information arrives when the organization does not install the risky arm.

Consider instead an alternative model in which the principal also has access to information that does not depend on R's installation. This adds an option value problem into the analysis. For instance, suppose that players also observe the diffusion process:[25]

$$dY_t = \theta dt + \sigma dZ_t, \text{ where } \{Z_t\}_{t \geq 0} \text{ is a standard Brownian motion}$$

Both in equilibrium and in the first-best, the organization will then install R whenever q_t is above an "installation" cutoff. This installation cutoff, moreover, is strictly higher in equilibrium than in the first-best. The latter is explained because the principal can alleviate her time inconsistency problem by decreasing the ex-ante uncertainty surrounding R's profitability. Hence, she has higher incentives to wait for a higher belief before installing R in equilibrium than in the first-best.

Thus the only difference between the baseline model and this richer setting is where the inefficiency shows up. In the baseline model, the incentive friction appears *either* in the agent's higher compensation or in the organization's inability to install R. In contrast, in the model involving the timing of adoption, the friction appears in the agent's higher compensation *and* in the greater delay relative to the efficient time of adoption.[26] Nevertheless, the results pertaining to the non-monotonicity of the principal's payoffs with respect to s and the importance of a slow resolution of the uncertainty for the dual moral hazard problem to emerge remain unchanged.

Monetary Transfers vs. Intrinsic Motivation.— Finally, even though this dual moral hazard problem appears to rely heavily on the existence of explicit monetary transfers, these are actually not essential. The same inefficiencies arise—and similar comparative statics are obtained—in a model where the principal pays the agent "in kind."

For example, suppose transfers are unfeasible, but that there is some discretion on how to "design" R (or there are multiple potential versions of R to install). Assume, furthermore, that players have conflicting preferences over R's design (as in, for example, Bonatti and Rantakari, 2016). If the principal must then compromise with the agent on R's design to incentivize him to install the arm (e.g., to increase the agent's "intrinsic" motivation), she will have incentives to return to S sooner

[24] In that version of the paper, I also showed that an alternative way to obtain Corollary 1(ii) is by letting $\sigma \to 0$ (instead of taking $\lambda \to \infty$ and $\Pi \to 0$ with $g \equiv \lambda \Pi$ constant).

[25] Note that, in contrast to the diffusion process given by (2), the informativeness of this diffusion process does not depend on whether R is installed nor the amount of resources allocated to the arm.

[26] The organization will always eventually install R in equilibrium since the incentive friction vanishes when the belief that $\theta = 1$ at the time of installation is sufficiently close to one (see Corollary 1).

than what is constrained efficient. Thus a similar double-sided moral hazard problem ensues and similar comparative statics with respect to s arise.

4 The Optimal Contract

I now characterize and discuss the optimal contract under different contractual assumptions. Unless explicitly stated otherwise, the setting I consider is identical to the baseline one, except that the set of available contracts is richer. I assume that the agent has no wealth and is protected by limited liability.

4.1 The Allocation of Resources is Contractible

4.1.1 Preliminaries

I begin by describing the feasible set of contracts and players' payoffs. A *resource allocation policy* is a stochastic process $k = \{k_t\}_{t\geq 0}$, where $k_t \in [0,1]$ is the fraction of the unit resource allocated to R at time t. A *compensation scheme* $c = \{c_t\}_{t\geq 0}$ is a nonnegative, nondecreasing process, where c_t denotes the cumulative transfers the principal has made to the agent up to, and including, time t.

With some abuse of notation, denote by R_t and S_t the cumulative profits delivered by R and S, respectively, up to time t, including t. Since $\{R_\tau\}_{\tau\leq t}$ is observationally equivalent to $\{N_\tau\}_{\tau\leq t}$ (recall that N_t denoted the number of successes observed up to, and including t), the contractible history up to t is given by:

$$h^t = \{N_\tau\}_{\tau\leq t} \cup \{S_\tau\}_{\tau\leq t} \cup \{k_\tau\}_{\tau\leq t} \cup \{c_\tau\}_{\tau\leq t}$$

That is, it includes (i) the number of successes with R (or, equivalently, the sequence of profits obtained from R), (ii) the sequence of profits obtained from S, (iii) the sequence of resources allocated to each arm, and (iv) the sequence of payments made by the principal to the agent.

A contract, which is offered and committed by the principal at $t = 0^-$, is defined as a resource allocation policy and a compensation scheme, i.e., it is a pair (k, c). The resource allocation policy may depend only on h^{t-}, while the compensation scheme only on $h^{t-} \cup \{(N_t, S_t, k_t)\}$. Players' expected payoffs at $t = 0$ given (k, c) and a_0, are then:

$$\begin{aligned} v_0 =& \mathbb{E}_0\left[\int_0^\infty re^{-rt}\big[a_0 k_t \lambda\theta\Pi + s(1-k_t)\big]dt - \int_0^\infty re^{-rt}dc_t\right] \\ u_0 =& \mathbb{E}_0\left[\int_0^\infty re^{-rt}dc_t\right] - ra_0 f \end{aligned}$$

The principal's problem is to choose (k, c) to maximize v_0. I refer to the solution to this problem as the "optimal contract."

For a given contract (k, c), define then $C^1(k, c)$ as the agent's expected (discounted) compensation if R is installed and $\theta = 1$, and $C^0(k, c)$ as the agent's expected (discounted) compensation when R

is not installed or $\theta = 0$.[27] That is:

$$C^1(k,c) = \mathbb{E}_0\left[\int_0^\infty re^{-rt}dc_t \;\middle|\; k, a_0\theta = 1\right] \qquad C^0(k,c) = \mathbb{E}_0\left[\int_0^\infty re^{-rt}dc_t \;\middle|\; k, a_0\theta = 0\right]$$

Note that u_0 can then be written as $u_0 = \bar{q}a_0C^1(k,c) + (1 - a_0\bar{q})C^0(k,c) - ra_0f$. Hence, for a given contract, the agent's equilibrium installation decision at $t = 0^+$ is:

$$a_0^*(k,c) = \begin{cases} 1 & \text{if } \bar{q}\big(C^1(k,c) - C^0(k,c)\big) \geq C^0(k,c) + rf \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

Equation (3) reveals a simple fact that is easy to miss: if the agent gets compensated when R is not installed (or when $\theta = 0$), i.e., $C^0(k,c) > 0$, this increases the cost of incentivizing the agent to install R and forces the principal to leave the agent with strictly positive rents. Intuitively, when $C^0(k,c) > 0$, the agent anticipates that he will be able to get some compensation irrespective of whether he installs R or not.

4.1.2 Independent Business Units and Commitment vs. Flexibility

When the allocation of resources is contractible, the principal's optimal contract implements the first-best outcome and allows her to achieve the first-best profits. This is formalized in the following proposition:

Proposition 3. *The principal can achieve the first-best payoffs, i.e., $v_0^* = w_0^{fb}$, by offering:*[28]

$$k_t^*(N_t = 0) = \begin{cases} 1 & \text{if } t \leq T^{fb}(\bar{q}) = (\ln\Omega(q^{fb}) - \ln\Omega(\bar{q}))/\lambda \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad k_t^*(N_t) = 1 \text{ for all } N_t \geq 1 \tag{4}$$

and a compensation scheme c^* such that $C^0(k^*,c^*) = 0$ and $C^1(k^*,c^*) = rf/\bar{q}$.

A first interpretation of the contract in Proposition 3 is the principal committing to a lockup period $T^{fb}(\bar{q})$, during which all of the organization's resources are directed towards R regardless of the outcome. Once the lockup period ends, the principal regains the flexibility to allocate resources as she wishes. An alternative interpretation is that of creating an independent business unit: transferring the agent the decision rights over a budget of $T^{fb}(\bar{q})$ units of resource over the entire horizon to develop the innovation.[29,30] In either case, the idea is for the principal to contractually establish ex-

[27] Note that, because the installation decision is not verifiable, the contract cannot distinguish whether R is not delivering profits because the arm is not installed or because it turns out to be unprofitable.

[28] Note that I am not claiming that the optimal contract is unique. As it will become clear in the next section, there are other contracts that also give the principal the first-best payoffs (e.g., a belief-contingent contract or contracts with damages payments linked to S's profits).

[29] The creation of independent business units to help established firms develop innovations has a long history in business (e.g., Gans, 2016, p. 86). The latter is despite the fact that formal economic analysis of the cost and benefits of them is relatively sparse (see Aghion and Tirole, 1994; Bresnahan, Greenstein and Henderson, 2011, for exceptions).

