Moral Hazard in Experiment Design: Implications for Financing Innovation[*]

Patrick Bolton
Imperial College London

Shannon Liu
Rotman School of Management

Ramana Nanda
Imperial College London

Savitar Sundaresan
Imperial College London

June 2024

Abstract

We develop a model of learning through experimentation in a principal-agent framework. Investors only observe an experiment's outcome, but entrepreneurs can impact the information contained in the outcome through the experiment's design. Investors prefer 'killer experiments' that are more likely to correctly identify true successes and failures, but entrepreneurs prefer to design experiments that are less likely to fail. We show that the ensuing moral hazard can create a market failure in financing the venture, which cannot be resolved through higher-powered incentives for the entrepreneur such as standard 'pay for performance' contracts. Our results speak to an important potential friction in the commercialization of innovations, particularly ones in areas such as 'Deep Tech' ventures based on fundamental science, that lack well-understood methodologies for investors to effectively validate the information contained in early experiments.

[*]We are grateful to Ajay Agrawal, Joshua Gans, Bill Janeway, Nicola Lacetera and Gustavo Manso for helpful discussions, and to seminar participants at HBS, IESE, Rotman, Berkeley for valuable feedback. This paper is part of a project that has received funding from the European Research Council (ERC) Horizon 2020 research and innovation programme, Grant Number 865127. All errors are our own.

1 Introduction

One of the most striking facts about the development of new technologies is how poor initial experiments can be at identifying whether or not those technologies will work at scale. In the pharmaceutical industry, for example, less than 10% of New Molecular Entities identified at the preclinical stage – following extensive lab experiments – progress through clinical trials to launch (Paul et al. 2010). Similar observations about the poor predictability of early lab experiments have been noted across several other industries (Siegmund et al. 2021; Greenwood et al. 2022).

The inherent uncertainty around radical innovations partly explains the low predictability of future success by early experiments. Nevertheless, given the billions of dollars invested each year into projects that ultimately go on to fail across these sectors, the ability to develop more effective early indicators of success goes to the heart of some first order concerns among Venture Capital (VC) investors, corporate R&D divisions and policy makers involved in providing financing for these innovations. Any potential frictions that might prevent the development of more predictive early experiments are worth investigating given the stakes.

In this paper, we focus on one such friction stemming from moral hazard in the *design* of early experiments with unfamiliar technologies. To emphasize the role that this friction plays relative to others that have been discussed in the literature (e.g., Bergemann and Hege (1998); Manso (2011)), we develop a model of learning through experimentation in a principal-agent framework that shuts down the informational frictions considered in this literature but introduces a novel one focused on experiment design. Specifically, all information about the project's value as well as priors about its likelihood of working is assumed to be public knowledge and is therefore known to both the investor and the entrepreneur. We motivate such a context using an example where an entrepreneur is commercializing a scientific invention such as a new cancer therapeutic, where the entrepreneur does not have any information advantage about the likelihood of it being successful.[1]

We focus attention on how an experiment's design can impact the information contained in the outcome of the experiment. Similar to the medical diagnostics literature where tests vary in their sensitivity (true positive rate) and specificity (true negative rate), we characterize an experiment's design by the degree to which it is likely to correctly identify viable/unviable technologies at the experimental stage. More robust experiments – with higher sensitivity and specificity – have fewer false positives and fewer false negatives.

In our model, the entrepreneur is assumed to not have sufficient funds to develop the invention himself, and depends on a VC to finance commercialization. The VC can finance an initial experiment to learn if the technology is likely to be viable, for example by examining whether a therapeutic candidate meets certain critical milestones related to in vivo experimental models in the lab. Based on the results from this experiment, she can choose to finance the development of the venture at the next stage or abandon funding its development. The VC will only fund the experiment and subsequent development if it is an NPV positive investment, but should she finance the venture's development, uncertainty about commercial viability is only resolved after her final investment has been made.

Asymmetry between the principal and the agent in our context stems from two sources. First, we assume that while the VC only cares about the NPV of the investment, the

[1]There is only one project that the entrepreneur works on in our model, and there is no diversion of effort or cash by the entrepreneur. In other words, to highlight the mechanism we emphasize, we shut down the typical sources of adverse selection and moral hazard articulated in principal-agent models on the financing of innovation.

entrepreneur/scientist also derives private benefits advancing science and working on the venture. The second source of asymmetry stems from the fact that although the VC can verify the outcome of an experiment – e.g. whether a particular technical milestone in the lab experiment has been met or not – the experiment's design, is largely a matter for the entrepreneur/scientist, who alone has the knowledge of how best to tweak the parameters of the experiment. All the VC can ascertain is the parameter set of possible experiment designs, but she cannot observe all the detailed tweaking that has gone into the experiment design. Since the experiment design impacts the degree to which the early experiment correctly identifies viable/unviable projects, it can impact the likelihood of meeting a given milestone. For example, an experiment that is worse at identifying unviable ventures is likely to generate more false positives and will therefore be more likely to pass the milestone.

This setup delivers several results. First, we show that although both the VC and the entrepreneur share the same prior about the project's likelihood of success and value if successful, they differ in the degree to which they want to *learn* from the experiment. Since the VC finances the experiment (and the subsequent scale up if the experiment shows promise), she is sensitive to 'throwing good money after bad' and is keen to identify unviable projects as soon as possible. She therefore prefers 'killer experiments' that minimize both false positives and false negatives as much as possible. On the other hand, the entrepreneur gets private benefits from continuation and does not bear the cost of financing an ultimately unviable venture. He therefore does not want the investor to shut down the project prematurely, so does not want the VC (or himself) to learn whether the project is likely to be unviable. Conditional on receiving funding to run the experiment, the entrepreneur prefers to design experiments that maximize the likelihood of generating a positive test outcome (and therefore false positives) – since such experiments minimize the chance that the project is shut down by the VC. It is important to emphasize that in this model, the entrepreneur does not have more information than the investor about whether the project will work, nor does he shirk or divert resources. Moral hazard arises from the fact that he has a differential incentive relative to the VC to *learn* if the project is unviable. Once the entrepreneur/scientist has secured funding to conduct an experiment, he cannot commit to designing the 'killer experiment' that the VC would like him to run.

Second, we show that this moral hazard in experiment design can lead to large inefficiencies. This is because the VC knows the entrepreneur's incentives, and in the absence of being able to verify the experiment's design, (correctly) assumes that if funded, he will design an experiment with the greatest potential false positives. Experiments with the most false positives are also the most inconclusive from the perspective of the VC. She cannot update her priors as much when she sees the experiment passing the milestone given its low information content. This makes it less likely for her to finance the development (or the initial experiment) and leads to a lower likelihood of the entrepreneur being funded relative to a benchmark without this friction. It can even lead to complete market failure for funding the venture. In instances where there is sufficient value for the VC to invest in spite of the friction, the venture will get funded, but would fail too often once implemented relative to a benchmark without this friction.

Our model therefore provides a theoretical rationale for both the low predictability (and high failure rate) of funded projects as well as the general lack of funding for 'deep tech' ventures building on fundamental science. While the challenges of financing such ventures has been the focus of a number of recent studies, our model provides a different, complementary, explanation for the low levels of VC funding available for these sectors. It also provides a theoretical underpinning for findings by Guedj and Scharfstein (2004) who show that startups are more likely than mature firms to advance from Phase I to

Phase II clinical trials, but conditional on passing, tend to have less promising clinical results in their Phase II trials and their Phase II drug candidates are also less likely to advance to Phase III and to receive FDA approval.

Third, we show that in most instances, the VC cannot align the entrepreneur's incentives to design a more informative experiment with standard pay-for-performance contracts. This is mainly because the entrepreneur does not bear the cost of financing the development of the venture; there is no downside risk for the entrepreneur. Therefore, being able to continue working on developing the project provides option value. He gets private benefits from working on the project and moreover, if the venture is viable, he gets a financial return. Providing him more 'skin in the game' only makes the option more valuable.

