



ELSEVIER

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb



Experimentation and project selection: Screening and learning [☆]

Renato Gomes ^a, Daniel Gottlieb ^{b,*}, Lucas Maestri ^c^a Toulouse School of Economics, University of Toulouse Capitole (CNRS), France^b Olin Business School, Washington University in St. Louis, United States^c FGV/EPGE – Escola Brasileira de Economia e Finanças, Brazil

ARTICLE INFO

Article history:

Received 11 November 2013

Available online 17 February 2016

JEL classification:

D21

J33

M12

O31

O34

Keywords:

Adverse selection

Experimentation

Bandit problem

Multi-dimensional screening

Entrepreneurship

ABSTRACT

Firms must strike a delicate balance between the *exploitation* of well-known business models and the *exploration* of risky, untested approaches. In this paper, we study financial contracting between an investor and a firm with private information about its returns from exploration and exploitation. The investor-optimal mechanism offers contracts with different tolerance for failures to screen returns from exploitation, and with different exposure to the project's revenues to screen returns from exploration. We derive necessary and sufficient conditions for private information about returns from exploration to have zero value to the firm. When these conditions fail, private information about exploration may even decrease the firm's payoff.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

To be successful, most firms need to strike a delicate balance between *exploiting* well-known business models and *exploring* new approaches (cf. Roberts and Weitzman, 1981). Exploration is costly and inherently risky. Most untested approaches fail, but some of them turn out to be exceptionally profitable.

Firms engaged in innovation typically require external financing to develop their ventures. Their choices of which projects to pursue, and the economic rents obtained from them, depend on the financing contracts signed with outside investors.

As long recognized by economists, asymmetric information is a key determinant of financial contracts. Firms often have an informational advantage (or expertise) on the potential of their innovative projects (exploration) – for example, because

[☆] We acknowledge funding support from the Mack Center for Technological Innovation at the Wharton School. Gottlieb also acknowledges financial support from the Dorinda and Mark Winkelman Distinguished Scholar Award. We benefited from discussions at the Toulouse–Northwestern IO Conference 2012, SAET 2013, and the Stony Brook Game Theory Conference 2013. We thank the Editor, the Advisory Editor, the anonymous referees, Thomas Mariotti and Niko Matouschek for detailed comments that greatly improved the paper.

* Corresponding author.

E-mail addresses: renato.gomes@tse-fr.eu (R. Gomes), dgottlieb@wustl.edu (D. Gottlieb), lucas.maestri@fgv.br (L. Maestri).

they have a better understanding of market conditions or the details of their technology.¹ They are also typically better informed about the profitability of on-going business ventures and the best alternative use of their human and physical assets (exploitation). The decision between exploiting and exploring (i.e., project selection) may also be affected by internal agency problems² and indirect costs of failure, such as stigma,³ that constrain the ability or willingness of firms to innovate.

In this paper, we study financial contracting in environments in which firms have private information about the payoff from *both* exploration and exploitation. Our analysis sheds light on three fundamental questions on the economics of innovation. First, what contractual solutions arise to mitigate adverse selection in the financing of innovation? Second, how does asymmetric information affect project selection and the frequency/duration of experimentation? Third, in light of the multiple sources of informational advantage held by firms, what determines the rents from engaging in innovation?

We employ a two-arm bandit problem to capture the firm's exploration/exploitation trade-off.⁴ In each period, the firm (or entrepreneur) either plays a *safe arm* with a known distribution of payoffs (capturing exploitation) or a *risky arm* (capturing exploration) that may be either "good" or "bad." Revenues are always zero if the risky arm is bad. If the risky arm is good, large revenues come at geometrically-distributed random times. Playing the risky arm therefore requires the firm to forego the payoff from the safe arm.

We model the financing of innovation by embedding the two-arm bandit problem described above into a principal-agent framework. The entrepreneur (agent) needs financing from an investor (principal) to play the risky arm. Without the investor's assistance, the entrepreneur is only able to play the safe arm. Both the investor and the entrepreneur *learn* about the risky arm by observing its revenues.

There is asymmetric information regarding the primitives of the bandit problem that the agent plays. The agent has private information about the probability that the risky arm is good (the *expected returns* from exploration) and about the expected payoff from playing the safe arm (the *opportunity cost* of exploration). We follow a mechanism design approach to derive the revenue-sharing rules and the exploration/exploitation strategies that maximize investor profits and identify the contractual features that optimally screen the entrepreneur's (two-dimensional) private information.

We start by analyzing two natural benchmarks where the informational asymmetry between the investor and the entrepreneur is one-dimensional. In the first benchmark, the payoff from the safe project is common knowledge but the entrepreneur has private information about the quality of the risky project. We show that the investor can attain the first-best profits in this case by reimbursing the entrepreneur's foregone rents from exploitation and collecting all the profit from the project. This result, which echoes the classic contributions of [Riordan and Sappington \(1988\)](#) and [Crémer and McLean \(1988\)](#), underscores the appropriability of innovation in arm's length relationships.

Next, we consider the *expert investor* benchmark, where the probability that the risky project is good is common knowledge, but the entrepreneur is privately informed about her opportunity cost of exploration. In the solution to this problem, the investor offers a menu of *stopping plans* to the entrepreneur, guaranteeing the financing of exploration until a critical period regardless of the output realizations. After this critical period, financing is continued if and only if high revenues were produced in at least one period. The investor screens the entrepreneur's opportunity cost by offering stopping plans with different durations. Entrepreneurs with high opportunity costs choose shorter plans, while those with unprofitable alternatives choose longer plans.

Because the investor needs to leave informational rents to the entrepreneur, the solution of this one-dimensional benchmark features under-experimentation (i.e., too few entrepreneurs experiment) and early termination (i.e., the risky project is halted too early) relative to the first-best outcome where innovation is self-financed. Moreover, the investor can implement the optimal mechanism using lump-sum payments (such as a golden handshake).

These two benchmarks lay the ground to our characterization of the optimal mechanism under two-dimensional private information. Our main result gives a necessary and sufficient condition for the solution to the expert investor benchmark to be implementable, in which case expertise about the risky project quality does not affect payoffs. We show that, with two-dimensional asymmetric information, the investor can implement the optimal mechanism using *linear contracts*, which are particularly simple and commonly used in project financing. A linear contract specifies the entrepreneur's remuneration as the sum of a lump-sum payment (such as a golden handshake) and a variable payment that is linear in the revenue from the project (such as equity). The investor uses the duration of the contract and the exposure to the revenue from exploration to screen the two dimensions of private information. Entrepreneurs with better projects pick higher-powered contracts (that is, contracts with larger revenue shares but smaller lump-sum payments). Entrepreneurs with higher opportunity costs pick contracts with smaller durations. The interaction between both dimensions of private information manifests itself in the use of variable payments that depend on the project's revenue (e.g., equity). Although expertise about the risky project's quality generates no additional informational rents to the entrepreneur, it determines the contractual form employed at the optimum.

¹ [Aboody and Lev \(2000\)](#) show that R&D is a major source of insider private information. [Gompers \(1995\)](#) shows that venture capitalists concentrate in early-stage companies and high technology industries, where informational asymmetries are more prevalent. [Gompers and Lerner \(2004, chapter 7\)](#) describe how staged investment helps venture capitalists screen asymmetric information about project quality.

² There is empirical evidence that, when deciding between exploration and exploitation, firms often behave as if the returns from their on-going ventures were larger than they actually are; as induced, for example, by career concerns of managers – see [Kerr et al. \(2014\)](#).

³ [Landier \(2006\)](#) argues that indirect costs such as the stigma from failure are important barriers to experimentation.

⁴ Two-arm bandit problems were first introduced in the economics literature by [Rothschild \(1974\)](#).

The necessary and sufficient condition mentioned above requires the reverse hazard rate of the distribution of payoffs from the safe project to be *concave*. Intuitively, this rate determines the impact of financing an entrepreneur with opportunity cost c on the informational rents that have to be left to all entrepreneurs with opportunity costs less than c . Concavity means that informational rents change at a decreasing rate (that is, the *marginal informational rent* is decreasing). Under this condition, holding fixed the probability θ that the risky arm is good, the *marginal firm* to be financed is the one with the highest opportunity cost c and, therefore, the one with the lowest marginal impact on informational rents. As a result, an increase in θ makes the expert investor willing to expand the range of opportunity costs that are financed, and the more so the higher is θ . At the optimum, this results in expected payments that are *convex* in θ .

Convex expected payments can be obtained by offering a menu of linear contracts where the variable part is increasing in the probability that the risky project is good, θ . Since incentive compatibility requires the variable pay to be increasing in θ , convexity of expected payments is the condition that allows the investor to use a menu of linear contracts to implement the expert investor benchmark. That is, the incentive compatibility constraint associated with the quality of the risky project is slack at the optimum if and only if the reverse hazard rate is concave.

It is well known that, in one-dimensional models, the principal's ability to screen types depends solely on the *slope* of informational rents. In our expert investor benchmark, for instance, the optimality of perfect screening is determined by the slope of the reverse hazard rate of the distribution of opportunity costs. In our two-dimensional model, for perfectly screening beliefs about project quality to be optimal, we have to take into account not only the *slope* of the reverse hazard rate but also its *curvature*. Because perfect screening occurs at no additional cost relative to the one-dimensional benchmark, the concavity of the reverse hazard rate also implies that the investor is able to fully extract the rents associated with the firm's private information about project quality.

When the concavity condition fails, the expert investor benchmark is not implementable under two-dimensional asymmetric information. While it is hard to obtain a complete (closed-form) characterization of this case, we are able to demonstrate two important facts. First, pooling is a robust feature of the optimal deterministic mechanism. Intuitively, since screening beliefs about project quality is incompatible with screening the entrepreneur's opportunity cost of exploration, the investor offers the same contract to a positive-measure set of entrepreneurs. Second, in the case where one round of experimentation perfectly distinguishes good and bad projects, we show that when the reverse hazard rate of the distribution of payoffs from the safe project is *strictly convex*, the investor offers the same contract to all types (*complete pooling*). Convexity magnifies the conflict between screening the payoffs from exploitation and exploration. When complete pooling is optimal, some entrepreneurs obtain *negative* informational rents from their private information about the risky project's returns (i.e., the entrepreneur's payoff would increase if the payoff from exploration were common knowledge). The reason is that, because of binding *global* incentive constraints, the investor denies financing to entrepreneurs with risky projects of intermediate expected quality as a way to reduce the informational rents of those with projects that are more likely to succeed.

Empirical implications. Our model has three main empirical implications: (1) projects with longer duration are offered greater protection (for example, through severance packages or golden parachutes) and are more likely to succeed; (2) projects with higher-powered long-term incentives are more likely to succeed; and (3) firms with higher payoffs from exploitation have shorter contracts.

The first and second implications are consistent with Lerner and Wulf (2007), who show that long-term incentives (such as stock options and restricted stock) are positively correlated with the number of patents, citations, and "originality" (measured by patent breadth). They are also consistent with Francis et al. (2009), who show that firms that offer golden parachutes to their CEOs are more innovative (measured by number of patents filed and awarded). We are not aware of any empirical work that tests the third implication.

Related literature

Many recent papers study agency issues related to experimentation. Most of them focus on the moral hazard aspect of how to motivate an agent to experiment. Manso (2011) presents a two-period model where the agent may shirk, follow a well-known approach, or explore a new approach. He shows that contracts that incentivize exploring new approaches are fundamentally different from standard pay-for-performance contracts, which are meant to incentivize well-known approaches. Ederer (2013) extends his framework to a situation where a principal faces multiple agents. Klein (2015) considers a continuous-time version of Manso's two-period model. In his model, the principal either cannot implement the innovative activity at all, or implements it at zero costs by making the exploitative activity sufficiently unappealing to the agent. Using a laboratory experiment, Ederer and Manso (2013) discuss the effectiveness of different contracts for incentivizing innovation.⁵

Bergemann and Hege (1998, 2005) and Hörner and Samuelson (2013) analyze models of moral hazard without commitment where the principal and the agent learn about the value of the investment over time.⁶ Bonatti and Hörner (2011)

⁵ Garfagnini (2011) considers a setting without transfers, and shows how the choice between a risky and a safe arm can be used to incentivize the agent to pick higher effort.

⁶ Bergemann and Hege (1998) also study a version of their model with commitment. In contrast to Manso (2011), there is moral hazard on the provision of effort, rather than on the choice between exploration and exploitation.

study experimentation in teams with unobservable actions, and [Campbell et al. \(2014\)](#) analyze how sharing information among team members affects experimentation.⁷

Our paper focuses on the adverse selection aspect of financing innovation. A few other papers also study experimentation in settings with adverse selection. [Heidhues et al. \(2015\)](#) analyze the free-riding aspect of strategic experimentation among multiple agents in a dynamic model with private payoffs. [Bouvard \(2014\)](#) studies a signaling model of financing. In his model, entrepreneurs have private information about the quality of their projects, which can be high or low. He shows that entrepreneurs signal the quality of their projects by distorting the timing of investments. Our model differs from his in three aspects. First, we consider a screening environment where financiers are the ones designing the contract. In our framework, the principal can always implement the first best if the only source of private information is the quality of the project. Second, we allow for private information about the entrepreneur's exploitation payoff. This additional source of private information prevents the principal from implementing the first best and introduces a rationale for screening through different contract durations. And third, we allow for a continuum of types whereas there are only two types in Bouvard's model.

[Halac et al. \(2015\)](#) also study optimal experimentation in a principal-agent relationship with a two-arm bandit problem. In their model, there is moral hazard regarding the agent's effort choice and adverse selection regarding the probability that the risky arm yields a success *conditional on being of good quality*.⁸ In the benchmarks with pure adverse selection or pure moral hazard, the first-best allocation is implementable. In the full-blown model, however, the interaction between learning, adverse selection and moral hazard makes it impossible to implement the first best allocation, resulting in under-experimentation by low-ability agents. Similarly, in the version of our model where there is only adverse selection on the agent's exploration payoffs, the first-best allocation is implementable.⁹ This is no longer the case when a second source of asymmetric information is present (regarding exploitation payoffs), and under-experimentation occurs at the optimum. One reason for this result is that there is no public signal available to the principal that is correlated with the agent's exploitation payoff.

[Guo \(2014\)](#) also embeds a two-arm bandit problem in a principal-agent framework but, instead, assumes that utility is not transferable. She shows that optimal policies place a cap on experimentation.¹⁰ Finally, [Fong \(2009\)](#) studies a model where the principal dynamically screens experts by offering scoring rules and threatening to end the relationship after a bad performance. Our models differ in that she assumes that there is no experimentation, private information is one dimensional, there is moral hazard, and side payments are not allowed.

It is often challenging to solve multi-dimensional screening problems, as one cannot determine from the outset the direction in which incentive constraints bind (see [Rochet and Stole, 2003](#) and the references therein). In our setting, the nature of the bandit problem played by the agent allows us to identify the exact condition under which a single dimension of private information is responsible for the agent's informational rents. In this case, we are able to show that incentive compatibility holds if and only if there are no profitable misreports *along each dimension* (in which case we can safely ignore two-dimensional deviations). By contrast, and similarly to the seminal contribution of [Rochet and Choné \(1998\)](#), profitable two-dimensional deviations cannot be ruled out when this condition does not hold. We then show that pooling is a robust feature of the optimal deterministic mechanism.

The remainder of the paper is organized as follows. Section 2 introduces the framework. Section 3 solves for the first-best and the two benchmark cases where the asymmetry of information is one-dimensional. Section 4 studies the optimal mechanism under two-dimensional asymmetric information. Section 5 concludes.

2. Model

There is one *investor* (principal) and one *innovative firm* or *entrepreneur* (agent). Both of them are risk neutral. Time $t \in \{1, \dots, T\}$ is discrete and finite, and both parties discount the future at rate $\delta \in (0, 1]$. At every time t , the firm undertakes either a risky but potentially profitable project (*risky arm*) or a safe project (*safe arm*).