[30] Note that the model, as currently written, requires the organization in the first-best to assign all of its resources to R early on. Hence, taken literally, the independent business unit interpretation of the optimal contract is calling

ante a minimum amount of resources to be allocated to R, to avoid the dual moral hazard problem that leads to the innovator's dilemma.

The agent's compensation is then chosen to reward the agent only if R is installed and turns out to be profitable (to avoid having to leave the agent with strictly positive rents), i.e., $C^0(k^*, c^*) = 0$, while keeping him indifferent in expectation between installing R or not, i.e., $\bar{q}C^1(k^*, c^*) = rf$. The latter can be achieved, for instance, by committing to a time-independent bonus B^* payable upon the first success (recall that $\mu = r/\lambda$):

$$C^1(k^*, c^*) = \lambda B^* \int_0^{T^{fb}(\bar{q})} re^{-(r+\lambda)t} dt \quad \text{and} \quad \bar{q}C^1(k^*, c^*) = rf$$
$$\Longrightarrow B^* = \frac{rf}{\bar{q}} \left(1 + \frac{1}{\mu}\right) \left[1 - \left(\frac{\Omega(\bar{q})}{\Omega(q^{fb})}\right)^{1+\mu}\right]^{-1}$$

Of course, Proposition 3 relies critically on the fact that with exponential bandits and conclusive news, the first-best experimentation policy is a stopping-time policy with a fixed stopping time $T^{fb}(\bar{q})$, that, furthermore, is perfectly known ex-ante. This implies that when the allocation of resources is contractible, the principal can perfectly commit at no cost to experiment with the risky arm efficiently. What would happen if that is not the case? If the efficient stopping time is not known at the time of contracting?

To explore this possibility, in the online Appendix I consider the following extension of the baseline model: If the risky arm is installed, then players observe a signal about its profitability before experimenting with the arm (say at a time $t = 0^{++}$). The latter changes the prior that $\theta = 1$ from $\bar{q}$ to $\phi \in [\underline{\phi}, \bar{\phi}]$, where ϕ distributes according to a cumulative distribution function $H(\phi)$ (so $\bar{q} = \int_{\underline{\phi}}^{\bar{\phi}} \phi dH(\phi)$ by Bayes' Rule). In this setting, the efficient experimentation policy is still a stopping-time policy with a fixed stopping time. However, the efficient stopping time is unknown ex-ante, as it depends on the realization of ϕ, i.e., $T^{fb}(\phi) = (\ln \Omega(q^{fb}) - \ln \Omega(\phi))/\lambda)$.

Following Amador, Werning and Angeletos (2006), the optimal contract can then be characterized by solving a principal-agent problem with one principal, the principal's "ex-ante" self, and two agents, the "true" agent, and the principal's "ex-post" self. More precisely, whenever the principal decides to incentivize the agent to install R, her ex-ante self commits to a direct mechanism $(k(\phi), c(\phi))$ such that the (true) agent installs R, and the principal's ex-post self truthfully reveals ϕ after the signal realization.

I then show, drawing heavily on the tools developed by Guo (2016), that if the resource allocation policy and the agent's compensation *in the event of a breakthrough* can be made contingent on the

---

for the division working on the new project (i.e., the agent working on the risky arm) to have the power of allocating all of the organization's resources for a period of time. At a first glance, this seems a lot more drastic than simply creating an independent business. In reality, however, it seems unlikely that such an extreme ("bang-bang") allocation of resources is efficient. Rather, a more likely scenario is that R's development requires a significant fraction (but not the totality) of the organization's resources per unit of time. The latter could be modeled, for instance, by assuming that only a fraction $\bar{\kappa} < 1$ of resources destined to R are useful (and that any amount above that is wasted). In such a model, the independent business unit interpretation makes more sense, as the principal does not need to transfer the agent all of the available resources per unit of time early on; she only commits to transfer a fraction κ of it per unit of time, with an overall cap or budget equal to $T^{fb}(\bar{q})$ over the entire horizon.

message ϕ (e.g., the principal can offer a message-contingent but time-independent bonus $B(\phi)$ payable upon a breakthrough), then the principal's optimal contract still allows her to obtain her ex-ante first-best payoffs.[31] However, when that is not the case, the principal starts facing a tradeoff between *commitment* and *flexibility*. On the one hand, she can mitigate her time inconsistency problem and reduce the agent's compensation by ex-ante restricting her ex-post options, i.e., by committing to experiment with R for a minimum pre-specified amount of time. On the other hand, the latter limits her ability to react to the new information that arrives once R is installed.

For a concrete example of this tradeoff, consider the optimal contract if the agent's compensation is restricted to entail a time and message-independent bonus B payable upon a breakthrough.[32] As it should be clear by now, the principal's ex-ante self would like the flexibility to observe ϕ to set the efficient stopping time $T^{fb}(\phi)$. However, once R is installed and due to bonus promised to the agent, the principal's ex-post self experiments with R only until $T^*(\phi;B) < T^{fb}(\phi)$, where:

$$T^*(\phi;B) = \frac{1}{\lambda}\ln\left(\frac{\Omega(q^*)}{\Omega(\phi)}\right) \quad \text{with} \quad q^* = \frac{\mu s}{\mu s + (1+\mu)(g-s) - \lambda\mu B}$$

In the online Appendix, I then show that under certain regularity conditions, the optimal contract in this case (which I refer to as the "(restricted) optimal contract") still involves stopping-time policies. However, now it additionally incorporates pooling at the bottom, i.e., for types $\phi \leq \hat{\phi}(B^*)$, where $\hat{\phi}(B^*)$ is always strictly above $\underline{\phi}$, and B^* denotes the optimal time-independent bonus promised by the principal.

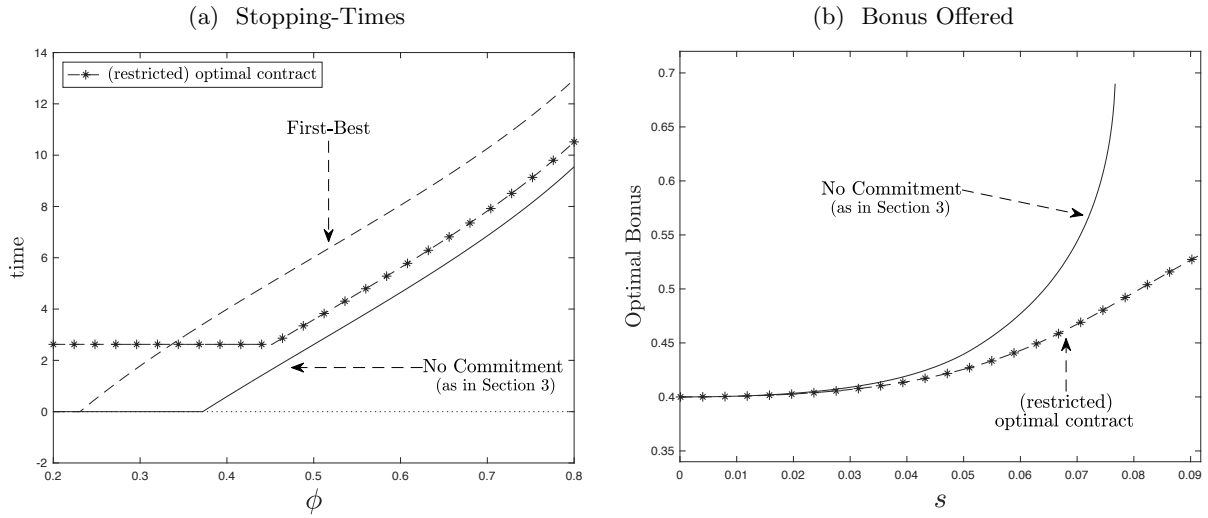
In particular, types below $\hat{\phi}(B^*)$ are pooled at $\hat{\phi}(B^*)$'s preferred stopping-time $T^*(\hat{\phi}(B^*);B^*)$, while types above $\hat{\phi}(B^*)$ are assigned to their preferred stopping-time $T^*(\phi;B^*)$. The latter can (again) be implemented with a lockup period of experimentation that lasts for $T^*(\hat{\phi}(B^*);B^*)$ or by setting an independent business unit with an overall budget of $T^*(\hat{\phi}(B^*);B^*)$ resources. The optimal bonus B^* is then chosen so to leave the agent indifferent between installing R or not:

$$\lambda B^*\left[\int_{\underline{\phi}}^{\hat{\phi}(B^*)}\phi\left(\int_0^{T^*(\hat{\phi}(B^*);B^*)} re^{-(r+\lambda)t}dt\right)dH(\phi) + \int_{\hat{\phi}(B^*)}^{\bar{\phi}}\phi\left(\int_0^{T^*(\phi;B^*)} re^{-(r+\lambda)t}dt\right)dH(\phi)\right] = rf$$

[31] The idea of the proof is as follows. Clearly, $v_0^* \leq w_0^{fb}$. Hence, if a contract allows the principal to achieve first-best payoffs, it must necessarily be optimal. Now, $v_0^* = w_0^{fb}$ only if the contract (i) implements the (ex-post) efficient resource allocation policy (which is a function of ϕ), and (ii) it induces the agent to install R without leaving him any rents. I then show that a direct mechanism $(k(\phi), c(\phi))$ satisfying (i) and (ii) is truthful only if $c(\phi)$ is such that $C^1(k(\phi), c(\phi)) = rf/\bar{q}$ for any $\phi \in [\underline{\phi}, \bar{\phi}]$. That is, only if the agent's expected (discounted) compensation is independent of ϕ. Since the efficient stopping time $T^{fb}(\phi)$ for a given ϕ is increasing in the latter, this requires the agent's reward in the event of a breakthrough to be decreasing in ϕ. Coincidentally, this is also the condition that guarantees that the agent will reveal ϕ truthfully if he were the only one observing the signal. Thus, whether one or both players observe the signal (and who observes it) is irrelevant. In either case, the optimal mechanism is the same, and players self-select accordingly ex-post.