Fourth, we show that expanding the contract space to reward 'proof of failure' can help address this challenge, but this expansion is very sensitive to the size of the payment, so will require an understanding of the entrepreneur's private benefit from continuing to work on the venture. If the VC does not pay the entrepreneur enough, the moral hazard is not solved. If the VC compensates the entrepreneur too highly, the entrepreneur is now incentivized to design experiments with excessive false negatives which can lead the VC to miss out on funding promising viable ventures. The challenge is particularly stark if the experiment's design is such that the sensitivity and specificity are substitutes and is attenuated in cases where they are complements.

Finally, we discuss several policy solutions, including the role of an organization that can validate the experiment design and thereby help address the market failure.

2 Relation to Literature

Our analysis builds on the theoretical literature on the financing of innovation and venture capital in a principal-agent framework (Aghion and Bolton 1992; Hellmann 1998), particularly the theories focusing on the entrepreneur's discretion to influence the learning of the project's value and hence investor's financing decisions (Bergemann and Hege 1998; Cornelli and Yosha 2003; Bergemann and Hege 2005). Specifically, Bergemann and Hege (1998) analyze a situation where the entrepreneur can divert funds (or effort) to her private ends instead of investing into experimenting with the project. Cornelli and Yosha (2003) address the window-dressing problem of performance signals by the entrepreneur to secure further funding. Bergemann and Hege (2005) study the dynamic agency conflicts surrounding the timing of terminating a research project under an infinite funding horizon and compare the overall efficiency of arm's length financing (actions are unobservable) and relationship financing (observable actions). In their analysis, the agency problem can generally be mitigated or resolved through 'skin-in-the-game' incentive contracting, such as staged financing (Bergemann and Hege 1998), convertible securities (Cornelli and Yosha 2003), and state-contingent control rights (Aghion and Bolton 1992; Hellmann 1998). We focus on a different type of moral hazard problem embedded in staged financing, arising from the ability of the entrepreneur to manipulate, in a general fashion, the learning technology that generates information indicative of the project's value. We also show that such agency problems cannot be resolved through higher-powered incentives for the entrepreneur.

Second, our theory is linked to the burgeoning entrepreneurship literature that treats entrepreneurial decision-making as strategic learning and experimentation (Gans, Stern, and Wu 2019; Camuffo et al. 2020; Agrawal, Gans, and Stern 2021; Camuffo et al. 2022).

In particular, Gans, Stern, and Wu (2019) argue that actors should design experiments that focus on information that is relevant to their particular decision. Agrawal, Gans, and Stern (2021) further highlights the potential tradeoff in the choice over the experiments when validating entrepreneurial ideas. Within this body of work, however, few models address the need to persuade investors to provide funding for the project and experiments, with the exception of Karp, Shelef, and Wuebker (2024). Using a Bayesian persuasion framework introduced by Kamenica and Gentzkow (2011), Karp, Shelef, and Wuebker (2024) finds that under a wide range of conditions, actors prefer to reduce the informativeness of the experiment to enhance credible 'cheap talk' persuasion even when the fully informative experiment is available. In this persuasion-based theory, the receiver is assumed to be fully informed about the designer's experimental strategy and the resulting information environment. On the contrary, what characterizes the essence of the agency problem in our setting is the lack of such understanding. In other words, the Bayesian persuasion approach endows the experiment designer with full commitment power of her choice over the information structure (Kolotilin 2015; Fréchette, Lizzeri, and Perego 2022). In contrast, our analysis stems from the absence of such commitment power, an assumption particularly suitable in ventures based on fundamental science and technology for which a huge knowledge gap exists between the founder and the venture capitalists.

Our work is also related to the empirical literature on the agency frictions in the financing of innovation (Gompers 1995; Hellmann and Puri 2000; Kaplan and Strömberg 2003; Guedj and Scharfstein 2004; Hall and Lerner 2010), around the challenge faced by deep tech ventures to get funded and scale up (Lerner and Nanda 2020; Fosfuri and Nagar 2023; Dalla Fontana and Nanda 2023). For instance, Dalla Fontana and Nanda (2023) document that VC financing accounts for a tiny share of all patents related to

Net Zero, and that the patenting focus of VC-backed firms has shifted away from "deep tech" in recent years. Fosfuri and Nagar (2023) show that startups at the frontier of science experience delays in VC funding, which impede its commercialization. Recent theories have been proposed to account for such difficulties. Arora, Fosfuri, and Roende (2022) reason that the allocation of costs across the different stages of the R&D process affects the division of innovative labour. As a result, startups facing both technology and market uncertainty are unable to find the required VC funding. Kremer, Levin, and Snyder (2022) provide a formal analysis of the effect of Advance Market Commitments in stimulating investment by suppliers of products to low-income countries with limited monopoly rents. We contribute to this discussion by studying an unexplored friction: moral hazard in experimental design that can be so severe that it leads to complete market failure in the financing of deep tech ventures. The implications for validating the experiments is well aligned with the notion of the Technology Readiness Level scale, proposed as an aid in better-informed decision-making regarding investments in several nascent tech industries including battery technologies, quantum, and machine learning (Lavin et al. 2022; Greenwood et al. 2022; Purohit et al. 2024).

Additionally, the paper speaks to the theoretical research on strategic interactions when an agent generates information through costly research to persuade a principal to approve an activity, an example being the regulatory process for drug approval (Di Tillio, Ottaviani, and Sørensen 2017; Henry and Ottaviani 2019; Bates et al. 2022; Balasubramanian, Pierce, and Cummings 2022; Bates et al. 2023). Specifically, Henry and Ottaviani (2019) compares organizations with different commitment power of informer and evaluator and shows that granting authority to the informer is socially optimal when information acquisition is sufficiently costly. Bates et al. (2022) discuss how the principal and agent can enter into a contract with payoffs based on statistical evidence that is robust to strategic action. Bates et al. (2023) view the agent as acting according to an implicit prior distribution and show how the principal can deduce information about this prior distribution from the agent's behavior. In all three instances, the underlying learning problem, typically a result of asymmetric information, exhibits statistical properties that can be exploited by the principal to make inferences. In our setting, however, the one-dimensional milestone tied to staged financing fundamentally limits such statistical inference. Moreover, Di Tillio, Ottaviani, and Sørensen (2017) analyze persuasion bias in randomized controlled trial design when the agent can use private information to manipulate the outcome of the experiment and discuss the welfare impact of three different types of strategic deviations. Chassang, Padró i Miquel, and Snowberg (2012) also study the design of randomized controlled experiments, with the distortion coming from experimental subjects' unobserved effort. In both cases, the structure of the learning problem differs from ours as given by the statistical regularities inherent in their specific settings.

Lastly, our work is closely related to medical research that surveys the overall performance of clinical trials and investigates specific ways to improve efficacy. In particular, the moral hazard problem identified in this paper is consistent with the presence of low success rates in subsequent clinical trial phases (Paul et al. 2010) . The economic reasoning offers a rationale as to why pivotal trials are often initiated with insufficient evidence (Kim et al. 2022) and the existence of a general lack of efficacy in the intended disease indication (Hingorani et al. 2019). Although our analysis does not directly address the question of how to improve clinical trial designs in a specific setting, the model highlights the two parameters - the type 1 (false-positive) and type 2 (false-negative) error rates of the experiments - and more crucially, their range of values as determinants of the potential efficacy of any tests in light of the agency problem.

3 The Model

We develop a contracting model for a venture between a risk-neutral venture capitalist (the investor) and a risk-neutral entrepreneur (the entrepreneur), where conditions for the venture's success can be assessed with an experiment.

3.0.1 Timing

The model is static, and broken into three periods. Period 0 is the *contracting period* when the investor makes a take-it-or-leave-it offer to the entrepreneur over an investment in the venture. Period 1 is the *experimentation period* when the entrepreneur conducts an experiment to determine whether necessary conditions for the venture's success are met. Below, to fix ideas, we consider the venture's technical feasibility as one such condition. Period 2 is the *implementation period* when the venture is fully developed, conditional on a follow-up investment based on the experiment results, and payoffs are realized. If the investor does not provide further investment, the venture is abandoned.