Playing the risky arm corresponds to experimenting something new, with an unknown distribution of returns. The innovative firm is not able to launch the risky project without the investor's assistance. For example, the firm may need access to assets or technology held by the investor in order to develop the project. It may also be financially constrained and therefore need the investor to help finance the project's operational costs. Undertaking the risky project requires a per-period outlay of $K \geq 0$.¹¹ The innovative firm has private information about the *quality of the risky project* θ (the *expected returns* from exploration).

⁷ [Canidio and Legros \(2015\)](#) analyze how the presence of career concerns shapes optimal experimentation contracts. [Board and Meyer-ter-Vehn \(2013\)](#) study the dynamics of reputation in a setting where firms make unobservable investments in quality and consumers learn about the quality of the good by experimentation.

⁸ By contrast, the agent in our model has private information about the prior probability that the risky arm is good.

⁹ Our mechanism is however different from that of [Halac et al. \(2015\)](#), as the two papers assume different forms of asymmetric information about the payoff distribution of the risky arm.

¹⁰ Because [Guo \(2014\)](#) studies an environment without side transfers, she is able to use Lagrangian methods to determine the optimal ranges of delegation. This method seems less promising in our environment, as the problem may not be concave.

¹¹ We set $K = 0$ if the investor's assistance consists in licensing intellectual property or providing expertise about how to develop and bring new projects to the market.

Playing the safe arm corresponds to exploiting business alternatives with a known distribution of returns. The expected payoff from the safe project c is also private information of the innovative firm. Beyond the pecuniary returns from exploitation, c may also reflect the inability or unwillingness of the firm to engage in experimentation (due to internal agency problems or indirect costs of failure, such as stigma), or the firm's cost of raising the capital needed to engage in the new project. We refer to c as the *opportunity cost* of exploration.

Geometric bandit returns. The risky project is of either *good* or *bad* quality; it is good with prior probability θ . A bad project always has zero revenues. A good project yields a positive revenue $\Delta > 0$ at geometrically-distributed random times. That is, in each period, a good project's revenue equals Δ with probability $\lambda \in (0, 1)$ and zero with complementary probability.¹² Therefore, the expected per-period profit (i.e., revenues net of costs) from a risky project of good quality is

$$\lambda \cdot \Delta - K,$$

which we assume to be strictly positive. The probability λ captures the underlying economic risk faced by agents in the economy (e.g., due to unpredictable business fluctuations) and is common knowledge between the investor and the innovative firm. The per period outlay K is also commonly known by both parties.

We refer to the event that the risky project generates revenue Δ as a *success*, and we refer to the event that it generates no revenue as a *failure*. Let $x_t \in \{s, f, \emptyset\} \equiv X$ denote the output of the project at time t , where $x_t = s$, denotes a success, $x_t = f$ denotes a failure, and $x_t = \emptyset$ means that the risky project is inactive at t . If the risky project is inactive, it produces no cash flow. We adopt the accounting convention that, whenever the risky project produces a success, the revenue Δ accrues to the investor.

Learning. An important feature of our model is that both the investor and the innovative firm learn about the quality of the risky project from its realized outputs, which are observed by both parties. We refer to the vector $x^t \equiv (x_1, \dots, x_t) \in X^t$ as the *output history* in period t . For future reference, and with a slight abuse of notation, let $S(x^t)$ and $F(x^t)$ denote the number of successes and failures associated with the output history x^t .

The firm and the investor update beliefs according to Bayes' rule. A success immediately reveals that the risky project is good, whereas a failure is only partial information about the quality of the risky project. A firm with initial belief θ which observes failures in all periods prior to t believes that the risky project is good with probability

$$\theta_t(\theta) \equiv \frac{\theta \cdot (1 - \lambda)^{t-1}}{\theta \cdot (1 - \lambda)^{t-1} + (1 - \theta)}.$$

Thus, at each time before a success is observed, either a new failure happens, which reduces the posterior probability to $\theta_{t+1}(\theta)$, or a success happens, which shifts the posterior probability to 1. The firm's belief that the risky project is good after t observations is a random variable with support on $\{\theta_t(\theta), 1\}$. We denote this random variable by $\tilde{\theta}_t(\theta)$.

Asymmetric information. The firm has private information about the probability that the project is good θ and about the payoff from the safe project c . From the investor's perspective, θ is drawn from a distribution G with support $[0, 1]$, while c is drawn from a distribution H with support $[\underline{c}, \bar{c}] \subseteq \mathbb{R}_{++}$. Both distributions are absolutely continuous with differentiable densities $g, h > 0$.

We make two technical assumptions. First, as usual in bandit models, the distributions of the payoffs from each arm, G and H , are independent. Second, as is standard in mechanism design, we require the distribution of safe project payoffs H to be log-concave, or, equivalently, the reverse hazard rate $\gamma(c) \equiv \frac{h(c)}{h(\bar{c})}$ to be weakly increasing.

To rule out uninteresting cases, we further assume that it is inefficient for firms with the highest c to undertake the risky project even when the project is known to be good, and that it is efficient for firms with the lowest c to undertake the risky project if it is known to be good: $\bar{c} \geq \lambda \cdot \Delta - K \geq \underline{c}$.

The firm's private information at $t = 1$ is described by the pair $\tau \equiv (\theta, c) \in \Gamma \equiv [0, 1] \times [\underline{c}, \bar{c}]$. We refer to τ as the firm's *type*.

Financing mechanisms

Since the output history is observable and verifiable, parties can write contracts contingent on the project's revenues. The investor makes a take-it-or-leave-it offer to the innovative firm in the form of a *mechanism*. No enforceability problems arise (for example, because of vesting provisions or large severance payments); i.e., there is commitment.¹³

A mechanism consists of an *action plan* ϕ and a *payment rule* p . The action plan ϕ determines when to undertake the risky project as a function of the firm's type and the output history. Using the Revelation Principle, we define an action plan as a collection of functions $\phi_t : \Gamma \times X^{t-1} \rightarrow [0, 1]$, $t = 1, \dots, T$, that map types and histories into the unit interval.¹⁴

¹² Keller et al. (2005) consider the continuous-time version of this stochastic process and call it an *exponential bandit* process, as the time before a success follows an exponential distribution. We adapt their nomenclature by noting that the geometric distribution is the discrete-time analog of the exponential distribution and therefore refer to this process as a *geometric bandit*. Two-armed geometric bandits are adopted, among many others, by Bergemann and Hege (1998, 2005), Hörner and Samuelson (2013), and Gerardi and Maestri (2012).

¹³ Kaplan and Stromberg (2003) document the widespread use of vesting provisions and severance payments in venture capital financing.

¹⁴ For national convenience, we set $X^0 \equiv \{\emptyset\}$ to be the output history before the relationship begins.

The function $\phi_t(\tau|x^{t-1})$ specifies, for each type τ , the probability of undertaking the risky project in period t when the output history up to that period is x^{t-1} . In this case, with probability $1 - \phi_t(\tau|x^{t-1})$, the risky arm is not played and the firm is allowed to pursue the safe project (reaping a payoff of c). Fixing the firm's type, we refer to the collection $\phi(\tau) \equiv \{\phi_t(\tau|\cdot) : t = 1, \dots, T\}$ as the *type- τ action plan*.

A payment rule is a function $p : \Gamma \times X^T \rightarrow \mathbb{R}$ that maps the firm's type τ and the final output history x^T into a payment to the firm. With no loss of generality, we evaluate payments in first period units. Although this normalization may seem counterintuitive (payments are conditional on the final history), it allows us to compare payments between projects that last for a different number of periods. A mechanism is therefore described by the pair $\mathcal{M} = (\phi, p)$.

For a given θ , an action plan ϕ uniquely defines a probability measure over the space of final output histories X^T , denoted by $\mu[\phi|\theta]$.¹⁵ Notice that $\mu[\phi|\theta]$ depends on the prior belief that the risky project is good, θ , but not on the safe project's payoff, c .

In order to describe the incentive compatibility constraints, it is useful to define the expected payment of a firm that reports type $\hat{\tau}$ when the prior probability that the risky project is good equals θ . This expectation is computed according to the measure $\mu[\phi(\hat{\tau})|\theta]$, which assigns probabilities to final output histories as induced by the type- $\hat{\tau}$ action plan $\phi(\hat{\tau})$ when the true probability that the risky project is good equals θ :

$$\tilde{P}(\hat{\tau}; \theta) \equiv \mathbb{E}^{\mu[\phi(\hat{\tau})|\theta]} \left[p(\hat{\tau}, \tilde{x}^T) \right]. \tag{1}$$

Let $P(\tau) \equiv \tilde{P}(\tau; \theta)$ denote the expected payment when the firm reports types truthfully.

If a firm reports type $\hat{\tau}$, the expected discounted duration of a project that is good with prior probability θ equals

$$\tilde{\Phi}(\hat{\tau}; \theta) \equiv \mathbb{E}^{\mu[\phi(\hat{\tau})|\theta]} \left[\sum_{t=1}^T \delta^{t-1} \cdot \phi_t(\hat{\tau}|x^{t-1}) \right]. \tag{2}$$

Let $\Phi(\tau) \equiv \tilde{\Phi}(\tau; \theta)$ denote the expected discounted duration under truthful reporting.

Using this notation, we can write the expected payoff of a firm of type $\tau = (\theta, c)$ which reports its type truthfully under the mechanism $\mathcal{M} = (\phi, p)$ as

$$U(\tau|\mathcal{M}) \equiv P(\tau) - \Phi(\tau) \cdot c. \tag{3}$$

A mechanism $\mathcal{M} = (\phi, p)$ is *incentive compatible (IC)* if the firm cannot benefit from deviating to a non-truthful reporting strategy:

$$\tilde{P}(\hat{\tau}; \theta) - \tilde{\Phi}(\hat{\tau}; \theta) \cdot c \leq U(\tau|\mathcal{M}) \tag{IC}$$

for all $\tau = (\theta, c)$ and $\hat{\tau}$ in Γ . A mechanism \mathcal{M} is *individually rational (IR)* if, for each $\tau = (\theta, c) \in \Gamma$, the firm is better off by accepting the investor's mechanism rather than undertaking the safe project in all periods:

$$U(\tau|\mathcal{M}) \geq 0. \tag{IR}$$

A mechanism $\mathcal{M} = (\phi, p)$ is *feasible* if it satisfies the (IC) and (IR) constraints. An action plan ϕ is *implementable* if there exists a payment rule p such that $\mathcal{M} = (\phi, p)$ is feasible. A mechanism \mathcal{M} is *optimal* if it maximizes the investor's ex-ante expected profit,

$$\Pi(\mathcal{M}) \equiv \mathbb{E}^{\theta, c} \left\{ \mathbb{E}^{\mu[\phi(\tau)|\theta]} \left[\sum_{t=1}^T \left(\delta^{t-1} \cdot \phi_t(\tau|x^{t-1}) \cdot \left(\tilde{\theta}_t(\theta) \cdot \lambda \cdot \Delta - K \right) \right) - p(\tau, x^T) \right] \right\}, \tag{4}$$

within the class of feasible mechanisms. To understand the expression above, note that the term inside brackets is the expected discounted revenue of the investor net of payments when the firm has type $\tau = (\theta, c)$. Taking expectations with respect to the firm's type gives the investor's ex-ante expected profits.

3. Preliminaries

We first consider three benchmarks that lay the ground to our characterization of the optimal mechanism.

¹⁵ To understand how to compute $\mu[\phi|\theta]$, let $T = 2$ and consider the action plan according to which a type- τ firm plays the risky arm in the first period, and, only in case of failure, switches to the safe arm at $t = 2$. The probability of history (s, s) , as induced by this action plan, is then the probability that the risky project is good and that it produces two consecutive successes: $\mu[\phi|\theta](s, s) = \theta\lambda^2$. By contrast, the measure $\mu[\phi|\theta]$ assigns zero probability to any final output history where the risky arm is not played at $t = 1$. A similar reasoning can be used to compute the probability of any other history.

3.1. First-best benchmark

In the first best, the investor observes the firm's type $\tau = (\theta, c)$, so we can disregard incentive constraints. A mechanism is *first-best optimal* if it maximizes the investor's ex-ante expected profit (4) subject to the (IR) constraint.

Because the investor wants to minimize payments to the firm, the (IR) constraint binds for every type $\tau \in \Gamma$. Using this fact, the principal's problem can be recast as that of choosing, for each $\tau \in \Gamma$, a type- τ action plan $\phi(\tau)$ that maximizes

$$\mathbb{E}^{\mu[\phi(\tau)|\theta]} \left[\sum_{t=1}^T \left(\delta^{t-1} \cdot \phi_t(\tau | \tilde{x}^{t-1}) \cdot \left(\tilde{\theta}_t(\theta) \cdot \lambda \cdot \Delta - K - c \right) \right) \right]. \tag{5}$$

As the objective function (5) shows, the principal faces a two-arm bandit problem with (i) a risky arm that gives positive revenues at geometrically-distributed random times (if the project is good) but costs $K + c$ at each round, and (ii) a safe arm with a zero payoff.

For each type $\tau = (\theta, c)$, the geometric bandit problem described above is the discrete-time analogue of the exponential bandit considered by Keller et al. (2005). As in its continuous-time version, the principal maximizes payoffs (5) by offering an action plan with the following threshold structure: If a success is observed in some period t , the risky project is continued until the end. If only failures occur, the project is terminated at some critical period $k(\tau)$. An action plan with this structure is called a *stopping plan*:

Definition 1 (*Stopping plans*). An action plan is a stopping plan if for every $\tau \in \Gamma$ there exists a number $k(\tau) \in \{1, \dots, T\}$ such that

1. the risky project is undertaken up to period $k(\tau)$ irrespective of its output history: $\phi_t(\tau|\cdot) = 1$ for all $t \leq k(\tau)$;
2. the risky project is terminated if no success occurs in the first $k(\tau)$ periods: $\phi_t(\tau|x^{t-1}) = 0$ for all $t > k(\tau)$ if $\mathbf{F}(x^{k(\tau)}) = k(\tau)$; and
3. the risky project is undertaken in all periods if at least one success occurs before period $k(\tau)$: $\phi_t(\tau|x^{t-1}) = 1$ for all $t \leq T$ if $\mathbf{S}(x^{k(\tau)}) \geq 1$.

In this case, we say that type τ tolerates $k(\tau)$ failures.

For a heuristic derivation of the first-best optimal action plan, fix the type $\tau = (\theta, c)$ and assume that $k - 1$ failures were observed in the first $k - 1$ periods. Consider the follow continuation strategy:

Pursue the risky project in period k . If a success happens in that period, continue until period T . Otherwise, terminate the project.

At time k , the expected payoff from this strategy is

$$V_k(\theta, c) \equiv \underbrace{\theta_k(\theta) \cdot \lambda \cdot \Delta - K - c}_{\text{flow payoff in period } k} + \mathbf{1}_{[k < T]} \cdot \theta_k(\theta) \cdot \lambda \cdot \delta \cdot \underbrace{\left[\frac{1 - \delta^{T-k}}{1 - \delta} \cdot (\lambda \cdot \Delta - K - c) \right]}_{\text{option value from a good project in period } k+1}, \tag{6}$$

where $\mathbf{1}_{[k < T]}$ is the indicator function that equals one if $k < T$ and zero if $k = T$. The first term is the expected flow payoff from continuing the risky project at time k . The second term is the expected option value of arriving at period $k + 1$ knowing that the risky project is good.

Since the expected payoff $V_k(\theta, c)$ is strictly decreasing in c , for every $k \in \{1, \dots, T\}$, we can define the *threshold value* $v_k^{FB}(\theta)$ as the safe project's payoff that makes the principal indifferent between pursuing the risky project at period k (and following the strategy above) and letting the firm undertake the safe project from period k on. That is, the threshold value $v_k^{FB}(\theta)$ is the solution to the following *Gittins Formula*:

$$V_k(\theta, v_k^{FB}(\theta)) = 0. \tag{7}$$

The next lemma describes the first-best action plan, which follows from Gittins (1979). For a parametric example, see Fig. 1.

