# Self-Discovery in Boom–Bust Investment

Christo A. Pirinsky Michael C. Tseng*

November 2026

**Abstract**

We analyze investment decisions when firms are uncertain about their own abilities to develop a new technology. Investing exposes firms to downside risk but generates self-knowledge about their fit with the technology, a process we refer to as introspective learning. This prospect of discovering their comparative advantage incentivizes firms to invest more and earlier. In the aggregate, learning about idiosyncratic fit gives rise to boom–bust investment patterns. The boom–bust is stronger when the technology engages broader skill sets and when firms are more impatient. Interestingly, external financing dampens rather than amplifies the boom–bust: because the benefits of discovery accrue to entrepreneurs, outside investors ration their capital.

*JEL classification*: D83; G31; O33.

*Keywords*: Introspective learning, technology, innovation, bubbles, matching.

*Christo A. Pirinsky is with the Department of Finance, University of Central Florida (cpirinsky@ucf.edu). Michael C. Tseng is with the Department of Economics, University of Central Florida (michael.tseng@ucf.edu). We are grateful to Vlad Gatchev, Indrajit Mitra, and participants at the 21st International Behavioral Finance Conference for helpful comments and discussions. All remaining errors are our own.

# 1 Introduction

The commercialization of new technologies is often characterized by boom–bust investment cycles: massive initial investment creates substantial excess capacity, followed by a sharp contraction. During the dot-com boom of the late 1990s and early 2000s, telecommunications companies invested hundreds of billions of dollars in laying millions of miles of fiber-optic cable, an estimated 85–95% of which initially remained unused. Similar patterns accompanied the railroad expansion of the 1870s and the electrification boom of the 1920s. The recent surge in investment in artificial intelligence (AI) may exhibit similar dynamics. These recurring episodes raise fundamental questions about the determinants of investment in new technologies, the optimality of firms' investment decisions, and their implications for social welfare.

Boom–bust investment in emerging technologies is often attributed to imperfections in external capital markets and investors' tendency to pursue speculative opportunities (e.g., Shiller (2000); Polk and Sapienza (2009); Smith et al. (2014)). Yet firms and entrepreneurs may have their own reasons to invest aggressively in unfamiliar technologies, independently of speculative incentives in capital markets. A long tradition in the literature on industry dynamics and experimentation recognizes that firms may not know ex ante whether they possess a comparative advantage in exploiting a new technological opportunity. Such advantages are often latent and can be discovered only through experimentation (Jovanovic (1982); Jovanovic and Lach (1989); Drugov and Macchiavello (2014)). By developing and testing new projects, firms learn whether their existing capabilities are suited to an emerging market and, if so, where their competitive advantage lies (Jovanovic and MacDonald (1994)). Thus, investment may serve not only to exploit an opportunity but also to discover whether the firm is particularly well positioned to do so. We refer to this motive as *introspective learning*.

We examine this self-discovery motive in a stylized model in which firms with heterogeneous bundles of know-how confront a new technology. Firms do not know ex ante whether, or where, their existing capabilities fit within the emerging set of investment opportunities. In AI, for example, these opportunities may include agentic service platforms, medical devices, drug-discovery platforms, and new outcome-based software businesses.[1] Firms can discover whether their capabilities are suited to a particular activity only by committing resources to it. If a firm's capabilities match the project, the investment generates a high return that can be replicated in subsequent periods. If they do not, the firm loses its investment but learns that its capabilities are poorly suited to that particular activity. Experimentation

[1] On these respective directions, see Dharani et al. (2025), FDA (2026), Zitnik (2025), and Khanna et al. (2025).

therefore has a dual payoff: it may generate immediate returns, but it also produces valuable information about the firm's own comparative advantage.

We show that introspective learning alone can generate boom–bust investment dynamics, characterized by a front-loaded investment surge followed by high failure and exit rates. Early in the adoption of a new technology, many firms invest to discover their individual fit within the emerging opportunity set. As experimentation unfolds, firms that remain in the market without having found a suitable match become an increasingly selected group, disproportionately composed of firms with no comparative advantage in the available opportunities. Meanwhile, firms that discover a favorable match early can scale up their investments and establish a strong position before unsuccessful experiments and capital depreciation erode the resources of other participants. The resulting dynamics combine broad initial participation with an endogenous shakeout and increasing concentration among successful firms.

The model also identifies technological characteristics that shape the intensity of these dynamics. In particular, the boom–bust pattern is stronger for more inclusive technologies—those that can be exploited using a broader range of existing skills. When barriers to experimentation are low, a larger population of firms has an incentive to participate early in the adoption process. This broad participation amplifies the initial investment boom but also generates a more pronounced subsequent shakeout as firms discover that their capabilities are not well suited to the available opportunities. By contrast, technologies requiring highly specialized expertise attract a narrower set of potential entrants and generate less extensive exploratory investment. Thus, the model suggests that skill-intensive technologies (such as AI) may be less susceptible to bubblelike overinvestment than technologies with lower skill barriers to entry (such as real estate).

The intensity of the boom–bust cycle also depends on firms' intertemporal incentives and access to capital. More impatient entrepreneurs place greater value on obtaining information sooner and therefore have stronger incentives to experiment early in the adoption cycle, making investment more front-loaded. Interestingly, greater access to external capital tends to mitigate this effect. By relaxing financing constraints, access to capital can reduce the pressure to accelerate experimentation in order to obtain early feedback. These results highlight how technological characteristics, entrepreneurial preferences, and capital-market conditions jointly shape the aggregate investment response to a new technology.

Our analysis also has implications for understanding the welfare consequences of technological investment booms. Aggregate experimentation can be socially valuable even when many individual projects fail. Each experiment may reveal information about the allocation of existing capabilities across emerging opportunities, thereby contributing to a broader process

of collective exploration. The relevant welfare question is therefore not simply whether individual investments generate positive immediate returns, but also how much information they produce about the economy's underlying comparative advantages. A period of extensive experimentation may generate substantial private losses while improving the allocation of resources across firms and activities. Whether the resulting investment boom is socially excessive or insufficient depends on the value of this information relative to the resources consumed in experimentation. We use this perspective to distinguish the private and social returns to exploratory investment.

Our learning mechanism differs fundamentally from conventional learning-by-doing. In learning-by-doing, agents become more productive at an activity through practice and repetition (see, e.g., Smith (1776); Arrow (1962); Dewey (1974); Pratten (1980); Lucas (1988); Lamoreaux et al. (1999); Thompson (2010); Morris (2020)). In our setting, by contrast, firms experiment to discover whether their existing capabilities are well matched to an activity. Whereas learning-by-doing emphasizes capability accumulation, introspective learning generates self-discovery. The distinction is important: firms need not become better at developing a technology to benefit from experimenting with it. They may benefit simply by discovering whether they are particularly well suited to the task.

Our setting also differs from models in which firms learn about a common external component of project payoffs. In Johnson (2007), for example, investors may optimally invest above the full-information level to learn about a new technology's returns to scale. In Pástor and Veronesi (2009), investors update their beliefs about another common parameter—the average productivity of a new technology—in a framework that focuses on asset prices and abstracts from real investment. By contrast, the object of introspective learning is not a common payoff primitive but each firm's individual fit with the technology. Investment reveals whether a firm has a comparative advantage in a particular activity, thereby sorting firms across the emerging opportunity set. This distinction has important aggregate implications: rather than generating a gradual adjustment to a common technological signal, introspective learning can produce broad early participation followed by an endogenous shakeout as unsuccessful firms exit.

The paper is related to the literature on industry dynamics and shakeouts. Entry followed by high failure and exit rates is a well-documented feature of industries adopting new technologies, and learning has been recognized as one of the mechanisms behind this pattern (Jovanovic (1982)). Jovanovic and Lach (1989) emphasize the advantage of late entrants, who can learn from the experiences of earlier participants. Jovanovic and MacDonald (1994) develop a competitive model explaining why the number of firms in an industry initially rises and subsequently declines as the industry evolves. Other explanations of excessive investment

emphasize relative-wealth concerns (Abel (1990); DeMarzo et al. (2008)) or behavioral biases interacting with capital-market frictions (De Long et al. (1990); Abreu and Brunnermeier (2003)). Our mechanism is distinct from these explanations. It does not rely on strategic interactions, competitive externalities, or behavioral distortions. Instead, we show that firms' individually rational incentives to discover their own fit with a new technology are sufficient to generate the familiar boom–bust pattern. We further characterize how the dynamics depend on technological characteristics, entrepreneurial preferences, and capital-market conditions.

Beyond aggregate industry dynamics, our analysis offers a complementary perspective on entrepreneurial entry, failure, and returns. A large empirical literature documents substantial entrepreneurial activity despite high failure rates and weak average ex post returns (Dunne et al. (1988); Moskowitz and Vissing-Jørgensen (2002)). A common interpretation attributes this pattern to entrepreneurial overconfidence (Camerer and Lovallo (1999); Bernardo and Welch (2001); Koellinger et al. (2007)). We highlight a rational motive for entrepreneurial entry: starting a venture generates information about whether an entrepreneur's existing capabilities are suited to developing and commercializing a particular opportunity. Even unsuccessful ventures can therefore create value by revealing where an entrepreneur's comparative advantage does not lie. High entry and subsequent exit need not reflect miscalibrated beliefs alone; they may also be the outcome of rational experimentation and self-discovery.

Our paper is also related to Drugov and Macchiavello (2014), who study entrepreneurs experimenting to discover their ability to undertake a new activity, including through a strategy of starting small. They show that financing experimentation can become more difficult precisely when experimentation is more profitable—for lower values of the known alternative and more optimistic priors. They also show that endogenous collateral requirements, such as those commonly observed in microcredit schemes, can be part of the optimal contract. Our analysis complements this work by focusing on the aggregate investment dynamics generated when heterogeneous firms experiment to discover their comparative advantages in a new technological opportunity set.

The rest of the paper is organized as follows. Section 2 discusses the broad notion of self-discovery and how it naturally applies to investment during new-technology episodes. Section 3 analyzes this self-discovery motive for investment in a stylized model and derives the resulting boom–bust investment dynamics. Section 4 turns to the cross-section, characterizing the sorting of firms over time into winners and losers. Section 5 introduces external financing and traces its effect on the boom–bust dynamics. Section 6 concludes. Proofs are provided in the Appendix.

# 2 Self-Discovery and Investment Behavior

Recent research across behavioral and neural sciences shows that living organisms need to explore their environment to decide where to forage, with whom to mate, or where to breed (Charnov (1976); Hills et al. (2015); Pirinsky (2026)). Humans face the same exploration–exploitation dilemma in virtually every decision: trying a new option to learn its payoff (exploration), or choosing the best among those whose payoffs are already known (exploitation).

This same logic carries over to firms. Experimenting with a new technology can reveal not only whether the technology is valuable in general, but also whether it is valuable *for this firm*, given its particular capabilities.[2] Firms carry heterogeneous bundles of know-how, organizational routines, and complementary assets, but those strengths are often *latent* until a new technological paradigm reveals where they apply. The mapping from "what we are good at" to "what the new technology rewards" is rarely transparent at the outset. Firms typically learn it only by building—that is, by exploring the technology first-hand. Doing so generally requires committing real resources before relative advantages and disadvantages are known. In that sense, early investment is an experiment that is informative precisely because it is risky.

Nvidia offers an immediate case in point. The capabilities behind its current AI advantage—massively parallel computation, programmable GPUs, and a deep developer ecosystem—were not originally designed for AI, but for graphics and gaming. At the outset, whether those strengths would travel to AI was itself unknown. The uncertainty was not just whether AI would matter, but whether Nvidia's existing capabilities would prove complementary to it. That complementarity could only have been discovered through costly engagement with the AI frontier—in the words of its CEO Jensen Huang, "if we didn't build it, we'd never know" (Salian (2024)).

The late-1990s Internet boom is the episode we examine empirically. During the expansion, firms were learning not only whether the Internet would prove economically important, but whether they themselves could turn the new technology into a viable business. Young companies with untested models entered in large numbers to find out. Many learned that they could not. The crash was therefore not merely a change in market sentiment, but a large-scale revelation that many entrants had no durable fit with the new commercial environment (Shiller (2000); Ofek and Richardson (2003); DeMarzo et al. (2007)).

[2]There is a long strand of economics literature on exploration, innovation, and learning-by-doing (Arrow (1962); March (1991); Aghion and Howitt (1998); Manso (2011); Lucas and Moll (2014)). While that literature studies external search, innovation, and capability accumulation, our focus is on the introspective margin.

This pattern is not confined to contemporary technology. During the British railway boom of the 1840s, promoters and investors were not simply betting on aggregate traffic demand. They were discovering, project by project, which engineering plans, financing structures, and operating organizations could actually work. Capital flowed in before firms knew whether they had the technical and organizational fit to execute particular lines successfully. The later crash reflected not only revised beliefs about demand, but the revelation that many ventures were poorly matched to the projects they had undertaken (Odlyzko (2010); Campbell and Turner (2012); Campbell (2013)).

Other episodes that display this same pattern of broad early entry, experimentation, and subsequent shakeout include the Florida land boom (Turner (2015)) and, going back still further, the Dutch tulip mania (Goldgar (2007)). Taken together, these episodes point to a simple but underemphasized mechanism: early investment in a new domain is often a diagnostic experiment about *fit*. The boom reflects broad participation by firms trying to learn whether their capabilities travel to the new technology; the bust reflects the subsequent realization that many do not.

# 3 Introspective Learning and Aggregate Investment

## 3.1 Model Setup

We now develop a model of firm self-discovery under a new technology. Firms face a range of possible activities (projects) through which they may learn whether their capabilities are aligned with that technology. "Projects" should be read broadly as lines of experimentation rather than as standalone undertakings in a narrow sense. A firm "matches" with a project type when undertaking that activity reveals a complementarity between that firm's capabilities and the new technology.

**Firms and Technology** We consider a unit mass of firms. Each firm has a type,

$$\theta \in \Theta := \{0, 1, \ldots, N\}, \tag{3.1}$$

which is ex ante unknown to the firm and identifies the project type through which the firm can discover its own fit with the technology. Type $\theta = 0$ is the no-fit type.

The cross-sectional distribution of types is

$$\mu_0(0) = 1 - p, \qquad \mu_0(i) = \frac{p}{N} \text{ for } i = 1, \ldots, N, \tag{3.2}$$

where $p \in (0,1)$. A share $1-p$ of firms has no fit with the technology, while the remaining share $p$ is distributed uniformly across the $N$ project types. Thus, $p$ measures the technology's economic *scope*—the breadth of firms able to leverage it. The parameter $N$ measures its *granularity*: as $N$ rises, the opportunity set is divided more finely, making any particular direction less likely to match a given firm.