3.0.2 Technology

The venture v has two states, $v \in V, 0$. It either succeeds and generates a value V, or fails, with the scrap value assumed to be 0 for simplicity. The venture requires K to develop the venture fully, with $K < V$. Ex-ante, the belief that the venture will be profitable is p_0, which is common knowledge. We also assume that

$$p_0 V - K < 0. \tag{1}$$

Failing any additional knowledge on the workability of this technology (a necessary condition for the venture to be profitable), the investor and entrepreneur will not pursue this venture. In other words, ex-ante, the venture has a negative NPV. More information can be gained by conducting an experiment.

3.0.3 Experiment

The entrepreneur has the exclusive capability to run an experiment (say, a lab test) that can generate information on the workability of the technology. The cost of the experiment is $C > 0$. The entrepreneur has the latitude to design the experiment to produce signals with varying false negative or false positive outcomes. This test can return two possible signals, s: F if the test fails the test, and P if the technology passes the test.

It may be helpful to begin by describing the benchmark of a perfect experiment. If technical feasibility is a sufficient and necessary condition for the venture's success, then an experiment that can conclusively determine the feasibility of the technology is a perfect experiment. If the technology passes the test in the lab, the entrepreneur knows for sure that the venture will be successful, and if it fails the test, the entrepreneur knows with equal certainty that the venture will flop. Given the prior p_0, the expected payoff from running the perfect experiment and implementing the technology if and only if a successful result is produced in the lab is:

$$p_0(V - K) - C.$$

We shall of course assume that

$$p_0(V - K) - C > 0 \tag{2}$$

The availability of a perfect experiment makes the venture positive NPV. Comparing conditions (1) and (2), it is immediately obvious what the value of running such an experiment is. By paying the cost of C, the investor eliminates the risk of sinking a large investment of K into an unprofitable venture, which, ex-ante, happens with probability $1 - p_0$.

In reality, however, verifying the existence of a necessary condition can reduce a venture's risk, but only to a certain extent. To simply the risks involved in developing a venture, we characterize the experiment by s_1, the probability that the test succeeds for a technology that is workable in practice and hence the venture can generate V, and s_2, the probability that the test fails when the technology is not workable in practice and hence the venture is a flop:

$$\begin{aligned} P(s = P | v = V) &= s_1 \\ P(s = F | v = 0) &= s_2 \end{aligned}$$

In other words, we define s_1 to be the parameter that reflects the *specificity* of the experiment, with $(1 - s_1)$ denoting the *rate of false negative* test outcomes. Similarly, we define s_2 to be the parameter that reflects the *sensitivity* of the experiment, with $(1 - s_2)$ denoting the *rate of false positive* test outcomes. The entrepreneur can jointly choose the parameter values (s_1, s_2) within a set S of possible test designs, $S = [\underline{s}_1, \overline{s}_1] \times [\underline{s}_2, \overline{s}_2]$, where $0 \leq \underline{s}_i < \overline{s}_i \leq 1$ for $i = 1, 2$.

For any given experiment characterized by s_1 and s_2, now the expected payoff from running such an experiment and investing K to implement the technology if and only if a pass signal P is generated is:

$$\pi_{s_1, s_2} = p_0 s_1 V - [p_0 s_1 + (1 - p_0)(1 - s_2)]K - C \tag{3}$$

Comparing conditions (1) and (3), in this imperfect world, the net gain from running such an experiment, $\Delta\pi_{s_1,s_2}$, is given by

$$\Delta\pi_{s_1,s_2} = (K - p_0 V - C) + p_0 s_1 (V - K) - (1 - p_0)(1 - s_2)K$$

Holding K, p_0, V, and C constant, and given that $V > K$, it is obvious that the net gain from the experiment is strictly increasing in both s_1 and s_2. This implies that the best experiment design in terms of maximizing value gained from information discovery is the one with the greatest specificity (the highest s_1) and the greatest sensitivity (the highest s_2). Note that a perfect experiment has $s_1 = 1$ and $s_2 = 1$, the greatest specificity and sensitivity possible, and therefore, if available, yields the highest gain.

The entrepreneur, in effect, faces a multitasking problem in designing the test (how specific and how sensitive to make the test). The contracting problem between the entrepreneur and the venture capitalist, therefore, has elements of a multitask Principal-Agent problem (Holmstrom and Milgrom, 1991). A key consideration in multitasking Principal-Agent problems is whether the tasks are independent, complementary, or substitutes. For the problem, we consider this issue boils down to the question of whether the entrepreneur can manipulate the specificity and sensitivity of the test independently, whether a more specific test is also more sensitive (complementary tasks), or whether greater specificity necessarily means less sensitivity (substitutable tasks). We consider each case in turn.

For simplicity, we assume that each test specification has the same cost $C > 0$. In the case of independent tasks, we assume that s_1 and s_2 can be chosen independently by the entrepreneur within the *test set* $S = [\underline{s}_1, \overline{s}_1] \times [\underline{s}_2, \overline{s}_2]$, where $0 \leq \underline{s}_i < \overline{s}_i \leq 1$ for $i = 1, 2$. In the case of substitute tasks, we consider the extreme case where specificity and sensitivity are perfect substitutes, which means that the sum of the parameter values $s_1$and s_2 always add up to a constant κ, so that $s_1 + s_2 = \kappa$, where $1 < \kappa < 2$. The test set in this case of perfect substitutes thus takes the form $S = \{(s_1, s_2) \mid s_1 + s_2 = \kappa, \underline{s}_1 \leq s_1 \leq \overline{s}_1, \text{and } \underline{s}_2 \leq s_2 \leq \overline{s}_2\}$. We assume $\overline{s}_1 + \overline{s}_2 > \kappa$. Hence, the entrepreneur can maximize specificity by setting $s_1 = \overline{s}_1$ or maximize sensitivity (setting $s_2 = \overline{s}_2$), but he cannot do both. At the frontier, a more sensitive test is inevitably a less specific test, and vice-versa. In the case of complementary tasks, we again consider an extreme case where sensitivity and specificity are perfect complements. This means that the test designs are all such that $s_2 = \lambda s_1$ where $\lambda \in (0, 1)$, so that the test set is given by $S = \{(s_1, s_2) \mid s_2 = \underline{s}_2 + \lambda s_1, \underline{s}_1 \leq s_1 \leq \overline{s}_1, \text{and } \underline{s}_2 \leq s_2 \leq \overline{s}_2\}$. Thus, in the perfect complements case, any improvements in test design improve both the sensitivity and specificity of the test.

To fix ideas, we could think of a setting where specificity and sensitivity are substitutes as one where the experimental design involves setting a temperature for experimentation on a particular chemical reaction. It might be that the chemical reaction is more likely to occur at higher temperatures, regardless of the quality of the technology being used, which would mean that there would a high number of true positives, but likewise a high number of false positives at higher temperatures. On the other hand there would be a high number of true and false negatives at lower temperatures. In this case, the specificity and the sensitivity of the experiment would negatively comove with each other.

A setting where specificity and sensitivity are complements could be one where the resolution of a camera used to identify the chemical reaction can be improved. At a higher resolution, false positives and false negatives would likely be lower meaning that the specificity and the sensitivity would positive comove with one another.

We are interested in situations where, absent any additional costs or frictions, it is socially desirable to run an experiment to test conditions such that a profitable venture can be realized in case of favourable test results, i.e. $\exists(s_1, s_2) \in S,$ such that $\pi_{s_1,s_2} > 0$. Given the above analysis, a necessary condition for this to be the case is that

$$p_0\overline{s}_1V - [p_0\overline{s}_1 + (1-p_0)(1-\overline{s}_2)]K - C > 0. \tag{4}$$

That is, under the most informative test with $s_1 = \overline{s}_1$ and $s_2 = \overline{s}_2$, it must be the case that running the experiment and developing the technology when the experiment is successful yields a positive net present value. We shall assume that condition (4) holds strictly.