[32] I restrict attention to this contract to make the comparison to the baseline model easier and because under some circumstances a time-contingent (but message-independent bonus) is sufficient to achieve the first-best (see the online Appendix for details).

Figure 3: Commitment versus Flexibility



Notes. Panel (a) has $s = 0.075$. Remaining parameters: $\Pi = 1$, $r = 1/5$, $\lambda = 1/5$, $f = 1/10$, $s = 0.075$, $\bar{q} = 1/2$, $\phi \sim \text{U}[0.2, 0.8]$. The value of s in Panel (a) was chosen so that the agent installs R even if the principal only commits to a time-independent bonus B without contractually specifying the allocation of resources (i.e., the problem without commitment does not unravel).

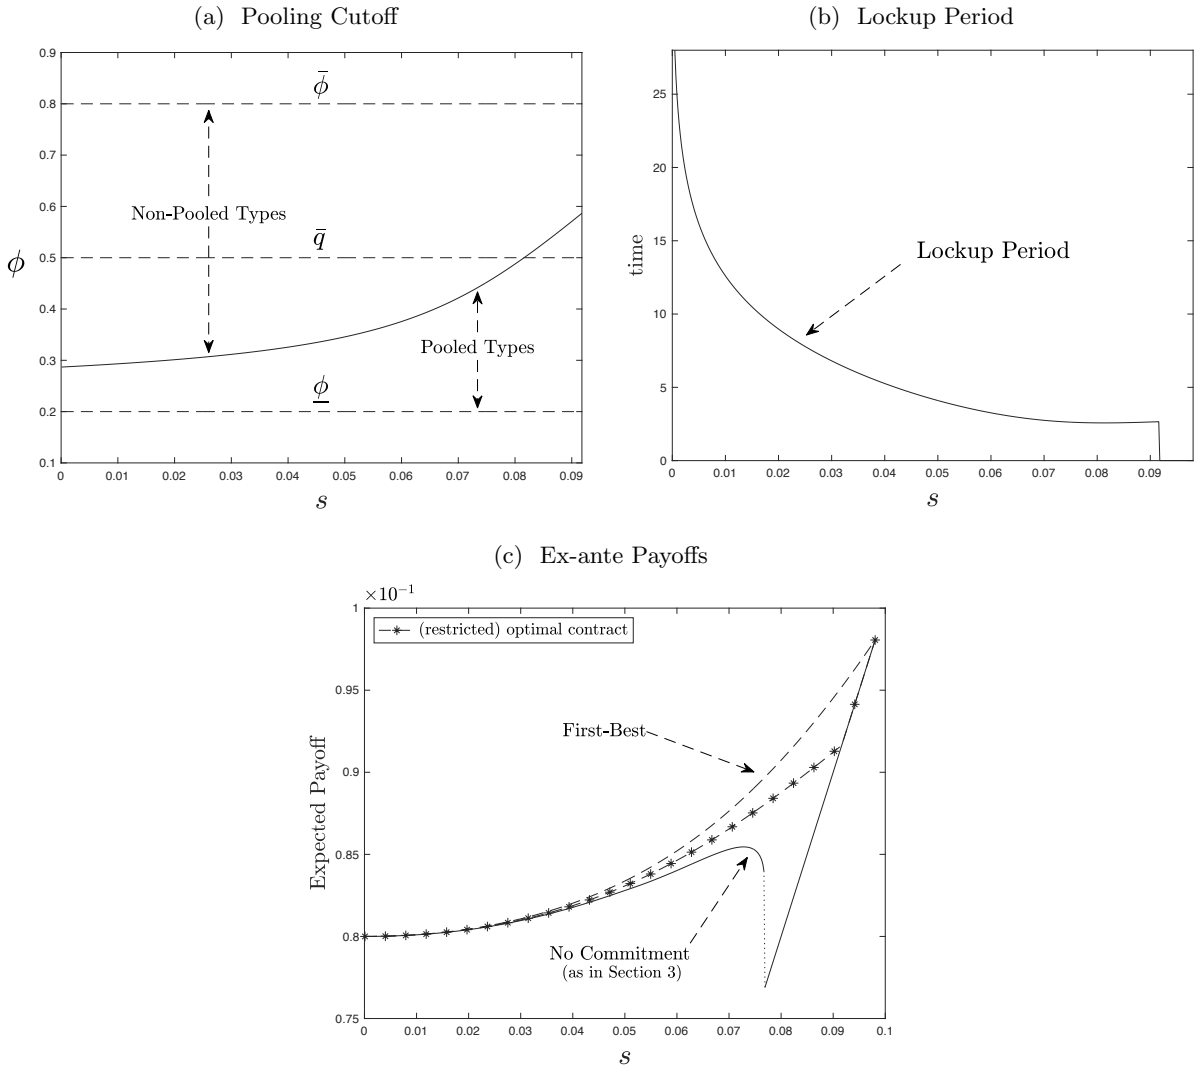
The resulting tradeoff between commitment and flexibility is illustrated in Figure 3. Panel (a) depicts the stopping times assigned to each type ϕ in the (restricted) optimal contract. For comparison purposes, that panel also illustrates the stopping times that each ϕ would follow in the first-best and in the case where the principal only commits to a time-independent bonus without contractually specifying the allocation of resources (the "no-commitment case," as in the baseline model of Section 3). Panel (b), in turn, plots the optimal bonus B^* offered in the (restricted) optimal contract and the equilibrium bonus that would ensue if the principal made no commitments on the allocation of resources, both as a function of s.

As the figure shows, setting a lockup period allows the principal to lower the agent's compensation compared to the no-commitment case. This explains why for all $\phi \geq \hat{\phi}(B^*)$, the stopping times in the (restricted) optimal contract are strictly higher than in the no-commitment case, even though the principal's ex-post self has the flexibility to choose her preferred stopping time in both cases. The flip side is that the lockup period sometimes forces the organization to experiment with R longer than it is efficient.

Figure 4, on the other hand, portrays the pooling cutoff and the duration of the lockup period $T^*(\hat{\phi}(B^*); B^*)$ as a function of the profits of the safe arm. It also illustrates the principal's ex-ante payoffs with the (restricted) optimal contract, comparing them with the principal's first-best payoffs and the payoffs she obtains in the no-commitment case.[33] As the figure shows, the set of types pooled increases with s (though the duration of the lockup period decreases with s). This is because the

[33]Players' ex-ante payoffs are continuous in s, though the shape of the contract exhibits a discontinuity at $s = 0$: $\hat{\phi}(B^*) \not\to \underline{\phi}$ as $s \to 0$, even though we would expect no pooling if $s = 0$. This discontinuity arises because when $s = 0$, the principal's ex-ante self is indifferent between no pooling, partial pooling, or full pooling, as the efficient stopping time is $T^{fb}(\phi, B) = \infty$ irrespective of the value of ϕ.

Figure 4: Optimal (Restricted) Contract as a Function of s



Notes. Remaining parameters: $\Pi = 1$, $r = 1/5$, $\lambda = 1/5$, $f = 1/10$, $\bar{q} = 1/2$, $\phi \sim \mathrm{U}[0.2, 0.8]$.

commitment problem on the principal's side worsens when the profits of the safe arm increase, so the optimal contract becomes less sensitive to the arrival of new information.[34]

The figure also shows that, although the (restricted) optimal contract does not implement the first-best outcome, the principal's payoffs are strictly increasing in the profits of the safe arm (as in the first-best). That is, there is no longer a dilemma between the organization's current success and its long-term prospects. Thus, the model supports the idea that creating an independent business unit might help established firms overcome the dilemma that arises when trying to develop certain types of innovations (those that require new organizational capabilities, and whose profitability is revealed only after ample time and testing).