**Capital and Investment** At each date $t = 1, \ldots, T$, a firm may choose to invest in a project, i.e., to undertake a new informative experiment. Capital depreciates at rate $\delta \in (0,1)$, and future payoffs are discounted by $\beta \in (0,1)$, where $\beta$ may vary across firms. Because learning about fit requires meaningful engagement with the technology, investment cannot be arbitrarily small. Any positive investment must therefore meet a minimum threshold, normalized to one. Each firm is endowed with an initial capital stock $k_1$. Capital then evolves according to

$$k_{t+1} = (1-\delta)\,(k_t - x_t), \tag{3.3}$$

where

$$x_t \in \{0\} \cup [1, k_t] \tag{3.4}$$

is the investment level at date $t$.

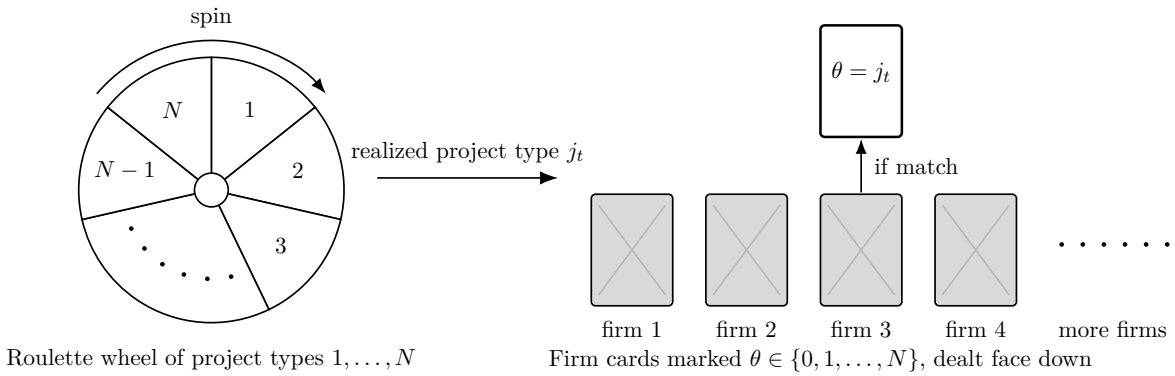



Figure 1: **The Roulette of New Technology**

**Project Draws: The Roulette of New Technology** Suppose successful development of the new technology can draw on $N$ distinct capabilities, skills, or organizational strengths. We index the corresponding development paths by project types $1, \ldots, N$. A firm's type $\theta$, defined in (3.1), indicates where, if anywhere, its capabilities fit within this opportunity set. In particular, $\theta = i$ means that undertaking a project of type $i$ reveals fit, while $\theta = 0$ means that the firm has no fit.

We visualize this economy-wide opportunity set as a roulette wheel, whose slots correspond to project types $1, \ldots, N$, as illustrated in Figure 1. In this picture, a firm's type is represented by a private card, dealt face down because firms do not know ex ante which project type, if any, will reveal their fit. A card marked $\theta = 0$ indicates that no such fit exists within the opportunity set.

**Matching and Cash Flow** As the new technology unfolds, each spin of the wheel lands on a project type $j$. Each firm may then choose whether to pursue that opportunity by placing at least the minimum stake. If a type-$\theta$ firm invests $x_t$ in a type-$j$ project, the one-period cash flow is

$$y_t = R x_t \, \mathbf{1}_{\{\theta = j\}}, \tag{3.5}$$

where $R > 0$ is the gross return on a unit of investment conditional on a successful match.

A positive payoff reveals that a firm's capabilities fit the new technology through the realized project type. Once matched, the firm can keep investing in that project type in subsequent periods, exploiting the fit it has uncovered. A zero payoff reveals that the firm did not find a fit through that project. Its underlying type remains unrevealed; in the roulette visualization, the firm's card stays face down.

We work in event time, suppressing repeated project realizations because they provide no new information and induce no new action. At early stages of a new technology, the remaining applications do not unfold according to a predictable roadmap. We therefore assume that the next informative economy-wide project arrival $j_t$—the next previously unseen project to land on the roulette wheel—follows

$$j_t \overset{d}{\sim} \text{Unif}(M_t), \tag{3.6}$$

where $M_t \subseteq \{1, \ldots, N\}$ is the set of unseen project types at the start of period $t$.

There is no loss in assuming that the horizon is no longer than the number of distinct project types the firm may test, i.e., $T \leq N$. After at most $N$ failures, a firm would have exhausted the relevant opportunity set and inferred that it has no fit.

**Firm's Problem** Each firm chooses a history-dependent investment policy

$$\pi = \{x_t, \; a_t(y_{t-1})\}_{t=1}^{T}, \quad \text{where} \quad a_t(y_{t-1}) \in \{\text{keep}, \text{redraw}\},$$

to maximize expected discounted cash flows

$$\max_{\pi} \ \mathbb{E}^{\pi}\left[\sum_{t=1}^{T} \beta^{t-1} y_t\right] = \max_{\pi} \ \mathbb{E}^{\pi}\left[\sum_{t=1}^{T} \beta^{t-1} R\, x_t\, \mathbf{1}_{\{\theta=j_t\}}\right], \tag{3.7}$$

subject to the capital accumulation equation (3.3) and the feasibility constraint (3.4). Here, $a_t(y_{t-1})$ is the firm's post-outcome decision at the start of period $t$: after observing the previous period's cash flow, it either continues exploiting a direction that has revealed a fit (keep) or returns to search and draws a new project type (redraw). $j_t$ is the project type at period $t$, whether kept from a revealed fit or newly drawn.

## 3.2 Investment Policy

Next we characterize the firm's optimal investment policy. Investment is not only capital deployment but also a form of exploration. Once fit is known, the problem reduces to pure exploitation. While fit remains unknown, however, the investment policy exhibits a sharp split: depending on time preference, firms either invest aggressively right away or proceed through cautious experimentation.

**Known-Match Regime** Suppose the firm has already matched with a project type. From that point on, there is no option value to waiting. Delaying investment only exposes capital to depreciation and yields no offsetting informational benefit. The firm's optimal investment policy in this regime is therefore to invest all available capital whenever investment is feasible:

$$x_t^*(k \mid \text{known match}) = \begin{cases} 0, & k < 1, \\ k, & k \geq 1, \end{cases}$$

and the associated continuation value $W_t(k)$ is

$$W_t(k) = \begin{cases} 0, & k < 1, \\ Rk, & k \geq 1. \end{cases} \tag{3.8}$$

**Searching Regime** Suppose instead that the firm is still searching for a fit. In this regime, investment is no longer pure exploitation. A successful experiment generates cash flow and reveals a direction in which the firm can scale. An unsuccessful experiment still has value because it rules out one remaining project type and narrows the firm's future search.

At date $t$, suppose the searching firm has capital $k$ and has not yet ruled out $m := |M_t|$ project types. Let $V_t(k, m)$ denote its continuation value in that state. The Bellman equation

for $V_t(k,m)$ is

$$V_t(k,m) = \max\Big\{\Phi_t^{\text{wait}}(k,m),\ \Phi_t^{\text{inv}}(k,m)\Big\}, \tag{3.9}$$

where $\Phi_t^{\text{wait}}(k,m)$ is the continuation value from waiting and $\Phi_t^{\text{inv}}(k,m)$ is the continuation value from investing.

*Waiting.* If the firm waits, capital simply depreciates and it learns nothing new about itself, so

$$\Phi_t^{\text{wait}}(k,m) = \beta\, V_{t+1}\Big((1-\delta)k,\ m\Big). \tag{3.10}$$

*Investing.* If instead the firm undertakes an experiment of size $x \in [1,k]$, then with probability $\mu(m)$ the project reveals a match, generates current cash flow $Rx$, and leaves the firm in the known-match regime next period. With the complementary probability, the project fails: the firm learns that this direction is not its fit, loses the capital committed to the experiment, and continues with one fewer remaining project type. Accordingly,

$$\Phi_t^{\text{inv}}(k,m) = \max_{x\in[1,k]}\Big[\mu(m)\Big(Rx+\beta\, W_{t+1}\Big((1-\delta)(k-x)\Big)\Big)+\Big(1-\mu(m)\Big)\beta\, V_{t+1}\Big((1-\delta)(k-x),\ m-1\Big)\Big], \tag{3.11}$$

where $\mu(m)$ is the conditional probability that the project draw reveals fit, given that the firm is still searching with $m$ surviving directions:

$$\mu(m) = \Pr\Big(\theta = j_t \,\Big|\, j_t \in M_t,\ |M_t| = m\Big) = \frac{\frac{p}{N}}{(1-p)+m\frac{p}{N}} \in (0,1). \tag{3.12}$$

**Proposition 3.1.** *In the searching regime* (3.9)*,*

$$\Phi_t^{\text{inv}}(k,m) > \Phi_t^{\text{wait}}(k,m) \quad \textit{for } k \geq 1.$$

*That is, whenever a firm still has enough capital to undertake an experiment, investing strictly dominates waiting.*

Proposition 3.1 formalizes a simple but key intuition: even a failed experiment is informative because it eliminates a dead-end direction and narrows the remaining search for fit. Waiting, by contrast, reveals nothing—because the uncertainty is introspective. Since fit can be learned only through experimentation, the value of information pulls investment forward in time. Consequently, when the firm is in the searching regime and positive investment is feasible, there is only the intensive margin: how much capital to risk in the experiment, trading off a larger experiment today against preserving capital for future experimentation.

To state the investment policy, it is useful to track the capital path generated by the most cautious form of active search. Suppose that, while still unmatched, the firm invests

the minimum feasible amount—one unit—in each period for as long as such experimentation remains feasible. Let

$$\tilde{k}_1 := k_1, \qquad \tilde{k}_{t+1} := (1-\delta)\,(\tilde{k}_t - 1) \tag{3.13}$$

denote the resulting capital path. Along this path, the last date at which a one-unit experiment remains feasible is

$$L := \min\{\, T,\ \ell(k_1)\,\}, \tag{3.14}$$

where

$$\ell(k_1) := \max\{\, t \geq 1 : \tilde{k}_t \geq 1\,\}.$$

**Theorem 3.2** (Optimal Investment Policy in Search)**.** *While searching for a fit, the firm's optimal policy is characterized by a sharp time-preference dichotomy:*

*(i)* (Impatient firms.) *If $\beta \leq \frac{1}{2(1-\delta)}$, the firm invests all available capital in the technology immediately.*

*(ii)* (Patient firms.) *If $\beta > \frac{1}{2(1-\delta)}$, the firm experiments cautiously, investing one unit at a time while its fit remains unknown and experimentation remains feasible, and committing all remaining capital only at $t = L$.*

**Impatient vs. Patient Firms** Theorem 3.2 shows that, under introspective learning, time preference does not merely tilt investment intensity, as it would in standard Bayesian learning. It determines the learning mode—through an all-in gamble or through sequential experimentation. While we abstract from risk aversion, lower-$\beta$ firms nonetheless behave as if they were more willing to gamble and invest all at once. Firms with higher $\beta$, by contrast, behave more cautiously, financing a sequence of small diagnostic experiments. Hereafter, we refer to firms with $\beta \leq \frac{1}{2(1-\delta)}$ as *impatient* and those with $\beta > \frac{1}{2(1-\delta)}$ as *patient*.

The factor of two in the cutoff is the return to patience. This can be seen by comparing the all-in investment with the simplest form of caution: test one unit today and, whatever the outcome, commit the rest next period. The deferred capital ends up behind the firm's fit by two routes: today's test succeeds, or it fails and the redraw succeeds. With $m$ project types remaining, the first route has probability $\mu(m)$ from (3.12); the second route is equally likely, because failure narrows the field and lifts the redraw odds to $\mu(m-1) = \mu(m)/(1-\mu(m))$, just offsetting the chance of having failed. Holding back thus doubles the odds on the deferred capital, which is worth a period of discounting and depreciation whenever $2\beta(1-\delta) > 1$.

**Remark 3.3** (Outside Option)**.** *Allowing firms to invest in an outside option—for example, an incumbent technology or a diversified outside investment—does not affect the analysis, provided the outside return $\bar{r}$ is below the ex ante return from experimentation. Specifically,*

*when $\bar{r} < \frac{p}{N}R$, the same investment policies as in Theorem 3.2 obtain. Impatient firms still invest immediately, and patient firms still experiment cautiously.*

As a corollary, we characterize the private value that brings firms into the search problem in the first place. This is the ex ante value of introspective learning before the firm knows whether it has a fit, and therefore captures the incentive behind broad early participation.

**Corollary 3.4** (Ex Ante Firm Value)**.** *The ex ante value $V_1(k_1, N)$ of a firm that begins with initial capital $k_1$ and faces $N$ possible project types—before discovering its fit—is:*

*(i) For an impatient firm,*

$$V_1(k_1, N) = R\,\frac{p}{N}\,k_1. \tag{3.15}$$

*(ii) For a patient firm,*

$$V_1(k_1, N) = R\,\frac{p}{N}\,\left[\sum_{t=1}^{L-1}\left(\beta^{t-1} + \beta^t \tilde{k}_{t+1}\right) + \beta^{L-1}\tilde{k}_L\right], \tag{3.16}$$

*where $(\tilde{k}_t)_{t\geq 1}$ is the feasible-testing capital path* (3.13)*.*

Corollary 3.4 shows that the value of self-discovery depends on the opportunity landscape opened by the new technology. The common factor $\frac{p}{N}$ in (3.15) and (3.16) is the ex ante probability that a given experiment uncovers fit. When the new technology has broader scope $p$, ex ante value rises because a firm is more likely to have capabilities that travel to the new domain. Higher granularity $N$, by contrast, spreads those potential fits over more projects. The opportunity may still be valuable, but the right match is harder to find: each experiment is less likely to be the one that uncovers a firm's edge.

**Discovery-Paced Investment** Under standard investment logic, there is no reason to hold capital back from a positive-NPV opportunity: a project worth taking is worth taking at scale (Jorgenson (1963)). In our model, an impatient firm faces such an opportunity (Corollary 3.4(i)) and invests as prescribed (Theorem 3.2(i)). Our mechanism sets apart the patient firm: its investment tracks its self-knowledge (Theorem 3.2(ii)), holding at the minimum feasible amount for as long as the firm does not know its own type, and stepping up to full investment when that knowledge arrives. In practice, the distinction is one of degree: most firms confronting a new technology experiment before they scale. While this tracking is exact in our stylized setting, the substantive point is more general: a firm's self-knowledge is itself a determinant of when and how much it invests, besides other factors the literature has already analyzed at length.