The Contracting Problem While the investor only cares about the monetary payoff from her investment, the entrepreneur derives a non-pecuniary utility Z per period from working on the venture. The entrepreneur has no money and requires funding from the investor to carry out the experiment. We also assume that the investor can only use equity shares (in a successful venture) to incentivize the entrepreneur. The outside option for both the investor and the entrepreneur is 0. There is also no discounting.

Suppose that the investor and entrepreneur agree on a contract in period 0, which commits the investor to pay C for the experiment in period 1 and the entrepreneur to undertake an experiment. The contract can also specify an ownership stake α for the investor in the venture should the venture go ahead at the end of period 1 and an ownership stake $(1-\alpha)$ for the entrepreneur. However, the contract cannot specify the design (s_1, s_2) of the experiment because this is not describable.

Note, to satisfy the investor's participation constraint, a necessary condition is

$$\alpha V \geq K$$

With that, the first question we want to answer is, do the entrepreneur and the investor have conflicting objectives regarding the test design?

For any given $\alpha \in [\frac{K}{V}, 1]$, the investor has a utility function from an experiment characterized by s_1 and s_2,

$$U_I = p_0 s_1 \alpha V - [p_0 s_1 + (1 - p_0)(1 - s_2)]K - C \tag{5}$$

Likewise, the entrepreneur's utility function if he pursues the venture with an experiment characterized by s_1 and s_2 is given by

$$U_E = p_0 s_1 (1 - \alpha)V + [p_0 s_1 + (1 - p_0)(1 - s_2)]Z + Z$$

3.1 Solving the Contracting Problem: Choice of Experiment

What experiment design will the entrepreneur choose under a contract that provides funding C to run an experiment and gives an ownership stake $\alpha \in [\frac{K}{V}, 1]$ to the investor (a typical arrangement found in practice for VC financing)? We begin our analysis of the entrepreneur's test design problem by considering first the case of independent tasks.

Independent Tasks. The entrepreneur's test set is given by $S = [\underline{s}_1, \overline{s}_1] \times [\underline{s}_2, \overline{s}_2]$, and the entrepreneur can independently choose any test design in this set.

Suppose to begin with that the venture will go ahead if the technology successfully passes the test in the lab, for any given α, then the entrepreneur's optimal experiment design problem becomes:

$$\max_{s_1, s_2 \in S} p_0 s_1(1-\alpha)V + [p_0 s_1 + (1-p_0)(1-s_2)]Z + Z$$

It is easy to see that the entrepreneur's utility is increasing in s_1 and decreasing in s_2. Hence, it immediately follows that

Lemma 1. *With independent tasks, $\forall \alpha \in [\frac{K}{V}, 1]$, the entrepreneur's optimal experiment sets $s_1 = \overline{s}_1$ and $s_2 = \underline{s}_2$*

Proof: See the discussion above. ■

Irrespective of how much skin the entrepreneur has in the game, and no matter how small Z is (in fact, Z can be ϵ), the entrepreneur's preferred experiment design maximizes the probability of continuation. Indeed, the design where $s_1 = \overline{s}_1$ and $s_2 = \underline{s}_2$ maximizes the probability that the technology passes the test in the lab and that the venture will be undertaken. The entrepreneur seeks success in the lab because this ensures that the venture will go ahead and that the venture will be fully developed. The entrepreneur wants to reduce the risk that he is unable to continue working on the venture. No amount of skin in the game will undo this objective because the entrepreneur is not committing any funding to the venture. All the funding comes from the investor. Moreover, if the entrepreneur has a stake $(1-\alpha) > 0$ in the venture, he will only profit if the venture is developed and successful, which is conditional on passing the test in the lab.

Understanding this, the investor will not issue equity shares in the contract and anticipates that the experiment design chosen will differ from her preferred experiment. The test chosen by the entrepreneur exposes the investor to a greater risk of the venture failing conditional on the continuation and can significantly diminish the expected investment payoff. In particular, if

$$p_0\overline{s}_1V - [p_0\overline{s}_1 + (1-p_0)(1-\underline{s}_2)]K - C < 0. \tag{6}$$

Then, the investor will be unwilling to fund the experiment.

Comparing the left-hand side of conditions (4) and (6), the agency problem results in an expected financial loss of $(1-p_0)(\overline{s}_2 - \underline{s}_2)K$. Denote the range of sensitivity in the available set $(\overline{s}_2 - \underline{s}_2)$ as Δ_{s_2}, then we have,

Proposition 1. *With independent tasks,*

- *If condition (6) does not hold, then the investor will fund the experiment and provide follow-up funding conditional on passing the experiment, with no equity share issued to the entrepreneur ($\alpha = 1$).*
  - *The entrepreneur will choose the experiment with $s_1 = \overline{s}_1$ and $s_2 = \underline{s}_2$.*
  - *Conditional on passing the experiment, however, the venture would fail more often than the first best scenario, resulting in a venture valuation loss of $(1-p_0)\Delta_{s_2}K$.*
- *If condition (6) holds, there is a market failure: the investor will not make a deal with the entrepreneur, and no experiment is conducted.*

Note, if $\underline{s}_2 = 0$, a market failure is certain, as it can easily be checked that, $\pi_{s_1,s_2=0}$ as defined in equation 3, the net present value of running an experiment with zero sensitivity is strictly negative, regardless of the level of specificity. Combining this with the condition 4 that the most informative test yield positive NPV, we have the following result,

Proposition 2. *Suppose condition (4) holds. For any venture with fixed parameters K, V, C, and p_0, and allow the experiment set/technology space S to vary, then*

- *$\exists\Delta^{\star}_{s_2}$, such that if $\Delta_{s_2} > \Delta^{\star}_{s_2}$, then we are in the scenario of market failure even though absent moral hazard problem, the venture has positive NPV. And if $\Delta_{s_2} \leq \Delta^{\star}_{s_2}$, the experiment is funded but is inefficiently conducted.*

- *If $\Delta_{s_2} \leq \Delta^{\star}_{s_2}$, the inefficiency cost resulting from the less conclusive experiment being chosen is proportional to Δ_{s_2}.*

Substitute Tasks. The entrepreneur's test set is now given by $S = \{(s_1, s_2) \mid 0 \leq s_i \leq \kappa \text{ and } s_1 + s_2 = \kappa\}$, with $\overline{s}_1 + \overline{s}_2 > \kappa$. Suppose again that the venture will go ahead if the technology successfully passes the test in the lab. For any given α, the entrepreneur's optimal experiment design problem is then given by:

$$\max_{s_1} s_1 p_0[(1-\alpha)V + Z] + (1-p_0)(1 - \max(\kappa - s_1, \underline{s_2}))Z + Z$$

The entrepreneur, a fortiori, now always maximizes s_1 so that we now have $s_1 = \overline{s}_1$ and $s_2 = \max(\kappa - \overline{s}_1, \underline{s_2})$, regardless of α.