[34] Additionally, it is not difficult to prove that the pooling cutoff is decreasing in λ while keeping $g \equiv \lambda\Pi$ fixed. This is consistent with Corollary 1: the easier it is to learn R's profitability, the less severe the dual moral hazard problem is and the optimal contract becomes more sensitive ϕ.

However, as it will become apparent in the next section, the key for this solution to work is for the allocation of resources to be contractible. The latter may explain the varying degrees of success independent business units have had in practice.[35]

4.2 The Allocation of Resources is Not Contractible

I now consider the case where the allocation of resources is not contractible. The setting is identical to the one described in Section 4.1.1, except that (i) a contract is now defined only as a compensation scheme $c = \{c_t\}_{t\geq 0}$, and (ii) the contractible history no longer includes the sequence of resources allocated to each arm, i.e.,:

$$h^t = \{N_\tau\}_{\tau\leq t} \cup \{S_\tau\}_{\tau\leq t} \cup \{c_\tau\}_{\tau\leq t}$$

The principal's problem is then to choose c to maximize v_0 subject to the $k = k^*(c)$, where $k^*(c)$ is the equilibrium resource allocation policy induced by c.

As in the case where the allocation of resources is contractible, the optimal contract can be characterized by solving a principal-agent problem between the principal's ex-ante and ex-post selves plus the true agent. The difference is that, if the principal's ex-ante self wants to induce the agent to install R, she now offers a compensation scheme $c = \{c_t\}_{t\geq 0}$ plus a "recommendation" policy $k = \{k_t\}_{t\geq 0}$ so that the true agent installs R, and the principal's ex-post self is willing to follow the recommendation.[36]

With this notation at hand, I begin by showing that, taking the model literally, the principal can again achieve the first-best payoff. I then discuss why this result is unlikely to hold in practice:

Proposition 4. *The principal can achieve the first-best payoffs, i.e., $v_0^* = w_0^{fb}$, by offering:*

$$dc_t^* = \begin{cases} \mathbb{1}_{N_t=1} B^* dN_t + \mathbb{1}_{dS_t>0} D^* & \text{if } t < T^{fb}(\bar{q}) \\ \mathbb{1}_{N_t=1} B^* dN_t & \text{otherwise} \end{cases} \tag{5}$$

where $D^* > s$, and:

$$B^* = \frac{rf}{\lambda \bar{q}} \left(1 + \frac{1}{\mu}\right) \left[1 - \left(\frac{\Omega(\bar{q})}{\Omega(q^{fb})}\right)^{1+\mu}\right]^{-1} \tag{6}$$

This compensation scheme then induces the principal to allocate resources as in the first-best policy given by (4), and it is such that $C^0(k^*(c^*), c^*) = 0$ and $C^1(k^*(c^*), c^*) = rf/\bar{q}$.

The idea behind Proposition 4 is to indirectly contract upon the allocation of resources by punishing the principal if S delivers profits before the efficient stopping time $T^{fb}(\bar{q})$. The latter is done by promising the agent a damages payment of $D^* > s$ every time such an event occurs. Note that,

[35] See Gans (2016, ch. 6) for an account of how creating a separate entity helped Quantum Corporation—a hard disk drive manufacturer in the 1980s and 1990s—lead the market of 3.5-inch drives—an innovation at the time. The same chapter also documents how the lack of commitment of IBM's upper management led to the failure (and subsequent absorption) of its Boca Raton Unit, an independent business unit created in response to the emergence of the personal computer.

[36] This recommendation-mechanism approach is similar in spirit to the one developed by Ely, Georgiadis and Rayo (2022) to characterize the optimal monetary and informational incentives in dynamic moral hazard settings.

crucially, D^* is never paid on- or off-path (i.e., if the agent fails to install R),[37] explaining why the agent obtains no rents despite the fact that his compensation involves payments that are not contingent on R's installation.

There are several reasons why Proposition 4 may fail to hold in practice. First, if S's profits were stochastic (as in, e.g., Bolton and Harris, 1999), then a damages clause as the one described above will be paid with positive probability off-path. As a result, incorporating it into the contract forces the principal to leave the agent with strictly positive rents, making their overall effect on the principal's ability to pre-commit ambiguous.[38] Hence, the principal herself might not have incentives to incorporate this clause ex-ante.[39]

Second, the scheme in Proposition 4 relies critically on the parties' ability to perfectly contract on S's profits. This assumption seems overly optimistic. For instance, it might be difficult to contractually specify what constitutes the "current business" exactly. Or, even if that is not the case, there could be ways to allocate resources to S to manipulate or backload profits (think of R&D investments to achieve incremental innovations in S). For all these reasons, it is then important to understand what the organization can achieve if contracting on S's profits is unfeasible, which is why I now analyze this case in more detail.

Restrict then the contractible history to just $h^t = \{N_\tau\}_{\tau \leq t} \cup \{c_\tau\}_{\tau \leq t}$. This implies that $c = \{c_t\}_{t \geq 0}$ can be made contingent on time, the number of successes with R, and the payments the agent has received in the past. However, since players are risk-neutral and the first success provides definitive evidence that $\theta = 1$, it is without loss to restrict attention to schemes involving a single time-contingent bonus $B_t \geq 0$ paid upon R's first success, i.e., $dc_t = \mathbb{1}_{N_t=1} B_t dN_t$ (see footnote 19).

To obtain some intuition about the issues that arise with time-contingent contracts, let me begin by considering the auxiliary case where the principal can make the bonus directly contingent on beliefs rather than on time, i.e., $B_t = B(q_t)$. The latter, of course, is not allowed in the current context; it would require keeping track of both time and the resources allocated to R, and the latter are not contractible by assumption.

Lemma 3. *By offering a single belief-contingent bonus payable upon a breakthrough $B(q_t)$, the principal can achieve the first-best payoffs, i.e., $v_0^* = w_0^{fb}$.*[40]

According to Lemma 3, even when the allocation of resources and S's profits are not contractible, it is still possible to make the efficient policy (4) incentive-compatible with the use of belief-contingent contracts. Intuitively, the idea is to exploit the dependence of $B(q_t)$ on beliefs to ensure that the principal has a sufficient stake in the R's profits to experiment efficiently ex-post. The latter requires that the agent's continuation value decreases sufficiently fast as players become more pessimistic about

[37] Given that $D^* > s$, the principal has incentives to allocate all the resources to R during $t \in (0, T^{fb}(\bar{q}))$ even if the agent fails to install R. This feature of the contract is important, since, otherwise, the agent would deviate and not install R to obtain the damages payments from the contract.

[38] On the one hand, the damages clause lowers the attractiveness of allocating resources to S, improving the dual moral hazard problem. On the other hand, by increasing the agent's compensation in the event of a breakthrough, they also further bias the principal away from R. The latter worsens the principal's commitment problem.

[39] In the numerical simulations I have conducted, this indeed appears to be the case (i.e., the principal shuns away from offering damages or equity on S's profits).

[40] Due to space constraints, the remaining lemmas and propositions are all proved in the online Appendix.

R's prospects, implying that the optimal bonus $B^*(q_t)$ must be—informally speaking—"increasing" in q_t.[41] An example of a bonus that works is the following:

$$B^*(q_t) = \frac{rf}{\lambda(w_0^{fb} - s)} \left[\left(1 + \frac{1}{\mu}\right)(g - s) - s\left(\frac{1 - q_t}{q_t}\right) \right]$$

Suppose now that the principal's ex-ante self tries to implement the efficient policy using a time-contingent rather than a belief-contingent bonus. Because in the absence of a breakthrough beliefs on-path evolve deterministically in time under the efficient resource policy, there is a one-to-one mapping between t and on-path beliefs q_t^* at this phase of the game, i.e.,:

$$q_t^* = \begin{cases} \frac{\bar{q}}{\bar{q} + (1-\bar{q})e^{\lambda t}} & \text{if } t < T^{fb} \\ q^{fb} & \text{otherwise} \end{cases}$$

Hence, by making $B_t^* = B^*(q_t^*)$, players' continuation values on-path are the same under both contracts. For instance, in the absence of a breakthrough, the principal's continuation values at a time $t \in (0, T^{fb})$ under $B^*(q_t^*)$ and B_t^* are, respectively:

$$v_t^*(\text{belief}) = \int_t^{T^{fb}} re^{-r(z-t)} q_z^* \left[g\left(1 + \frac{\lambda}{r}\right) - \lambda B^*(q_z^*) \right] \left(\frac{1 - q_t^*}{1 - q_z^*}\right) dz + se^{-r(T^{fb}-t)} \left(\frac{1 - q_t^*}{1 - q_{T^{fb}}^*}\right)$$

$$v_t^*(\text{time}) = \int_t^{T^{fb}} re^{-r(z-t)} q_z^* \left[g\left(1 + \frac{\lambda}{r}\right) - \lambda B_z^* \right] \left(\frac{1 - q_t^*}{1 - q_z^*}\right) dz + se^{-r(T^{fb}-t)} \left(\frac{1 - q_t^*}{1 - q_{T^{fb}}^*}\right)$$

As a result, one might be tempted to conclude that the principal can achieve the same with a time-contingent bonus B_t than with a belief-contingent bonus $B(q_t)$. This logic, however, turns out to be incorrect.