Discovery-paced investment reverses the usual real-options intuition. There, the option is to wait and see, holding investment back until more is known (Dixit and Pindyck (1994)). Here, the capital held back waits on the firm's own experiments, not on the passage of time: the option value of self-discovery derives from active search, with waiting made even less attractive by depreciating capital.

## 3.3 Self-Discovery and Aggregate Investment

We now turn from individual investment policy to aggregate investment dynamics. The key question is how introspective learning about firm-specific fit maps into the time path of investment at the economy-wide level. When firms invest not only to exploit a new technology but also to learn whether their own capabilities are suited to it, adoption becomes naturally front-loaded: many invest early to test fit, and aggregate investment later contracts as mismatch is revealed and the learning motive is gradually exhausted. To allow for heterogeneity in the urgency of this self-discovery motive, we consider a population consisting of both impatient firms, which invest aggressively up front, and patient firms, which experiment before scaling.

We continue to write $(\tilde{k}_t)_{t=1}^{L}$ for the feasible-testing capital path and $L$ for the associated feasible experimentation horizon, as defined in (3.13) and (3.14). Let $\pi_{\text{imp}} \in [0,1]$ be the population share of impatient firms, and $\pi_{\text{pat}} \coloneqq 1 - \pi_{\text{imp}}$ the complementary share of patient firms.

**Proposition 3.5.** *Fix initial capital $k_1 \geq 1$. For each period $t \geq 1$, let $\text{Inv}_{\text{imp},t}$ and $\text{Inv}_{\text{pat},t}$ denote aggregate investment by impatient and patient firms, respectively. Then the aggregate investment paths are:*

*(i) Impatient firms invest all capital immediately:*

$$\text{Inv}_{\text{imp},t} = \pi_{\text{imp}}\, k_1\, \mathbf{1}_{\{t=1\}}.$$

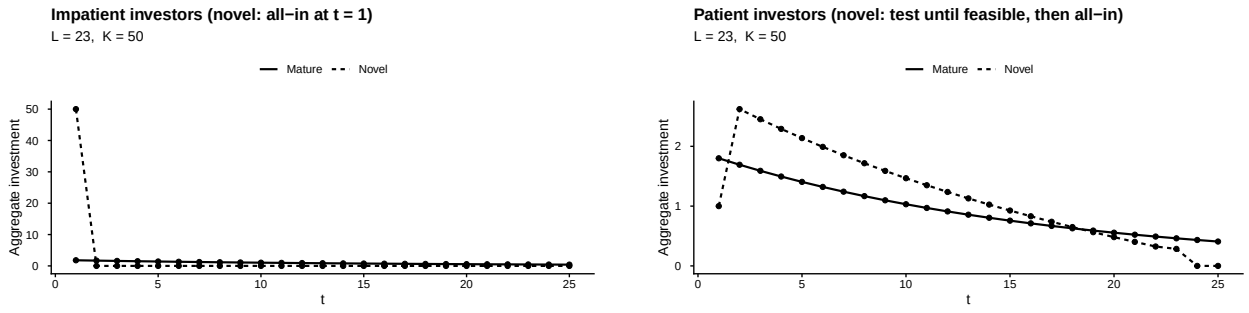
*(ii) Patient firms experiment and scale up if matched, implying*

$$\text{Inv}_{\text{pat},t} = \begin{cases} \pi_{\text{pat}}, & t = 1, \\ \pi_{\text{pat}}\left(S_t + \dfrac{p}{N}\tilde{k}_t\right), & t = 2, \ldots, L-1, \\ \pi_{\text{pat}}\left(S_L + \dfrac{p}{N}\right)\tilde{k}_L = \pi_{\text{pat}}\left(1 - \dfrac{p}{N}(L-2)\right)\tilde{k}_L, & t = L, \\ 0, & t > L. \end{cases} \tag{3.17}$$

*Here,* $S_t = 1 - \frac{p}{N}(t-1)$*, for* $1 \leq t \leq L$*, is the mass of patient firms still searching at the start of period* $t$*.*

Aggregate investment (3.17) by patient firms comes from two groups. A mass $S_t$ of firms is still searching, each staking the minimum feasible unit. A mass $p/N$ found fit one period earlier, each now committing all its remaining capital. Investment from both groups dwindles as the technology ages: the search pool drains each period, and each successive cohort of winners reaches scale-up with less capital left to deploy.

**Aggregate Investment Path** Figure 2 compares the investment path in Proposition 3.5 with a mature-technology benchmark in which firms know their type and invest only when the project type matching their capabilities arrives. This comparison isolates the self-discovery motive behind the front-loaded boom–bust pattern. In the mature benchmark, firms already know where their capabilities fit. Under a new technology, they must invest to find out.




(a) Impatient firms.

(b) Patient firms.

Figure 2: **Boom–bust investment under a new technology.** The dashed line plots aggregate investment when firms must invest to learn their fit with the new technology. The solid line plots the mature-technology benchmark, where fit is already known. Panel (a) shows impatient firms, for which learning produces an immediate investment spike at $t = 1$. Panel (b) shows patient firms, for which new-technology investment is spread over the experimentation phase but remains front-loaded relative to the mature benchmark and declines as the learning motive is exhausted.

Thus, firm self-discovery generates a boom–bust investment pattern without aggregate shocks or financial frictions. Under a new technology, investment is not merely a scale choice but also a diagnostic experiment. Firms risk capital before knowing whether their existing capabilities travel to the new domain. The return to that experiment is highest early on, when uncertainty about fit is still greatest. At that stage, a successful experiment reveals a direction worth scaling, while a failed one rules out a dead end. Investment is pulled toward the front end of the technology's life cycle. The initial *boom* reflects broad participation

by firms trying to discover their match. The subsequent *bust* follows from the endogenous exhaustion of the same learning motive as experiments accumulate.

**The Effect of Time Preference** Time preference shapes the boom's timing. Visually, a larger impatient share would front-load aggregate investment toward the impatient path in Figure 2. The same front-loading effect applies to the comparative-static paths in Figures 3 and 4 below. This comparative static is most relevant when capital is evaluated over short horizons—for example, benchmarked institutional capital or fund managers sensitive to near-term performance. Such capital compresses self-discovery into the first wave of experimentation.

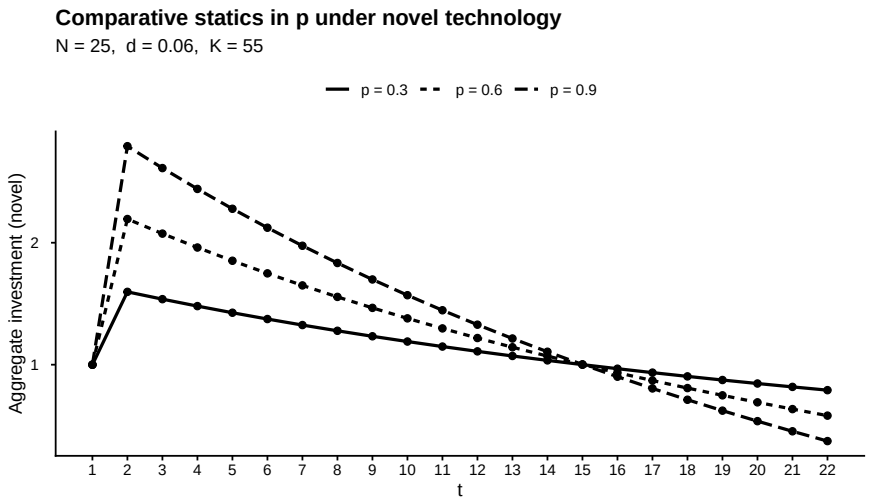



Figure 3: **Technology scope and boom–bust investment.** Each line plots patient-firm aggregate investment for a different scope $p = \Pr(\theta \neq 0)$. A larger scope $p$ raises the early boom because fit is more likely and more experiments turn into scale-up. Investment then declines as fewer firms remain in search.

**The Effect of Technology Scope** Figure 3 illustrates the investment path of patient firms under a new technology for different values of the scope $p = \Pr(\theta \neq 0)$. A higher $p$ means that the technology is accessible to a larger share of firms, so each experiment is more likely to uncover a productive match. As shown in Corollary 3.4, this raises the ex ante value of experimentation. Therefore, more inclusive technologies generate larger booms—early participation is broad, and more experiments reveal scalable uses. A more specialized technology offers fewer plausible routes to fit, and the boom is correspondingly muted.

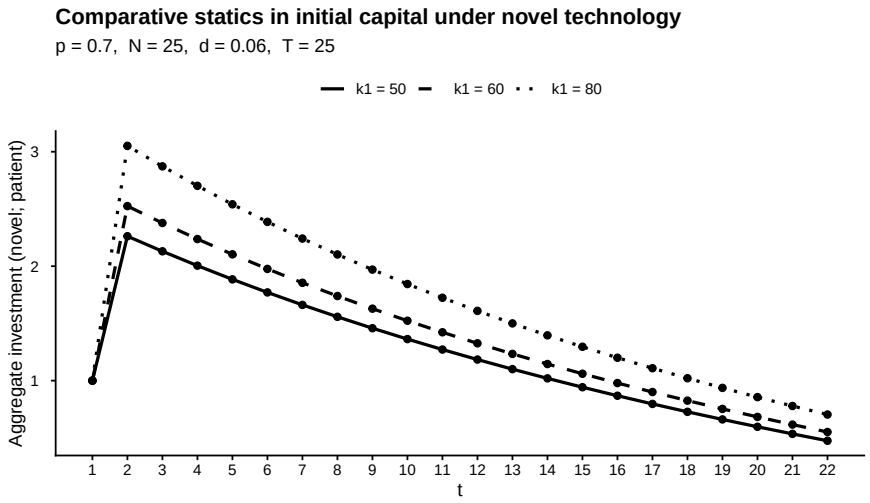



Figure 4: **Access to capital and boom–bust investment.** Each line plots patient-firm aggregate investment for a different initial capital stock $k_1$. Higher $k_1$ raises and prolongs the boom because firms can keep searching longer and successful firms scale with more capital left to deploy.

**The Effect of Access to Capital** More initial capital $k_1$ extends the boom phase because firms can keep searching longer, and it makes the eventual scale-up larger because successful firms reach that stage with more capital left to deploy. When $k_1$ is small, experimentation ends sooner and the boom–bust cycle is compressed. Figure 4 visualizes this effect by plotting the aggregate investment path of patient firms for different values of $k_1$. Higher $k_1$ lifts the boom and holds it up longer.

This is distinct from the traditional view that financial constraints prevent firms from undertaking positive-NPV projects (Fazzari et al. (1988); Hall and Lerner (2010); Kerr and Nanda (2015)). Here, the issue is not only whether firms can finance projects whose payoffs are already understood, but whether they can survive the costly trials needed to discover such projects in the first place. Access to finance lets firms absorb zero-payoff experiments long enough for fit to become visible.[3]

**The Effect of Creative Destruction** In many new-technology episodes, adoption also erodes the value of existing business models (see Schumpeter (1942); Aghion and Howitt (1992); Grossman and Helpman (1991); Caballero and Hammour (1996); Hobijn and Jovanovic (2001); Klette and Kortum (2004)). In our setting, this corresponds to letting the outside return in Remark 3.3 decline as the new technology advances. Such incumbent displacement

[3]For example, in its pivot to deep learning, Nvidia "invested billions of dollars and years of engineering resources" before knowing how far deep learning would scale (Salian (2024)).

would make the boom–bust investment dynamics sharper still—as the incumbent alternative becomes less attractive, the opportunity cost of delay increases and firms have stronger incentives to experiment early.

# 4 Sorting into Winners and Losers

The aggregate investment path reflects an underlying firm-level sorting process: broad entry, failure, exit, and the eventual concentration of activity among a smaller set of scaled survivors. We now characterize these dynamics.

## 4.1 Firm Failure Rates

Early in a new-technology episode, failure is common because firms have learned little about where their capabilities fit. These failures, however, are informative. Each unsuccessful attempt closes off one route to fit, so the firms still searching face a narrower set of unresolved directions.

Proposition 4.1 records the probability of another no-match outcome for patient firms. Impatient firms make a single all-in diagnostic investment at $t = 1$, so their failure probability is the corresponding first-period value, $1 - \frac{p}{N}$.

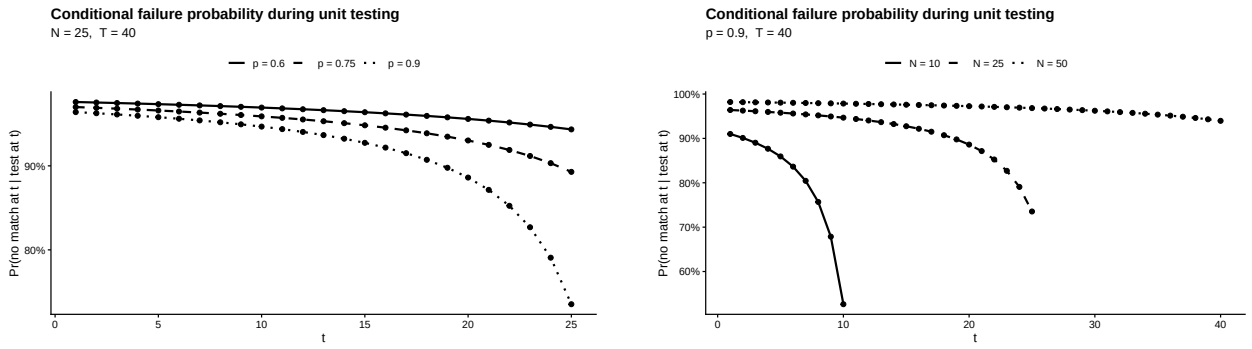
**Proposition 4.1** (Failure Rate While Searching for Fit)**.** *Consider a patient firm that remains in search at the start of period $t < L$ and undertakes one more unit experiment. The probability that the experiment does not reveal fit is*

$$\Pr(\textit{no match at } t \mid \textit{still searching at } t) = 1 - \frac{\frac{p}{N}}{1 - \frac{p}{N}(t-1)}. \tag{4.1}$$

Figure 5 plots the probability (4.1) that a firm still searching for fit receives another no-match outcome in period $t$, under selected parameter values. The common downward slope reflects the selection induced by introspective learning.