Lemma 1 (*First best*). The first-best action plan ϕ^{FB} is a stopping plan. An innovative firm with type $\tau = (\theta, c)$ tolerates $k^{FB}(\tau) \in \{1, \dots, T\}$ failures if and only if

$$v_{k^{FB}(\tau)+1}^{FB}(\theta) < c \leq v_{k^{FB}(\tau)}^{FB}(\theta),$$

where $v_k^{FB}(\cdot)$ is the threshold function given by the Gittins Formula (7).




Fig. 1. The first-best optimal action plan when $[\underline{c}, \bar{c}] = [1, 2]$ under the following parameterization: $T = 3$, $\lambda = \frac{1}{2}$, $K = 1$, $\Delta = 6$ and $\delta = 1$. Areas 1, 2 and 3 refer to the set of types that play the risky arm at $t = 1$ and tolerate, respectively, 1, 2 and 3 failures according to ϕ^{FB} .

3.2. Symmetric information about the safe project

In this subsection, we study the problem of an investor who knows the payoff from the safe project c but does not know the probability that the risky project is good θ (which is private information of the firm). In this case, the investor can implement the first-best action plan and extract the firm’s entire rents by compensating it for its forgone payoff from the safe project. To understand how to achieve this, consider the mechanism consisting of the first-best action plan ϕ^{FB} and the payment rule

$$p^{share}(\tau, x^T) = \alpha \cdot \pi(x^T) + \sum_{\{t: x_t \neq \emptyset\}} \delta^{t-1} \cdot c,$$

where $\alpha \in (0, 1)$, and

$$\pi(x^T) \equiv \sum_{t=1}^T \delta^{t-1} (\mathbf{1}_{[x_t=s]} \cdot \Delta - \mathbf{1}_{[x_t \neq \emptyset]} \cdot (K + c))$$

are the discounted profits generated by the risky project at the final history x^T . The payment rule p^{share} fully reimburses the firm for its foregone rents from the safe project and offers a fraction α of the discounted net profit $\pi(x^T)$ generated by the risky project. Note that the payment rule p^{share} is invariant to the belief θ reported by the firm.

We now verify that the mechanism (ϕ^{FB}, p^{share}) is incentive compatible. Because the payoff from the safe project c is commonly known, the firm can only lie about θ . That is, a firm of type $\tau = (\theta, c)$ can only send reports of the form $\hat{\tau} = (\hat{\theta}, c)$. We can therefore use formulas (1) and (2) to obtain

$$\tilde{P}(\hat{\tau}; \theta) - \tilde{\Phi}(\hat{\tau}; \theta) \cdot c = \alpha \cdot \mathbb{E}^{\mu[\phi^{FB}(\hat{\tau})|\theta]} \left[\sum_{t=1}^T \left(\delta^{t-1} \cdot \phi_t^{FB}(\hat{\tau}|\tilde{x}^{t-1}) \cdot (\tilde{\theta}_t(\theta) \cdot \lambda \cdot \Delta - K - c) \right) \right].$$

Notice that the firm’s expected payoff from reporting $\hat{\theta}$ is proportional to the investor’s payoff from following the action $\phi^{FB}(\hat{\theta}, c)$ in the first-best benchmark when the true type is (θ, c) . Therefore, for any $\alpha > 0$, the firm has a strict incentive to pursue the action plan $\phi^{FB}(\theta, c)$ (in which case we say that ϕ^{FB} is strictly implemented). Intuitively, the profit-sharing agreement p^{share} perfectly aligns the preferences of both parties, since the firm becomes the residual claimant of the risky project’s proceeds.

By choosing α close to zero, the investor is able to strictly implement the first-best action plan while obtaining rents arbitrarily close to those obtained under complete information.¹⁶ This discussion is summarized in the following lemma.

Lemma 2 (Symmetric information about the safe project). *When there is symmetric information about the payoff from the safe project, the investor is able to strictly implement the first-best action plan ϕ^{FB} , while enjoying expected net profits arbitrarily close to first-best.*

Similarly to Cr mer and McLean (1988) and Riordan and Sappington (1988) (for short, CM and RS), the construction above implements the first-best leaving no rents to the agent. The ideas behind their constructions are, however, significantly different from ours. In CM and RS, the principal uses a publicly observed signal to devise a lottery that has zero expected value upon truthful reporting, but a negative expected value if the agent lies. In these papers, the possibility of negative payoffs is crucial to discipline the agent to truthfully report his type. By contrast, in our model the principal gives

¹⁶ When the firm’s profit-share is exactly zero, the firm is indifferent between any two reports on the belief θ or, equivalently, between following any two action plans. In this case, full rent extraction is exact, but the first-best action plan is only weakly implemented.

the agent (an arbitrarily small) share of the project’s profit. This guarantees that the agent has the correct incentives to take the right action, while receiving monetary transfers that are close to zero with certainty (provided the agent’s share is low enough).

The reason for why there exists a simpler mechanism that implements the first-best in our setting lies in how the parameter of private information affects the agent’s payoff. In our model, the parameter θ stands for the agent’s prior belief about the quality of the risky project and only affects the agent’s expected payoffs indirectly through the expectation of payments. By appropriately devising the payment rule p^{share} (setting $\alpha \approx 0$), the principal is able to make the impact of θ on the agent’s payoff arbitrarily small, thus capturing (approximately) all informational rents while generating strict incentives for truth-telling. By contrast, in CM and RS, the agent’s parameter of private information describes his preferences over allocations, and it directly enters the agent’s payoff. This feature creates the need for a more complex implementation than the profit-sharing mechanism described above.

3.3. Symmetric information about the risky project

In this last benchmark, we assume that the investor is capable of evaluating the risky project as accurately as the firm but is unable to determine the firm’s opportunity cost of experimentation c . For example, the investor may be a venture capitalist who specializes in the same sector as the firm’s technology, but who is unsure about the firm’s best alternative use of human and physical capital or its cost of raising funding.

As in the previous subsection, the firm’s incentive constraints are one-dimensional since they only require truthful revelation of the payoff from the safe project. Accordingly, we say that a mechanism $\mathcal{M} = (\phi, p)$ satisfies *incentive compatibility on the safe arm (ICS)* if no type (θ, c) benefits from reporting some other type (θ, \hat{c}) :

$$\tilde{P}(\theta, \hat{c}; \theta) - \tilde{\Phi}(\theta, \hat{c}; \theta) \cdot c \leq U(\tau | \mathcal{M}) \tag{ICS}$$

for all $\tau = (\theta, c) \in \Gamma$ and $\hat{c} \in [\underline{c}, \bar{c}]$.

A mechanism \mathcal{M} is *feasible for the expert investor* (or *e-feasible*, for short) if it satisfies individual rationality (IR) and incentive compatibility on the safe arm (ICS). A mechanism \mathcal{M} is *optimal for the expert investor* (or *e-optimal*, for short) if it maximizes the investor’s expected profit (4) within the class of e-feasible mechanisms. The next lemma characterizes the e-optimal mechanism.

Lemma 3 (Symmetric information about the risky project). *The e-optimal mechanism uses a stopping plan ϕ^e described by the threshold functions $v_k^e(\theta)$ that solve the following Virtual Gittins Formula:*

$$V_k(\theta, v_k^e(\theta) + \gamma(v_k^e(\theta))) = 0. \tag{8}$$

A firm with type $\tau = (\theta, c)$ tolerates $k^e(\tau) \in \{1, \dots, T\}$ failures if and only if

$$v_{k^e(\tau)+1}^e(\theta) < c \leq v_{k^e(\tau)}^e(\theta)$$

and receives an expected payment of

$$P^e(\tau) = \Phi_1(\theta) \cdot v_1^e(\theta) + \sum_{k=2}^{k^e(\tau)} (\Phi_k(\theta) - \Phi_{k-1}(\theta)) \cdot v_k^e(\theta), \tag{9}$$

where $\Phi_k(\theta)$ is the expected discounted duration of a project which is good with probability θ , for a type that tolerates k failures:

$$\Phi_k(\theta) = \frac{1 - \delta^k}{1 - \delta} + \mathbf{1}_{[k < T]} \cdot \delta^k \cdot \theta \cdot (1 - (1 - \lambda)^k) \cdot \frac{1 - \delta^{T-k}}{1 - \delta}. \tag{10}$$

The investor must leave informational rents in order to induce the firm to report the payoff from the safe project truthfully. These rents are captured by the reverse hazard rate $\gamma(c)$. The lemma above shows that the expert investor optimum coincides with the first-best action plan, except that the payoff from the safe project c is replaced by its virtual counterpart $c + \gamma(c)$, which incorporates the informational rents paid to the firm. The Virtual Gittins Formula (8) is the virtual counterpart of the Gittins Formula (7): it equates the expected flow payoff from continuing the project at period k to the expected option value from arriving at $k + 1$ with a good project *after adjusting for the firm’s informational rents*.

For each θ , the tolerance for failure $k^e(\theta, c)$, which is decreasing in c , screens firms with different safe projects. Firms with *more valuable safe projects* pick contracts with a *lower tolerance* for failure. Let c^{max} solve

$$\lambda \cdot \Delta - K = c^{max} + \gamma(c^{max}). \tag{11}$$

It follows from (8) that a firm with a payoff from the safe project greater than c^{max} will never play the risky arm, regardless of θ . Because $\bar{c} \geq \lambda \cdot \Delta - K \geq c$, it follows that $c^{max} \in [\underline{c}, \bar{c}]$. We refer to $[\underline{c}, c^{max}]$ as the *relevant range*.

Equation (9) determines the payments needed to implement the action plan ϕ^e . In order to prevent upward deviations regarding the payoff from the safe project, a firm with type $\tau = (\theta, c)$ has to receive the foregone rents that would arise from

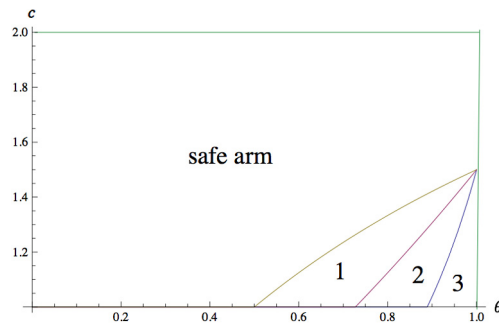


Fig. 2. The e -optimal action plan when $c \sim U[1, 2]$ under the following parameterization: $T = 3$, $\lambda = \frac{1}{2}$, $K = 1$, $\Delta = 6$ and $\delta = 1$. Areas 1, 2 and 3 refer to the set of types that play the risky arm at $t = 1$ and tolerate, respectively, 1, 2 and 3 failures according to ϕ^e .

such an inflated report. These foregone rents take into account the different expected durations associated with different reports (since firms reporting a higher c receive contracts that allow for less experimentation).

Notice that, in the case of an expert investor, optimality pins down the firm's expected payments $P^e(\tau)$, but not the actual payment rule $p(\tau, x^T)$ that induces $P^e(\tau)$. Because there is commitment and the investor observes θ , $P^e(\tau)$ can be achieved by multiple payments rules. For instance, it can be achieved by a lump-sum transfer $P^e(\tau)$ or by more complicated output-contingent payment schemes (with transfers that depend on the realized output history of the risky project). As we will see in the next section, asymmetric information about *both* the safe and the risky project introduces additional constraints on optimal payment rules, as contingent payment schemes are needed to screen both dimensions of the firm's type.

Fig. 2 illustrates the e -optimal action plan when $c \sim U[1, 2]$ under the same parameterization of Fig. 1. Two patterns emerge from comparing Figs. 1 and 2: (i) fewer types play the risky arm in the e -optimal mechanism than in the first best, and (ii) the first-best mechanism tolerates more failure from all types. The next corollary shows that these properties hold beyond the uniform case.

Corollary 1 (*Distortions*). *Relative to the first-best, the e -optimal action plan features:*

1. **Under-experimentation:** $\{\tau : k^{FB}(\tau) = 0\}$ is a proper subset of $\{\tau : k^e(\tau) = 0\}$.
2. **Early termination:** $k^e(\tau) \leq k^{FB}(\tau)$ for all τ such that $k^e(\tau) \geq 1$, with strict inequality for a set of types with positive measure.

Both under-experimentation and early termination are natural consequences of the information asymmetry between the expert investor and the innovative firm: Since the investor needs to leave informational rents to the firm, she is less willing to undertake the risky project than in the first-best benchmark and, when undertaking it, she has a lower the tolerance for failure.

4. Investor-optimal financing mechanism

We now turn to the investor's optimal mechanism when the innovative firm has private information about both the safe and the risky project. The analysis of this more realistic case will help clarify the source of rents from innovation and flesh out important features of the optimal financing contract between investors and innovative firms.

First, we examine how the mechanisms derived in the one-dimensional benchmarks from the previous section perform under two-dimensional private information. Recall that when the investor observes the payoff from the safe project c , she is able to implement the first-best action plan and extract all the firm's rents. This is accomplished by reimbursing the firm's payoff from the safe project c , thus making it indifferent between both projects. This mechanism is no longer incentive compatible when the firm has private information about c , since all firms would then claim to have higher payoffs from the safe project.

More broadly, with two-dimensional private information, full rent extraction is not achievable by any mechanism. The reason is that, when both θ and c are private information, the principal cannot do better than in the expert investor optimum (as the e -optimal problem is a relaxation of the two-dimensional screening problem studied below). Because the e -optimal mechanism generates positive informational rents to the agent, it follows that full rent extraction is not achievable (as it would lead to higher profits for the principal).¹⁷ Specifically, the mechanisms proposed by Crémer and McLean (1988) and Riordan and Sappington (1988) do not work in our setting. The reason is that the principal does not have access to

¹⁷ Yet, echoing results in Mezzetti (2004), the first-best allocation is implementable. This is possible by making the agent the residual claimant of all revenues accrued to the project and charging him K every time he chooses the risky arm. This construction is ruled out in Jehiel and Moldovanu (2001), where payment rules cannot be output-contingent. While achieving the first best, this mechanism leaves large rents to the agent.

any public signal correlated with the agent’s opportunity cost of exploration, what makes the construction of these papers inapplicable.

Since full rent extraction is no longer achievable, it is natural to ask if the investor can still obtain the e -optimal profits with two-dimensional private information. To fix ideas, consider the e -optimal mechanism $\mathcal{M}^e = (\phi^e, \tilde{p}^e)$ according to which a firm of type τ is paid the expected amount $P^e(\tau)$ in a lump-sum fashion (that is, $\tilde{p}^e(\tau, x^T) = P^e(\tau)$ for all final output histories x^T). This mechanism is not incentive compatible under two-dimensional asymmetric information: holding fixed the tolerance for failure, the lump-sum payment $P^e(\tau)$ strictly increases with the reported belief θ . Firms will then inflate their reported beliefs.

This observation does not preclude the possibility that the e -optimal action plan ϕ^e can be implemented by some *output-contingent* payment rule that adjusts the firm’s exposure to the project’s revenue as a function of the reported belief. We explore this possibility in the next subsection.

4.1. Linear contracts and the irrelevance of expertise

We will show that, under appropriate conditions, the e -optimal stopping plan can be implemented by an important class of payment rules: *linear contracts*.

Definition 2 (*Linear contracts*). A payment rule p takes the form of a *menu of linear contracts* $\{a(\tau), b(\tau)\}_{\tau \in \Gamma}$ if for every type $\tau = (\theta, c) \in \Gamma$, the payments received by the innovative firm are given by

$$p(\tau, x^T) = a(\tau) + \sum_{\{t: x_t = s\}} \delta^{t-1} \cdot b(\tau).$$

That is, payments to the firm are linear in the number of successes. We refer to $a(\tau)$ as the lump-sum payment and to $b(\tau)$ as the success bonus.

