# Self-Discovery in Real Investment Decisions

Christo A. Pirinsky Michael C. Tseng*

August 2026

**Abstract**

We explore the idea that real investment generates information about each firm's own ability to develop a new technology. In a stylized model, we show that this self–discovery motive gives rise to boom–bust investment cycles, characterized by an initial surge in investment followed by elevated rates of failure and exit. These dynamics are amplified when the technology draws on broader skill sets, investors are more impatient, and capital is more abundant. Using comprehensive venture capital data, we find that financing activity during the 1990s Internet boom exhibits patterns consistent with the model's predictions.

*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

*"You can't connect the dots looking forward. You can only connect them looking backwards, so you have to trust that the dots will somehow connect in your future."*
*—Steve Jobs, Stanford University Commencement Address, June 12, 2005*

Investment decisions are large-scale, long-term, and often irreversible. Not surprisingly, investors dedicate substantial resources to assessing the likelihood of investment success. Standard models of investment under uncertainty focus predominantly on external unknowns—demand, prices, policy, or other common sources of payoff uncertainty—while taking firms' knowledge of their own capabilities as given. But firms may not know ex ante whether they have a comparative advantage in exploiting a new technological opportunity. Often, that latent comparative advantage can be learned only through experimentation.[1] As firms build and test projects, they discover whether they can participate successfully in the emerging market and, if so, where their niche lies. Only in retrospect do the dots connect and reveal a firm's real edge.

We bring out this self-discovery logic of investment in a stylized model of *introspective learning*, in which firms with different bundles of know-how confront a new technology. Firms do not know ex ante where, if anywhere, their own capabilities fit within that emerging opportunity set. In AI, for example, such project opportunities may include agentic service platforms, medical devices, drug-discovery platforms, and new outcome-based software businesses.[2] Firms can learn whether their capabilities are suited to a particular project only by committing capital to it. The uncertainty is not external but introspective—a firm's idiosyncratic fit cannot be learned by observing other firms' outcomes. If their capabilities match the project, the investment yields a high return that can be replicated in subsequent periods. If not, the firm loses the invested amount but learns that it is not well suited to that type of project.

We show that introspective learning alone is sufficient to generate boom–bust investment dynamics—a front-loaded surge followed by high failure and exit rates. Early in the adoption wave, many firms invest to resolve their own idiosyncratic match problems within the opportunity set. As this experimentation unfolds, the pool still searching shrinks and becomes increasingly tilted toward firms with no match. Firms that discover fit earliest become the largest winners, because they scale up before failed experiments and depreciation

[1] Indeed, a growing body of research spanning the behavioral and neural sciences suggests that people and organizations explore their environment not only to search for external opportunities, but also to discover their own capabilities and limitations; see Section 2.

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

erode their capital.

We also show that this boom–bust pattern is stronger for more inclusive technologies—those accessible to a broader range of skills. When barriers to experimentation are low, early participation is broad, and the same breadth that fuels the initial boom makes the eventual bust more severe. By contrast, technologies that require highly specialized expertise are less prone to bubbles, because exclusivity dampens the exploration benefits of early investment. This suggests that skill-intensive technologies (such as AI) may be less susceptible to bubble-like overinvestment than technologies with lower skill barriers to entry (such as real estate).

The boom–bust pattern also intensifies with investor impatience. More impatient investors want feedback sooner, so they invest earlier in the adoption cycle, making the boom more front-loaded. This prediction connects naturally to the literature suggesting that institutional investors may be more impatient than individual investors, and that money-management incentives can induce short horizons focused on benchmark performance and capital flows (e.g., Brown et al. (1996); Bushee (1998); Bailey et al. (2011); Jovanovic and Szentes (2013)). To the extent that institutional capital is indeed more impatient, the rise in institutional ownership may have changed not only how much capital enters new technologies, but also when it arrives—making some technology booms more front-loaded and the subsequent shakeouts sharper.

Our setting departs in a central way from traditional Bayesian learning frameworks for investment under uncertainty, in which firms learn about a common external payoff component. When observations across firms all bear on the same systematic component, they accumulate into a common posterior, so aggregate investment can move only as that posterior moves—gradually, toward the full-information benchmark. In Johnson (2007), for example, investors may optimally invest above the full-information level to learn the new technology's returns to scale, but the resulting investment path remains one of gradual adjustment as information accumulates. In Pástor and Veronesi (2009), investors update on another common parameter—the technology's average productivity—in a setting that focuses on asset prices and abstracts from real investment.