To see why, suppose that there has been no breakthrough before $t \in (0, T^{fb})$, and that the principal ex-post self deviates to $k_t' = 0$ for $[t, t + dt)$ while returning to the efficient policy recommendation thereafter. Denoting by $v_t'(\text{belief})$ and $v_t'(\text{time})$ the principal's deviation payoffs under the two different compensation schemes, I have that:

$$v_t'(\text{time}) - v_t'(\text{belief}) = \int_{t+dt}^{T^{fb}} re^{-r(z-t)} q_{z-dt}^* \lambda\big(B^*(q_{z-dt}^*) - B_z^*\big) \left(\frac{1 - q_t^*}{1 - q_{z-dt}^*}\right) dz$$
$$+ se^{-r(T^{fb}-t)} \left(\frac{1 - q_t^*}{1 - q_{T^{fb}-dt}^*}\right) \left[1 - \left(\frac{1 - q_{T^{fb}-dt}^*}{1 - q_{T^{fb}}^*}\right)\right]$$

Hence, even though $B_t^* = B^*(q_t^*)$ gives players' the same on-path payoffs, the principal's incentives to deviate are different under the two contracts. In particular, with a time-contingent contract, the principal can "freeze beliefs" while keeping the clock on B_t running, something that is not possible when the bonus depends directly on beliefs.

How does the optimal time-contingent contract look like then? And what is the best the principal

[41] As I show in the online Appendix, $B^*(q)$ must be increasing in q on "average;" it does not need to be monotone necessarily.

can achieve? To describe the equilibrium outcome in this case, let:

$$\Gamma(x,s) \equiv \bar{q}(g-s)\left[1-\left(\frac{\Omega(\bar{q})}{\Omega(x)}\right)^{1+\mu}\right]+\frac{\mu s(1-\bar{q})}{1-\mu}\left[1-\left(\frac{\Omega(\bar{q})}{\Omega(x)}\right)^{\mu-1}\right]$$

and define $\hat{s}$ and $\bar{s}$ as the (unique) solutions to $\Gamma(q^{fb}(\hat{s}),\hat{s})=rf$ and $\max_x \Gamma(x,\bar{s})=rf$ (where $q^{fb}(s)$ is defined as in Lemma 1). It is not difficult to prove that $\hat{s}<\bar{s}<s^{fb}$. The following proposition then characterizes the optimal time-contingent contract:

Proposition 5. *Under the optimal time-contingent contract, the principal can incentivize the agent to install R only when $s<\bar{s}$. If that is the case, the optimal contract implements a stopping-time policy, i.e.:*

$$k_t^*(N_t=0)=\begin{cases}1 & \text{if } t\leq T^{**}\\ 0 & \text{otherwise}\end{cases} \quad \text{and} \quad k_t^*(N_t)=1 \text{ for all } N_t\geq 1$$

with a fixed stopping time equal to:

$$T^{**}=\begin{cases}T^{fb} & \text{if } s\in(0,\hat{s}]\\ \dfrac{1}{\lambda}\ln\left(\dfrac{\Omega(q^{**})}{\Omega(\bar{q})}\right) & \text{if } s\in(\hat{s},\bar{s}]\end{cases}$$

where $q^{**}=\inf\{x\in[q^{fb},\bar{q}]:\Gamma(x,s)=rf\}>q^{fb}$. Denoting by v_0^{**} and u_0^{**}, respectively, the principal's and the agent's equilibrium expected payoff at $t=0$ under the optimal time-contingent contract, then:

$$v_0^{**}=\begin{cases}w_0^{fb} & \text{if } s\in(0,\hat{s}]\\ g\bar{q}+(s-gq^{**})\left(\dfrac{1-\bar{q}}{1-q^{**}}\right)\left(\dfrac{\Omega(\bar{q})}{\Omega(q^{**})}\right)^{\mu}-rf & \text{if } s\in(\hat{s},\bar{s}]\\ s & \text{if } s\in(\bar{s},s^{fb})\end{cases}$$

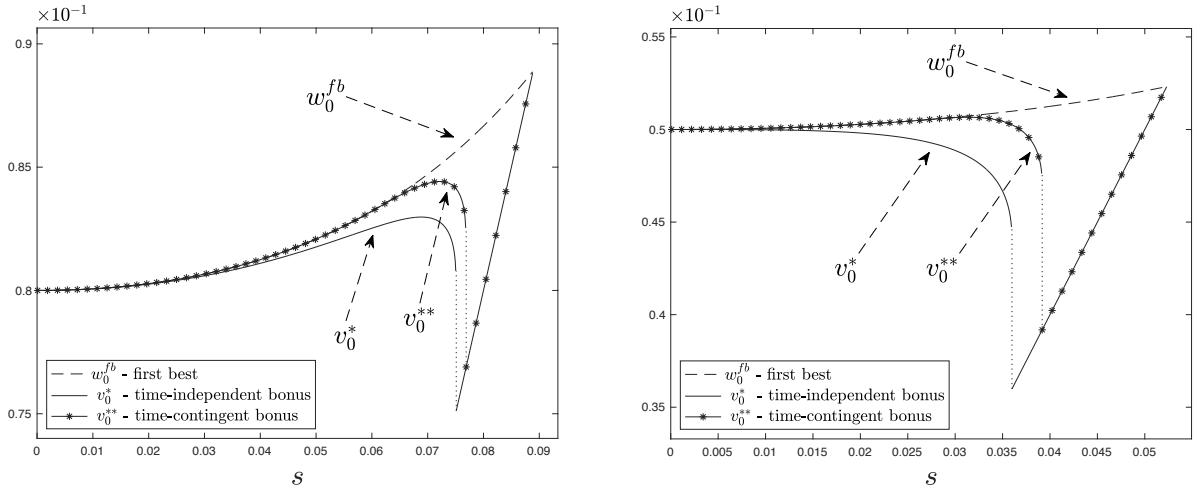
while $u_0^{**}=0$ for all $s\in(0,s^{fb})$.

Figure 5, which builds upon Figure 2, depicts the principal's ex-ante payoffs under the optimal time-contingent contract, v_0^{**}, as well as the first-best payoff, w_0^{fb}, and the principal's payoffs under a time-independent bonus of Section 3, v_0^*. As the figure shows, while time-contingent contracts allow the principal to improve upon time-independent contracts (allowing her, for instance, to achieve the first-best payoffs when $s\leq\hat{s}$), they still lead to a dilemma between currents profits and long-term prospects: there is a range of values of s over which v_0^{**} is strictly decreasing in that parameter. Moreover, similarly to the case where the principal offers a time-independent bonus, for s sufficiently high the incentive problem also unravels in this case.

Intuitively, the principal's ability to freeze beliefs while keeping B_t running, bounds how fast the agent's continuation value can decrease if a breakthrough does not materialize. As a result, when the safe arm is too profitable (so the principal's time-inconsistency problem is significant), the principal is unable to pre-commit not to distort the allocation of resources ex-post. The dual moral hazard

Figure 5: Time-Contingent Contracts Do Not always Solve the Innovator's Dilemma

(a) Welfare ($f = 1/4$) (b) Welfare ($f = 1/10$)



Notes. Remaining parameters: $\Pi = 1$, $r = 1/5$, $\lambda = 1/5$, $\bar{q} = 1/2$.

problem described in Section 3 then claws its way back into the scene and the innovator's dilemma, unraveling included, emerges again.