Panel 5a varies technology scope $p$. Larger scope shifts the failure path downward. This is the same inclusiveness margin behind Figure 3, viewed through firm-level failures rather than aggregate investment. A more inclusive technology offers more plausible routes to fit. Larger scope amplifies the investment boom by moving more firms into scale-up, while also making continued search less likely to end in another failed attempt.

Panel 5b varies technology granularity $N$. Greater granularity makes self-discovery slower and more failure-prone by spreading fit across a larger set of possible directions. This is the




(a) Technology scope $p$. Larger scope lowers failure rates.

(b) Technology granularity $N$. Higher $N$ keeps failure rates elevated.

Figure 5: **Failure rates while searching for fit.** Each path plots the probability that a still-searching firm receives another no-match outcome. Failure rates decline over time because continuing firms have already ruled out some dead ends. Larger scope $p$ lowers failure rates, while greater granularity $N$ spreads fit across more possible directions and prolongs the period of elevated failure.

dispersion margin in Corollary 3.4, viewed through firm-level failures. When the opportunity set is divided into more distinct directions, a given experiment is less likely to reveal a match.

Thus the same early participation that generates the investment boom also produces high early failure. Failure is not only evidence of mistaken entry. It is also the information produced by entry.

## 4.2 Winner–Loser Sorting over Time

The conditional failure rates above describe the experience of firms that remain in search. In the cross-section, the same dynamics appear as sorting: each period a fraction of firms discovers fit and moves into exploitation, the remaining search pool becomes increasingly tilted toward no-fit firms, and firms that discover fit earlier enter exploitation at larger scale. This firm-level separation is absent under Bayesian learning about common payoff primitives.

**Proposition 4.2** (Cross-Sectional Sorting and Shakeout)**.** *With $\tilde{k}_t$ again denoting the capital path under one-unit experimentation and $L$ the corresponding feasible experimentation horizon, as in* (3.13) *and* (3.14)*, the cross-sectional composition of patient firms evolves as follows.*

*(i) At the start of period $t \leq L$, the mass of firms still searching is*

$$S_t = 1 - \frac{p}{N}(t-1).$$

*By the end of period $t$, the cumulative mass of firms that have discovered fit is*

$$W_t = \frac{p}{N}t.$$

*(ii) The remaining search pool becomes increasingly selected toward no-fit firms:*

$$\Pr(\theta = 0 \mid \textit{still searching at } t) = \frac{1-p}{1-\frac{p}{N}(t-1)}.$$

*(iii) If two firms first discover fit in periods $t < t' < L$, then $\tilde{k}_{t+1} > \tilde{k}_{t'+1}$. Thus, earlier winners enter exploitation at larger scale.*

Impatient firms sort immediately: the all-in experiment at $t = 1$ creates winner mass $p/N$ and loser mass $1 - p/N$. Patient firms sort gradually: each round peels off a cohort of winners, and the remaining search pool shrinks and tilts toward firms that were never a good fit.

The sorting on both the extensive margin (who becomes a winner) and the intensive margin (how large each winner becomes) produces the well-documented pattern of entrepreneurial activity: broad early participation, severe attrition, and eventual concentration among a narrower set of scaled survivors. In our setting, this sequence arises from optimizing behavior alone—high entry and subsequent exit need not reflect miscalibration.

# 5 External Financing

The firms of Section 3 fund their projects with their own capital. In practice, entrepreneurs are funded by venture investors, and a project then has two parties whose incentives diverge—the entrepreneur stands to discover her own fit with the technology, which belongs to her, by deploying the investor's capital. We now consider the implications of this friction for the boom–bust dynamics.

## 5.1 Entrepreneurs and Investors

**Agents** Each firm is now a pair—an *entrepreneur* (she) and an *investor* (he). She carries a potential fit with the technology—her type $\theta \in \Theta$ of (3.1)—but now has no capital. He has capital and no access to the technology on his own.

At date 1 he chooses how much capital to give the firm. The most he can commit to any one firm is $\bar{K}$.[4] His capital earns the outside return $\bar{r}$ of Remark 3.3. As before, an

[4]In practice, venture funds are closed-end and operate under partnership agreements that cap exposure to

investment of size $x_t \in \{0\} \cup [1, k_t]$ in project direction $j_t$ returns $Rx_t$ if $\theta = j_t$ and nothing otherwise.

**Control** While the investor supplies the capital, the entrepreneur directs its use: faced with the date-$t$ project $j_t$, she chooses the size of the investment, $x_t$, from the capital remaining at that date. If she finds fit, she may be tempted to take what she learned elsewhere and leave the firm.[5] Her leverage at that point determines how much of the scale-up surplus she can bargain for, a participation constraint that binds the moment she discovers her value-added. We summarize this leverage by her *Nash bargaining weight* $\eta \in (0, 1)$. For example, when the two parties have equal bargaining power, the standard symmetric Nash solution $\eta = \frac{1}{2}$ splits the surplus equally between them.

If she has found fit and the firm has remaining capital $k'$, she then receives the fraction $\eta$ of the scale-up surplus $(R - \bar{r})k'$ over the outside return. Her share of the scale-up return $R$ is therefore

$$h \coloneqq \eta\Big(1 - \frac{\bar{r}}{R}\Big), \tag{5.1}$$

and $1 - h$ is the investor's share.

Variation in $\eta$ reflects whether the value created by the discovery is due to the entrepreneur or to an asset—"jockey or horse" (Kaplan et al. (2009))—which can depend on the technology (e.g., AI vs. tulips). A higher $\eta$ means that more of that value is attributable to the entrepreneur and thus harder for the firm to keep: leading to weaker non-competition covenants, less intellectual property assigned to the firm, etc. This dimension is distinct from the scope $p$ of (3.2): e.g., a technology can be easy to fit and hard to protect.

**Contract** The contract $(k_1, \alpha)$ sets the investor's capital commitment $k_1 \leq \bar{K}$ and the entrepreneur's equity share $\alpha \in [0, 1]$ of the firm's earnings. The contract is written before anyone knows whether she has a fit, so the equity split governs what is earned before that is known.

To isolate the friction that comes from the option value of self-discovery accruing asymmetrically to one side, we take the investor to be competitive; an investor with market power would strengthen the results.[6] For the friction to have bite, the outside return $\bar{r}$ must lie

any one company. In our setting the exposure cap $\bar{K}$ replaces the capital endowment $k_1$ of Section 3: with constant returns some bound on the scale of a venture is needed, and the question here is how much of it the investor uses.

[5] See Hart and Moore (1994), Anton and Yao (1994), Gompers et al. (2005), Klepper and Sleeper (2005), and Franco and Filson (2006).

[6] Such an investor could push the entrepreneur's share $\alpha$ to zero, and a claim paid entirely out of the scale-up would make even the most impatient entrepreneur stage, making the dampening of the initial boom

between the investor's secure share of a project's return and that project's full expected return.

**Assumption 1.** $(1-h)\,R\frac{p}{N} < \bar{r} < R\frac{p}{N}$.

A project of unknown fit is thus worth financing at face value, as in Remark 3.3, but not out of the investor's secure share of it alone. Below the lower bound in Assumption 1 the hold-up would never bind, and above the upper bound no project would be financed. The question is therefore how much capital and what equity share the investor will give the entrepreneur.

**Self-Financed Benchmark** The self-financed benchmark is the setting of Section 3 with the firm owning $\bar{K}$ outright and keeping every cash flow. It is also what external finance delivers when the entrepreneur has no leverage, $\eta = 0$, and the investor can be promised everything.

## 5.2 The Discoverer's Share and Pledgeable Value

**Her Reservation Share** Given a contract $(k_1, \alpha)$, the entrepreneur's claim is $\alpha$ of the experiments and $h$ of the scale-ups. Her problem is therefore the firm's problem of Section 3.2 but with the known-match value (3.8) scaled by $h/\alpha$. Extending Theorem 3.2 under external financing characterizes her optimal investment policy.

**Proposition 5.1.**

*(i) Given a contract, the entrepreneur carries out unit experiments—as in Theorem 3.2(ii)—if and only if her return to patience exceeds the cost of one round's delay,*

$$\beta(1-\delta)\left(1+\frac{h}{\alpha}\right) > 1. \tag{5.2}$$

*(ii) Equivalently, she carries out unit experiments if her share $\alpha$ is below her **reservation share**,*

$$\underbrace{\hat{\alpha}(\beta)}_{\text{reservation share}} \;\coloneqq\; \underbrace{h}_{\text{her part of the scale-up}} \cdot \underbrace{\frac{\beta(1-\delta)}{1-\beta(1-\delta)}}_{\text{value of waiting for it}}. \tag{5.3}$$

*Otherwise, she commits all capital at once—as in Theorem 3.2(i).*

---

stronger still.

The factor $1 + h/\alpha$ in (5.2) replaces the factor of two in Theorem 3.2—which is the self-financing case $\alpha = h = 1$. For the self-financed firm, the return to patience was two because deferred capital reaches the firm's fit by two routes—either today's experiment succeeds or it fails and the redraw succeeds—and the firm keeps the proceeds of both. The externally financed entrepreneur keeps $h$ of the first route and $\alpha$ of the second, so her return to patience is $(h + \alpha)/\alpha$.

A contract that pays the entrepreneur mainly out of the scale-up leads her to experiment one unit at a time; one that pays her mainly out of the experiments leads her to commit everything at once. Her reservation share $\hat{\alpha}(\beta)$ of (5.3) increases in her patience $\beta$: a patient entrepreneur values the option of self-discovery more, and so it takes a larger share of the firm's pre-discovery earnings to entice her to commit all capital now.

By Proposition 5.1, the entrepreneur's investment policy $x = (x_t)$ is thus either experimenting one unit at a time or gambling all capital at once. Under a policy $x$, the firm's expected discounted proceeds $V(x; k_1)$ split into those from the experimentation phase, $V^{\mathrm{e}}(x; k_1)$, and those from the scale-up, $V^{\mathrm{s}}(x; k_1)$. By Corollary 3.4, when $x$ is the experimenting policy,

$$V^{\mathrm{e}}(x; k_1) = R\,\frac{p}{N}\Big[\sum_{t=1}^{L-1} \beta^{t-1} + \beta^{L-1}\tilde{k}_L\Big], \qquad V^{\mathrm{s}}(x; k_1) = R\,\frac{p}{N}\sum_{t=1}^{L-1} \beta^{t}\,\tilde{k}_{t+1}, \tag{5.4}$$

with $\tilde{k}_t$ the feasible-testing path generated by $k_1$ and $L$ its horizon, as in (3.13) and (3.14).[7] When $x$ is the gambling policy there is no scale-up: $V^{\mathrm{e}}(x; k_1) = R\frac{p}{N}k_1$ and $V^{\mathrm{s}}(x; k_1) = 0$.

**The Investor's Contract Problem** The competitive investor chooses the contract $(k_1, \alpha)$ that maximizes the entrepreneur's payoff subject to his break-even (BE) and her incentive-compatibility (IC) constraints:

$$\max_{k_1 \le \bar{K},\, \alpha \in [0,1],\, x} \quad \alpha V^{\mathrm{e}}(x; k_1) + h V^{\mathrm{s}}(x; k_1) \tag{5.5}$$

$$\text{subject to} \quad x \in \arg\max_{x' \in \{\mathrm{e},\,\mathrm{g}\}} \ \alpha V^{\mathrm{e}}(x'; k_1) + h V^{\mathrm{s}}(x'; k_1), \tag{IC}$$

$$(1-\alpha) V^{\mathrm{e}}(x; k_1) + (1-h) V^{\mathrm{s}}(x; k_1) \ \ge\ \bar{r} k_1, \tag{BE}$$

where in (IC) the policy $x'$ is either experimenting (e) or gambling (g).

**Pledgeable Value of Self-Discovery** The investor's contract problem (5.5) is equivalent to choosing which policy to implement—experimenting or gambling—and how much capital

[7] The final round at $L$ is an experiment, so its proceeds sit in $V^{\mathrm{e}}$.

to commit, so as to maximize the surplus subject to his outlay being covered:

$$\max_{x\in\{\mathrm{e},\,\mathrm{g}\},\; k_1\le \bar{K}} \quad V(x;k_1)-\bar{r}k_1 \tag{5.6}$$

$$\text{subject to} \quad P(x;k_1)\;\ge\;\bar{r}k_1, \tag{5.7}$$

where $P(x;k_1)$, the *pledgeable income*, is the most payoff he can obtain among all of her equity shares that are incentive-compatible with the investment policy $x$:

$$P(x;k_1)\coloneqq \max_{\alpha\in[0,1]} (1-\alpha)V^{\mathrm{e}}(x;k_1)+(1-h)V^{\mathrm{s}}(x;k_1) \quad \text{s.t.} \quad x\in \arg\max_{x'\in\{\mathrm{e},\,\mathrm{g}\}} \alpha V^{\mathrm{e}}(x';k_1)+hV^{\mathrm{s}}(x';k_1). \tag{5.8}$$

The objective in (5.6) is the entrepreneur's payoff when (BE) binds: the investor is paid just his outlay, and she keeps the surplus $V-\bar{r}k_1$ over the cost of capital.[8]

**Lemma 5.1** (Pledgeable Income)**.**

*(i) Under experimenting, the investor's pledgeable income is attained at $\alpha=0$, the entrepreneur receiving nothing before discovery:*

$$P(\mathrm{e};k_1)=V^{\mathrm{e}}(\mathrm{e};k_1)+(1-h)\,V^{\mathrm{s}}(\mathrm{e};k_1). \tag{5.9}$$

*(ii) Under gambling, it is attained at the entrepreneur's reservation share $\hat{\alpha}(\beta)$ of* (5.3)*:*

$$P(\mathrm{g};k_1)=\Big(1-\hat{\alpha}(\beta)\Big)\,R\,\frac{p}{N}\,k_1. \tag{5.10}$$

In (5.9), the entrepreneur can pledge the experiments' proceeds in full: they are realized before she has anything to hold up. The scale-ups are pledgeable only up to $1-h$: she cannot pledge the option value of *her* self-discovery, and under the experimenting policy that is where most of the firm's value sits.

## 5.3 Capital Rationing

How much capital the investor commits depends on the financing constraint (5.7)—in particular, how the pledgeable income $P(x;k_1)$ grows with $k_1$. Under gambling, the financing constraint is scale-free (Lemma 5.1(ii)): $P(\mathrm{g};k_1)$ is proportional to $k_1$, so it covers $\bar{r}k_1$ at

[8](BE) need not bind at the optimum of (5.5), but this is immaterial. If under experimenting the (competitive) investor is covered at every share below the entrepreneur's reservation share $\hat{\alpha}(\beta)$ (5.3), he gives her that share and keeps a rent, so her payoff falls short of the surplus $V-\bar{r}k_1$; but the maximizer $(x,k_1)$ of (5.6) is still that of (5.5).

every commitment or at none—therefore, a gambling entrepreneur (if she can be financed at all) is financed at the full amount $\bar{K}$.