The main result of this section identifies a necessary and sufficient condition for the e -optimal action plan to be implementable with two-dimensional private information. Under this condition, the investor can elicit information about the quality of the project at no additional cost (relative to the expert investor benchmark). To fix terminology, consider the following definition:

Definition 3 (*Irrelevance of expertise*). We say that expertise (about the quality of the risky project) is irrelevant if the action plan ϕ^e associated with the e -optimal mechanism is implementable under asymmetric information over the two-dimensional type (θ, c) .

Before stating our main result, we introduce the following condition, which is satisfied by many standard distributions, such as the uniform, the mirrored generalized Pareto, and the beta distributions (for a wide range of parameters):

Condition C. The reverse hazard rate $\gamma(c)$ is weakly concave on the relevant cost range $[\underline{c}, c^{max}]$.

The next proposition establishes that **Condition C** is both necessary and sufficient for expertise to be irrelevant. This is, therefore, the “minimal” condition that makes screening information about the risky arm “compatible” with the e -optimal solution. Moreover, under **Condition C**, the stopping plan ϕ^e can be implemented by a menu of linear contracts.

Proposition 1 (*Optimal mechanism*). Expertise is irrelevant if and only if **Condition C** holds. In this case, the optimal action plan coincides with the e -optimal action plan: $\phi^* = \phi^e$. Moreover, ϕ^* is implementable by a menu of linear contracts p^* , with lump-sum payments and success bonuses given by

$$a^*(\tau) = P^e(\tau) - b^*(\tau) \cdot \Lambda_{k^e(\tau)}(\theta) \quad \text{and} \quad b^*(\tau) = \frac{\frac{\partial}{\partial \theta} [P^e(\tau)]}{\frac{d}{d\theta} [\Lambda_{k^e(\tau)}(\theta)]}, \tag{12}$$

where $\Lambda_k(\theta)$ is the expected discounted number of successes of a firm which tolerates exactly k failures with a project which is good with probability θ :

$$\Lambda_k(\theta) = \theta \cdot \lambda \cdot \frac{1 - \delta^k}{1 - \delta} + \mathbf{1}_{[k < T]} \cdot \delta^k \cdot \theta \cdot \lambda \cdot \left(1 - (1 - \lambda)^k\right) \cdot \frac{1 - \delta^{T-k}}{1 - \delta}. \tag{13}$$

Before presenting its formal proof, it is useful to intuitively describe the main ideas behind **Proposition 1**.

Let us investigate if and how one can implement the e -optimal action plan ϕ^e by a menu of linear contracts $p = \{a(\tau), b(\tau)\}_{\tau \in \Gamma}$. **Lemma 3** implies that, whenever the mechanism (ϕ^e, p) is feasible, the expected payments received by a

firm of type τ have to equal $P^e(\tau) + \Psi$, for some nonnegative constant Ψ . Therefore, if the menu of linear contracts p implements ϕ^e , the expected payoff of type $\tau = (\theta, c)$ is

$$P^e(\tau) + \Psi - \Phi_{k^e(\tau)}(\theta) \cdot c = a(\tau) + b(\tau) \cdot \Lambda_{k^e(\tau)}(\theta) - \Phi_{k^e(\tau)}(\theta) \cdot c, \quad (14)$$

where the equality follows from the linearity of p .

Equation (14) shows that the firm's payoff under p satisfies two important monotonicity properties. First, there are decreasing differences in the contract's expected duration and the payoff from the safe arm, so that firms with higher payoffs from the safe project are more tempted to pick contracts with shorter durations. This is why incentive compatibility requires the number of failures tolerated by the stopping plan to be weakly decreasing in the safe project payoff. The fact that the reverse hazard rate $\gamma(c)$ is weakly increasing in c , as usually assumed in mechanism design, guarantees that this monotonicity requirement is slack at the e -optimal mechanism.

Second, there are increasing differences in the success bonus and the belief that the risky project is good: firms with better risky projects θ prefer contracts with more exposure to the revenue from the project. Thus, any linear contract p that implements the e -optimal action plan ϕ^e must offer a higher bonus to projects that are more likely to be of good quality. As we will see next, this second monotonicity requirement is slack at the optimum if and only if the reverse hazard rate $\gamma(c)$ is weakly concave in c . Although not standard in one-dimensional screening environments, this condition is essential to the two-dimensional screening problem studied in this paper, as it determines how informational rents change *at the margin*.¹⁸

To understand why marginal information rents are important in our context, differentiate both sides of (14) with respect to θ . Using the envelope theorem on the right-side of (14) and manipulating, we obtain:

$$\frac{\partial}{\partial \theta} [P^e(\tau)] = b^*(\tau) \cdot \frac{d}{d\theta} [\Lambda_{k^e(\tau)}(\theta)].$$

This equation shows how the success bonus $b^*(\tau)$ (i.e., the firm's exposure to the risky project's revenue) varies with the belief that the risky project is good. Solving for $b^*(\tau)$ gives the expression in the statement of Proposition 1. Because $\Lambda_{k^e(\tau)}(\theta)$ is piecewise linear in θ , the bonus $b^*(\tau)$ is weakly increasing in θ if and only if the expected payment $P^e(\tau)$ is convex in θ .

The marginal informational rents determine whether expected payments $P^e(\tau)$ are convex in θ . Holding fixed the maximal number of failures before termination, the *marginal* firm to be financed at each θ is the one with the highest payoff from the safe project c , which, by Condition C, is the one with the lowest marginal impact on informational rents among all firms that are financed. Therefore, as θ increases, the expert investor is willing to finance a positive and increasing number of firms, resulting in expected payments $P^e(\tau)$ that are convex in θ .

By the same reasoning it follows that, if Condition C does not hold, it is possible to find a type $\tau = (\theta, c)$ where $P^e(\tau)$ is locally strictly concave in θ . In this case, the success bonus in (12) is strictly decreasing in θ in a neighborhood of τ , making the e -optimal stopping plan not implementable by a menu of linear contracts. That is, Condition C is not only sufficient but also necessary for the success bonus in (12) to be weakly increasing in θ .

This heuristic argument considers only a subset of the incentive constraints that need to be satisfied for the e -optimal action plan to be implementable under two-dimensional asymmetric information. Namely, it only considers deviations along each dimension, where the innovative firm misreports either θ or c but not both. The proof of Proposition 1 formalizes this heuristic argument and verifies that the remaining incentive constraints are satisfied, ruling out two-dimensional deviations where the firm simultaneously misreports both θ and c . It also establishes that the investor cannot gain by using non-linear payment rules.

Notice that, even when the e -optimal mechanism can be implemented, private information about the quality of the risky project places additional constraints on the optimal mechanism. Recall that with symmetric information about the quality of the risky project, the optimal mechanism could be implemented with lump-sum transfers. With two-dimensional private information, however, because the expected payment $P^e(\tau)$ must be increasing and convex in θ , payments must depend on the revenue from the project. Therefore, although under Condition C the entrepreneur's expertise is irrelevant for payoffs, the interaction of both dimensions of private information is a key determinant of the optimal contractual form.

Before presenting the proof, we consider an example illustrating the results of Proposition 1 in closed form. The example also illustrates a remarkable robustness property of the optimum: the menu $(a^*(\tau), b^*(\tau))$ is invariant to the distribution of beliefs, $G(\theta)$.

¹⁸ A similar condition plays an important role in the work of Beaudry et al. (2009). Their paper studies optimal redistribution in an environment in which workers divide their time between working on the formal and informal sectors. The government observes the worker's productivity and the time allocated to the formal sector. Each worker can imitate those with lower productivities by shirking, but they cannot fake being more productive. A reverse hazard rate condition similar to Condition C is used to control for the effect of correlation between the productivities in the two sectors. This condition is not needed in the model of Beaudry et al. when the productivities are independently distributed. Because we cannot rule out both upward and downward deviations, incentive compatibility requires tighter conditions in our model (even under independence).

Example 1 (Optimal mechanisms with mirrored generalized Pareto distributions). Let the payoff from the safe project be distributed according to

$$H(c) = \left(\frac{c - \underline{c}}{\bar{c} - \underline{c}} \right)^\eta, \quad \text{where } \eta > 0.$$

For this distribution, $\gamma(c) = \frac{c - \underline{c}}{\eta}$, **Condition C** holds, and the firm’s expertise is irrelevant. The optimal action plan is implementable by a menu of linear contracts with lump-sum payments and success bonuses:

$$a^*(\tau) = \frac{\Phi_{k^*(\tau)}(\theta)}{\eta + 1} \cdot (\underline{c} - \eta \cdot K) \quad \text{and} \quad b^*(\tau) = \Delta \cdot \frac{\eta}{\eta + 1}.$$

The optimal payment rule then features a menu of linear contracts with identical success bonuses and lump-sum payments that increase as the innovative firm agrees to tolerate more failures before terminating the project.

Notice that parameter η orders the distributions of the safe project payoff in terms of first-order stochastic dominance (that is, larger values of η are associated with higher safe project payoffs). The optimal success bonus offered to all types is increasing in η . As $\eta \rightarrow 0$, the distribution H approaches a mass-point distribution at the lower limit \underline{c} . In this case, the investor gets the entire cash flow from the risky project and reimburses the firm for the foregone rents from the safe project. When costs are uniformly distributed ($\eta = 1$), the optimal contract features an equal sharing rule: $\beta^*(\tau) = \frac{\Delta}{2}$. Finally, as $\eta \rightarrow +\infty$, the distribution H approaches a mass-point distribution at the upper limit \bar{c} . In this case, the firm gets the entire cash flow from the project and compensates the investor for the per-period outlay K .

Remark 1 (Limited liability and equity-based contracts). Although in reality firms are protected by limited liability, we abstracted from it in the analysis above. When the firm has no pledgeable income, limited liability requires that payments cannot fall below zero for any output history:

$$p(\tau, x^T) \geq 0 \quad \forall \tau \in \Gamma \quad \text{and} \quad \forall x^T \in X^T. \tag{LL}$$

Whenever this condition is satisfied, we can construct equity-based contracts specifying (i) a lump-sum transfer (*golden handshake*) $\alpha(\tau) \geq 0$ paid by the investor to the innovative firm; and (ii) a *revenue (or equity) share* $\beta(\tau) \in (0, 1)$ that determines how the investor and the firm share revenues accrued from the risky project. Equity-based contracts are remarkably simple and bear close resemblance to the actual practice of venture capital financing (see, for example, [Kaplan and Stromberg, 2003](#)).

In the working paper version of this article, we show that the optimal menu $(a^*(\tau), b^*(\tau))$ from [Proposition 1](#) satisfies (LL) if and only if $\frac{P^e(\theta, c)}{\theta}$ is weakly decreasing in θ for every c in the relevant range $[\underline{c}, c^{max}]$. Although not directly expressed in terms of primitives, this monotonicity condition can be easily verified. For instance, in the setting of [Example 1](#), it holds if and only if $\underline{c} \geq \eta \cdot K$. In this case, the optimal action plan is implementable by a menu of equity-based contracts with equity shares $\beta^*(\tau) = \frac{\eta}{\eta + 1}$ and golden handshake $\alpha^*(\tau) = a^*(\tau)$.

4.2. Proof of [Proposition 1](#)

We now prove [Proposition 1](#), formalizing the ideas from the heuristic derivation above. Readers interested in the less technical aspects of our analysis may prefer to skip this subsection and proceed to [Subsection 4.3](#), where we investigate optimal financing mechanisms when **Condition C** does not hold.

Proof of Proposition 1. First note that the expert investor problem is a relaxation of the investor problem when there is two-dimensional asymmetric information. As a consequence, the optimal action plan ϕ^* coincides with the e -optimal action plan ϕ^e provided that the latter is implementable.

We structure this proof in two steps. In the first step, we will show that there exists a menu of linear contracts that implements ϕ^e if and only if **Condition C** holds. In the second step, we will show that if there is no menu of linear contracts that implements ϕ^e , then ϕ^e is not implementable.

Step 1. There exists a menu of linear contracts that implements ϕ^e if and only if **Condition C** holds. Moreover, under this condition, there exists a unique such menu that leads to expected payments $P^e(\tau)$.

By ([ICS](#)), any payment rule that implements ϕ^e leads to expected payments $P^e(\tau) + \Psi$, where Ψ is a weakly positive constant. By defining the lump-sum payment $a^*(\tau)$ according to $a^*(\tau) = P^e(\tau) - b^*(\tau) \cdot \Lambda_{k^e(\tau)}(\theta)$, we assure by construction that

$$\begin{aligned} P^*(\tau) &\equiv \mathbb{E}^{\mu[\phi^*(\tau)|\theta]} \left[a^*(\tau) + \sum_{\{t: X_t = s\}} \delta^{t-1} \cdot b^*(\tau) \right] \\ &= \mathbb{E}^{\mu[\phi^e(\tau)|\theta]} \left[a^*(\tau) + \sum_{\{t: X_t = s\}} \delta^{t-1} \cdot b^*(\tau) \right] \\ &= P^e(\tau) - b^*(\tau) \cdot \Lambda_{k^e(\tau)}(\theta) + b^*(\tau) \cdot \Lambda_{k^e(\tau)}(\theta) = P^e(\tau), \end{aligned}$$

where the equality in the second line follows from $\phi^* = \phi^e$.

We will now derive the unique success bonus $b^*(\tau)$ that, together with $a^*(\tau)$, implements the action plan ϕ^e . To this end, consider a mechanism $\mathcal{M} = (\phi, p)$ such that ϕ is a stopping plan. Consider the following incentive constraints regarding non-truthful reporting about the risky arm (denoting $\tau = (\theta, c)$ and $\hat{\tau} = (\hat{\theta}, \hat{c})$):

$$\tilde{P}(\hat{\tau}; \theta) - \Phi_{k(\hat{\tau})}(\theta) \cdot c \leq U(\tau|\mathcal{M}) \quad \forall \tau, \hat{\tau} \in \Gamma \text{ with } \theta \neq \hat{\theta} \text{ and } k(\hat{\tau}) = k(\tau), \tag{ICR_1}$$

and

$$\tilde{P}(\hat{\tau}; \theta) - \Phi_{k(\hat{\tau})}(\theta) \cdot c \leq U(\tau|\mathcal{M}) \quad \forall \tau, \hat{\tau} \in \Gamma \text{ with } \theta \neq \hat{\theta} \text{ and } k(\hat{\tau}) \neq k(\tau). \tag{ICR_2}$$

(ICR₁) rules out non-truthful reports on θ (and potentially on c as well) such that $k(\hat{\tau}) = k(\tau)$ (i.e., non-truthful reports that retain the same tolerance to failure as in the truthful report). In turn, (ICR₂) rules out non-truthful reports on θ (and potentially on c as well) such that $k(\hat{\tau}) \neq k(\tau)$ (i.e., non-truthful reports that change the tolerance to failure). Clearly, for a stopping plan ϕ , the mechanism $\mathcal{M} = (\phi, p)$ satisfies condition (IC) if and only if it satisfies conditions (ICS), (ICR₁) and (ICR₂).

By construction, $\phi^* = \phi^e$, and, by the choice of $a^*(\tau)$, $P^*(\tau) = P^e(\tau)$. Lemma 3 then implies that the optimal mechanism $\mathcal{M}^* = (\phi^*, p^*)$ satisfies condition (ICS), in which case the firm’s payoff is $U(\theta, c|\mathcal{M}^*) = P^*(\tau) - \Phi_{k^e(\tau)}(\theta) \cdot c = P^e(\tau) - \Phi_{k^e(\tau)}(\theta) \cdot c$. Note that $k^e(\tau)$ has zero partial derivative with respect to θ almost everywhere (and is not differentiable along the path $(\theta, v_{k^e(\tau)}^e(\theta))$). Therefore,

$$\frac{\partial U(\theta, c|\mathcal{M}^*)}{\partial \theta} = \frac{\partial}{\partial \theta} [P^e(\tau)] - \frac{d}{d\theta} \{ \Phi_{k^e(\tau)}(\theta) \} \cdot c \quad \text{almost everywhere.} \tag{15}$$

We will now turn to condition (ICR₁). Observe that the firm’s payoff satisfies strict increasing differences in the success bonus and the reported belief. Building on this observation, we can use standard techniques to characterize the set of mechanisms that satisfy (ICR₁). In particular, the next lemma delivers an alternative envelope formula that depends explicitly on the schedule of bonuses $b(\tau)$. Coupled with (15), this new envelope formula pins down the unique candidate schedule of bonuses that implement $\phi^* = \phi^e$.