Understanding that, the investor will not issue equity shares and hence $\alpha = 1$. Now the investor's utility function conditional on funding the experiment with s_1 and $s_2 = \kappa - s_1$ becomes,

$$U_I = p_0 s_1 V - [p_0 s_1 + (1-p_0)(1 - \max(\kappa - s_1, \underline{s_2})]K - C \tag{7}$$

In this case, the investor's preferred experiment depends on whether the substitution bites at the corner solution. Specifically,

Proposition 3. *With substitute tasks, $\forall\alpha$, the entrepreneur's optimal experiment sets $s_1 = \overline{s}_1$ and $s_2 = max(\kappa - \overline{s}_1, \underline{s_2})$, while the investor's prefered experiment varies:*

- *If $\kappa - \overline{s}_1 \leq \underline{s}_2$, the investor's preferred experiment is the same as that of the entrepreneur's, that is, $s_1 = \overline{s}_1$ and $s_2 = \underline{s}_2$*

- *If $\kappa - \overline{s}_1 > \underline{s}_2$, the investor's preferred experiment vastly differs from that of the entrepreneur's, that is, $s_1 = \underline{s}_1$ and $s_2 = \kappa - \underline{s}_1$*

Proof: Basic algebra. ■

Intuition: the investor and the entrepreneur are fundamentally aligned on the parameter choice of specificity. High specificity per se increases both the likelihood of passing the test and the expected payoff. However, the investor and the entrepreneur are fundamentally misaligned on the parameter choice of sensitivity, as demonstrated by the independent tasks case. This is because higher sensitivity increases the value of the experiment and hence the expected payoff from investment by reducing the false positive rate and decreasing the likelihood of passing the test. Therefore, choosing the highest likelihood of a pass signal does not come at a huge cost if the available specificity is sufficiently high and on the margin, increasing specificity does not lower sensitivity.

However, if increasing specificity comes at the cost of lowering sensitivity on the margin, then the investor would prefer a more sensitive test to screen off bad apples, while the entrepreneur, as analyzed before, would always prefer a less sensitive one. The investor is even more worse off in this situation than in the independent tasks case.

Complementary Tasks. The entrepreneur's test set is now given by $S = \{(s_1, s_2) \mid \underline{s}_1 \leq s_1 \leq \overline{s}_1 \text{ and } s_2 = \underline{s}_2 + \lambda s_1, \underline{s}_2 \leq s_2 \leq \overline{s}_2\}$, where $\lambda \in (0, 1)$. For simplicity, suppose $\underline{s}_2 + \lambda s_1 \leq \overline{s}_2$. If the venture goes ahead conditional on the technology successfully passing the test in the lab, for any given α, now the entrepreneur's optimal experiment design problem becomes:

$$\max_{s_1} s_1 p_0[(1-\alpha)V + Z] + (1-p_0)Z(1-\underline{s}_2 - \lambda s_1) + Z$$

Differentiating with respect to s_1, we obtain that the entrepreneur maximizes s_1 (and therefore also s_2 as much as possible) if and only if

$$p_0[(1-\alpha)V + Z] - (1-p_0)\lambda Z > 0 \tag{8}$$

When $(1-\alpha) = 0$, this condition reduces to

$$p_0 Z - (1-p_0)\lambda Z > 0.$$

Note that if $p_0 < 1/2$ and λ is close to 1, this condition is violated, so that the entrepreneur is willing to choose a maximally sensitive (and specific) test only if he has sufficient skin in the game.

We summarize this discussion in the proposition below:

Proposition 4. *When (s_1, s_2) are complementary tasks, the entrepreneur's optimal experiment sets $s_1 = \overline{s}_1$ and $s_2 = \underline{s}_2 + \lambda\overline{s}_1$ if and only if the entrepreneur has sufficient skin in the game that*

$$1-\alpha \geq \frac{Z(\lambda(1-p_0) - p_0)}{p_0 V}. \tag{9}$$

Proof: See the discussion above. ■

It follows from condition (9) that the more the entrepreneur values doing science (the higher is Z) the more the entrepreneur must be financially rewarded to design a more conclusive test (such that $s_1 = \overline{s}_1$ and $s_2 = \underline{s}_2 + \lambda\overline{s}_1$). A more conclusive test may reveal that the technology the entrepreneur has studied is a dead end (at least as far as the application the VC is interested in is concerned). This would put any future funding to do more science at risk. Under low-powered financial incentives the entrepreneur therefore prefers to design a test that is more likely to pass, and would assure the continuation of his future scientific endeavours, than take the risk that a more stringent (but more conclusive) test might fail. Similarly, the more conclusive the test is (as measured by a higher λ) the more the entrepreneur must be rewarded financially.

However, it can be profitable for the VC to finance the experiment for some parameter values when the two tasks are complementary, whereas such investment will not be worth it if s_1 and s_2 are independent or substitute tasks.

Proposition 5. *When (s_1, s_2) are complementary tasks, the investor is willing to fund the experiment if*

$$\overline{s}_1(p_0V - (Z - K)(\lambda(1 - p_0) - p_0)) \geq [((1 - p_0)(1 - \underline{s}_2)]K + C \tag{10}$$

Proof: The investor must provide a share of the final value of the venture $(1 - \alpha)$ to the entrepreneur such that condition (9) holds. Assuming that this condition is binding, and substituting for α in the condition below

$$p_0s_1\alpha V - [p_0s_1 + (1 - p_0)(1 - s_2)]K \geq C,$$

we obtain condition (10), which ensures that with minimum skin in the game for the entrepreneur (so that he has an incentive to choose a maximally conclusive test design), the investor at least breaks even in expectation by funding the experiment, and subsequently the venture, should the experiment produce a positive test outcome. ■

3.2 Paying for Failed Test

One reason why there is a deep tech market failure is that the financial compensation structure of the entrepreneur is inadequate. In essence, the entrepreneur is rewarded if the venture is implemented. He gets a share of the realized value (even if only a small one) if the technology is successful and he faces no downside risk, as he does not put up any money to fund the venture. What is more, the entrepreneur derives private benefits from doing science and gets rewarded with more private benefits if he can continue to do science. In contrast, the investor faces all the downside risk and has a risky upside, particularly if the lab test is not very conclusive.

How can the entrepreneur be given better incentives to design a more conclusive experiment? We show next that the entrepreneur can be given incentives to choose the most conclusive experiment design $s_1 = \overline{s}_1$ and $s_2 = \overline{s}_2$ if he gets compensated for a failed test rather than for passing a test in the lab.

Independent Tasks. Consider first the case where s_1 and s_2 are independent tasks. If compensation is based on *proof of failure*, then the entrepreneur's objectives could be better aligned with those of the investor. Concretely, suppose that the contract between the entrepreneur and the investor, based on Proposition 1, includes a payment $X > 0$ conditional on the outcome $s = F$. The entrepreneur's best response in the experiment design problem in period 1 is then the solution to the following maximization problem:

$$\max_{s_1, s_2 \in S} p_0 s_1 Z + (1 - p_0)(1 - s_2)Z + [p_0(1 - s_1) + (1 - p_0)s_2)]X + Z \tag{11}$$

Differentiating with respect to s_1 and s_2, the following proposition immediately obtains.

Lemma 2. *The entrepreneur chooses $s_1 = \overline{s}_1$ and $s_2 = \overline{s}_2$ if $X \geq Z$.*

Proof: Collecting terms we obtain that the entrepreneur's objective function (11) can be written as

$$s_1 p_0[Z - X] + s_2(1 - p_0)[X - Z] + p_0 X + (1 - p_0)Z + Z \tag{12}$$

In effect, the payment $X \geq Z$ for producing the outcome $s = F$ fully compensates for the opportunity cost of getting a negative test result. Therefore, by rewarding a failed test, the VC can ensure that the entrepreneur is no longer distorted away from the most conclusive test possible. The investor may be willing to finance the experiment under these terms, and fund the venture should the technology successfully pass the test.

Proposition 6. *When (s_1, s_2) are independent tasks, the investor is willing to fund an experiment with compensation for failed test $X = Z$ if*

$$p_0[\bar{s}_1(V - K) - (1 - \bar{s}_1)Z)] - (1 - p_0)[\bar{s}_2 Z + (1 - \bar{s}_2)K] - C \geq 0 \tag{13}$$

Proof: When condition (13) holds, the investor at least breaks even in expectation when financing the entrepreneur who will, according to Lemma 2, conduct the most informative experiment when compensated for failed test. ■

Comparing condition (13) with condition (4), by compensating for a failed test, the investor is in expectation spending $[p_0(1 - \bar{s}_1) + (1 - p_0)\bar{s}_2]Z$ more than the first best case. Given $Z << Y$, the private benefit Z to the entrepreneur from doing science is likely to be significantly smaller than the net present value in the first best scenario, therefore condition (13) can almost be satisfied for free.