**Collateral Value of Experiments** Under experimenting, the financing constraint is no longer scale-free (Lemma 5.1(i)): with a small capital commitment, a marginal dollar of capital goes to the experiments, whose proceeds are pledgeable in full; with a large commitment, it goes to the scale-up, where only $1-h$ of its return can be pledged. The commitment is therefore determined by the *marginal value of commitment* $V'(\mathrm{e};k_1)$, what a marginal unit of capital adds to the value of the firm, and the *marginal pledgeable income* $P'(\mathrm{e};k_1)$, the part of $V'(\mathrm{e};k_1)$ that can be pledged to the investor. Because the entrepreneur captures the surplus, she wants the marginal capital when $V'(\mathrm{e};k_1)$ exceeds the cost of capital $\bar{r}$. On the other hand, the investor is paid for that marginal capital only when $P'(\mathrm{e};k_1)$ exceeds $\bar{r}$. In between—where a marginal unit is worth having but cannot pay for itself out of what the investor can be promised—the commitment grows until the pledgeable income of the inframarginal units just covers the total outlay $\bar{r}k_1$.

**Lemma 5.2.**

*(i) Under the experimenting policy, firm value and pledgeable income are piecewise affine in the commitment* $k_1$:[9]

$$V(\mathrm{e};k_1) = \underbrace{V'(\mathrm{e};k_1)}_{\substack{\textit{marginal value}\\ \textit{of commitment}}}\, k_1 \;+\; V^0(\mathrm{e};k_1), \tag{5.11}$$

$$P(\mathrm{e};k_1) = \underbrace{P'(\mathrm{e};k_1)}_{\substack{\textit{marginal}\\ \textit{pledgeable income}}}\, k_1 \;+\; \underbrace{P^0(\mathrm{e};k_1)}_{\substack{\textit{collateral value}\\ \textit{of experiments}}}. \tag{5.12}$$

*(ii) Where the two parties' returns to patience disagree,* $\frac{1}{2} < \beta(1-\delta) < 1/(2-h)$*, the marginal value of commitment* $V'(\mathrm{e};k_1)$ *rises with the commitment* $k_1$ *while the marginal pledgeable income* $P'(\mathrm{e};k_1)$ *falls: as firm value grows with the capital committed, more of it accrues to the entrepreneur's finding her fit, and she is less willing to pledge that value.*

The part of pledgeable income (5.12) that does not scale with the commitment (its intercept) is the **collateral value of experiments**

$$P^0(\mathrm{e};k_1) \coloneqq P(\mathrm{e};k_1) - P'(\mathrm{e};k_1)\,k_1, \tag{5.13}$$

[9]The slopes $V'$, $P'$ and the intercepts $V^0$, $P^0$ are the same for all commitments that afford the same number of experimentation rounds $L$ of (3.14), and change only when the commitment affords one more.

their proceeds net of the pledgeable value of the capital they draw down. The investor's financing constraint (5.7) therefore holds if and only if

$$P^0(\mathrm{e}; k_1) \;\geq\; \Big(\bar{r} - P'(\mathrm{e}; k_1)\Big)\, k_1,$$

that is, if the collateral value of the experiments covers the shortfall on the capital committed.[10]

Under external financing, therefore, the entrepreneur gambles if and only if her reservation share $\hat{\alpha}(\beta)$ leaves the investor enough to cover his cost of capital. If her reservation share leaves him short, she experiments, and how much capital he gives her depends on whether a marginal unit of reserve pays for itself and, if not, on how far the collateral value of the experiments stretches.

Denote by $\rho$ the cost of capital in units of a project's expected return,

$$\rho \coloneqq \frac{\bar{r}}{R\,\frac{p}{N}} \qquad \text{(normalized cost of capital)}, \tag{5.14}$$

so that Assumption 1 reads $1 - h < \rho < 1$, and define

$$\underline{\beta}\,(1-\delta) \coloneqq \frac{1-\rho}{1-\rho+h}, \qquad \bar{\beta}\,(1-\delta) \coloneqq \frac{\rho}{\rho+1-h}.$$

The band $(\underline{\beta}, \bar{\beta})$ straddles the patience cutoff $\frac{1}{2(1-\delta)}$ of Theorem 3.2.

**Theorem 5.2** (Equilibrium Contract: Capital Rationing)**.** *The equilibrium contract $(k_1^*, \alpha^*)$ depends on the entrepreneur's patience $\beta$ as follows.*

*(i)* (Gambling at full scale.) *If $\beta \leq \underline{\beta}$, she gambles on the full commitment $k_1^* = \bar{K}$. Her equity share is the project's return net of the cost of capital, $\alpha^* = 1 - \rho$.*

*(ii)* (Rationed experimenting.) *If $\underline{\beta} < \beta < \bar{\beta}$, she experiments on a rationed capital commitment $k_1^* < \bar{K}$,*

$$k_1^* = \frac{P^0(\mathrm{e}; k_1^*)}{\bar{r} - P'(\mathrm{e}; k_1^*)}, \tag{5.15}$$

*at which the collateral value of experiments $P^0$ of (5.13) just covers the shortfall $\bar{r} - P'$ on each unit of capital. She receives none of the pre-discovery earnings, $\alpha^* = 0$, and is paid only out of the scale-up. For an impatient entrepreneur, this rationing can be so severe that $k_1^*$ allows her to invest for only one round.*[11]

[10] In this sense, $P^0(\mathrm{e}; k_1)$ plays the same role as the inside wealth in Holmström and Tirole (1997).

[11] In such cases of severe rationing, $k_1^*$ is just short of the $1 + \frac{1}{1-\delta}$ that would afford a second round; there is no capital for post-discovery scale-up, and she receives the share $\alpha^* = 1 - \rho$ of (i) instead of $\alpha^* = 0$.

*(iii)* (Experimenting at full scale.) *If $\beta \geq \bar{\beta}$, she experiments on the full commitment $k_1^* = \bar{K}$, and her equity share is*

$$\alpha^* = \min\left\{1 - \frac{\bar{r}\bar{K} - (1-h)\,V^{\mathrm{s}}(\mathrm{e};\bar{K})}{V^{\mathrm{e}}(\mathrm{e};\bar{K})}\,,\ \hat{\alpha}(\beta)\right\}.$$

*When the first term is smaller, the investor just breaks even; otherwise she gets her reservation share $\hat{\alpha}(\beta)$, and he keeps a rent.*

**Impatient Entrepreneurs Who Experiment** Among the rationed entrepreneurs of Theorem 5.2(ii) are impatient ones, who would have gambled with their own capital—those with $\beta$ below the cutoff $\frac{1}{2(1-\delta)}$ of Theorem 3.2(i). Under external financing their reservation share leaves the investor short, so the gamble cannot be financed; the only contract the two parties agree on loads the entrepreneur's claim onto the scale-up, which leads her to experiment, on the rationed commitment $k_1^*$ of (5.15). (Either this, or the rationing is so severe that she can invest for only one round—no capital held back for a second round would pay for itself.) This conversion from gambling to experimenting begins with the least impatient entrepreneurs, those nearest the patience cutoff.

**The Capital-Rationed Band** Above the upper boundary $\bar{\beta}$ there is no rationing: the firm is patient enough that the investor's share of the deferred scale-up covers his cost of capital on a marginal unit of commitment. Below the lower boundary $\underline{\beta}$, rationing is also absent but for the opposite reason: the entrepreneur is impatient enough that her reservation share $\hat{\alpha}(\beta)$ is small, and paying it still leaves the investor his (normalized) cost of capital $\rho$ of (5.14), so the gamble is financed at full scale.

In the rationed band $(\underline{\beta}, \bar{\beta})$, the potential hold-up drives a wedge between the pace of investment each would prefer: she, paid out of the scale-up, holds capital back to wait for discovery; he, held up at the scale-up, would prefer his claim be already secure upon discovery. From his perspective, a marginal unit of capital no longer pays for itself, and his commitment is limited to the amount on which the collateral value of the experiments (5.13) covers the shortfall.

**Theorem 5.3.** *The rationed band $(\underline{\beta}, \bar{\beta})$ widens with the cost of capital $\rho$ and with the entrepreneur's leverage $h$: the lower boundary $\underline{\beta}$ falls and the upper boundary $\bar{\beta}$ rises in both $\rho$ and $h$.*

More costly capital rations more entrepreneurs. The investor must capture more of each project's return, so fewer entrepreneurs can be paid their reservation share and still be

financed as gamblers. The band widens at both ends: more impatient gamblers are turned into experimenters, and more patient experimenters are rationed. As the cost of capital approaches a project's expected return, $\rho \to 1$, every impatient entrepreneur experiments on rationed capital.

A more portable discovery also leads to more rationing of capital, through her claim rather than his cost. The more of the scale-up she can keep, the larger the reservation share it takes for her to commit all capital at once, so fewer gambles can be financed. On the investor side, the less of the scale-up is left to him, the lower the marginal pledgeable income at every commitment, so a unit of reserve capital pays for itself for fewer experimenters. Rationing is thus more extensive when the discovery belongs more to the discoverer and less to the asset.

## 5.4 Dampened Boom–Bust

We now return to the population of Section 3.3, now with each entrepreneur financed on the terms of Theorem 5.2.

**Theorem 5.4.** *Relative to the self-financed benchmark with the same $\bar{K}$:*

*(i)* (A lower, shorter boom.) *Rationed experimenters, $\frac{1}{2(1-\delta)} < \beta < \bar{\beta}$, follow the benchmark's experimenting path on a smaller commitment: they run fewer rounds and scale up less at each, so their boom is lower and shorter in the sense of Figure 4.*

*(ii)* (The spike cut down.) *In addition to being rationed, the impatient entrepreneurs in the band, $\underline{\beta} < \beta \leq \frac{1}{2(1-\delta)}$, no longer commit everything at $t = 1$; their reduced capital arrives over the rounds of the experimenting path—or the rationing is so severe that it is enough for only one round.*

*(iii)* (Self-limiting capital abundance.) *An increase in the fund size $\bar{K}$ raises aggregate investment only through the entrepreneurs outside the rationed band, and leaves the paths of those inside it unchanged.*

The benchmark's comparative static in access to capital, Figure 4, thus survives external finance only in part. For those in the rationed band, the commitment is set by the collateral value of their experiments, not by the size of the fund, and more capital in the fund changes nothing. Capital abundance is therefore self-limiting (Theorem 5.4(iii)). External finance, on this account, does not add a speculative layer on top of the self-discovery motive. It disciplines that motive, because the discovery belongs to the discoverer and not to the firm.

**The Conventional View** Accounts in which investors fuel booms share a common assumption: capital supplied is capital invested, at a time of the investor's choosing. More

money chasing deals raises what is funded (Sahlman and Stevenson (1985); Gompers and Lerner (2000); Gompers et al. (2008); Nanda and Rhodes-Kropf (2013)), and investors with short horizons bring it forward (Brown et al. (1996); Bushee (1998); Jovanovic and Szentes (2013)). Theorem 5.4 suggests the opposite view: for many ventures funding is limited by how much of the discovery the investor can secure, and an investor who wants results sooner can only commit less capital.

**Remark 5.5** (Stage Financing)**.** *The experimenting contracts of Theorem 5.2 are equivalent to the staged contracts observed in venture capital (Sahlman (1990); Gompers (1995); Neher (1999)): the investor releases the capital a round at a time, and the balance when the entrepreneur finds fit. Rationing thus precedes staging: the rationed commitment determines how much there is to stage. Theorem 5.3 then speaks to staging as well as rationing. Where the entrepreneur has more leverage, staging is more extensive: more gamblers are turned into experimenters, whose capital is staged, and on smaller commitments. This is consistent with the finding of Gompers (1995) that capital is staged most tightly in industries rich in intangible assets and research.*

# 6 Conclusion

Firms invest not only for expected cash flows, but also to learn where their own capabilities apply. This observation is familiar from the entrepreneurial experience of building at the frontier but underexplored in the literature—comparative advantage is often understood only retrospectively, after firms have tried to build with the technology themselves. Investment is therefore not merely capital allocation under external uncertainty, but also an act of introspective learning through which firms discover whether their own capabilities fit the new domain. We have formalized this motive in a stylized model and traced its predictions in the dot-com boom.

In addition to the recognizable boom–bust patterns of new-technology episodes, the model further predicts where these dynamics should be strongest: in technologies accessible to a broader range of firms, among more impatient investors, and when capital is abundant. That this full set of dynamics arises from introspective learning alone—without externalities or irrationality—suggests that the self-discovery motive may be a first-order driver of boom–bust episodes.

# A Appendix

## A.1 Proof of Proposition 3.1

Fix a searching state $(k, m)$ at date $t$ with $k \geq 1$, and consider any feasible continuation policy $\pi$ that begins with $x_t = 0$.

If $\pi$ never invests at any later date, then its payoff is zero. But since $k \geq 1$, the firm can instead undertake a unit experiment immediately, which yields expected payoff at least $\mu(m)R > 0$. Hence any policy that never invests is strictly dominated.

Now suppose $\pi$ does invest at some future date. Let $r \geq 1$ denote the number of initial waiting periods under $\pi$, so that

$$x_t = x_{t+1} = \cdots = x_{t+r-1} = 0, \qquad x_{t+r} > 0.$$

Because no information arrives while waiting, after these $r$ periods of pure delay the state becomes $\Big((1-\delta)^r k,\, m\Big)$.

Construct a new policy $\hat{\pi}$ that shifts this entire post-wait plan forward by one period: it waits only $r-1$ periods, reaches state $\Big((1-\delta)^{r-1}k,\, m\Big)$ at date $t+r-1$, and from that point onward takes exactly the same investment actions as $\pi$ would take one period later after the same sequence of success/failure outcomes.