Lemma 4. Consider a mechanism $\mathcal{M} = (\phi, p)$, where ϕ is a stopping plan and $p = \{a(\tau), b(\tau)\}_{\tau \in \Gamma}$ is a menu of linear contracts. The mechanism \mathcal{M} satisfies condition (ICR₁) if and only if the following conditions jointly hold:

1. For every $\hat{k} \in \{1, \dots, T\}$, the success bonus $b(\tau)$ restricted to support $\{\tau : k(\tau) = \hat{k}\}$ is weakly increasing in θ ,
2. The envelope formula holds almost everywhere:

$$\frac{\partial U(\theta, c|\mathcal{M})}{\partial \theta} = b(\tau) \cdot \frac{d}{d\theta} [\Lambda_{k(\tau)}(\theta)] - \frac{d}{d\theta} \{ \Phi_{k(\tau)}(\theta) \} \cdot c. \tag{16}$$

Proof of Lemma 4. See appendix. □

Setting the envelope formula (16) evaluated at \mathcal{M}^* equal to equation (15) pins down the optimal bonus $b^*(\tau)$ described in the statement of the current proposition. To conclude that the optimal mechanism \mathcal{M}^* satisfies condition (ICR₁), it only remains to be shown that the bonus $b^*(\tau)$ satisfies the monotonicity condition 1 from Lemma 4. The next lemma establishes that Condition C is both a necessary and sufficient condition for the latter property to hold.

Lemma 5. The following statements are equivalent:

1. For any $\hat{k} \in \{1, \dots, T\}$, the bonus $b^*(\tau)$ restricted to the support $\{\tau : k^e(\tau) = \hat{k}\}$ is weakly increasing in θ .
2. For any $\hat{k} \in \{1, \dots, T\}$, $P^e(\tau)$ restricted to the support $\{\tau : k^e(\tau) = \hat{k}\}$ is weakly convex in θ .
3. Condition C holds.

Proof of Lemma 5. Formula (13) determines the expected discounted number of successes of a firm that tolerates exactly k failures with a project that is good with probability θ , $\Lambda_k(\theta)$. Note that, for any k , $\Lambda_k(\theta)$ is a linear function of θ , in which case $\frac{d}{d\theta} [\Lambda_k(\theta)]$ is a constant. Therefore, by the definition of $b^*(\tau)$ in equation (12), it follows that conditions 1 and 2 are equivalent.

We will now establish the equivalence between conditions 2 and 3. To clarify ideas, let us start with the simpler case where $T = 1$. By equation (9), the payments received by an active type (i.e., $k^*(\tau) = 1$) equal $P^*(\tau) = P^e(\tau) = v_1^e(\theta)$. As a consequence, $P^e(\tau)$ is weakly convex in θ if and only if $v_1^e(\theta)$ is weakly convex. In turn, the Virtual Gittins Formula (8) implies that $v_1^e(\theta)$ is given by

$$\theta \cdot \lambda \cdot \Delta - K = v_1^e(\theta) + \gamma (v_1^e(\theta)). \tag{17}$$

Because the left-hand side of the equation above is linear, it follows that $v_1^e(\theta)$ is weakly convex if and only if $\gamma(c)$ is weakly concave in every $c \in [\underline{c}, c^{max}]$.¹⁹

Now consider $T > 1$ and fix some $\hat{k} \in \{1, \dots, T\}$. Analogously to the case where $T = 1$, weak concavity of $\gamma(c)$ in every $c \in [\underline{c}, c^{max}]$ is both a necessary and sufficient condition for $P^e(\tau)$ to be weakly convex in θ . To establish this claim, differentiate $P^e(\tau)$ with respect to θ and employ the Virtual Gittins Formula (8). The details of this computation can be found in the appendix. \square

In light of Lemmas 4 and 5, it follows that that mechanism $\mathcal{M}^* = (\phi^*, p^*)$ satisfies condition (ICR₁) if and only if Condition C holds. In order to complete Step 1, it only remains to be shown that \mathcal{M}^* satisfies condition (ICR₂). To see why this is true, fix some type $\tau = (\theta, c)$ that tolerates $k^*(\tau)$ failures, and consider his payoff from misreporting some type $\hat{\tau} = (\hat{\theta}, \hat{c})$ that tolerates a different number of failures $k^*(\hat{\tau}) \neq k^*(\tau)$. It follows from the thresholds in the Virtual Gittins Formula (8) that for any type $\tau = (\theta, c)$ either (i) there exists some \tilde{c} such that $k^*(\hat{\tau}) = k^*(\theta, \tilde{c})$, or (ii) there exists some $\tilde{\theta}$ such that $k^*(\hat{\tau}) = k^*(\tilde{\theta}, c)$ (or both). Assume for a moment that (i) is true. Then

$$\begin{aligned} U(\theta, c | \mathcal{M}^*) &\geq \tilde{P}^*(\theta, \tilde{c}; \theta) - \Phi_{k^*(\hat{\tau})}(\theta) \cdot c \\ &= P^*(\theta, \tilde{c}) - \Phi_{k^*(\hat{\tau})}(\theta) \cdot c \\ &\geq \tilde{P}^*(\hat{\theta}, \tilde{c}; \theta) - \Phi_{k^*(\hat{\tau})}(\theta) \cdot c \\ &= \tilde{P}^*(\hat{\tau}; \theta) - \Phi_{k^*(\hat{\tau})}(\theta) \cdot c, \end{aligned}$$

where the inequality in the first line follows from (ICS), the equality from the first to the second line follows from the invariance of $P^*(\cdot)$ with respect to c within the support $\{\check{\tau} : k^*(\hat{\tau}) = k^*(\check{\tau})\}$, the inequality from the second to the third line follows from (ICR₁), and the equality from the third to the fourth line follows from the invariance of $P^*(\cdot)$ with respect to c within the support $\{\check{\tau} : k^*(\hat{\tau}) = k^*(\check{\tau})\}$. The case where (ii) is true is analogous (we just invoke (ICS) and (ICR₁) in the opposite order). This completes Step 1.

Step 2. If there is no menu of linear contracts that implements ϕ^e , then ϕ^e is not implementable.

The proof of this claim establishes that for any payment rule p , there exists a menu of linear contracts that lead to the same expected payments as p . The details appear in the appendix. \square

4.3. Pooling when expertise is relevant

On theoretical grounds, Proposition 1 helps assessing the robustness of payment schemes that lead to full rent extraction to environments with multi-dimensional private information. In these environments, full rent extraction in one dimension of private information (in our case, the belief θ) is only possible if the informational rents of the other dimension (in our case, the payoff c) have the “right” behavior at the margin. Here, because the firm’s payoff has increasing differences in (b, θ) , expected payments $P^e(\theta, c)$ have to be convex in θ . Whether this requirement is slack depends on the curvature (rather than on the slope) of the reverse hazard rate, as described by Condition C.

When Condition C fails, the e -optimal mechanism is no longer feasible, and the interaction of both dimensions of private information impacts the investor’s payoff (relative to the expert investor benchmark, where private information is one-dimensional). In this case, obtaining a complete characterization of the optimal mechanism becomes a very complex task. First, it is hard to rule out the optimality of stochastic mechanisms.²⁰ It is also difficult to rule out complex action plans in which experimentation stops for some periods and restarts later.

We take a more modest approach and focus on optimal deterministic stopping-plan mechanisms. Beyond tractability, we restrict attention to deterministic mechanisms because stochastic mechanisms are rarely seen in practice. In turn, stopping plans arise naturally when it is too costly to restart the experimentation process after it has been stopped. Indeed, for most firms, shutting down is a very expensive – if not entirely irreversible – decision.

To simplify matters, we consider strictly convex reverse hazard rates, which are the “polar-opposite” of Condition C. Then, pooling is a general property in the sense that, in every optimal deterministic mechanism, there is a positive measure of innovative firms which choose the same contract²¹:

Proposition 2 (Pooling). *Let the reverse hazard rate $\gamma(c)$ be strictly convex. Then, in every optimal deterministic mechanism, there is a positive measure of firms which choose the same contract.*

¹⁹ Indeed, because $\gamma(c)$ is weakly increasing, it follows that $v_1^e(\theta) = \sigma^{-1}(\theta \cdot \lambda \cdot \Delta - K)$ where the function $\sigma(c)$ is defined by $\sigma(c) \equiv c + \gamma(c)$. Note that $\sigma(c)$ is strictly increasing, and is weakly concave if and only if Condition C holds. This implies that Condition C is necessary and sufficient for $v_1^e(\theta)$ to be convex.

²⁰ On this point, see Strausz (2006).

²¹ In the online appendix we show that an optimal mechanism exists.

For simplicity, we illustrate the result from Proposition 2 with an example where $\lambda = 1$, so that one period of experimentation is enough to determine the quality of the risky project. The proof of the general case is presented in the appendix.

Example 2 (Full pooling). Suppose the reverse hazard rate $\gamma(c)$ is strictly convex and let $\lambda = 1$. Since one round of experimentation perfectly distinguishes good a bad projects in this case, a risky project that succeeds in the first period is undertaken in every future period; a project that fails is abandoned forever.

For a given θ , the expert investor lets all types with payoffs from the safe project $c \leq v^e(\theta)$ experiment. When θ is not observable, the investor offers a menu of contracts to elicit the quality of the risky project. Without loss of generality, we can focus on a menu of linear contracts. When a firm with type $\tau = (\theta, c)$ accepts the linear contract $(a(\tau), b(\tau))$, it gets payoff

$$a(\tau) + \theta b(\tau) - c + \theta \left(\sum_{t \geq 1} \delta^t \cdot (b(\tau) - c) \right).$$

Let $v(\theta)$ denote the frontier of types that experiment, so that (θ, c) tries the risky project whenever $c \leq v(\theta)$. Then, for $\underline{c} < v(\theta) < \bar{c}$, the frontier is given by the indifference condition:

$$a(\tau) + \theta b(\tau) - v(\theta) + \theta \left(\sum_{t \geq 1} \delta^t \cdot (b(\tau) - v(\theta)) \right) = 0.$$

Because experimentation lasts for at most one period, the investor cannot use the length of experimentation to screen the payoff from the safe project c among firms that experiment. Therefore, we abuse notation and write $(a(\theta, c), b(\theta, c)) = (a(\theta), b(\theta))$ for the contract offered to all firms with belief θ who engage in experimentation.

Following a standard argument from mechanism design (used, for example, in the proof of the revenue equivalence result in auction theory), the investor’s payoff depends only on the payoff of types with the lowest beliefs (which is determined by the participation constraint) and the experimentation cutoff $v(\theta)$. Since the investor’s payoff is a quasi-concave function of the experimentation cutoff $v(\theta)$ for a given θ , the investor would like to choose a contract that sets $v(\theta)$ as close to $v^e(\theta)$ as possible.

By the envelope theorem,

$$b(\theta) = \frac{d}{d\tilde{\theta}} \left(\frac{v(\tilde{\theta}) + \tilde{\theta} \sum_{t \geq 1} \delta^t \cdot v(\tilde{\theta})}{\sum_{t \geq 0} \delta^t} \right)_{\tilde{\theta}=\theta}, \tag{18}$$

for almost all θ for which $v(\theta) \in (\underline{c}, \bar{c})$. Moreover incentive compatibility requires $b(\theta)$ to be increasing. Thus, the mechanism will satisfy incentive compatibility if $\xi^v(\theta) = \left(\frac{v(\theta) + \theta \sum_{t \geq 1} \delta^t \cdot v(\theta)}{\sum_{t \geq 0} \delta^t} \right)$ is a convex function of θ .

When the reverse hazard rate is strictly convex, the experimentation curve for an expert investor $\xi^{v^e}(\theta)$ is a strictly concave function of θ . Therefore, the solution of the relaxed program (which ignores the convexity constraint) is strictly concave, whereas incentive compatibility requires it to be convex. Then, as in Guesnerie and Laffont (1984), the optimal contract is “non-responsive” (i.e., it has complete pooling).

To understand the intuition for this non-responsiveness result, first notice that, by equation (16), intervals of pooling correspond to intervals where ξ^v is affine. Assume towards a contradiction that there is an interval (θ_1, θ_2) in which $\xi^v(\theta_1) = \xi^{v^e}(\theta_1)$ and $\xi^v(\theta_2) = \xi^{v^e}(\theta_2)$ and assume that the restriction of ξ^v to this interval is not affine. As the left-side panel of Fig. 3 shows, the slope of ξ^v is lower than the slope of ξ^{v^e} for θ close to θ_1 and higher for θ close to θ_2 . Therefore, the investor could obtain higher profits by offering types $\theta \in (\theta_1, \theta_2)$ the experimentation curve $\xi^{\bar{v}}$ defined by the line segment passing through $(\theta_1, \xi^v(\theta_1))$ and $(\theta_2, \xi^v(\theta_2))$. In fact, the investor could induce this curve by offering to every type $\theta \in (\theta_1, \theta_2)$ the same contract (\bar{a}, \bar{b}) , making types $(\theta_i, v(\theta_i))$ indifferent between $(a(\theta_i), b(\theta_i))$ and (\bar{a}, \bar{b}) , where $i = 1, 2$. Therefore, there must be pooling in the interval (θ_1, θ_2) . Applying a similar argument to the intervals $(0, \theta_1)$ and $(\theta_2, 1)$ shows that all types are pooled, as illustrated in the right-side panel of Fig. 3.

The right-side panel of Fig. 3 plots the expected payments (as a function of θ) of all types that play the risky arm in the example above. Firms with projects of either low or high qualities get positive informational rents from their private information about θ . More precisely, they obtain higher rents when θ is their private information than when it is also observed by the investor. However, firms with medium-quality projects (i.e., those in the interval $(\hat{\theta}_1, \hat{\theta}_2)$) obtain *negative informational rents* from their private information about θ , since their payoffs are higher in the expert investor solution. These firms would benefit if they were able to provide hard evidence regarding the quality of their project. Thus, unlike in one-dimensional screening models, private information can make the agent strictly worse (even when the dimensions of private information are uncorrelated, as c and θ are here). This is due to the multidimensionality of types – with one dimensional types, agents are always weakly better off when information is asymmetric.




Fig. 3. The left-side panel illustrates why pooling occurs when the reverse hazard rate $\gamma(c)$ is strictly convex. It plots the experimentation curve of the expert investor problem, $\xi^{v^e}(\theta)$ (dotted), and a candidate experimentation curve under two-dimensional asymmetric information, $\xi^v(\theta)$ (full). The right-side panel plots the expected payments (as a function of θ) on the expert investor optimum, $P^e(\theta)$ (dotted), and on the optimal deterministic mechanism, $P^*(\theta)$ (full).

When experimentation lasts for more than one period, it is more complex to verify the monotonicity constraint, as it would require the bonus of each stopping plan to be increasing. One can construct examples in which a continuum of different contracts are offered. Nevertheless, as established in Proposition 2, when the reverse hazard rate $\gamma(c)$ is strictly convex, there is always some pooling.

5. Conclusion

In this paper, we study optimal contracting between an investor and an innovative firm. Our model embeds a geometric two-arm bandit problem into a principal-agent framework. The agent has private information about his expected payoff from exploitation, as well as about the value of exploration. We construct a two-step procedure to solve for the investor’s profit-maximizing mechanism: In the first stage, we consider a relaxed (one-dimensional screening) problem in which private information is confined to the firm’s exploitation payoff. Having solved this relaxed program, we derive a necessary and sufficient condition for the solution of this program to be implementable in the setting with two-dimensional private information.