With regards to the optimal bonus, when $s \in (0, \hat{s}]$ there is a myriad of different time-contingent bonuses that allow the principal to obtain the first-best payoff. In contrast, when $s \in (\hat{s}, \bar{s}]$, the optimal bonus B_t^* is unique (among the class of single time-contingent bonuses paid upon a breakthrough) and is given by:[42]

$$B_t^* = \begin{cases} \dfrac{1}{\lambda}\left(1+\dfrac{1}{\mu}\right)(g-s) - \dfrac{se^{\lambda t}}{\lambda}\left(\dfrac{1-\bar{q}}{\bar{q}}\right)\left(\dfrac{\mu - e^{-\lambda(\mu-1)(T^{**}-t)}}{\mu-1}\right) & \text{if } t < T^{**} \\ \bar{B} & \text{otherwise} \end{cases} \tag{7}$$

where $\bar{B}$ must be set sufficiently high (e.g., $\bar{B} \to \infty$) so that, if the principal tries to freeze beliefs while keeping the clock on B_t^* running, she is unable to extend the experimentation with R beyond T^{**}.

The key feature of B_t^* as given by (7) is that the principal's ex-post self is always indifferent (on-path) between playing $k_t^* = 1$ during $t \in (0, T^{**})$ and deviating to $k_t' = 0$ for $[t, t+dt)$ for a small dt (while returning to the on-path recommendation thereafter). That is, B_t^* is pinned down by the local incentive-compatibility constraint on the principal's ex-post self side, plus the fact that the agent is always made ex-ante indifferent between installing R or not.

[42]This expression is valid as long as $\mu \neq 1$. To obtain the optimal bonus when $\mu = 1$, one needs to take the limit of this expression as $\mu \to 1$.

5 Conclusions

In this paper, I proposed an explanation for why an organization's success in exploiting its current business can undermine its ability to explore new opportunities. My theory provides testable implications about which innovations will be more difficult for successful firms to adopt and supports the idea that setting up an independent business unit can help successful firms be more effective at innovating.

Appendix A The First-Best Outcome

A.1 Proof of Lemma 1

If $a_0 = 0$, then it is clearly optimal to always allocate resources to S, as R will never deliver any profits. If $a_0 = 1$, on the other hand, then the organization faces a single-player experimentation problem with exponential bandits and conclusive news. The solution to this problem is characterized in Keller, Rady and Cripps (2005, Section 3). □

A.2 Proof of Proposition 1

Installing R gives $w(\bar{q}, a_0 = 1; s) - rf$, while not installing gives s. Hence, it is optimal to install R whenever $w(\bar{q}, a_0 = 1; s) - rf - s \geq 0$. The term $w(\bar{q}, a_0 = 1; s) - rf - s$ is strictly decreasing in s. Furthermore, as $s \to 0$, it converges to $\bar{q}g - rf > 0$, and as $s \to g$, to $-rf < 0$ (recall that $g \equiv \lambda\Pi$). Hence, it is optimal to install R whenever $s \leq s^{fb}$, where $s^{fb} \in (0, g)$ is the unique solution to $w(\bar{q}, a_0 = 1; s^{fb}) - rf = s^{fb}$. □

Appendix B The Equilibrium

B.1 Proof of Lemma 2

The Principal's Value Function and her Resource Allocation Problem.— If the agent does not install R, then it is clear that $k_t^*(q_t, a_0 = 0, B; s) = 0$ for all $q_t \in [0, 1]$, given that R never delivers profits. Hence, the principal's value in this case is $v(q_t, a_0 = 0, B; s) = s$.

On the other hand, if the agent installs R, by the principle of optimality, the principal's value function satisfies:

$$v(q) = \max_{k\in[0,1]} \Big\{ r\left[(1-k)s + \lambda qk(\Pi - B)\right] dt + e^{-rdt}\mathbb{E}\left[v(q+dq) \,|\, q, k\right] \Big\} \tag{8}$$

where, to avoid cluttering notation, I am expressing v only as a function of the state variable q. The fact that (i) the instantaneous probability of a breakthrough is λqk, (ii) that if a breakthrough occurs the value function jumps to $v(1) = \lambda\Pi \equiv g$, and that (iii) $dq = \lambda kq(1-q)dt$, implies that:

$$\mathbb{E}\left[v(q+dq) \,|\, q, k\right] = g\lambda qkdt + (1 - \lambda kqdt)v(q - \lambda kq(1-q)dt)$$

Hence, evaluating the latter on (8) and taking its continuous-time limit, $v(q)$ must satisfy:

$$rv(q) = \max_{k\in[0,1]} \left\{ r\left[(1-k)s + qk(g - \lambda B)\right] + \lambda kq(g - v(q)) - \lambda kq(1-q)v'(q) \right\} \tag{9}$$

I then conjecture that the principal follows a cutoff strategy (as in the first-best):

$$k^*(q, a_0 = 1, B; s) = \begin{cases} 1 & \text{if } q > q^* \\ 0 & \text{otherwise} \end{cases}$$

If so, the differential equation (9), plus the value-matching and the smooth-pasting conditions $v(q^*) = s$ and $v'(q^*) = 0$, imply that:

$$q^* = \frac{\mu s}{\mu s + (1+\mu)(g-s) - \lambda \mu B}$$

and that:

$$v(q, a_0 = 1, B; s) = gq + (s - gq^*)\left(\frac{1-q}{1-q^*}\right)\left(\frac{\Omega(q)}{\Omega(q^*)}\right)^{\mu} - \lambda B q \left(\frac{\mu}{1+\mu}\right)\left[1 - \left(\frac{\Omega(q)}{\Omega(q^*)}\right)^{1+\mu}\right]$$

if $q > q^*$ and $v(q, a_0 = 1, B; s) = s$, otherwise (recall that $\mu \equiv r/\lambda$). Optimality of this cutoff strategy then follows from standard verification arguments. □

The Agent's Value Function.— If the agent does not install R, then the principal plays $k_t^*(q_t, a_0 = 1, B; s) = 0$ for all $q_t \in [0, 1]$. Thus, the agent's value function in this case is $u(q_t, a_0 = 0, B; s) = 0$.

Suppose instead that the agent installs R. If $q_t \leq q^*$, then given the principal's equilibrium resource allocation policy, we have that $u(q_t, a_0 = 1, B; s) = 0$. In contrast, if $q_t > q^*$, then following similar reasoning as in the principal's case, the agent's value function must satisfy:

$$ru(q) = \lambda q(rB - u(q)) - \lambda q(1-q)u'(q)$$

where I have omitted the dependence of u on a_0, B, and s, to avoid cluttering notation. This differential equation has solution $u(q) = \lambda B q \mu/(1+\mu) + (1-q)\Omega(q)^{\mu} C$, where C is a constant. Using then the fact that $u(q^*) = 0$ (value matching) we can solve for C. Substituting back into $u(q)$ and rearranging, we obtain the agent's value function as stated in the lemma. □

B.2 Proof of Proposition 2

Preliminaries.— I begin with two preliminary claims that will be used in the proof of this proposition.

Claim 1. *Let $\hat{B}(s) \in \arg\max_{B \geq 0} u(\bar{q}, a_0 = 1, B; s)$. For any $s \in (0, s^{fb})$, $\hat{B}(s)$ exists and is unique, and $u(\bar{q}, a_0 = 1, \hat{B}(s); s)$ is strictly decreasing in s.*

Proof. Note that when $s = s^{fb}$, then $q^{fb} < \bar{q}$ (otherwise, $w(\bar{q}, a_0 = 1; s^{fb}) = s^{fb}$ and it would not be optimal to install R). By continuity, this implies that when $s < s^{fb}$, there exists an $B > 0$ such that $q^* < \bar{q}$. Moreover, if B is such that $q^* \geq \bar{q}$, then $u(\bar{q}, a_0 = 1, B; s) = 0$. In contrast, if B is such that $q^* < \bar{q}$, then $u(\bar{q}, a_0 = 1, B; s) > 0$. Hence, $\hat{B}(s)$ must be such that $q^* < \bar{q}$. If so, then:

$$\frac{\partial^2 u}{\partial B^2} = -\frac{\lambda^2 \mu^2 q \left(2(1+\mu)(g-s) + \lambda \mu^2 B\right)}{\left((1+\mu)(g-s) - \lambda \mu B\right)^2}\left(\frac{\Omega(\bar{q})}{\Omega(q^*)}\right)^{1+\mu} < 0$$

Thus, in the relevant region, $u(\bar{q}, a_0 = 1, B; s)$ is strictly concave, and therefore, $\hat{B}(s)$ is unique. Finally, since $u(\bar{q}, a_0 = 1, B; s)$ is strictly decreasing in s when B is such that $q^* < \bar{q}$, we have that in particular, $u(\bar{q}, a_0 = 1, \hat{B}(s); s)$ is strictly decreasing in s. □

Claim 2. *Define s^* as a solution of $u(\bar{q}, a_0 = 1, \hat{B}(s^*); s^*) = rf$. s^* exists, is unique and is strictly between 0 and s^{fb}.*

Proof. By Claim 1, $u(\bar{q}, a_0 = 1, \hat{B}(s); s)$ is strictly decreasing in s. Hence, $u(\bar{q}, a_0 = 1, \hat{B}(s); s)$ cuts rf at most once. Furthermore, when $s \to 0$, then $u(\bar{q}, a_0 = 1, \hat{B}(s); s) \to \bar{q}g > rf$. Thus, to show the claim it suffices to show that $u(\bar{q}, a_0 = 1, B(s^{fb}); s^{fb}) < rf$.