The first shifted investment is feasible because under $\pi$,

$$x_{t+r} \leq (1-\delta)^r k < (1-\delta)^{r-1} k,$$

so the same investment can be undertaken one period earlier under $\hat{\pi}$. More generally, fix any realized sequence of subsequent success/failure outcomes, and let $\hat{k}_s$ denote the capital under $\hat{\pi}$ at date $s$, while $k^{\pi}_{s+1}$ denotes the capital under $\pi$ one period later along the corresponding history. After the first shifted investment,

$$\hat{k}_{t+r} = (1-\delta)\Big((1-\delta)^{r-1}k - x_{t+r}\Big) > (1-\delta)\Big((1-\delta)^r k - x_{t+r}\Big) = k^{\pi}_{t+r+1}.$$

If at some subsequent date $\hat{k}_s > k^{\pi}_{s+1}$, then because $\hat{\pi}$ and $\pi$ choose the same next investment $x$ along the corresponding history,

$$\hat{k}_{s+1} = (1-\delta)(\hat{k}_s - x) > (1-\delta)(k^{\pi}_{s+1} - x) = k^{\pi}_{s+2}.$$

By induction, the shifted policy always has strictly more capital at each corresponding future date. Hence every action taken by $\pi$ after its initial waiting spell remains feasible under $\hat{\pi}$

one period earlier.

Along any fixed sequence of success/failure realizations, $\hat{\pi}$ generates the same sequence of investment payoffs as $\pi$, but shifted one period earlier. Therefore the discounted payoff under $\hat{\pi}$ is weakly larger pathwise. It is strictly larger on any history in which one of those shifted experiments succeeds, because the same positive cash flow is then received one period earlier. Such a history has strictly positive probability, since the first positive experiment under $\pi$ succeeds with conditional probability $\mu(m) > 0$. Therefore $\hat{\pi}$ yields strictly higher expected payoff than $\pi$.

Since $\pi$ was arbitrary, every feasible continuation policy with initial choice $x_t = 0$ is strictly dominated. Hence waiting is never optimal when $k \geq 1$. This proves the proposition.

## A.2 Proof of Theorem 3.2

Since the right-hand side of (3.11) is piecewise affine and convex in $x$, the optimal policy must be bang-bang: either investing minimally, $x = 1$ (a unit experiment), or investing all available capital, $x = k$ (an all-in gamble). Accordingly, (3.11) further reduces to

$$V_t(k,m) = \max\{\underbrace{E_t(k,m)}_{\text{experiment}}, \underbrace{G_t(k,m)}_{\text{gamble}}\}, \tag{A.1}$$

where

$$\begin{aligned} E_t(k,m) &:= \mu(m)\Big(R + \beta\, W_{t+1}\Big((1-\delta)(k-1)\Big)\Big) \\ &\quad + \Big(1-\mu(m)\Big)\beta\, V_{t+1}\Big((1-\delta)(k-1),\, m-1\Big), \\ G_t(k,m) &:= \mu(m)\, Rk. \end{aligned} \tag{A.2}$$

We solve the Bellman equation (A.1) by making the following ansatz. Fix $(k,m)$ with $m \geq 1$ and $k \geq 1$, and consider the restricted plan: *test now* ($x = 1$) and then, if still searching next period, *gamble all-in* (i.e., invest $x = (1-\delta)(k-1)$ at $t+1$). Let $\widehat{E}_t(k,m)$ denote the expected payoff from this plan. When $(1-\delta)(k-1) \geq 1$ (so an additional unit test would be feasible next period), this plan yields

$$\widehat{E}_t(k,m) = \mu(m)\Big(R + \beta\, R(1-\delta)(k-1)\Big) + \Big(1-\mu(m)\Big)\beta\Big(\mu(m-1)\, R(1-\delta)(k-1)\Big), \tag{A.3}$$

because success at $t$ implies a known match at $t+1$ with payoff $R(1-\delta)(k-1)$, while failure at $t$ leaves the firm searching with $m-1$ remaining types and leads it to gamble at $t+1$, succeeding with probability $\mu(m-1)$.

Using

$$\mu(m-1) = \frac{\mu(m)}{1-\mu(m)},$$

the expression (A.3) for $\widehat{E}_t(k,m)$ simplifies to

$$\widehat{E}_t(k,m) = \mu(m)\Big(R + 2\beta\, R(1-\delta)(k-1)\Big). \tag{A.4}$$

Now the difference between this restricted-plan payoff and the all-in gamble payoff is

$$\widehat{E}_t(k,m) - G_t(k,m) = \mu(m)R\Big(1 + 2\beta(1-\delta)(k-1) - k\Big) = \mu(m)R\,(k-1)\,\Big(2\beta(1-\delta) - 1\Big).$$

Hence, for all $m \geq 1$ and $k > 1$,

$$\begin{cases} \widehat{E}_t(k,m) < G_t(k,m), & \text{if } \beta(1-\delta) < \frac{1}{2}, \\ \widehat{E}_t(k,m) > G_t(k,m), & \text{if } \beta(1-\delta) > \frac{1}{2}, \\ \widehat{E}_t(k,m) = G_t(k,m), & \text{if } \beta(1-\delta) = \frac{1}{2}. \end{cases}$$

This comparison yields the claimed characterization of optimal policy:

- If $\beta(1-\delta) < \frac{1}{2}$, then along the feasible-testing capital path $\{\tilde{k}_t\}_{t=1}^{L}$ we have $\widehat{E}_t(\tilde{k}_t, m) < G_t(\tilde{k}_t, m)$ whenever testing is feasible. Backward induction along the feasible-testing path therefore implies that the agent gambles all-in immediately at $t = 1$.
- If $\beta(1-\delta) > \frac{1}{2}$, then $\widehat{E}_t(k,m) > G_t(k,m)$ for all $k > 1$ implies $E_t(k,m) \geq \widehat{E}_t(k,m) > G_t(k,m)$. Thus gambling is never optimal at any $t < L$. The agent tests in every feasible searching period and gambles only at the last feasible date $t = L$ if still searching.
- If $\beta(1-\delta) = \frac{1}{2}$, then $\widehat{E}_t(k,m) = G_t(k,m)$ and the agent is indifferent between gambling and experimenting whenever both are feasible.

## A.3 Proof of Corollary 3.4

Statement (i) follows immediately from Theorem 3.2(i). We therefore consider case (ii), where $\beta(1-\delta) > \frac{1}{2}$.

By Theorem 3.2, the optimal policy $\pi_L$ is: while searching, choose $x_t = 1$ for all $t < L$; if still searching at $t = L$, choose the all-in gamble $x_L = \tilde{k}_L$ (absorbing). Along this feasible unit-experiment capital path, the number of remaining project types is given by $m_t = N - t + 1$, for $t = 1, \ldots, L$.

Consider the searching-state continuation value under $\pi_L$ along this feasible-testing path,

$$V_t^{\pi_L}(\tilde{k}_t, m_t), \qquad t = 1, \ldots, L.$$

In particular, $V_1^{\pi_L}(\tilde{k}_1, m_1) = V_1(k_1, N)$. Evaluating the Bellman equation (3.11) along this optimal policy path yields, for $t = 1, \ldots, L-1$,

$$V_t^{\pi_L}(\tilde{k}_t, m_t) = \mu(m_t)\Big(R + \beta R\tilde{k}_{t+1}\Big) + \Big(1 - \mu(m_t)\Big)\beta\, V_{t+1}^{\pi_L}(\tilde{k}_{t+1}, m_{t+1}),$$

with terminal condition

$$V_L^{\pi_L}(\tilde{k}_L, m_L) = \mu(m_L)\, R\, \tilde{k}_L.$$

Iterating this recursion forward gives

$$\begin{aligned} V_1(k_1, N) &= V_1^{\pi_L}(\tilde{k}_1, m_1) \\ &= \sum_{s=1}^{L-1} \beta^{s-1} \left(\prod_{u=1}^{s-1}\Big(1 - \mu(m_u)\Big)\right) \mu(m_s)\Big(R + \beta R\tilde{k}_{s+1}\Big) \\ &\qquad + \beta^{L-1} \left(\prod_{u=1}^{L-1}\Big(1 - \mu(m_u)\Big)\right) \mu(m_L)\, R\, \tilde{k}_L, \end{aligned} \tag{A.5}$$

where, by convention, $\prod_{u=1}^{0}(\cdot) = 1$.

Now fix $s \le L$ and let $S_s$ denote the event that the agent is still searching at the start of period $s$. Then

$$\left(\prod_{u=1}^{s-1}\Big(1 - \mu(m_u)\Big)\right) \mu(m_s) = \Pr(S_s)\ \Pr(\text{success at } s \mid S_s) = \Pr(\text{first success occurs at } s).$$

Along the feasible-testing path, $\Pr(S_s) = (1-p) + m_s \frac{p}{N}$ and $\Pr(\text{success at } s \mid S_s) = \mu(m_s) = \frac{\frac{p}{N}}{(1-p) + m_s \frac{p}{N}}$. Hence

$$\left(\prod_{u=1}^{s-1}\Big(1 - \mu(m_u)\Big)\right) \mu(m_s) = \Big((1-p) + m_s \tfrac{p}{N}\Big) \cdot \frac{\frac{p}{N}}{(1-p) + m_s \frac{p}{N}} = \frac{p}{N}.$$

Substituting into the unrolled Bellman expression (A.5) yields

$$V_1(k_1, N) = \frac{pR}{N}\left[\sum_{s=1}^{L-1}\Big(\beta^{s-1} + \beta^s \tilde{k}_{s+1}\Big) + \beta^{L-1}\tilde{k}_L\right],$$

which proves the claim.

## A.4 Proof of Proposition 3.5

The impatient-firm case is immediate. We consider the patient firms. Let $S_t$ denote the event "still searching at the start of period $t$." Along the feasible experimentation path, define

$$m_t \coloneqq N - t + 1.$$

Then

$$\Pr(S_t) = (1-p) + m_t \frac{p}{N} = 1 - \frac{p}{N}(t-1), \qquad t = 1, \dots, L.$$

Moreover, for $s \leq L-1$ (while the agent is still experimenting),

$$\Pr(\text{first success at } s) = \frac{p}{N}.$$

In period $t = 1$, all patient agents carry out one unit experiment, so

$$\text{Inv}_{\text{pat},1} = \pi_{\text{pat}} \cdot 1.$$

In periods $t = 2, \dots, L-1$, a mass $\Pr(S_t)$ of patient agents is still searching at the start of period $t$ and therefore invests 1 in that period. In addition, a mass $\Pr(\text{first success at } t-1) = \frac{p}{N}$ succeeded in period $t-1$ and thus invests all remaining capital $\tilde{k}_t$ in period $t$. Hence,

$$\text{Inv}_{\text{pat},t} = \pi_{\text{pat}} \left[ \Pr(S_t) \cdot 1 + \frac{p}{N} \tilde{k}_t \right], \qquad t = 2, \dots, L-1.$$

In the final feasible period $t = L$, those who first succeed at $L-1$ invest $\tilde{k}_L$ in period $L$, and those still searching at the start of period $L$ also invest $\tilde{k}_L$ as the terminal all-in action. Therefore,

$$\text{Inv}_{\text{pat},L} = \pi_{\text{pat}} \left[ \left( \frac{p}{N} + \Pr(S_L) \right) \tilde{k}_L \right] = \pi_{\text{pat}} \left[ \left( 1 - \frac{p}{N}(L-2) \right) \tilde{k}_L \right].$$

This establishes the claimed per-period investment expressions for patient agents.

## A.5 Proof of Proposition 4.2

Fix $t \leq L$. Let $S_t$ denote the event that a firm is still searching at the start of period $t$, and let $F_t$ denote the event that it finds fit for the first time in period $t$. Under Theorem 3.2(ii), an unmatched firm runs a unit experiment in every feasible searching period, and each funded failure removes one distinct candidate type. Hence, after $t-1$ unsuccessful tests, the number

of remaining candidate types is

$$|M_t| = N - t + 1.$$

A firm is still searching at the start of period $t$ if and only if either $\theta = 0$ or $\theta \in M_t$. Using the prior in (3.2),

$$\Pr(S_t) = (1 - p) + |M_t|\frac{p}{N} = 1 - \frac{p}{N}(t - 1).$$

This is the search-pool formula in part (i).

Conditional on still searching, the probability that the current draw matches the firm's type is the match probability in (3.12). Therefore

$$\Pr(F_t \mid S_t) = \mu(|M_t|) = \frac{p/N}{1 - \frac{p}{N}(t-1)}.$$

Multiplying by $\Pr(S_t)$ gives

$$\Pr(F_t) = \Pr(S_t)\Pr(F_t \mid S_t) = \frac{p}{N}.$$

Thus each period produces the same mass $p/N$ of new winners. Summing over $s = 1, \ldots, t$ yields

$$\sum_{s=1}^{t} \Pr(F_s) = \frac{tp}{N},$$

which proves part (i).

The same expression for $\Pr(F_t \mid S_t)$ shows that the success hazard rises over time. Equivalently, the conditional failure probability declines over time, consistent with Proposition 4.1. At the same time, Bayes' rule gives

$$\Pr(\theta = 0 \mid S_t) = \frac{1 - p}{1 - \frac{p}{N}(t-1)},$$

which is increasing in $t$. Hence the firms that remain in search are increasingly tilted toward true no-fit types. This proves part (ii).

Finally, if a firm first discovers fit earlier, it reaches the known-match regime with more remaining capital. Along the unit-experiment path, remaining capital evolves according to (3.13), which is decreasing over time. Hence early winners scale sooner and at larger size than late winners, proving part (iii).

## External Financing: The Investor's Capital Held in Reserve

In the proofs for Section 5 that follow, let $b \coloneqq \beta(1-\delta)$ denote what a unit of capital carried one round is worth today. $\tilde{k}_t$ and $L$ are still the feasible-testing path (3.13) and horizon (3.14) generated by the commitment $k_1$. At each date $t < L$ the entrepreneur spends one unit of the investor's commitment on the experiment and holds back the rest, $\tilde{k}_t - 1$: the *reserve* at $t$, which carried to the next round is the remaining capital $k'$ of Section 5.1. Its discounted sum over the rounds it is held is

$$D(k_1) \coloneqq \sum_{t=1}^{L-1} \beta^{t-1}\Big(\tilde{k}_t - 1\Big). \tag{A.6}$$

Since $\beta^t\,\tilde{k}_{t+1} = b\,\beta^{t-1}(\tilde{k}_t - 1)$ by (3.13), the two parts of (5.4) are

$$V^{\mathrm{s}}(\mathrm{e}; k_1) = R\,\frac{p}{N}\,b\,D(k_1), \qquad V^{\mathrm{e}}(\mathrm{e}; k_1) = R\,\frac{p}{N}\Big[k_1 - (1-b)\,D(k_1)\Big], \tag{A.7}$$

the second because $\sum_{t=1}^{L-1}\left(\beta^{t-1}\tilde{k}_t - \beta^t\tilde{k}_{t+1}\right) = k_1 - \beta^{L-1}\tilde{k}_L$. A unit held in reserve gives up a unit of today's experiment and comes back, worth $b$, by each of the two routes behind Theorem 3.2: as scale-up capital if today's experiment succeeds, as next round's capital if it fails. Gambling holds no reserve.