Our analysis delivers four main insights. First, we identify the contractual features that optimally screen the firm’s private information. The payoffs from exploitation are screened by financing contracts of different durations. In turn, the prospects from exploration are screened by different exposures to the risky project’s proceeds (e.g., in the form of revenue shares or equity).

Second, we show that asymmetric information causes fewer firms to experiment and, the ones who do, experiment for fewer periods than they would under symmetric information.

Third, we derive a necessary and sufficient condition under which the firm’s expertise about its exploration prospects generates no informational rents. This condition controls for how the informational rents from exploitation vary at the margin, implying restrictions on the curvature of the appropriate reverse hazard rate (rather than on its slope, as usual in one-dimensional screening problems).

Fourth, we show that when this condition fails, screening the firm’s payoff from exploitation is incompatible with screening its explorations prospects. As a result, the investor has to offer the same contract to a positive-measure set of firms. In some cases, the incompatibility of screening each dimension of private information is so extreme that the optimal mechanism offers the same contract to all types. In those cases, global incentive constraints imply that some firms obtain *negative* informational rents from their expertise on the quality of risky project. These firms would be better off if the investor knew the prospect of the risky project, or if they could provide hard evidence about it.

Our results have the following implications regarding the firms’ incentives to invest in R&D. When expertise on exploration is irrelevant for payoffs, only firms with pockets deep enough to self-finance their experimentation would invest in R&D, since the rents from their expertise are easily appropriable by investors. Moreover, if experimentation is too costly to be self-financed, firms might gain from reducing the informational asymmetries on the quality of their projects. This can be accomplished by R&D partnerships or early-stage (small-scale) market experiments monitored by potential investors.²²

Finally, notice that our results rely on the investor’s ability to observe whether the entrepreneur is experimenting. To see this, consider the case in which limited liability does not bind. In this case, the entrepreneur with the highest cost who experiments obtains a zero payoff. Hence, he would prefer to take the fixed payment and choose the safe arm, obtaining a positive payoff. Extending our model to the case in which the entrepreneur’s action is unobservable (in the spirit of Manso, 2011) is an interesting and challenging extension, which we leave for future work.

²² The management literature has long recognized the detrimental effects of informational asymmetries for the financing of innovation. See, for example, Ries (2011).

Appendix A. Omitted proofs

The proofs omitted in the main text are presented below.

Proof of Lemma 3. The expert investor chooses a mechanism $\mathcal{M} = (\phi, p)$ to maximize expected profits (4) subject to incentive compatibility on the safe arm (ICS). Note that the firm’s payoff (3) satisfies strictly decreasing differences in (Φ, c) , where Φ is the expected duration of the action plan. Therefore, we can use standard arguments to prove that:

Lemma 6. (ICS) A mechanism $\mathcal{M} = (\phi, p)$ satisfies incentive compatibility on the safe arm (ICS) if and only if the following conditions jointly hold:

1. For every $\theta \in [0, 1]$, the expected discounted duration of the project, $\Phi(\theta, c)$, is weakly decreasing in c ,
2. For every $\theta \in [0, 1]$, the expected payment of type $\tau = (\theta, c)$, $P(\tau)$, satisfies the envelope formula:

$$P(\tau) = \Phi(\tau) \cdot c + \int_c^{\hat{c}} \Phi(\theta, \hat{c}) d\hat{c}. \tag{19}$$

Plugging the envelope formula above into the objective function (4) leads to:

$$\mathbb{E}^c \left\{ \mathbb{E}^{\mu[\phi(\tau)|\theta]} \left[\sum_{t=1}^T \left(\delta^{t-1} \cdot \phi_t(\tau, \tilde{x}^{t-1}) \cdot \left(\tilde{\theta}_t \cdot \Delta - K - c - \gamma(c) \right) \right) \right] \right\}. \tag{20}$$

For now, ignore the monotonicity constraint on $\Phi(\tau)$ and consider the relaxed problem \mathcal{P}^e of choosing an action plan ϕ to maximize (20) point-wise in c :

$$\mathcal{P}^e : \quad \max_{\phi} \quad \mathbb{E}^{\mu[\phi(\tau)|\theta]} \left[\sum_{t=1}^T \left(\delta^{t-1} \cdot \phi_t(\tau, \tilde{x}^{t-1}) \cdot \left(\tilde{\theta}_t \cdot \Delta - K - c - \gamma(c) \right) \right) \right].$$

The objective function above is similar to the first-best objective function; the only difference being that the safe project payoff c is replaced by its virtual counterpart, $c + \gamma(c)$. We can therefore apply Lemma 1 to conclude that the solution of problem \mathcal{P}^e is the stopping plan ϕ^e according to which an innovative firm with type $\tau = (\theta, c)$ tolerates exactly $k \in \{1, \dots, T\}$ failures if and only if $v_{k+1}^e(\theta) < c \leq v_k^e(\theta)$, where the functions $v_k^e(\theta)$ are determined by equation (8).

Consider now the types $\tau = (\theta, c)$ and $\hat{\tau} = (\theta, \hat{c})$ such that $\hat{c} > c$. Because H is log concave, $\gamma(c)$ is increasing in c . Therefore, under ϕ^e , an innovative firm of type τ tolerates a weakly greater number of failures than a firm of type $\hat{\tau}$. This implies that the expected discounted duration $\Phi(\tau)$ is weakly decreasing in c for every $\theta \in [0, 1]$. We can therefore conclude that the solution of the expert investor’s problem employs the action plan ϕ^e . Evaluating the envelope formula (19) at the action plan ϕ^e leads to (9). \square

Proof of Corollary 1. The function $V_k(\theta, c)$, defined in (6), is strictly decreasing in c for every $k \in \{1, \dots, K\}$. Therefore

$$0 = V_k(\theta, v_k^e(\theta) + \gamma(v_k^e(\theta))) < V_k(\theta, v_k^e(\theta)),$$

what implies that $v_k^{FB}(\theta) > v_k^e(\theta)$ for all $k \in \{1, \dots, K\}$. The result then follows from the characterization of Lemma 1 and Lemma 3. \square

Proof of Lemma 4. Condition (ICR1) is equivalent to saying that for every $\hat{k} \in \{1, \dots, T\}$

$$(a(\tau), b(\tau)) \in \arg \max_{(\hat{a}, \hat{b})} \left\{ \hat{a} + \hat{b} \cdot \Lambda_k(\theta) - \Phi_k(\theta) \cdot c \right\} \text{ s.t. } (\hat{a}, \hat{b}) \in \left\{ (a(\tau), b(\tau)) : k(\tau) = \hat{k} \right\}.$$

Notice that, because $\Lambda_k(\theta)$ is strictly increasing in θ , the objective above satisfy strict increasing differences in (\hat{b}, θ) . Milgrom’s Constraint Simplification theorem (Milgrom, 2004, pp. 105) then implies that $(a(\tau), b(\tau))$ is a maximizing schedule if and only if the envelope condition (16) holds and $b(\tau)$ restricted to the support $\{\tau : k(\tau) = \hat{k}\}$ is weakly increasing in θ . \square

The proof of the Lemma 5 uses the following convenient mathematical fact.

Lemma 7. Consider an interval $(\theta_1, \theta_2) \in \mathbb{R}_+$ and let A_1, A_2, B_1 , and B_2 be scalars such that $B_1 + B_2 \cdot \theta > 0$ for all $\theta \in (\theta_1, \theta_2)$. Let the function $\zeta : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ be twice continuously differentiable with $\zeta'(x) \geq 0$ for all x . Let $x(\theta)$ be implicitly defined by

$$A_1 + A_2 \cdot \theta = (B_1 + B_2 \cdot \theta) \cdot (x(\theta) + \zeta(x(\theta))). \tag{21}$$

Then $\frac{d^2}{d\theta^2} [(B_1 + B_2\theta)x(\theta)] \geq 0$ if and only if $\zeta''(x(\theta)) \leq 0$.

Proof of Lemma 7. Let us first compute the following derivatives:

$$\begin{aligned} \frac{d}{d\theta} [(B_1 + B_2 \cdot \theta) x(\theta)] &= (B_1 + B_2 \cdot \theta) \cdot x'(\theta) + B_2 \cdot x(\theta) \\ \frac{d}{d\theta} [(B_1 + B_2 \cdot \theta) \cdot \zeta(x(\theta))] &= (B_1 + B_2 \cdot \theta) \cdot \zeta'(x(\theta)) \cdot x'(\theta) + B_2 \cdot \zeta(x(\theta)) \\ \frac{d^2}{d\theta^2} [(B_1 + B_2 \cdot \theta) \cdot x(\theta)] &= (B_1 + B_2 \cdot \theta) \cdot x''(\theta) + 2 \cdot B_2 \cdot x'(\theta) \\ \frac{d^2}{d\theta^2} [(B_1 + B_2 \cdot \theta) \cdot \zeta(x(\theta))] &= (B_1 + B_2 \cdot \theta) \cdot \zeta''(x(\theta)) \cdot [x'(\theta)]^2 \\ &\quad + (B_1 + B_2 \cdot \theta) \cdot \zeta'(x(\theta)) \cdot x''(\theta) + 2 \cdot B_2 \cdot \zeta'(x(\theta)) \cdot x'(\theta). \end{aligned}$$

From the third and fourth equalities above it follows that

$$\frac{d^2}{d\theta^2} [(B_1 + B_2 \cdot \theta) \cdot \zeta(x(\theta))] = (B_1 + B_2 \cdot \theta) \cdot \zeta''(x(\theta)) \cdot [x'(\theta)]^2 + \zeta'(x(\theta)) \cdot \frac{d^2}{d\theta^2} [(B_1 + B_2 \cdot \theta) \cdot x(\theta)]. \tag{22}$$

In turn, because of (21), we know that

$$\frac{d^2}{d\theta^2} [(B_1 + B_2 \cdot \theta) \cdot (x(\theta) + \zeta(x(\theta)))] = 0. \tag{23}$$

Putting (22) and (23) together leads to

$$\frac{d^2}{d\theta^2} [(B_1 + B_2 \cdot \theta) \cdot x(\theta)] = - \left[\frac{(B_1 + B_2 \cdot \theta) \cdot [x'(\theta)]^2}{(1 + \zeta'(x(\theta)))} \right] \cdot \zeta''(x(\theta)),$$

proving the result. \square

Proof of Lemma 5. With a slight abuse of notation, we write $\Phi_0(\theta) = 0$. Therefore, from equation (9), it follows that the expected payments $P^e(\tau)$ restricted to the support $\{\tau : k^e(\tau) = \hat{k}\}$ is weakly convex in θ for every \hat{k} provided that for all $k \in \{1, \dots, T\}$

$$\frac{d^2}{d\theta^2} \{(\Phi_k(\theta) - \Phi_{k-1}(\theta)) \cdot v_k^e(\theta)\} \geq 0. \tag{24}$$

Consider the Virtual Gittins formula (8) at some $k \in \{1, \dots, T\}$. Multiplying both sides by $\theta \cdot (1 - \lambda)^{t-1} + (1 - \theta)$ leads to

$$\begin{aligned} &\theta \cdot (1 - \lambda)^{k-1} \cdot (\lambda \cdot \Delta - K - v_k^e(\theta) - \gamma(v_k^e(\theta))) - (1 - \theta) \cdot (K + v_k^e(\theta) + \gamma(v_k^e(\theta))) \\ &= \mathbf{1}_{[k < T]} \cdot \theta \cdot (1 - \lambda)^{k-1} \cdot \lambda \cdot \delta \cdot \left[\frac{1 - \delta^{T-k}}{1 - \delta} \cdot (\lambda \cdot \Delta - K - v_k^e(\theta) - \gamma(v_k^e(\theta))) \right]. \end{aligned}$$

Rearranging the equation above leads to

$$\begin{aligned} &\theta \cdot (1 - \lambda)^{k-1} \cdot \left(\lambda \cdot \Delta - K - \mathbf{1}_{[k < T]} \cdot \lambda \cdot \delta \cdot \left[\frac{1 - \delta^{T-k}}{1 - \delta} \cdot (\lambda \cdot \Delta - K) \right] \right) - (1 - \theta) \cdot K \\ &= \left(\theta \cdot (1 - \lambda)^{k-1} \cdot \left[1 + \mathbf{1}_{[k < T]} \cdot \lambda \cdot \delta \cdot \frac{1 - \delta^{T-k}}{1 - \delta} \right] + (1 - \theta) \right) (v_k^e(\theta) + \gamma(v_k^e(\theta))) \\ &= (\Phi_k(\theta) - \Phi_{k-1}(\theta)) \cdot (v_k^e(\theta) + \gamma(v_k^e(\theta))), \end{aligned} \tag{25}$$

where the last equality follows from (10).

The first line of (25) is affine in θ . Moreover, the expression for $\Phi_k(\theta) - \Phi_{k-1}(\theta)$ is also affine in θ and strictly positive for every θ . From Lemma 7 it then follows that Condition C, which guarantees the concavity of γ in the relevant range, implies that (24) holds for all $k \in \{1, \dots, T\}$. Therefore, Condition C implies that statement 2 holds.

Finally, it follows from the argument in the text for $T = 1$ that if Condition C fails for some $c \in [\underline{c}, c^{max}]$, then $P^e(\tau)$ restricted to the support $\{\tau : k^e(\tau) = 1\}$ is not weakly convex in θ . This establishes the equivalence between statements 2 and 3. \square

Proof of Step 2 of Proposition 1. The result follows from the following lemma.

Lemma 8. For any deterministic mechanism (ϕ, p) , where ϕ is a stopping plan, there exists a menu of linear contracts $p_l : \Gamma \times X^T \rightarrow \mathbb{R}$ such that $P(\hat{\tau}; \theta) = P_l(\hat{\tau}; \theta)$ for any $\hat{\tau} \in \Gamma$ and any $\theta \in [0, 1]$.

Proof. Define the subset of terminal nodes $B \equiv \{x^T : x_t \neq s \forall t \in \{1, \dots, T\}\}$. Take a payment rule $p : \Gamma \times X^T \rightarrow \mathbb{R}$ and, without any loss of generality, set $p(\hat{\tau}, x^T) = 0$ to all terminal nodes x^T that are not reachable with positive probability under the mechanism (ϕ, p) . It follows by definition that

$$\begin{aligned} \tilde{P}(\hat{\tau}; \theta) &= \theta \cdot \sum_{x^T} \left\{ p(\hat{\tau}, x^T) \cdot \left(\lambda^{S(x^T)} \cdot (1 - \lambda)^{F(x^T)} \right) \right\} + (1 - \theta) \cdot \sum_{x^T \in B} \left\{ p(\hat{\tau}, x^T) \right\} \\ &= \theta \cdot \left[\sum_{x^T} \left\{ p(\hat{\tau}, x^T) \cdot \left(\lambda^{S(x^T)} \cdot (1 - \lambda)^{F(x^T)} \right) \right\} - \sum_{x^T \in B} \left\{ p(\hat{\tau}, x^T) \right\} \right] + \sum_{x^T \in B} \left\{ p(\hat{\tau}, x^T) \right\} \\ &= \Lambda_{k(\hat{\tau})}(\theta) \cdot \left[\frac{\sum_{x^T} \left\{ p(\hat{\tau}, x^T) \cdot \left(\lambda^{S(x^T)} \cdot (1 - \lambda)^{F(x^T)} \right) \right\} - \sum_{x^T \in B} \left\{ p(\hat{\tau}, x^T) \right\}}{\frac{d}{d\theta} [\Lambda_{k(\hat{\tau})}(\theta)]} \right] + \sum_{x^T \in B} \left\{ p(\hat{\tau}, x^T) \right\}, \end{aligned}$$

where the last equality follows from the fact that $\frac{\Lambda_{k(\hat{\tau})}(\theta)}{\frac{d}{d\theta} [\Lambda_{k(\hat{\tau})}(\theta)]} = \theta$ for any stopping plan ϕ (with tolerance for failures described by $k(\cdot)$).