To do this, note that when $w(\bar{q}, a_0 = 1; s) > s$, then $w(\bar{q}, a_0 = 1; s) = G(q^{fb}(s), s)$, where $q^{fb}(s) \equiv \mu s((1+\mu)(g-s) + \mu s)^{-1}$ and $G(x, s)$ is defined as:

$$G(x, s) \equiv \bar{q}g + (s - xg)\left(\frac{1-\bar{q}}{1-x}\right)\left(\frac{\Omega(\bar{q})}{\Omega(x)}\right)^{\mu}$$

Note, further, that $q^{fb}(s) = \arg\max_{x \in [0,\bar{q}]} G(x, s)$, and that whenever $B \leq \hat{B}(s)$, then $v(\bar{q}, a_0 = 1, B; s) + u(\bar{q}, a_0 = 1, B; s) = G(q^*(B, s), s)$, where:

$$q^*(B, s) \equiv \frac{\mu s}{\mu s + (1+\mu)(g-s) - \lambda \mu B}$$

.

Suppose then by contradiction that $u(\bar{q}, a_0 = 1, \hat{B}(s^{fb}); s^{fb}) \geq rf$, and, to simplify notation, let $\hat{B}^{fb} \equiv \hat{B}(s^{fb})$. Since $f > 0$, this implies that $q^*(\hat{B}^{fb}, s^{fb}) < \bar{q}$, i.e., the equilibrium cutoff belief q^* when the bonus is $\hat{B}^{fb}$ and S's profits s^{fb} must be strictly less than the prior $\bar{q}$ (otherwise, there would be no experimentation with R and $u(\bar{q}, a_0 = 1, \hat{B}^{fb}; s^{fb}) = 0$).

Now, using the definition of s^{fb}, the premise that $u(\bar{q}, a_0 = 1, \hat{B}^{fb}; s^{fb}) \geq rf$ implies that $u(\bar{q}, a_0 = 1, \hat{B}^{fb}; s^{fb}) \geq G(q^{fb}(s^{fb}), s^{fb}) - s^{fb}$. However, if so, then:

$$\underbrace{[u(\bar{q}, a_0 = 1, \hat{B}^{fb}; s^{fb}) + v(\bar{q}, a_0 = 1, \hat{B}^{fb}; s^{fb})]}_{G(q^*(\hat{B}^{fb}, s^{fb}), s^{fb})} \geq G(q^{fb}(s^{fb}), s^{fb}) + \underbrace{[v(\bar{q}, a_0 = 1, \hat{B}^{fb}; s^{fb}) - s^{fb}]}_{>0 \text{ given that } q^*(\hat{B}^{fb}, s^{fb}) < \bar{q}}$$

which implies that $G(q^*(\hat{B}^{fb}, s^{fb}), s^{fb}) > G(q^{fb}(s^{fb}), s^{fb})$. This contradicts contradicts the fact that $G(q^{fb}(s^{fb}), s^{fb}) = \max_{x \in [0,\bar{q}]} G(x, s^{fb})$, so $u(\bar{q}, a_0 = 1, \hat{B}^{fb}; s^{fb}) < rf$ necessarily. □

Proof of Proposition 2.— Recall that the principal solves:

$$\max_{B \geq 0} \left\{ s + a_0^*(B; f, s)\left[v(\bar{q}, a_0 = 1, B; s) - s\right] \right\} \text{ where } a_0^*(B; f, s) = \begin{cases} 1 & \text{if } rf \leq u(\bar{q}, a_0 = 1, B; s) \\ 0 & \text{otherwise} \end{cases}$$

Define then $s^* \in (0, s^{fb})$ as in Claim 2. Clearly, motivating the agent to install R is impossible when $s > s^*$. Hence, in this case, any B is optimal.

Consider then the case $s \leq s^*$. Note that in this case the principal always finds it optimal to incentivize the agent to exert effort. Indeed, for the agent to install R, it must be that $u(\bar{q}, a_0 = 1, B; s) \geq rf$. However, since $f > 0$, this requires that B must be such that $q^* < \bar{q}$, implying that $v(\bar{q}, a_0 = 1, B) > s$.

Now, $v(\bar{q}, a_0 = 1, B; s)$ is strictly decreasing in B; thus, the principal offers the agent the minimum B that satisfies $u(\bar{q}, a_0 = 1, B; s) = rf$. That is:

$$B^* = \underline{B}(f, s) \equiv \inf\{B \geq 0 : u(\bar{q}, a_0 = 1, B; s) = rf\}$$

which always exists given that $s \leq s^*$. Consequently, the principal expected payoffs at $t = 0$ is:

$$v_0^* = \begin{cases} v(\bar{q}, a_0 = 1, \underline{B}(f, s); s) > s & \text{if } s \leq s^* \\ s & \text{otherwise} \end{cases}$$

Since the agent is always indifferent between installing and not installing R, then his expected payoffs at $t = 0$ is $u_0^* = 0$. □

B.3 Proof of Corollary 1

Part (i).— First, note that $w(\bar{q}, a_0 = 1; s) \to g$ as $\bar{q} \to 1$, so $s^{fb} \to g - rf$ when we take such a limit. Thus, $w_0^{fb} \to \max\{g - rf, s\}$ as $\bar{q} \to 1$. Now, $v(\bar{q}, a_0 = 1, B; s) \to g - \lambda\mu B/(1 + \mu)$ and $u(\bar{q}, a_0 = 1, B; s) \to \lambda\mu B/(1 + \mu)$ as $\bar{q} \to 1$, so in the limit the principal solves $\max_{B \geq 0}\{s + \mathbb{1}_{\lambda\mu B/(1+\mu) \geq rf}\left[g - \lambda\mu B/(1 + \mu) - s\right]\}$, where $\mathbb{1}_X$ is the indicator function of the event X. The solution to this problem is $\lambda\mu B^*/(1+\mu) = rf$ if $g - rf > s$ and $B^* = 0$ otherwise, so $v_0^* \to \max\{g - rf, s\}$ as $\bar{q} \to 1$ also. □

Part (ii).— Note that $w(\bar{q}, a_0 = 1; s) \to g\bar{q} + s(1 - \bar{q})$ as $\lambda \to \infty$ and $\Pi \to 0$ keeping $g \equiv \lambda\Pi$ constant, so $w_0^{fb} \to \max\{g\bar{q} + s(1 - \bar{q}) - rf, s\}$. Now, $v(\bar{q}, a_0 = 1, B; s) \to \bar{q}(g - \lambda\mu B/(1 + \mu)) + s(1 - \bar{q})$ and $u(\bar{q}, a_0 = 1, B; s) \to \bar{q}\lambda\mu B/(1 + \mu)$ as $\lambda \to \infty$ and $\Pi \to 0$ keeping $g \equiv \lambda\Pi$ constant, so in the limit the principal solves:

$$\max_{B \geq 0}\{s + \mathbb{1}_{\bar{q}\lambda\mu B/(1+\mu) \geq rf}\left[\bar{q}(g - \lambda\mu B/(1 + \mu)) + s(1 - \bar{q}) - s\right]\}$$

The solution to this problem is $\bar{q}\lambda\mu B^*/(1 + \mu) = rf$ if $\bar{q}(g - s) > rf$ and $B^* = 0$ otherwise, so v_0^* also converges to $\max\{g\bar{q} + s(1 - \bar{q}) - rf, s\}$ as $\lambda \to \infty$ and $\Pi \to 0$ keeping $g \equiv \lambda\Pi$ constant. □

Appendix C The Optimal Contract

C.1 Proof of Proposition 3

Note that the principal's expected ex-ante payoff can be written as:

$$v_0 = \mathbb{E}_0\left[\int_0^\infty re^{-rt}\big[a_0 k_t \lambda\theta\Pi + s(1 - k_t)\big]dt\right] - u_0 - ra_0 f$$

where $u_0 = \bar{q}a_0 C^1(k, c) + (1 - a_0\bar{q})C^0(k, c) - ra_0 f$. If the principal offers the contract specified in the statement of Proposition 3, then the agent installs R and obtains an ex-ante payoff of zero, i.e., $a_0^* = 1$ and $u_0^* = 0$, and:

$$\mathbb{E}_0\left[\int_0^\infty re^{-rt}\big[k_t\lambda\theta\Pi + s(1 - k_t)\big]dt\right] = w(\bar{q}, a_0 = 1; s)$$

Hence, $v_0^* = w(\bar{q}, a_0 = 1; s) - rf = w_0^{fb}$. □.