Over the commitments that afford the same horizon $L$, $\tilde{k}_t - 1 = (1-\delta)^{t-1}k_1 - \sum_{s=0}^{t-1}(1-\delta)^s$, so $D(k_1)$ is affine,

$$D(k_1) = D'_L\,k_1 - D^0_L, \qquad D'_L \coloneqq \sum_{t=1}^{L-1} b^{\,t-1}, \qquad D^0_L \coloneqq \sum_{t=1}^{L-1}\beta^{t-1}\sum_{s=0}^{t-1}(1-\delta)^s. \tag{A.8}$$

At the commitment that first affords an $(L+1)$-th round, $\tilde{k}_{L+1} = 1$ and $D(k_1)$ jumps up by $\beta^{L-1}(\tilde{k}_L - 1) = \beta^{L-1}/(1-\delta)$. $D(k_1)$ is therefore nondecreasing in the commitment, and $D(k_1) \le D'_L\,k_1 < k_1/(1-b)$.

## Proof of Proposition 5.1

Under the contract the entrepreneur receives $\alpha$ of each experiment's proceeds and $h$ of each scale-up's. For $\alpha > 0$, her Bellman equation is (3.9)–(3.11) with $R$ replaced by $\alpha R$ in the current-period term and the known-match value (3.8) replaced by $hW$; dividing through by $\alpha$ gives the benchmark problem with $W$ scaled by $\lambda \coloneqq h/\alpha$. Waiting remains dominated, by the argument of Proposition 3.1. Since $(1-\mu(m))\,\mu(m-1) = \mu(m)$ by (3.12), the value function of the scaled problem factors as $V_t(k,m) = \mu(m)\,R\,v_t(k)$, where $v_t(k) = 0$ for $k < 1$,

$v_T(k) = k$ for $k \geq 1$ because nothing follows date $T$, and for $t < T$ and $k \geq 1$

$$v_t(k) = \max_{x \in [1,k]} \Big\{ x + \beta \Big[ \lambda\, k'\, \mathbf{1}_{\{k' \geq 1\}} + v_{t+1}(k') \Big] \Big\}, \qquad k' = (1-\delta)(k-x). \tag{A.9}$$

*If* $(1+\lambda)\, b \leq 1$, then $v_t(k) = k$ at every date, attained at $x = k$: if $v_{t+1}(k') = k'$ for $k' \geq 1$, the bracket in (A.9) is at most $(1+\lambda)\, k'$, and the maximand is at most $x + (1+\lambda)\, b\, (k-x) \leq k$. She commits all capital at once.

*If* $(1+\lambda)\, b > 1$, suppose that $v_{t+1}(k'') - v_{t+1}(k') \geq k'' - k'$ whenever $k'' \geq k' \geq 1$, as holds at $T$. The bracket in (A.9) then equals $1+\lambda$ at $k' = 1$ and rises by at least $1+\lambda$ per unit of $k'$ beyond it, so raising $x$ by a unit lowers the maximand by at least $(1+\lambda)\, b - 1 > 0$ for as long as the capital carried forward affords another experiment. When $(1-\delta)(k-1) \geq 1$, the unit experiment thus beats every larger one that leaves another experiment feasible; and it beats those that do not, each worth $x \leq k$, because it is worth at least $1 + (1+\lambda)\, b\, (k-1) > k$. When $(1-\delta)(k-1) < 1$, no reserve is of use and she commits everything. The hypothesis passes to $v_t$, which equals $k$ below $1 + 1/(1-\delta)$, the least capital that affords a second round, jumps up there, and rises by at least $(1+\lambda)\, b$ per unit above it. By induction she invests one unit at a time while another experiment remains feasible and commits the remainder at $t = L$, the policy of Theorem 3.2(ii).

*At* $\alpha = 0$ she is paid only out of the scale-ups: an experiment larger than one unit earns her nothing and costs her reserve, so she experiments one unit at a time, and (5.2) holds with its left side infinite. This proves (i).

(ii) Rearranging, $b\,(1 + h/\alpha) > 1$ if and only if $\alpha < h\, b/(1-b) = \hat{\alpha}(\beta)$. In terms of the reserve, by (A.7), experimenting on $k_1$ pays her

$$\alpha V^{\mathrm{e}}(\mathrm{e}; k_1) + h V^{\mathrm{s}}(\mathrm{e}; k_1) = R\, \frac{p}{N} \Big[ \alpha\, k_1 + \Big( (\alpha + h)\, b - \alpha \Big) D(k_1) \Big], \tag{A.10}$$

and gambling on it pays her $R \frac{p}{N}\, \alpha\, k_1$. The difference has the sign of $\hat{\alpha}(\beta) - \alpha$; at $\alpha = \hat{\alpha}(\beta)$ she is indifferent.

## Proof of Lemma 5.1

By Proposition 5.1, experimenting satisfies (IC) if and only if $\alpha \leq \hat{\alpha}(\beta)$, and gambling if and only if $\alpha \geq \hat{\alpha}(\beta)$. Under either policy the investor's receipts $(1-\alpha)V^{\mathrm{e}} + (1-h)V^{\mathrm{s}}$ fall in $\alpha$, so (5.8) is attained at the smallest share that induces the policy. Under experimenting that share is $\alpha = 0$, which gives (5.9). Under gambling it is $\alpha = \hat{\alpha}(\beta)$, where $V^{\mathrm{e}} = R \frac{p}{N} k_1$ and $V^{\mathrm{s}} = 0$, which gives (5.10). A commitment below $1 + 1/(1-\delta)$ affords a single round, on which the two policies coincide and (5.9) applies; and when $\hat{\alpha}(\beta) > 1$ no share induces

the gamble, so that (5.10), then negative, fails (5.7) as it should.

By (A.7),

$$V(\mathrm{e};k_1)=R\,\frac{p}{N}\Big[k_1+(2b-1)\,D(k_1)\Big],\qquad P(\mathrm{e};k_1)=R\,\frac{p}{N}\Big[k_1-\Big(1-(2-h)\,b\Big)\,D(k_1)\Big]. \tag{A.11}$$

Against the unit of experiment it replaces, a unit of reserve is worth $2b$ to the firm and $(2-h)\,b$ to the investor.

## Proof of Lemma 5.2

(i) By (A.8) and (A.11), over the commitments that afford $L$ rounds $V(\mathrm{e};k_1)$ and $P(\mathrm{e};k_1)$ are affine, with slopes

$$\begin{aligned} V'(\mathrm{e};k_1)&=R\,\frac{p}{N}\Big[1+(2b-1)\,D'_L\Big]=R\,\frac{p}{N}\Big[\sum_{t=1}^{L-1}b^t+b^{L-1}\Big],\\ P'(\mathrm{e};k_1)&=R\,\frac{p}{N}\Big[1-\Big(1-(2-h)\,b\Big)\,D'_L\Big]=R\,\frac{p}{N}\Big[(1-h)\sum_{t=1}^{L-1}b^t+b^{L-1}\Big]. \end{aligned} \tag{A.12}$$

A marginal unit of commitment leaves the unit experiments as they are: it waits, depreciating, and is deployed in the scale-up that follows the first success—the sums—or in the final round—the last terms.

(ii) $D'_{L+1}-D'_L=b^{L-1}$, and $L$ is nondecreasing in $k_1$. Each further round the commitment affords therefore moves $V'$ by $R\frac{p}{N}\,b^{L-1}(2b-1)$ and $P'$ by $R\frac{p}{N}\,b^{L-1}\Big((2-h)\,b-1\Big)$: up and down, respectively, when $\frac{1}{2}<b<1/(2-h)$.

(iii) By (i), $P^0(\mathrm{e};k_1)=R\frac{p}{N}\Big(1-(2-h)\,b\Big)D^0_L$. Write $P'_j$ for the marginal pledgeable income of a commitment that affords $j$ rounds, so that $R\frac{p}{N}-P'_j=R\frac{p}{N}\Big(1-(2-h)\,b\Big)D'_j$. Collecting the terms $\beta^{t-1}(1-\delta)^s=\beta^{t-s-1}\,b^s$ of $D^0_L$ by the date $\tau=t-s$ at which a unit is drawn down gives $D^0_L=\sum_{\tau=1}^{L-1}\beta^{\tau-1}D'_{L-\tau+1}$, hence

$$P^0(\mathrm{e};k_1)=\sum_{\tau=1}^{L-1}\beta^{\tau-1}\Big[R\,\frac{p}{N}-P'_{L-\tau+1}\Big].$$

The experiment at date $\tau$ brings in its expected proceeds $R\frac{p}{N}$, pledgeable in full, and draws down one unit of capital that, with $L-\tau+1$ rounds remaining, would have been worth $P'_{L-\tau+1}$ to the investor. Every term is positive when $(2-h)\,b<1$. The equivalence stated in the lemma is (5.13) substituted into (5.7).

## Proof of Theorem 5.2

*The share, given the policy and the commitment.* The entrepreneur's payoff $\alpha V^{\mathrm{e}} + hV^{\mathrm{s}}$ rises in $\alpha$, so the contract carries the largest share that satisfies (IC) and (BE).

Under gambling, (IC) requires $\alpha \geq \hat{\alpha}(\beta)$ and (BE) requires $(1-\alpha)R\frac{p}{N} \geq \bar{r}$, that is, $\alpha \leq 1-\rho$. The gamble can be financed if and only if $\hat{\alpha}(\beta) \leq 1-\rho$, which rearranges to $\beta \leq \underline{\beta}$, and then at every commitment; her share is $1-\rho$, and her payoff $(1-\rho)R\frac{p}{N}k_1$ is largest at $\bar{K}$.

Under experimenting, (IC) requires $\alpha \leq \hat{\alpha}(\beta)$, and (BE) requires $\alpha \leq \alpha^{\mathrm{BE}}(k_1) \coloneqq 1 - \big[\bar{r}k_1 - (1-h)V^{\mathrm{s}}(\mathrm{e};k_1)\big]/V^{\mathrm{e}}(\mathrm{e};k_1)$, which is nonnegative if and only if (5.7) holds. Her share is $\min\{\alpha^{\mathrm{BE}}(k_1),\, \bar{\alpha}\}$, with $\bar{\alpha} \coloneqq \min\{\hat{\alpha}(\beta), 1\}$, and her payoff is the smaller of

$$V(\mathrm{e};k_1) - \bar{r}k_1 \qquad \text{and} \qquad \bar{\alpha}V^{\mathrm{e}} + hV^{\mathrm{s}} = R\,\frac{p}{N}\Big[\bar{\alpha}\,k_1 + \Big((\bar{\alpha}+h)\,b - \bar{\alpha}\Big)D(k_1)\Big], \tag{A.13}$$

the first when (BE) binds, the second when she is paid $\bar{\alpha}$ and the investor keeps a rent. The second rises with the commitment, since $\bar{\alpha} \leq \hat{\alpha}(\beta)$, and by (A.10) it never exceeds $\hat{\alpha}(\beta)R\frac{p}{N}k_1$, what gambling the same commitment would pay her at her reservation share. By (A.11), with $\rho$ of (5.14),

$$\begin{aligned} V(\mathrm{e};k_1) - \bar{r}k_1 &= R\,\frac{p}{N}\Big[(1-\rho)\,k_1 + (2b-1)\,D(k_1)\Big], \\ P(\mathrm{e};k_1) - \bar{r}k_1 &= R\,\frac{p}{N}\Big[(1-\rho)\,k_1 - \Big(1-(2-h)\,b\Big)\,D(k_1)\Big]. \end{aligned} \tag{A.14}$$

Assumption 1, $1-h < \rho < 1$, places the patience cutoff inside the band: $\underline{\beta}\,(1-\delta) < \frac{1}{2} < \bar{\beta}\,(1-\delta) < 1/(2-h)$.

(i) $\beta \leq \underline{\beta}$. Experimenting on any commitment pays her at most $\hat{\alpha}(\beta)R\frac{p}{N}k_1 \leq (1-\rho)R\frac{p}{N}\bar{K}$, the payoff of the gamble at full scale. The contract is $(\bar{K},\, 1-\rho)$, and she gambles.

(ii) $\underline{\beta} < \beta < \bar{\beta}$. The gamble cannot be financed on a commitment that affords a second round, and $(2-h)\,b < 1$: the reserve costs the investor.

*Rationing.* On the last bracket, $L = T$, the slack $P(\mathrm{e};k_1) - \bar{r}k_1$ has slope $P'_T - \bar{r}$, with $P'_j$ as in the proof of Lemma 5.2, and $P'_T$ falls to $R\frac{p}{N}\,(1-h)\,b/(1-b) < \bar{r}$ as the horizon lengthens. For a horizon long enough that $P'_T < \bar{r}$, the commitments on which experimenting can be financed are bounded; let $\kappa$ be their supremum, and let the fund exceed it. Then $k_1^* \leq \kappa < \bar{K}$, whatever the size of the fund.

*Where the constraint binds.* A patient entrepreneur, $2b > 1$, takes the largest such commitment, $k_1^* = \kappa$, because both amounts in (A.13) rise with it. Where the financing constraint binds at $\kappa$, $\alpha^{\mathrm{BE}}(\kappa) = 0$: she receives none of the pre-discovery earnings, $\alpha^* = 0$,

and $P^0(\mathrm{e};\kappa) + P'(\mathrm{e};\kappa)\,\kappa = \bar{r}\kappa$ rearranges to (5.15). If instead $\kappa$ is a commitment that first affords another round, the constraint is slack just below it, and the contract sits there on the share $\alpha^{\mathrm{BE}} > 0$ that takes up the slack.

*A single round.* A commitment below $1 + 1/(1-\delta)$ affords one round, on which the two policies coincide: (IC) is idle, (BE) reads $\alpha \leq 1 - \rho$, and her payoff $(1-\rho)R\frac{p}{N}k_1$ rises with the commitment. Such a contract therefore sits just short of $1 + 1/(1-\delta)$, on the share $\alpha^* = 1 - \rho$ of (i), and leaves no capital for a scale-up. It is the equilibrium contract in two cases. The first is that no commitment affording a second round can be financed, so that $\kappa = 1 + 1/(1-\delta)$; this can happen at any patience in the band. The second is that an impatient entrepreneur prefers it to every larger commitment she could have.