Now consider the menu of linear contracts $p_l : \Gamma \times X^T \rightarrow \mathbb{R}$ with lump-sum payments $a(\hat{\tau})$ and success bonuses $b(\hat{\tau})$. By definition,

$$\tilde{P}_l(\hat{\tau}; \theta) = a(\hat{\tau}) + \Lambda_{k(\hat{\tau})}(\theta) \cdot b(\hat{\tau}).$$

Now set

$$a(\hat{\tau}) = \sum_{x^T \in B} \left\{ p(\hat{\tau}, x^T) \right\}$$

and

$$b(\hat{\tau}) = \frac{\sum_{x^T} \left\{ p(\hat{\tau}, x^T) \cdot \left(\lambda^{S(x^T)} \cdot (1 - \lambda)^{F(x^T)} \right) \right\} - \sum_{x^T \in B} \left\{ p(\hat{\tau}, x^T) \right\}}{\frac{d}{d\theta} [\Lambda_{k(\hat{\tau})}(\theta)]}.$$

Note from equation (13) that $\frac{d}{d\theta} [\Lambda_{k(\hat{\tau})}(\theta)]$, for each tolerance level $k(\hat{\tau})$, is a constant that does not depend on θ . This implies that the payment rule p_l is well-defined. By construction, $\tilde{P}(\hat{\tau}; \theta) = \tilde{P}_l(\hat{\tau}; \theta)$ for any $\hat{\tau} \in \Gamma$ and any $\theta \in [0, 1]$, as we wanted to show.

By Lemma 8, a deterministic mechanism (ϕ, p) , where ϕ is a stopping plan, is feasible if and only if there exists a deterministic mechanism (ϕ, p_l) , where p_l is a menu of linear contracts, that is also feasible. This concludes the proof of Step 2. \square

Proof of Proposition 2. In the proof below, we assume that there exists an optimal deterministic mechanism. We relegate the proof of this claim to the online appendix. Consider a stopping plan where the number of failures born by agents is described by the set $\{r_1, \dots, r_L\} \subset \{1, \dots, T\}$, where $r_1 < r_2 < \dots < r_L$. This mechanism leads to L threshold curves $v_{r_l} : \Theta \rightarrow [\underline{c}, \bar{c}]$, $l \in \{1, \dots, L\}$. Let us follow the convention that $v_{r_{L+1}}(\theta) = \underline{c}$. Hence, for a fixed θ , the set of types (θ, c) who choose a stopping plan with tolerance for r_l failures are the ones with $c \in [v_{r_{l+1}}(\theta), v_{r_l}(\theta)]$ if $v_{r_{l+1}}(\theta) < v_{r_l}(\theta)$ (if $v_{r_{l+1}}(\theta) = v_{r_l}(\theta)$ then this set has at most one type).²³

Now, for $l \in \{1, \dots, L\}$, consider the threshold curve $\theta' \mapsto v_{r_l}^e(\theta')$ as the curve of the expert investor who is restricted to offer the stopping plan where the number of failures born by agents is described by the set $\{r_1, \dots, r_L\}$. Define $\theta_{r_l}^e$ by $\theta_{r_l}^e \equiv \inf \left\{ \theta' : v_{r_l}^e(\theta') > \underline{c} \right\}$ for $l \in \{1, \dots, L\}$.

We say that pooling is a generic property of the optimum if in every optimal mechanism there is a positive measure of firms choosing the same contract. Suppose towards a contradiction that there exists an optimal stopping plan mechanism (henceforth mechanism) \mathcal{M} in which pooling is not generic. For $l \in \{1, \dots, L\}$ let

$$\theta_{r_l} \equiv \inf \left\{ \theta' : v_{r_l}(\theta') - v_{r_{l+1}}(\theta') > 0 \right\},$$

and let $\underline{\theta} \equiv \min_{l \in \{1, \dots, L\}} \theta_{r_l}$. For concreteness we assume that $\underline{\theta} = \theta_{r_1}$. As we will see, the modification of our mechanism will impact types arbitrarily close to $\underline{\theta}$. Hence, if we had $\underline{\theta} = \theta_{r_l}$ and $\underline{\theta} < \theta_{r_k}$ for all $k < l$ then the same analysis could be performed with the stopping plans $\{l, \dots, L\}$. We have to deal with 3 cases.

²³ We ignore the choices at the zero measure sets given by the curves $(\theta, v_l(\theta))$.

Case 1: $\theta_{r_1}^e > \theta_{r_1}$.

In this case, the principal obtains a negative expected payoff from the experimentation of every type $(\theta, c) \in (\theta_{r_1}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1})) \times [\underline{c}, \bar{c}]$. Let k^* be the largest k for which $\theta_{r_k} \in [\theta_{r_1}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1})]$. First assume that $k^* > 1$. Recall that all types in

$$B \equiv \left\{ (\theta, c) : \theta \in \left(\theta_{r_{k^*}}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}) \right), c < v_{r_{k^*}}(\theta) \right\}$$

choose a stopping plan with tolerance for r_{k^*} failures. Let $(a_{k^*}(\tau), b_{k^*}(\tau))$ be the linear contract offered to types $\tau \in B$, where $a_k(\tau)$ stands for the firm's lump-sum payment and $b_{k^*}(\tau)$ for the success bonus.²⁴ The contradiction assumption asserts that whenever $(\theta, c), (\theta', c') \in B, \theta' > \theta$ implies $b_{k^*}(\theta', c') > b_{k^*}(\theta, c)$.

For each $\varepsilon \in (\theta_{r_{k^*}}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}))$ we propose a new mechanism \mathcal{M}^ε . For $\varepsilon \in (0, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}) - \theta_{r_{k^*}})$ let \mathcal{C}^ε be the set of contracts chosen by the types $(\theta, c) \in B$ for which $\theta \in (\theta_{r_{k^*}}, \theta_{r_{k^*}} + \varepsilon)$. Consider a new mechanism \mathcal{M}^ε which differs from the putative optimal mechanism \mathcal{M} only in that it suppresses the contracts \mathcal{C}^ε .²⁵ We focus on the case in which $v_{r_{l+1}}(\theta) < v_{r_l}(\theta)$ for all $(\theta, l) \in (\theta_{r_{k^*}}, \theta_{r_{k^*}} + \varepsilon) \times \{1, \dots, k^*\}$ (the argument for other cases requires minor changes and heavier notation; we omit it for brevity). It is easy to verify that for ε small enough the mechanism \mathcal{M}^ε satisfies the following properties.

a) Almost every type $(\theta, c) \in B^c$ chooses the same allocation in \mathcal{M} and \mathcal{M}^ε .

b) A positive measure of types of B choose a stopping plan with a tolerance for r_{k^*-1} failures in \mathcal{M}^ε .

We claim that reducing the tolerance for failures for a type $\tau \in B$ improves the principal's profit. The principal's payoff conditional on types with prior $\theta \in (\theta_{r_{k^*}}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}))$ can be written as:

$$\begin{aligned} & H(v_{r_1}(\theta)) \left[\sum_{t=1}^{r_1} \delta^{t-1} (\lambda \theta \Delta - K - v_{r_1}(\theta)) + \theta (1 - (1 - \lambda)^{r_1}) \sum_{t=r_1+1}^T \delta^{t-1} (\lambda \Delta - K - v_{r_1}(\theta)) \right] \\ & + H(v_{r_2}(\theta)) [(1 - \theta) + \theta \cdot (1 - \lambda)^{r_1}] \left[\sum_{t=r_1+1}^{r_2} \delta^{t-1} (\lambda \theta_{r_1+1}(\theta) \Delta - K - v_{r_2}(\theta)) \right. \\ & \quad \left. \theta_{r_1+1}(\theta) (1 - (1 - \lambda)^{r_2-r_1}) \sum_{t=r_2+1}^T \delta^{t-1} (\lambda \Delta - K - v_{r_2}(\theta)) \right] \\ & + \dots \\ & + \left[\begin{array}{c} H(v_{r_{k^*}}(\theta)) \\ \times [(1 - \theta) + \theta \cdot (1 - \lambda)^{r_{k^*-1}}] \end{array} \right] \left[\sum_{t=r_{k^*-1}+1}^{r_{k^*}} \delta^{t-1} (\lambda \theta_{r_{k^*-1}+1}(\theta) \Delta - K - v_{r_{k^*}}(\theta)) \right. \\ & \quad \left. \theta_{r_{k^*-1}+1}(\theta) (1 - (1 - \lambda)^{r_{k^*}-r_{k^*-1}}) \sum_{t=r_{k^*}+1}^T \delta^{t-1} (\lambda \Delta - K - v_{r_{k^*}}(\theta)) \right]. \quad (26) \end{aligned}$$

Notice that the first $k^* - 1$ terms in (26) are independent of $v_{r_{k^*}}(\theta)$ and hence the principal's payoff as a function of $v_{r_{k^*}}(\theta)$ is proportional to

$$H(v_{r_{k^*}}(\theta)) \cdot \left[\sum_{t=r_{k^*-1}+1}^{r_{k^*}} \delta^{t-1} (\lambda \theta_{r_{k^*-1}+1}(\theta) \Delta - K - v_{r_{k^*}}(\theta)) \right. \\ \left. \theta_{r_{k^*-1}+1}(\theta) \cdot (1 - (1 - \lambda)^{r_{k^*}-r_{k^*-1}}) \sum_{t=r_{k^*}+1}^T \delta^{t-1} (\lambda \Delta - K - v_{r_{k^*}}(\theta)) \right].$$

The expression above can be written as:

$$\begin{aligned} & H(v_{r_{k^*}}(\theta)) \cdot \left[\delta^{r_{k^*-1}+1} (\lambda \theta_{r_{k^*-1}+1}(\theta) \Delta - K - v_{r_{k^*}}(\theta)) \right. \\ & \quad \left. + \theta_{r_{k^*-1}+1}(\theta) \lambda \sum_{t=r_{k^*-1}+1}^T \delta^{t-1} (\lambda \Delta - K - v_{r_{k^*}}(\theta)) \right] \\ & + H(v_{r_{k^*}}(\theta)) \cdot (1 - \theta_{r_{k^*-1}+1}(\theta) \cdot \lambda) \cdot \left[\delta^{r_{k^*-1}+2} (\lambda \theta_{r_{k^*-1}+2}(\theta) \Delta - K - v_{r_{k^*}}(\theta)) \right. \\ & \quad \left. + \theta_{r_{k^*-1}+2}(\theta) \lambda \sum_{t=r_{k^*-1}+2}^T \delta^{t-1} (\lambda \Delta - K - v_{r_{k^*}}(\theta)) \right] \end{aligned}$$

²⁴ Recall from Lemma 8 that it is without loss to restrict attention to linear contracts.

²⁵ There is a multiplicity of such mechanisms. However, all of them lead almost surely to the same allocation.

$$\begin{aligned}
 &+ \dots \\
 &+ H(v_{r_{k^*}}(\theta)) \cdot \prod_{k=1}^{r_{k^*}-1} (1 - \theta_{r_{k^*}-1+k}(\theta) \cdot \lambda) \cdot \left[\begin{aligned} &\delta^{r_{k^*}} (\lambda \theta_{r_{k^*}}(\theta) \Delta - K - v_{r_{k^*}}(\theta)) \\ &+ \theta_{r_{k^*}}(\theta) \lambda \sum_{t=r_{k^*}}^T \delta^{t-1} (\lambda \Delta - K - v_{r_{k^*}}(\theta)) \end{aligned} \right].
 \end{aligned}$$

Notice that, since $\theta \in (\theta_{r_{k^*}}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}))$ and H is log-concave, each one of the terms above is strictly decreasing in $v_{r_{k^*}}(\theta)$. Therefore a) and b) imply that \mathcal{M}^ε is preferred to \mathcal{M} , which proves the result for the case that $k^* > 1$.

Next, assume that $k^* = 1$. Notice that all types in

$$D \equiv \left\{ (\theta, c) : \theta \in \left(\theta_{r_1}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}) \right), c < v_{r_1}(\theta) \right\}$$

choose the stopping plan with tolerance for r_1 failures. Take $\varepsilon \in (0, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}) - \theta_{r_1})$ and let \mathcal{C}^ε be the set of contracts chosen by types $(\theta, c) \in B$ for which $\theta \in (\theta_{r_1}, \theta_{r_1} + \varepsilon)$. Consider a new mechanism \mathcal{M}^ε which differs from the putative optimal mechanism \mathcal{M} only in that it suppresses the contracts \mathcal{C}^ε .²⁶ It is easy to verify that for ε small enough the mechanism \mathcal{M}^ε satisfies the following properties.

- a) Almost every type $(\theta, c) \in B^c$ chooses the same allocation in \mathcal{M} and \mathcal{M}^ε .
- b) A positive measure of types of B choose no action plan.

The principal payoff from a (fixed) type $\theta \in (\theta_{r_1}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}))$ can be written as:

$$H(v_{r_1}(\theta)) \left[\sum_{t=1}^{r_1} \delta^{t-1} (\lambda \cdot \theta \Delta - K - v_{r_1}(\theta)) + \theta \cdot (1 - (1 - \lambda)^{r_1}) \cdot \sum_{t=r_1+1}^T \delta^{t-1} (\lambda \Delta - K - v_{r_1}(\theta)) \right]. \tag{27}$$

The payoff in (27) is strictly decreasing in $v_{r_1}(\theta)$ since $\theta < \theta_{r_1}^e$. Therefore a) and b) imply that \mathcal{M}^ε is preferred to \mathcal{M} , which proves this case.

Case 2: $\theta_{r_1}^e = \theta_{r_1}$.

First assume that $\theta_{r_k} \in [\theta_{r_1}^e, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_2}^e)]$ for some $\tilde{k} > 1$. In this case, the analysis is almost identical to the one from the first part of Case 1: We propose a new mechanism in which a positive measure of types τ with $\theta \in (\theta_{r_{\tilde{k}}}, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_2}^e))$ migrate from a contract with a tolerance for r_{k^*} failures to another with a tolerance for r_{k^*-l} failures (for some $l > 1$). The calculation above implies that this change increases the principal's payoff. Thus we assume that $k^* = 1$ for the remainder of this case.

Before continuing with the proof, we set up some notation that will be used for remainder of this proof. Recall that for a fixed $\theta \in (\theta_{r_1}^e, \frac{1}{2}(\theta_{r_1}^e + \theta_{r_2}^e))$ all types (θ, c) such that $c < v_{r_1}(\theta)$ receive the same contract and hence we may abuse notation and write $(a(\theta), b(\theta))$ for this contract. The curve $v_{r_1}(\theta)$ is determined (a.e.) by:

$$\begin{aligned}
 a(\theta) + D(\theta) \cdot b(\theta) - (E + F(\theta)) \cdot v_{r_1}(\theta) &= 0 \\
 D'(\theta) \cdot b(\theta) - (E + F(\theta)) \cdot v'_{r_1}(\theta) - F'(\theta) \cdot v_{r_1}(\theta) &= 0,
 \end{aligned}$$

where

$$\begin{aligned}
 D(\theta) &\equiv \left[\sum_{t=1}^{r_1} \delta^{t-1} \lambda + (1 - (1 - \lambda)^{r_1}) \cdot \sum_{t=r_1+1}^T \delta^{t-1} \lambda \right] \cdot \theta \\
 E &\equiv \sum_{t=1}^{r_1} \delta^{t-1} \\
 F(\theta) &\equiv \theta \cdot (1 - (1 - \lambda)^{r_1}) \cdot \sum_{t=r_1+1}^T \delta^{t-1}.
 \end{aligned}$$

The power of the contract $b(\theta)$ satisfies (a.e.)

$$b(\theta) = \left[\frac{(E + F(\theta)) \cdot v'_{r_1}(\theta) + F'(\theta) \cdot v_{r_1}(\theta)}{D'(\theta)} \right]$$

²⁶ There is a multiplicity of such mechanisms. However, all of them lead almost surely to the same allocation.

$$= \left[\frac{\frac{d}{d\theta} (E + F(\theta)) \cdot v_{r_1}(\theta)}{D'(\theta)} \right]$$

Incentive compatibility implies that $b(\theta)$ is increasing. It is convenient to change variables and define $\xi(\theta) \equiv (E + F(\theta)) \cdot v_{r_1}(\theta)$ and thus incentive compatibility requires that $\xi'(\theta)$ is increasing.