C.2 Proof of Proposition 4

To prove the proposition, it suffices to show that under the stated contract the principal follows the first-best experimentation strategy.

For $t < T^{fb}(\bar{q})$.— It is easy to see that the principal plays $k_t^* = 1$. Indeed, if $N_t = 1$, then $q_t = 1$ and the bonus B^* was already disbursed. Hence $k_t^* = 1$ as $\lambda\Pi > s$. If $N_t = 0$, on the other hand, then $k_t^* = 1$, since $D^* > s$ and the principal's continuation payoff in the event of a breakthrough is $r(\Pi - B^*) + \lambda\Pi$, which is always strictly greater than zero given the value of B^* and the fact that $s < s^{fb}$.

For $t \geq T^{fb}(\bar{q})$.— If $N_t = 1$, then $k_t^* = 1$, as $q_t = 1$, the bonus B^* was already disbursed, and $\lambda\Pi > s$. If $N_t = 0$, then, $k_t^* = 0$ given that the principal no longer has to pay damages for S's profits, and $q_t \leq q^{fb}$ (given that $k_t^* = 1$ for all $t < T^{fb}(\bar{q})$, as we showed above).

□

References

Acs, Zoltan J., and David B. Audretsch. 1988. "Innovation in Large and Small Firms: An Empirical Analysis." *American Economic Review*, 78(4): 678–690.

Aghion, Philippe, and Jean Tirole. 1994. "The Management of Innovation." *The Quarterly Journal of Economics*, 109(4): 1185–1209.

Almazan, Andres, and Javier Suarez. 2003. "Entrenchment and Severance Pay in Optimal Governance Structures." *The Journal of Finance*, 58(2): 519–547.

Amador, Manuel, Iván Werning, and George-Marios Angeletos. 2006. "Commitment vs. Flexibility." *Econometrica*, 74(2): 365–396.

Arrow, Kenneth. 1962. "The Rate and Direction of Inventive Activity." , ed. R. Nelson, Chapter Economic Welfare and the Allocation of Resources for Invention. Princeton, NJ: Princeton University Press.

Arrow, Kenneth. 1974. *The Limits of Organization.* New York: Norton.

Awaya, Yu, and Vijay Krishna. forthcoming. "Startups and Upstarts: Disadvantageous Information in R&D." *Journal of Political Economy.*

Bergemann, Dirk, and Ulrich Hege. 1998. "Venture Capital Financing, Moral Hazard, and Learning." *Journal of Banking & Finance*, 22(6): 703 – 735.

Bergemann, Dirk, and Ulrich Hege. 2005. "The Financing of Innovation: Learning and Stopping." *The RAND Journal of Economics*, 36(4): 719–752.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica*, 67(2): 349–374.

Bolton, Patrick, Markus K. Brunnermeier, and Laura Veldkamp. 2012. "Leadership, Coordination, and Corporate Culture." *The Review of Economic Studies*, 80(2): 512–537.

Bonatti, Alessandro, and Heikki Rantakari. 2016. "The Politics of Compromise." *American Economic Review*, 106(2): 229–59.

Bresnahan, Timothy F., Shane Greenstein, and Rebecca M. Henderson. 2011. "Schumpeterian Competition and Diseconomies of Scope: Illustrations from the Histories of Microsoft and IBM." *The Rate and Direction of Inventive Activity Revisited*, 203–271. University of Chicago Press.

Brown, DG., Wobst HJ., Kapoor A., Kenna LA., and Southall N. forthcoming. "Clinical Development Times for Innovative Drugs." *Nature Reviews Drug Discovery.*

Chassang, Sylvain. 2010. "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts." *American Economic Review*, 100(1): 448–65.

Christensen, Clayton M. 1993. "The Rigid Disk Drive Industry: A History of Commercial and Technological Turbulence." *The Business History Review*, 67(4): 531–588.

Christensen, Clayton M. 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail.* Boston, MA: Harvard Business School Press.

Crémer, Jacques, Luis Garicano, and Andrea Prat. 2007. "Language and the Theory of the Firm." *The Quarterly Journal of Economics*, 122(1): 373–407.

DiMasi, Joseph A., Henry G. Grabowski, and Ronald W. Hansen. 2016. "Innovation in the pharmaceutical industry: New estimates of R&D costs." *Journal of Health Economics*, 47: 20–33.

Ely, Jeffrey, George Georgiadis, and Luis Rayo. 2022. "Incentive Compatibility in Dynamic Information Design." Working Paper.

Futures Group. 1984. "Characterization of Innovations Introduced on the U.S. Market in 1982." Nat. Tech. Information Service Document PB-84-212067.

Gans, Joshua. 2016. *The Disruption Dilemma.* Cambridge, MA: The MIT Press.

Garicano, Luis, and Luis Rayo. 2016. "Why Organizations Fail: Models and Cases." *Journal of Economic Literature*, 54(1): 137–92.

Gellman Research Associates. 1976. "Indicators of International Trends in Technological Innovation." Final Report, National Science Foundation, Nat. Tech. Information Service Document PB-263-738.

Gellman Research Associates. 1982. "The Relationship between Industrial Concentration, Firm Size and Technological Innovation." Final Report, US Small Bus. Admin.

Guo, Yingni. 2016. "Dynamic Delegation of Experimentation." *American Economic Review*, 106(8): 1969–2008.

Halac, Marina, Navin Kartik, and Qingmin Liu. 2017. "Contests for Experimentation." *Journal of Political Economy*, 125(5): 1523–1569.

Henderson, Rebecca. 1993. "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Equipment Industry." *The RAND Journal of Economics*, 24(2): 248–270.

Henderson, Rebecca M., and Kim B. Clark. 1990. "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms." *Administrative Science Quarterly*, 35(1): 9–30.

Holmstrom, Bengt. 2016. "Nobel Prize Lecture: Pay for Performance and Beyond." Retrieved from Nobelprize.org. (Accessed: 17 January 2021).

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

Igami, Mitsuru. 2017. "Estimating the Innovator's Dilemma: Structural Analysis of Creative Destruction in the Hard Disk Drive Industry, 1981—1998." *Journal of Political Economy*, 125(3): 798–847.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica*, 73(1): 39–68.

Klepper, Steven, and Sally Sleeper. 2005. "Entry by Spinoffs." *Management Science*, 51(8): 1291–1306.

Leland, Hayne E. 1994. "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure." *The Journal of Finance*, 49(4): 1213–1252.

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

Markides, Constantinos, and Paul A. Geroski. 2005. *Fast Second: How Smart Companies Bypass Radical Innovation to Enter and Dominate New Markets.* San Francisco, CA: John Wiley & Sons.

Moroni, Sofia. 2020. "Experimentation in Organizations." Working Paper.

Murphy, Kevin J. 1999. "Chapter 38: Executive Compensation." In . Vol. 3 of *Handbook of Labor Economics*, 2485 – 2563. Elsevier.

Nelson, Richard R., and Sidney G. Winter. 1982. *An Evolutionary Theory of Economic Change.* Cambridge, MA: Belknap Press of Harvard University Press.

O'Reilly, Charles A., and Michael L. Tushman. 2016. *Lead and Disrupt: How to Solve the Innovator's*

Dilemma. Stanford, CA: Stanford University Press.

Pavitt, K., M. Robson, and J. Townsend. 1987. "The Size Distribution of Innovating Firms in the UK: 1945-1983." *The Journal of Industrial Economics*, 35(3): 297–316.

Prendergast, Canice. 1993. "The Role of Promotion in Inducing Specific Human Capital Acquisition." *The Quarterly Journal of Economics*, 108(2): 523–534.

Prendergast, Canice. 1999. "The Provision of Incentives in Firms." *Journal of Economic Literature*, 37(1): 7–63.

Rotemberg, Julio J., and Garth Saloner. 1993. "Leadership Style and Incentives." *Management Science*, 39(11): 1299–1318.

Rotemberg, Julio J., and Garth Saloner. 1994. "Benefits of Narrow Business Strategies." *American Economic Review*, 84(5): 1330–1349.

Scherer, F. M. 1965. "Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions." *American Economic Review*, 55(5): 1097–1125.

Van den Steen, Eric. 2005. "Organizational Beliefs and Managerial Vision." *Journal of Law, Economics, & Organization*, 21(1): 256–283.