If we compare condition (13) with condition (6), the gross gain for the investor from compensating for the failure, is total inefficiency cost $(1 - p_0)(\bar{s}_2 - \underline{s}_2)K$. If $Z << K$, the net gain is large, thus, compensating for the failed test might just be what tilts the venture from a negative to a positive NPV venture for the investor.

Substitute Tasks. When tasks are perfect substitutes, it is not possible to provide incentives to the entrepreneur to both maximize s_1 and s_2. It is one or the other. This is the typical multitask moral hazard problem.

As Holmstrom and Milgrom (1991) have argued, providing low-powered incentives to the entrepreneur may be the best solution for the investor. If $\underline{s}_1$ is high enough, it may indeed pay he VC not to give any skin in the game to the entrepreneur, but to reward the entrepreneur for proof of failure. This would induce the entrepreneur to maximize the *sensitivity* of the experiment (giving up on improving the specificity of the experiment), which would benefit the VC by reducing the risk of false positives.

Proposition 7. *When (s_1, s_2) are substitute tasks, the investor is willing to fund an experiment with a reward for proof of failure $X = Z$ if*

$$p_0[\underline{s}_1(V - K) - (1 - \underline{s}_1)Z)] - (1 - \xi_0)\pi_0[\overline{s}_2 Z + (1 - \overline{s}_2)K] - C \geq 0 \tag{14}$$

Proof: When condition (14) holds the VC at least breaks even in expectation when funding a conclusive test that rewards the entrepreneur for proof of failure and maximizes the *sensitivity* of the experiment. ■

Complementary Tasks. When tasks are perfect complements it is possible for the VC to give incentives to the entrepreneur to maximize both s_1 and s_2 by providing either sufficient skin in the game or by rewarding proof of failure. Providing sufficient skin in the game requires condition (9) to hold, which could be onerous for the VC. Alternatively, if the VC rewards proof of failure, the entrepreneur has sufficient incentives to maximize both s_1 and s_2 if $X = Z$, which could be much cheaper for the VC.

Proposition 8. *When (s_1, s_2) are complementary tasks, it is cheaper to reward the entrepreneur for proof of failure if*

$$p_0(1 - \overline{s}_1) + (1 - p_0)(1 - \underline{s}_2 - \lambda\overline{s}_1) \leq \lambda(1 - p_0) - p_0 \tag{15}$$

Proof: When $X = Z$ the entrepreneur is indifferent between any $(s_1, s_2) \in S$ since irrespective of the outcome of the experiment the entrepreneur obtains Z. If the technology passes the test, the entrepreneur can continue to do science and obtains Z in kind. If the technology fails the test, the entrepreneur is rewarded financially the amount Z for proof of failure. When indifferent, the entrepreneur can be assumed to choose the test design that is best for the VC. Under this test design the VC pays the entrepreneur Z with probability $p_0(1 - \overline{s}_1) + (1 - p_0)(1 - \underline{s}_2 - \lambda\overline{s}_1)$. If the VC instead provides skin in the game incentives, then she must grant the entrepreneur a share of the value of the venture

$$1 - \alpha = \frac{Z(\lambda(1 - p_0) - p_0)}{p_0 V},$$

which is worth ex-ante

$$(1 - \alpha)p_0 V = p_0 V[\frac{Z(\lambda(1 - p_0) - p_0)}{p_0 V}] = Z(\lambda(1 - p_0) - p_0).$$

It is straightforward to verify that when condition (15) holds, this is more expensive for the VC. ■

4 Policy Responses

When optimal contracting between the VC and the entrepreneur/scientist is unable to overcome the moral hazard problem in experiment design, third-party institutional interventions can be useful. In this section, we analyze several plausible interventions with independent tasks only.

4.1 The Third Party (Planner's) Problem

We have seen that the VC prefers the most informative experiment design, while the entrepreneur/scientist prefers experiment designs more likely to result in a positive test outcome. What is the planner's preferred experiment design and incentive contract with the entrepreneur/scientist and VC? The social planner's objective is to maximize expected social surplus from the experiment:

$$U_{SP} = p_0 s_1 V - [p_0 s_1 + (1 - p_0)(1 - s_2)](K - Z) - C + Z.$$

If the social planner could pick s_1 and s_2 freely, he would therefore set $s_1 = \overline{s}_1$ and $s_2 = \overline{s}_2$ - the same choice the VC would make. It follows that when the VC incentivizes the entrepreneur to choose her preferred design by X to pay for validation, she induces the first-best outcome. This outcome could also be implemented as a conditional transfer made by the social planner to the entrepreneur. A payment to the entrepreneur of Z (the smallest payment to induce an informative experiment) would have a PV of $(p_0(1-\overline{s}_1)+(1-p_0)\overline{s}_2)Z \equiv Z_{PV}$. The social planner could also instead of making a conditional transfer of Z upon a failed experiment, incentivize the entrepreneur by making a conditional transfer to increase V in case of success, or an unconditional transfer to reduce the cost of the experiment C.

Setting the PV of the three possibilities equal to each other, we can measure the impact of spending Z_{PV} in three different ways:

1. Impact via validation: If the planner makes a conditional payment of Z_{PV} to the entrepreneur if the technology fails the experiment, the social benefit is $(1-p_0)(\overline{s}_2-\underline{s}_2)K$

2. Impact via C reduction: If the planner subsidizes the cost of running the experiment by Z_{PV}, the social benefit is also Z_{PV}.

3. Impact via V increase: If the planner subsidizes the benefit of the technology, shoul dit prove viable by Z_{PV}, the social benefit is $p_0 s_1 Z_{PV}$.

Reducing cost strictly dominates increasing reward, as the cost must be paid, and the reward may never arrive. For sufficiently small values of Z, paying for validation has an even bigger impact on welfare than subsidizing the experiment.

4.2 The Role of Universities as Venture Incubators

In this subsection, we explore the role an intermediary, such as a university, can play in reducing moral hazard, absent the option discussed in the previous subsection.

The moral hazard problem arises from the knowledge gap between the investor and the entrepreneur around the informativeness of the experiment. When the investor does not have the capability to effectively validate the information contained in the experiments conducted by the entrepreneur, moral hazard can result in a complete market failure or at least vastly reduce the learning efficiency.

A potential solution to this problem is to involve a third-party organization that is able and willing to verify the informativeness of the experiment. The university is a natural candidate for playing such a crucial role. A university houses experts across diverse fields who can provide peer reviews, ensuring that the experiment's methodology, data analysis, and conclusions meet high standards. Universities also have state-of-the-art laboratories and equipment that allow for precise replication and verification of experiments. By leveraging academic expertise, resources, and collaborative frameworks, universities have the capability to ensure that scientific experiments are thoroughly vetted and validated. As implied by Proposition 2, the problem of market failure can be resolved, if the validation provided by the university can reduce the wide range of possible sensitivity measures $\Delta_{s_2} = \bar{s}_2 - \underline{s}_2$ from a value above $\Delta^{\star}_{s_2}$ to a value below. Moreover, the more effective the validation (the higher reduction of $\Delta^{\star}_{s_2}$), the higher the valuation of the venture.

Why is it in the university's interest to provide such objective, high-quality validation while the individual entrepreneur/scientist cannot commit to doing so? The universities compete with each other to attract funding for cutting-edge research and the commercialization of innovations. Hence, a university values the credibility of being able to effectively validate early experiments, which directly makes the university more competitive in the market for financial resources and, at the same time, reflects the academic excellence of the institution, strengthening its academic reputation. The effect is particularly pronounced around areas such as "Deep Tech" where such moral hazard problem is most severe (Δ_{s_2} is high, hence the odds are market failure absent validation). Another reason why the university can and is willing to develop a credible reputation while an individual entrepreneur/scientist cannot is because the university is and will be engaged in repeated interaction with the financiers, while an individual entrepreneur/scientist likely will interact only once or limited times with the market.