Notice that the curve $v_{r_1}^e(\theta)$ satisfies:

$$v_{r_1}^e(\theta) \equiv \arg \max_c H(c) [D(\theta)\Delta - (E + F(\theta)) \cdot c].$$

Hence $v_{r_1}^e(\theta)$ must satisfy

$$D(\theta)\Delta = \left[\left(\frac{H(v_{r_1}^e(\theta))}{h(v_{r_1}^e(\theta))} \right) + v_{r_1}^e(\theta) \right] \cdot (E + F(\theta)).$$

Defining $\xi^e(\theta) \equiv (E + F(\theta)) \cdot v_{r_1}^e(\theta)$, and recalling that γ is the reverse hazard rate, we have:

$$D(\theta)\Delta = \gamma \left(\frac{\xi^e(\theta)}{E + F(\theta)} \right) \cdot (E + F(\theta)) + \xi^e(\theta).$$

Since we are assuming that $\gamma(c)$ is strictly convex, the function $\xi^e(\theta)$ is concave, while monotonicity implies that $\xi(\theta)$ is convex. Furthermore, our contradiction assumption implies that $\xi(\theta)$ is strictly convex. Recall also that we have $\xi(\theta_{r_1}^e) = \xi^e(\theta_{r_1}^e)$. We must deal with two cases. First assume that we can find $\varepsilon \in (0, \frac{1}{2}(\theta_{r_1} + \theta_{r_2}) - \theta_{r_1})$ such that $\xi(\theta') < \xi^e(\theta')$ for all $\theta' \in (\theta_{r_1}, \theta_{r_1} + \varepsilon)$. Let $\bar{\xi}$ satisfy

$$\bar{\xi} = \frac{\xi(\theta_{r_1} + \varepsilon) - \xi(\theta_{r_1})}{\varepsilon}.$$

And define a new curve $\hat{\xi}$ by

$$\hat{\xi}(\theta) \equiv \begin{cases} \xi(\theta_{r_1}) + \bar{\xi} \cdot (\theta - \theta_{r_1}) & \text{if } \theta \in (\theta_{r_1}, \theta_{r_1} + \varepsilon) \\ \xi(\theta) & \text{otherwise.} \end{cases}$$

By construction, we have $\hat{\xi}(\theta) \in (\xi(\theta), \xi^e(\theta))$ for all $\theta \in (\theta_{r_1}, \theta_{r_1} + \varepsilon)$. Therefore, a new mechanism $\hat{\mathcal{M}}$ which implements the curve $\hat{\xi}$ leads to a strictly higher profit than the original mechanism \mathcal{M} , contradicting the optimality of \mathcal{M} . This is because that the experimentation region is approaching (one can take the Hausdorff topology to measure the distance between two sets) the experimentation region given by $\xi^e(\theta)$.²⁷ The argument follows because the profit function is strictly quasi-concave.

Next, assume that $\xi(\theta') > \xi^e(\theta')$ for all $\theta' \in (\theta_{r_1}, \theta_{r_1} + \varepsilon)$. Notice that, since ξ is convex, the sub-differentials $\partial\xi(\theta_{r_1})$ and $\partial\xi(\theta_{r_1} + \varepsilon)$ are compact and nonempty. Let $\xi_1 \equiv \max \partial\xi(\theta_{r_1})$ and $\xi_2 \equiv \max \partial\xi(\theta_{r_1} + \varepsilon)$. The strict convexity of ξ implies that $\xi_1 < \xi_2$. Next, consider the following lines: $L_1(\theta) \equiv \xi(\theta_{r_1}) + \xi_1 \cdot (\theta - \theta_{r_1})$ and $L_2(\theta) \equiv \xi(\theta_{r_1} + \varepsilon) + \xi_2 \cdot (\theta - \theta_{r_1} + \varepsilon)$. Notice that $\xi_1 < \xi_2$ implies that there is a unique $\hat{\theta} \in (\theta_{r_1}, \theta_{r_1} + \varepsilon)$ such that $L_1(\hat{\theta}) = L_2(\hat{\theta})$. Therefore, we define a new (convex) curve $\hat{\xi}$ by:

$$\hat{\xi}(\theta) \equiv \begin{cases} L_1(\theta) & \text{if } \theta \in (\theta_{r_1}, \hat{\theta}) \\ L_2(\theta) & \text{if } \theta \in [\hat{\theta}, \theta_{r_1} + \varepsilon) \\ \xi(\theta) & \text{otherwise.} \end{cases}$$

Notice that, by construction, $\hat{\xi}(\theta) \in (\xi^e(\theta), \xi(\theta))$ for all $\theta \in (\theta_{r_1}, \theta_{r_1} + \varepsilon)$. Thus, as above, we can construct a new mechanism $\hat{\mathcal{M}}$ which leads to a strictly higher payoff to the principal.

Case 3: $\theta_{r_1}^e < \theta_{r_1}$.

Define $\theta \mapsto \underline{\xi}(\theta)$ by $\underline{\xi}(\theta) = \underline{c}(E + D(\theta))$. In this case, take $\varepsilon > 0$ such that for all $\theta' \in (\theta_{r_1}, \theta_{r_1} + \varepsilon)$ we have $\xi^e(\theta) > \xi(\theta)$.

Consider the line passing through the points $(\theta_{r_1}^e, \underline{\xi}(\theta_{r_1}^e))$ and $(\theta_{r_1} + \varepsilon, \xi(\theta_{r_1} + \varepsilon))$:

$$L(\theta) \equiv \underline{\xi}(\theta_{r_1}^e) + \left(\frac{\xi(\theta_{r_1} + \varepsilon) - \underline{\xi}(\theta_{r_1}^e)}{\theta_{r_1} + \varepsilon - \theta_{r_1}^e} \right) \cdot (\theta - \theta_{r_1}^e).$$

²⁷ For $i \in \{1, 2, 3\}$ take 3 functions $\theta \rightarrow \psi^i(\theta)$. We say that ψ^2 is closer to ψ^3 than ψ^1 if for all θ such that $\psi^2(\theta) \leq \psi^3(\theta)$ (resp. $\psi^2(\theta) \geq \psi^3(\theta)$) we have $\psi^1(\theta) \leq \psi^2(\theta)$ (resp. $\psi^1(\theta) \geq \psi^2(\theta)$). The construction above guarantees that the function $\hat{\xi}$ is closer to ξ^e than ξ .

Define the convex curve $\hat{\xi}$ by:

$$\hat{\xi}(\theta) := \begin{cases} L(\theta) & \text{if } \theta \in (\theta_{r_1}^e, \theta_{r_1} + \varepsilon) \\ \xi(\theta) & \text{otherwise.} \end{cases}$$

Consider the mechanism $\hat{\mathcal{M}}$ in which the curve $v_{r_1}(\theta)$ is defined by $\hat{\xi}$. Notice that this mechanism may in addition change the curves $v_{r_1}(\theta)$ for $(l, \theta) \in \{2, \dots, L\} \times (\theta_{r_1}, \theta_{r_1} + \varepsilon)$. This new mechanism $\hat{\mathcal{M}}$ clearly increases the principal's payoff for all $\theta \in (\theta_{r_1}^e, \theta_{r_1})$, while it may decrease the principal's payoff in a subset of $(\theta_{r_1}, \theta_{r_1} + \varepsilon)$. Finally, the mechanism $\hat{\mathcal{M}}$ does not change the principal's payoff for $\theta \notin (\theta_{r_1}^e, \theta_{r_1} + \varepsilon)$. Let $\underline{g} \equiv \min g(\theta)$, $\bar{g} \equiv \max g(\theta)$, $\underline{h} \equiv \min h(\theta)$ and $\bar{h} \equiv \max h(\theta)$. Notice that, for ε sufficiently small, there exists a constant $A > 0$ such that the principal's payoff increase by at least A for every type from $\{(\theta, c) : \theta \in (\frac{1}{2}(\theta_{r_1}^e + \theta_{r_1}), \theta_{r_1}), c < \hat{\xi}(\theta)\}$ who plays the risky arm (below we will integrate over this set). Notice also that there exists a constant $B > 0$ such that the absolute value of the payoff obtained by the principal in the mechanism $\hat{\mathcal{M}}$ from the experimentation of any type $\{(\theta, c) : \theta \in (\theta_{r_1}^e, \theta_{r_1} + \varepsilon), c < \hat{\xi}(\theta)\}$ is at most B (below we will integrate over this set). Therefore, a lower bound to the difference of the principal's payoff from $\hat{\mathcal{M}}$ and \mathcal{M} is:

$$\left(A \underline{g} \underline{h}\right) \int_{\frac{1}{2}(\theta_{r_1}^e + \theta_{r_1})}^{\theta_{r_1}} (\hat{\xi}(\theta) - \underline{\xi}(\theta)) d\theta - \left(B \bar{g} \bar{h}\right) \int_{\theta_{r_1}}^{\theta_{r_1} + \varepsilon} (\hat{\xi}(\theta) - \underline{\xi}(\theta)) d\theta. \tag{28}$$

Notice that for ε small we have $\hat{\xi}(\theta) - \underline{\xi}(\theta) > \left(\frac{\hat{\xi}(\theta_{r_1} + \varepsilon) - \underline{\xi}(\theta_{r_1} + \varepsilon)}{4}\right)$ for all $\theta \in (\theta_{r_1}^e, \theta_{r_1} + \varepsilon)$. Hence, (28) is at least:

$$\left[\hat{\xi}(\theta_{r_1} + \varepsilon) - \underline{\xi}(\theta_{r_1} + \varepsilon)\right] \left[\left(\frac{A \underline{g} \underline{h}}{4}\right) \left(\frac{\theta_{r_1} - \theta_{r_1}^e}{2}\right) - \varepsilon \left(B \bar{g} \bar{h}\right)\right],$$

which is strictly positive whenever $\varepsilon < \left(\frac{A \underline{g} \underline{h}}{4 B \bar{g} \bar{h}}\right) \left(\frac{\theta_{r_1} - \theta_{r_1}^e}{4}\right)$. Therefore the mechanism $\hat{\mathcal{M}}$ leads to a strictly higher profit, a contradiction. \square

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <http://dx.doi.org/10.1016/j.geb.2016.02.001>.

References

Aboody, D., Lev, B., 2000. Information asymmetry, R&D, and insider gains. *J. Finance* 55, 2747–2766.
 Beaudry, P., Blackorby, C., Szalay, D., 2009. Taxes and employment subsidies in optimal redistribution programs. *Amer. Econ. Rev.* 99, 216–242.
 Bergemann, D., Hege, U., 1998. Venture capital financing, moral hazard, and learning. *J. Banking Finance* 22, 703–735.
 Bergemann, D., Hege, U., 2005. The financing of innovation: learning and stopping. *RAND J. Econ.* 36, 719–752.
 Board, S., Meyer-ter-Vehn, M., 2013. Reputation for quality. *Econometrica* 81, 2381–2462.
 Bonatti, A., Hörner, J., 2011. Collaborating. *Amer. Econ. Rev.* 101, 632–663.
 Bouvard, M., 2014. Real option financing under asymmetric information. *Rev. Finan. Stud.* 27 (1), 180–210.
 Canidio, A., Legros, P., 2015. The value of entrepreneurial failures: task allocation and career concerns. Unpublished.
 Campbell, A., Ederer, F., Spinnewijn, J., 2014. Delay and deadlines: free-riding and information revelation in partnerships. *Amer. Econ. J., Microecon.* 6 (2), 163–204.
 Crémer, J., McLean, R., 1988. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica* 56, 1247–1257.
 Ederer, F., 2013. Incentives for parallel innovation. Unpublished.
 Ederer, F., Manso, G., 2013. Is pay-for-performance detrimental to innovation? *Manage. Sci.* 59 (7), 1496–1513.
 Fong, K., 2009. Evaluating skilled experts: optimal scoring rules for surgeons. Unpublished.
 Francis, B., Hasan, I., Sharma, Z., 2009. Do incentives create innovation? Evidence from CEO compensation contracts. Unpublished.
 Garfagnini, U., 2011. Delegated experimentation. Unpublished.
 Gerardi, D., Maestri, L., 2012. A principal-agent model of sequential testing. *Theoretical Econ.* 7, 425–463.
 Gittins, J.C., 1979. Bandit processes and dynamic allocation indices. *J. Roy. Statistical Soc., Ser. B Methodological* 41, 148–177.
 Gompers, P., 1995. Optimal investment, monitoring, and the staging of venture capital. *J. Finance* 50, 1461–1489.
 Gompers, P.A., Lerner, J., 2004. *The Venture Capital Cycle*. MIT Press, Cambridge.
 Guesnerie, R., Laffont, J.-J., 1984. A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm. *J. Public Econ.* 25, 329–369.
 Guo, Y., 2014. Dynamic delegation of experimentation. Unpublished.
 Halac, M., Kartik, N., Liu, Q., 2015. Optimal contracts for experimentation. Unpublished.
 Heidhues, P., Rady, S., Strack, P., 2015. Strategic experimentation with private payoffs. *J. Econ. Theory* 159 (A), 531–551.
 Hörner, J., Samuelson, L., 2013. Incentives for experimenting agents. *RAND J. Econ.* 44 (4), 632–663.
 Jehiel, P., Moldovanu, B., 2001. Efficient design with interdependent valuations. *Econometrica* 69 (5), 1237–1259.
 Kaplan, S., Stromberg, P., 2003. Financial contracting theory meets the real world: an empirical analysis of venture capital contracts. *Rev. Econ. Stud.* 70, 281–315.

- Keller, G., Rady, S., Cripps, M., 2005. Strategic experimentation with exponential bandits. *Econometrica* 73, 39–68.
- Klein, N., 2015. The importance of being honest. *Theoretical Econ.*, forthcoming.
- Kerr, W., Nanda, R., Rhodes-Kropf, M., 2014. Entrepreneurship as experimentation. *J. Econ. Perspect.* 28, 25–48.
- Landier, A., 2006. Entrepreneurship and the stigma of failure. Unpublished.
- Lerner, J., Wulf, J., 2007. Innovation and incentives: evidence from corporate R&D. *Rev. Econ. Statist.* 89, 634–644.
- Manso, G., 2011. Motivating innovation. *J. Finance* 66, 1823–1860.
- Mezzetti, C., 2004. Mechanism design with interdependent valuations: efficiency. *Econometrica* 72 (5), 1617–1626.
- Milgrom, P., 2004. *Putting Auction Theory to Work*. Cambridge University Press, New York.
- Ries, E., 2011. *The Lean Startup: How Today Entrepreneurs Use Continuous Innovation to Create Radically Successful Business*. Crown Business Publishing.
- Riordan, M., Sappington, D., 1988. Optimal contracts with public ex post information. *J. Econ. Theory* 45, 189–199.
- Roberts, K., Weitzman, M., 1981. Funding criteria for research, development, and exploration projects. *Econometrica* 49, 1261–1288.
- Rochet, J.C., Choné, P., 1998. Ironing, sweeping and multidimensional screening. *Econometrica* 66 (4), 783–826.
- Rochet, J.C., Stole, L., 2003. The economics of multidimensional screening. In: *Advances in Economics and Econometrics: Theory and Applications. Eight World Congress*. In: *Econometric Society Monographs*, vol. 36. Cambridge University Press, pp. 783–826.
- Rothschild, M., 1974. A two-armed bandit theory of market pricing. *J. Econ. Theory* 9, 185–202.
- Strausz, R., 2006. Deterministic versus stochastic mechanisms in principal-agent models. *J. Econ. Theory* 127 (1), 306–314.