*Impatient entrepreneurs.* For $2b \leq 1$ the reserve costs the firm value as well. At her reservation share the investor would receive $R\frac{p}{N}\big[(1-\hat{\alpha}(\beta))\,k_1 - (1-2b)\,D(k_1)\big] < \bar{r}k_1$, so (BE) binds first at every commitment, and her payoff is the surplus $V(\mathrm{e};k_1) - \bar{r}k_1$ of (A.14). That surplus rises with the commitment on a bracket only if $V'(\mathrm{e};k_1) \geq \bar{r}$, and it drops where the commitment first affords another round. Among the commitments that afford a second round, $V'(\mathrm{e};k_1)$ is largest on those that afford exactly two, where it equals $2b\,R\frac{p}{N}$. When $2b \leq \rho$, then, no capital held back for a second round pays for itself: the surplus falls on every commitment past a single round's stake, and that stake is her contract. When $2b > \rho$, she takes the financeable commitment on which the surplus is largest: $\kappa$ on the terms above, a smaller one on the share $\alpha^* > 0$ at which the investor just breaks even, or again the single round.

(iii) $\beta \geq \bar{\beta}$. The gamble cannot be financed, and $2b > 1$. The financing constraint (5.7) holds at every commitment and every horizon: by (A.14) this is immediate when $(2-h)\,b \geq 1$, and otherwise $D(k_1) \leq k_1/(1-b)$ gives

$$P(\mathrm{e};k_1) - \bar{r}k_1 \;\geq\; R\,\frac{p}{N}\,\frac{(1-h)\,b - \rho\,(1-b)}{1-b}\,k_1 \;\geq\; 0.$$

Both amounts in (A.13) rise with the commitment, so $k_1^* = \bar{K}$, on the share

$$\alpha^* = \min\Big\{\alpha^{\mathrm{BE}}(\bar{K}),\ \hat{\alpha}(\beta),\ 1\Big\}.$$

*The reduced problem.* In every case $(x, k_1^*)$ also solves (5.6) subject to (5.7). In (i) the gamble at full scale carries the largest surplus, because an impatient firm is worth less experimenting. In (ii) and (iii) only experimenting can be financed, and the surplus is largest where her payoff is: both rise with the commitment for a patient entrepreneur, and the two coincide for an impatient one. The equivalence rests on Assumption 1, which keeps $\underline{\beta}$ below

the patience cutoff: no entrepreneur whose surplus is larger under experimenting can be financed as a gambler.

## Proof of Theorem 5.3

Both boundaries have the form $x/(x+y)$, increasing in $x$ and decreasing in $y$: $\underline{\beta}\,(1-\delta)$ with $x = 1-\rho$ and $y = h$, and $\bar{\beta}\,(1-\delta)$ with $x = \rho$ and $y = 1-h$. The first falls and the second rises in both $\rho$ and $h$. A higher outside return $\bar{r}$ raises $\rho$ and lowers $h = \eta\,(1-\bar{r}/R)$, and the band still widens at both ends: $d\underline{\beta}/d\bar{r}$ has the sign of $-\eta\left(1-\frac{p}{N}\right)$, and $d\bar{\beta}/d\bar{r}$ that of $1-\eta$.

## Proof of Theorem 5.4

(i) Along the experimenting path, $\tilde{k}_t = (1-\delta)^{t-1}k_1 - \sum_{s=1}^{t-1}(1-\delta)^s$ increases in $k_1$ at every $t$, and $L$ is nondecreasing in $k_1$. On $k_1^* < \bar{K}$ the entrepreneur therefore runs no more rounds than on $\bar{K}$, and every scale-up term $\frac{p}{N}\tilde{k}_t$ of (3.17) is smaller.

(ii) By Theorem 5.2(ii), the gamble of these entrepreneurs is not financed at scale, and their commitment is $k_1^* < \bar{K}$. Where $k_1^*$ affords more than one round, the spike $\bar{K}$ at $t = 1$ of Proposition 3.5(i) gives way to the path (3.17) generated by $k_1^*$, which opens with a single unit. Where it affords only one, all of $k_1^* < 1 + 1/(1-\delta)$ is invested at $t = 1$, in place of $\bar{K}$.

(iii) Inside the band, $k_1^*$ is determined by the financing constraint (5.7) and the surplus, in neither of which $\bar{K}$ appears; the path is unchanged when the fund grows. Outside it $k_1^* = \bar{K}$, and the gamblers' spike and the experimenters' scale-ups grow with the fund.

# References


Abel, A. B. (1990, May). Asset prices under habit formation and catching up with the Joneses. *American Economic Review 80*(2), 38–42.

Abreu, D. and M. K. Brunnermeier (2003). Bubbles and crashes. *Econometrica 71*(1), 173–204.

Aghion, P. and P. Howitt (1992, March). A model of growth through creative destruction. *Econometrica 60*(2), 323–351.

Aghion, P. and P. Howitt (1998). *Endogenous Growth Theory*. Cambridge, MA: MIT Press.

Anton, J. J. and D. A. Yao (1994). Expropriation and inventions: Appropriable rents in the absence of property rights. *American Economic Review 84*(1), 190–209.

Arrow, K. J. (1962). The economic implications of learning by doing. *Review of Economic Studies 29*(3), 155–173.

Bernardo, A. E. and I. Welch (2001). On the evolution of overconfidence and entrepreneurs. *Journal of Economics & Management Strategy 10*(3), 301–330.

Brown, K. C., W. V. Harlow, and L. T. Starks (1996). Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry. *Journal of Finance 51*(1), 85–110.

Bushee, B. J. (1998). The influence of institutional investors on myopic R&D investment behavior. *Accounting Review 73*(3), 305–333.

Caballero, R. J. and M. L. Hammour (1996, August). On the timing and efficiency of creative destruction. *Quarterly Journal of Economics 111*(3), 805–852.

Camerer, C. and D. Lovallo (1999). Overconfidence and excess entry: An experimental approach. *American Economic Review 89*(1), 306–318.

Campbell, G. (2013). Deriving the Railway Mania. *Financial History Review 20*(1), 1–27.

Campbell, G. and J. D. Turner (2012). Dispelling the myth of the naive investor during the British Railway Mania, 1845–1846. *Business History Review 86*(1), 3–41.

Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology 9*(2), 129–136.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990). Positive feedback investment strategies and destabilizing rational speculation. *Journal of Finance 45*(2), 379–395.

DeMarzo, P., R. Kaniel, and I. Kremer (2007). Technological innovation and real investment booms and busts. *Journal of Financial Economics 85*(3), 735–754.

DeMarzo, P., R. Kaniel, and I. Kremer (2008). Relative wealth concerns and financial bubbles. *Review of Financial Studies 21*(1), 19–50.

Dewey, J. (1974). *John Dewey on Education: Selected Writings*. Chicago: University of Chicago Press.

Dharani, N., J. Pineda, A. Andeyro, N. Hamming, Y. Kim, N. Misra, and A. Bhatia (2025, August). Rethinking B2B software pricing in the agentic AI era. Article.

Dixit, A. K. and R. S. Pindyck (1994). *Investment under Uncertainty*. Princeton, NJ: Princeton University Press.

Drugov, M. and R. Macchiavello (2014). Financing experimentation. *American Economic Journal: Microeconomics 6*(1), 315–349.

Dunne, T., M. J. Roberts, and L. Samuelson (1988). Patterns of firm entry and exit in U.S. manufacturing industries. *RAND Journal of Economics 19*(4), 495–515.

Fazzari, S. M., R. G. Hubbard, and B. C. Petersen (1988). Financing constraints and corporate investment. *Brookings Papers on Economic Activity 1988*(1), 141–206.

FDA (2026, March). Artificial intelligence-enabled medical devices. Content current as of March 4, 2026.

Franco, A. M. and D. Filson (2006). Spin-outs: Knowledge diffusion through employee mobility. *RAND Journal of Economics 37*(4), 841–860.

Goldgar, A. (2007). *Tulipmania: Money, Honor, and Knowledge in the Dutch Golden Age*. Chicago: University of Chicago Press.

Gompers, P., A. Kovner, J. Lerner, and D. Scharfstein (2008). Venture capital investment cycles: The impact of public markets. *Journal of Financial Economics 87*(1), 1–23.

Gompers, P. and J. Lerner (2000). Money chasing deals? the impact of fund inflows on private equity valuations. *Journal of Financial Economics 55*(2), 281–325.

Gompers, P., J. Lerner, and D. Scharfstein (2005). Entrepreneurial spawning: Public corporations and the genesis of new ventures, 1986 to 1999. *Journal of Finance 60*(2), 577–614.

Gompers, P. A. (1995). Optimal investment, monitoring, and the staging of venture capital. *Journal of Finance 50*(5), 1461–1489.

Grossman, G. M. and E. Helpman (1991). *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.

Hall, B. H. and J. Lerner (2010). The financing of R&D and innovation. In B. H. Hall and N. Rosenberg (Eds.), *Handbook of the Economics of Innovation*, Volume 1, pp. 609–639. Elsevier.

Hart, O. and J. Moore (1994). A theory of debt based on the inalienability of human capital. *Quarterly Journal of Economics 109*(4), 841–879.

Hills, T. T., P. M. Todd, D. Lazer, A. D. Redish, and I. D. Couzin (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences 19*(1), 46–54.

Hobijn, B. and B. Jovanovic (2001, December). The information-technology revolution and the stock market: Evidence. *American Economic Review 91*(5), 1203–1220.

Holmström, B. and J. Tirole (1997). Financial intermediation, loanable funds, and the real sector. *Quarterly Journal of Economics 112*(3), 663–691.

Johnson, T. C. (2007). Optimal learning and new technology bubbles. *Journal of Monetary Economics 54*(8), 2486–2511.

Jorgenson, D. W. (1963). Capital theory and investment behavior. *American Economic Review 53*(2), 247–259.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica 50*(3), 649–670.

Jovanovic, B. and S. Lach (1989). Entry, exit, and diffusion with learning by doing. *American Economic Review 79*(4), 690–699.

Jovanovic, B. and G. M. MacDonald (1994, April). The life cycle of a competitive industry. *Journal of Political Economy 102*(2), 322–347.

Jovanovic, B. and B. Szentes (2013). On the market for venture capital. *Journal of Political Economy 121*(3), 493–527.

Kaplan, S. N., B. A. Sensoy, and P. Strömberg (2009). Should investors bet on the jockey or the horse? evidence from the evolution of firms from early business plans to public companies. *Journal of Finance 64*(1), 75–115.

Kerr, W. R. and R. Nanda (2015, November). Financing innovation. *Annual Review of Financial Economics 7*(1), 445–462.

Khanna, M., N. Mittal, A. Wu, and M. Castillo (2025, September). Upgrading software business models to thrive in the AI era. Article.

Klepper, S. and S. Sleeper (2005). Entry by spinoffs. *Management Science 51*(8), 1291–1306.

Klette, T. J. and S. Kortum (2004, October). Innovating firms and aggregate innovation. *Journal of Political Economy 112*(5), 986–1018.

Koellinger, P., M. Minniti, and C. Schade (2007). I think I can, I think I can: Overconfidence and entrepreneurial behavior. *Journal of Economic Psychology 28*(4), 502–527.

Lamoreaux, N. R., D. M. G. Raff, and P. Temin (Eds.) (1999). *Learning by Doing in Markets, Firms, and Countries*. Chicago: University of Chicago Press.

Lucas, Jr., R. E. (1988). On the mechanics of economic development. *Journal of Monetary Economics 22*(1), 3–42.

Lucas, Jr., R. E. and B. Moll (2014). Knowledge growth and the allocation of time. *Journal of Political Economy 122*(1), 1–51.

Manso, G. (2011). Motivating innovation. *Journal of Finance 66*(5), 1823–1860.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science 2*(1), 71–87.

Morris, T. H. (2020). Experiential learning—a systematic review and revision of Kolb's model. *Interactive Learning Environments 28*(8), 1064–1077.

Moskowitz, T. J. and A. Vissing-Jørgensen (2002). The returns to entrepreneurial investment: A private equity premium puzzle? *American Economic Review 92*(4), 745–778.

Nanda, R. and M. Rhodes-Kropf (2013). Investment cycles and startup innovation. *Journal of Financial Economics 110*(2), 403–418.

Neher, D. V. (1999). Staged financing: An agency perspective. *Review of Economic Studies 66*(2), 255–274.

Odlyzko, A. (2010). Collective hallucinations and inefficient markets: The British Railway Mania of the 1840s. Technical report, University of Minnesota. Preliminary version.

Ofek, E. and M. Richardson (2003). DotCom mania: The rise and fall of internet stock prices. *Journal of Finance 58*(3), 1113–1137.

Pástor, Ľ. and P. Veronesi (2009, September). Technological revolutions and stock prices. *American Economic Review 99*(4), 1451–1483.

Pirinsky, C. A. (2026). *A Theory of Dynamic Preferences: Using Economics and Behavioral Sciences to Explore Bias, Irrationality, and Choice*. Palgrave Macmillan.

Polk, C. and P. Sapienza (2009). The stock market and corporate investment: A test of catering theory. *Review of Financial Studies 22*(1), 187–217.

Pratten, C. F. (1980). The manufacture of pins. *Journal of Economic Literature 18*(1), 93–96.

Sahlman, W. A. (1990). The structure and governance of venture-capital organizations. *Journal of Financial Economics 27*(2), 473–521.

Sahlman, W. A. and H. H. Stevenson (1985). Capital market myopia. *Journal of Business Venturing 1*(1), 7–30.

Salian, I. (2024, June). "Believe in Something Unconventional, Something Unexplored," NVIDIA CEO tells Caltech grads. Accessed May 1, 2026.

Schumpeter, J. A. (1942). *Capitalism, Socialism and Democracy*. New York: Harper & Brothers.

Shiller, R. J. (2000). *Irrational Exuberance*. Princeton University Press.

Smith, A. (1776). *An Enquiry into the Nature and Causes of the Wealth of Nations*. London: W. Strahan and T. Cadell. First published 1776.

Smith, A., T. Lohrenz, J. King, P. R. Montague, and C. F. Camerer (2014). Irrational exuberance and neural crash warning signals during endogenous experimental market bubbles. *Proceedings of the National Academy of Sciences 111*(29), 10503–10508.

Thompson, P. (2010). Learning by doing. In B. H. Hall and N. Rosenberg (Eds.), *Handbook of the Economics of Innovation*, Volume 1, pp. 429–476. Elsevier.

Turner, G. M. (2015). *The Florida Land Boom of the 1920s*. Jefferson, NC: McFarland & Company, Inc., Publishers.

Zitnik, M. (2025, July). AI-enabled drug discovery reaches clinical milestone. *Nature Medicine 31*, 2490–2491.