4.3 Learning Spillovers and Sequential Learning

In our main analysis, we have assumed for simplicity that the venture is a standalone project, and the experiment result is relevant only to the venture itself. In practice, however, learning from an experiment aiming to test the technical feasibility of one project may have a knowledge spillover effect for a host of other projects.

We model this more general problem in the following way. Suppose that the feasibility of the venture's technology relies on the success of two separate and independent components: a component G, which is a *general purpose component* and can be common to many different startup enterprises, and a component r which is a *application specific component* and is idiosyncratic to the entrepreneur's chosen venture. We assume that each component can take the value 1 if it is workable in practice and 0 if it is not. Specifically, we assume that $0 < P(G = 1) \equiv \pi_0 < 1$ and $0 < P(r = 1) \equiv \xi_0 < 1$. Therefore, the technology's development is successful with probability $p_0 \equiv \pi_0 \xi_0$.

Each experiment is characterized again, by respectively s_1 the probability that the test succeeds for a technology that is workable in practice, and s_2 the probability that the test fails when the technology is not workable in practice, with the following experimental technology:

$$
\begin{aligned}
P(s = P|G = 1, r = 1) &= s_1 \\
P(s = F|G = 1, r = 0) &= s_2 \\
P(s = F|G = 0) &= 1
\end{aligned}
$$

Note that s_1 and s_2 only depend on the idiosyncratic component, r. If $G = 0$, the experiment fails.

This setup captures the idea that new technologies developed for one project often have other potential applications. In the presence of multiple potential projects, an experiment can reveal information about both the promise of a new technology for a specific project and the more general applicability to other projects.

The entrepreneur can jointly choose the parameter values (s_1, s_2) within a set S of possible test designs. The best experiment design in terms of information discovery is the one with the greatest specificity (the highest s_1) and the greatest sensitivity (the highest s_2).

Accordingly, suppose a planner is considering investments in potentially $N+1$ different projects, all based on the same general purpose component G. The specific applications of each project have the same ξ_0 and the same learning technology. Hence, the projects, though different, are ex-ante identical from a learning perspective. The planner can learn sequentially about the viability of each project. For simplicity, assume that the planner can randomly pick one entrepreneur to run their experiment first and then have the remaining N entrepreneurs run their experiments simultaneously in the next period.

Conditional on a certain s_1 and s_2, the planner's posterior distribution on G after a pass or fail signal is:

$$
\begin{aligned}
P(G = 1|s = P) &= 1 \\
P(G = 1|s = F) &= \frac{(\xi_0(1 - s_1) + (1 - \xi_0)s_2)\pi_0}{(\xi_0(1 - s_1) + (1 - \xi_0)s_2)\pi_0 + (1 - \pi_0)} \equiv \pi_1 < \pi_0
\end{aligned}
$$

Effectively, a pass signal indicates that the general component *must* be workable as the experiment would have failed otherwise, while a fail signal could indicate that the general component was not workable, or that the entrepreneur's application was faulty. This contrasts with the setup of the previous section, as a pass signal now has broader usefulness than just reflecting the viability of the entrepreneur's application.

The planner's posterior distribution on r after a pass or fail signal is:

$$\begin{aligned} P(r=1|s=P) &= \frac{\pi_0 s_1 \xi_0}{\pi_0\xi_0 s_1 + \pi_0(1-\xi_0)(1-s_2)} \\ P(r=1|s=F) &= \frac{((1-\pi_0)+\pi_0(1-s_1))\xi_0}{(\xi_0(1-s_1)+(1-\xi_0)s_2)\pi_0 + (1-\pi_0)} \end{aligned}$$

As before, a pass signal does not completely validate the entrepreneur's tech nor does a fail signal completely invalidate it, but now the presence of the general component reduces the ability of all parties to infer the idiosyncratic quality of the tech.

Suppose the social planner could pick s_1 and s_2 directly. The social planner's utility takes the form:

$$\begin{aligned} U_P \;=\; & \underbrace{U(\pi_0, s_1, s_2)}_{\text{Utility from first venture}} + \underbrace{((\xi_0(1-s_1)+(1-\xi_0)s_2)\pi_0 + (1-\pi_0))}_{\text{Probability of a `fail' signal}} N\max(U(\pi_1, s'_{1F}, s'_{2F}), 0) \\ & + \underbrace{(\pi_0\xi_0 s_1 + \pi_0(1-\xi_0)(1-s_2))}_{\text{Probability of a `pass' signal}} NU(1, s'_{1P}, s'_{2P}) \end{aligned}$$

If the initial venture's experiment yields a 'pass' signal, the general component is validated, and enters subsequent experiments with surety (probability of 1 instead of π_0). If the initial venture's experiment yields a 'fail' signal, the general component might be unworkable, or it might not, so the VC has a choice of funding the other N projects, or abandoning them.

Deriving U_P with respect to s_1 we get:

$$\begin{aligned} U'_P &= p_0(V-K)+\xi_0\pi_0 N(U_I(1,s'_{1P},s'_{2P})-U_I(\pi_1,s'_{1F},s'_{2F})) \\ &\quad +\frac{(\pi_0 s_1+(1-\xi_0)(1-s_2))K-\xi_0 s_1\alpha V}{(\xi_0(1-s_1)+(1-\xi_0)s_2)\pi_0+(1-\pi_0)}N\xi_0\pi_0((1-\pi_0)) \end{aligned}$$

The above expression is always positive, so $s_1=\overline{s}_1$. The social planner has no incentive to reduce the true positive rate. Having a higher s_1 not only improves the informational content of the experiment as a whole, but also allows for a better inference on the general component of the technology. Differentiating U_P with respect to s_2 we get:

$$\begin{aligned} U'_P &= (1-\xi_0)\pi_0 K+(1-\xi_0)\pi_0 N(U_I(\pi_1,s'_{1P},s'_{2P})-U_I(1,s'_{1F},s'_{2F})) \\ &\quad -\frac{(\pi_0 s_1+(1-\xi_0)(1-s_2))K-\xi_0 s_1\alpha V}{(\xi_0(1-s_1)+(1-\xi_0)s_2)\pi_0+(1-\pi_0)}(1-\pi_0)(1-\xi_0)\pi_0 NK \end{aligned}$$

This expression is decreasing in N and will be negative for N sufficiently large.

Proposition 9. *$\exists N^*$ such that $\forall N \leq N^*$, the planner wants $s_1=\overline{s}_1$ and $s_2=\overline{s}_2$, and $\forall N > N^*$, the planner wants $s_1=\overline{s}_1$ and $s_2=\underline{s}_2$. The entrepreneur's optimal $s_1=\overline{s}_1$ and $s_2=\underline{s}_2$ is invariant in N.*

Proof: See the discussion above. ■

Intuitively, having a higher false positive rate reduces the informational content of the experiment, which consequently reduces the planner's ability to infer the quality of the specific application to the first project being considered, but a positive signal necessarily means that the general purpose technology is workable, so even a *false* positive, which obfuscates the quality of the idiosyncratic experiment, is valuable for determining the viability of the general component. Therefore if N is larger, the planner has a more congruent objective with the entrepreneur's. If N is smaller, the planner's objective is more congruent with the VC's. The value of an informative signal about the common component is increasing in the number of projects that depend on that common component.

5 Conclusion

Given the huge uncertainty in the outcomes associated with new technologies, investors financing these innovations engage in staged-financing which is equivalent to financing a sequence of experiments over time. It is striking, however, that some of these early experiments appear to be particularly poor at predicting whether a technology will work at scale. The poor predictability of these early experiments also seems correlated with a lack of venture capital financing, despite great potential societal need.

In this paper, we characterize early experiments in terms of their ability to capture true positives and true negatives, and conversely the degree to which they lead to false negatives or false positives. We highlight a novel source of moral hazard for the entrepreneur, that leads them to design experiments that are more inconclusive or have lower learning efficacy. We show this moral hazard can lead to large inefficiencies, including a lower likelihood of getting funded and among those that do get funded, leading to greater failure (relative to a benchmark without this friction) at later stages. From a theoretical perspective, we show that the nature of the moral hazard we identify cannot be easily addressed through 'skin in the game' that is able to align incentives in many such principal-agent models. It requires the principal to pay for failure, although this solution is fragile as it can lead to 'you get what you pay for' with the entrepreneur designing experiments that have too many false negatives that lead the experiment to fail even if the project is viable.

Our model provides a complementary explanation for the lack of funding for deep tech ventures building on fundamental science, where it is harder for investors to validate experiment designs. It also creates an understanding of how universities might commit to being the bodies that validate the experiment designs from early de-risking experiments, and how this can help alleviate the market failure arising from the friction.

References

Aghion, Philippe and Patrick Bolton. 1992. "An incomplete contracts approach to financial contracting." *The review of economic Studies* 59 (3):473–494.

Agrawal, Ajay, Joshua S Gans, and Scott Stern. 2021. "Enabling entrepreneurial choice." *Management Science* 67 (9):5510–5524.

Arora, Ashish, Andrea Fosfuri, and Thomas Roende. 2022. "Caught in the middle: the bias against startup innovation with technical and commercial challenges." Tech. rep., National Bureau of Economic Research.

Balasubramanian, Parasuram, Lamar Pierce, and Trey Cummings. 2022. "Research Validity Across Organizational Forms: Evidence from Phase 2 Oncology Clinical Trials." *Technical Article* .

Bates, Stephen, Michael I Jordan, Michael Sklar, and Jake A Soloff. 2022. "Principal-agent hypothesis testing." *arXiv preprint arXiv:2205.06812* .

———. 2023. "Incentive-Theoretic Bayesian Inference for Collaborative Science." *arXiv preprint arXiv:2307.03748* .

Bergemann, Dirk and Ulrich Hege. 1998. "Venture capital financing, moral hazard, and learning." *Journal of Banking & Finance* 22 (6-8):703–735.

———. 2005. "The financing of innovation: Learning and stopping." *RAND Journal of Economics* :719–752.

Camuffo, Arnaldo, Alessandro Cordova, Alfonso Gambardella, and Chiara Spina. 2020. "A scientific approach to entrepreneurial decision making: Evidence from a randomized control trial." *Management Science* 66 (2):564–586.

Camuffo, Arnaldo, Alfonso Gambardella, Fabio Maccheroni, Massimo Marinacci, and Andrea Pignataro. 2022. "Microfoundations of low-frequency high-impact decisions." *CEPR Discussion Paper No. DP17392* .

Chassang, Sylvain, Gerard Padró i Miquel, and Erik Snowberg. 2012. "Selective trials: A principal-agent approach to randomized controlled experiments." *American Economic Review* 102 (4):1279–1309.

Cornelli, Francesca and Oved Yosha. 2003. "Stage financing and the role of convertible securities." *Review of Economic Studies* 70 (2):1–32.

Dalla Fontana, Silvia and Ramana Nanda. 2023. "Innovating to Net Zero: Can Venture Capital and Start-Ups Play a Meaningful Role?" *Entrepreneurship and Innovation Policy and the Economy* 2 (1):79–105.

Di Tillio, Alfredo, Marco Ottaviani, and Peter Norman Sørensen. 2017. "Persuasion bias in science: can economics help?"

Fosfuri, Andrew and Jay Prakash Nagar. 2023. "timing is key: navigating venture capital funding for science-based startup." ().

Fréchette, Guillaume R, Alessandro Lizzeri, and Jacopo Perego. 2022. "Rules and commitment in communication: An experimental analysis." *Econometrica* 90 (5):2283–2318.

Gans, Joshua S, Scott Stern, and Jane Wu. 2019. "Foundations of entrepreneurial strategy." *Strategic Management Journal* 40 (5):736–756.

Gompers, Paul. 1995. "Optimal Investment, monitoring, and the staging of venture capital." *Journal of Finance* 50:1461–1490.

Greenwood, Matthew, Jens Matthies Wrogemann, Richard Schmuch, Hwamyung Jang, Martin Winter, and Jens Leker. 2022. "The Battery Component Readiness Level (BC-RL) framework: A technology-specific development framework." *Journal of Power Sources Advances* 14:1–16.

Guedj, Ilan and David S Scharfstein. 2004. "Organizational scope and investment: Evidence from the drug development strategies and performance of biopharmaceutical firms."

Hall, Bronwyn H and Josh Lerner. 2010. "The financing of R&D and innovation." In *Handbook of the Economics of Innovation*, vol. 1. Elsevier, 609–639.

Hellmann, Thomas. 1998. "The allocation of control rights in venture capital contracts." *RAND Journal of Economics* 29:57–76.

Hellmann, Thomas and Manju Puri. 2000. "The interaction between product market and financing strategy: The role of venture capital." *Review of Financial Studies* 13:959–984.

Henry, Emeric and Marco Ottaviani. 2019. "Research and the approval process: The organization of persuasion." *American Economic Review* 109 (3):911–955.

Hingorani, Aroon D, Valerie Kuan, Chris Finan, Felix A Kruger, Anna Gaulton, Sandesh Chopade, Reecha Sofat, Raymond J MacAllister, John P Overington, Harry Hemingway et al. 2019. "Improving the odds of drug development success through human genomics: modelling study." *Scientific reports* 9 (1):18911.

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian persuasion." *American Economic Review* 101 (6):2590–2615.

Kaplan, Steven N and Per Strömberg. 2003. "Financial contracting theory meets the real world: An empirical analysis of venture capital contracts." *The review of economic studies* 70 (2):281–315.

Karp, Rebecca, Orie Shelef, and Robert Wuebker. 2024. "Business Experiments As Persuasion." ().

Kim, C Kwon, Yin Rui Lee, Lynnett Ong, Michael Gold, Amir Kalali, and Joydeep Sarkar. 2022. "Alzheimer's disease: key insights from two decades of clinical trial failures." *Journal of Alzheimer's Disease* 87 (1):83–100.

Kolotilin, Anton. 2015. "Experimental design to persuade." *Games and Economic Behavior* 90:215–226.

Kremer, Michael, Jonathan Levin, and Christopher M Snyder. 2022. "Designing advance market commitments for new vaccines." *Management Science* 68 (7):4786–4814.

Lavin, Alexander, Ciarán M Gilligan-Lee, Alessya Visnjic, Siddha Ganju, Dava Newman, Sujoy Ganguly, Danny Lange, Atílím Güneş Baydin, Amit Sharma, Adam Gibson et al. 2022. "Technology readiness levels for machine learning systems." *Nature Communications* 13 (1):6039.

Lerner, Josh and Ramana Nanda. 2020. "Venture Capital's Role in Financing Innovation: What We Know and How Much We Still Need to Learn." *Journal of Economic Perspectives* 34 (3).

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

Paul, Steven M., Daniel S. Mytelka, Christopher T. Dunwiddie, Charles C. Persinger, Bernard H. Munos, Stacy R. Lindborg, and Aaron L. Schacht. 2010. "How to improve

R&D productivity: the pharmaceutical industry's grand challenge." *Nature Reviews Drug Discovery* 9:203–14.

Purohit, Abhishek, Maninder Kaur, Zeki Can Seskir, Matthew T Posner, and Araceli Venegas-Gomez. 2024. "Building a quantum-ready ecosystem." *IET Quantum Communication* 5 (1):1–18.

Siegmund, Daniel, Sebastian Metz, Volker Peinecke, Terence E. Warner, Carsten Cremers, Anna Grevé, Tom Smolinka, Doris Segets, and Ulf-Peter Apfel. 2021. "Crossing the Valley of Death: From Fundamental to Applied Research in Electrolysis." *JACS Au* 1 (5):527–535.