By contrast, the object of introspective learning is not a common payoff primitive but a firm's individual fit. The uncertainty is idiosyncratic rather than systematic, and its resolution is binary rather than incremental. By revealing match or mismatch, investment sorts firms across the opportunity set. Aggregate investment is therefore no longer tied to the gradual adjustment logic of Bayesian learning. Instead, it exhibits the familiar pattern of new-technology episodes—a broad early-participation boom that gives way to an endogenous shakeout as unsuccessful firms exit.

Our learning channel is likewise distinct from learning-by-doing. In learning-by-doing, agents become better 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 in order to learn whether their existing capabilities are well matched to the activity. Whereas learning-by-doing points toward capability accumulation, introspective learning generates self-discovery.

Our mechanism does not rely on the cross-agent forces or departures from rationality emphasized in other accounts of boom–bust dynamics. In that literature, overinvestment is often attributed to relative-wealth concerns (Abel (1990); DeMarzo et al. (2008)) or to behavioral biases interacting with capital-market frictions (De Long et al. (1990); Abreu and Brunnermeier (2003)). We show that, even in an environment without externalities or irrationality, firms' intrinsic motive to learn their own fit with a new technology generates the same boom–bust pattern.

Accordingly, the model abstracts from cross-agent effects. Introducing externalities leaves the self-discovery logic intact and may further reinforce the results. For example, in a fixed market with a race for dominant share, strategic pressure to move early would amplify the front-loading dynamics we identify.

We take our model to data on the financing rounds of venture-capital-backed firms during the dot-com boom of the late 1990s. Venture capital (VC) finances firms at precisely the stage where uncertainty about fit is greatest: the companies are young, their business models are untested, and their founders have yet to learn whether they can turn a new technology into a viable business. Venture capitalists, moreover, are not passive providers of funds; they sit on boards, shape strategy, and release capital round by round. As a result, the financing record of this market offers a window onto the exploratory element of investment decisions.

The data bear out three distinct predictions from the model. First, the arrival of a new technology sets off a pronounced boom in aggregate investment, followed by a bust. Specifically, we find that the share of aggregate VC investment directed toward Internet companies rose from 6 percent in 1995 to nearly 30 percent by the end of 1999, before falling back to roughly its mid-1990s level by the end of 2003. Second, during the run-up, Internet companies were financed through significantly more rounds than other venture-backed companies. Third, Internet companies that discovered their match during the expansion period accelerated investment faster than firms in older industries.

The dot-com episode is a canonical case in the literature on financial bubbles. While theories in that tradition can explain some of the empirical patterns above, they are unlikely to account for all of them. For example, in asymmetric-information models, investors ride overheated markets to exploit superior information (e.g., Brunnermeier and Nagel (2004)); in models of heterogeneous beliefs with short-sale constraints, investors buy overvalued assets expecting to resell them to still more optimistic buyers (Miller (1977); Harrison and Kreps (1978); Hong et al. (2006)). Bubbles of either kind hinge on the ability to trade dynamically—to build a position while sentiment rises and to exit before it turns. These conditions are largely absent in the highly illiquid venture capital market, where capital is returned only when portfolio companies are sold or funds are wound down. Speculative mechanisms based on bubble-riding are therefore unlikely to explain the investment behavior we document. We discuss these and other alternative explanations in Section 5.2.

In addition to aggregate market outcomes, our analysis provides insight into entrepreneurial entry, failure, and returns. A large empirical literature documents high entrepreneurial activity despite severe attrition and weak average ex post returns (Dunne et al. (1988); Moskowitz and Vissing-Jørgensen (2002)). A common interpretation emphasizes overconfidence (Camerer and Lovallo (1999); Bernardo and Welch (2001); Koellinger et al. (2007)). We highlight a complementary rational margin. Starting a venture generates information about whether an entrepreneur is capable of developing and commercializing the underlying opportunity, or whether their effort is better directed elsewhere. Even failure can be valuable through the self-knowledge it produces—high entry and subsequent exit need not reflect miscalibration alone.

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 takes the model to data on venture capital investment during the dot-com boom. 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).

A key benefit of exploration is that it enables *introspective* learning. Economic agents often know far less about their own dispositions than standard models assume. The environment may be uncertain, but a deeper uncertainty concerns our compatibility with it. The reward system of living organisms encodes important information about the evolutionary advantages of the options they face. It consists of brain structures and neural pathways that associate beneficial activities with positive sensory and emotional responses, and harmful activities with negative responses (Berridge and Kringelbach (2015)). But living organisms can learn the reward of an activity only through engagement in the activity itself (Wilson and Gilbert (2005)). Exploration is therefore a mechanism for discovering one's own preferences and comparative strengths, not merely for sampling unknown payoffs.

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.[3] 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)).

This pattern is not confined to contemporary technology. During the British railway boom

[3] 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.

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,

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

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

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

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 $t$.

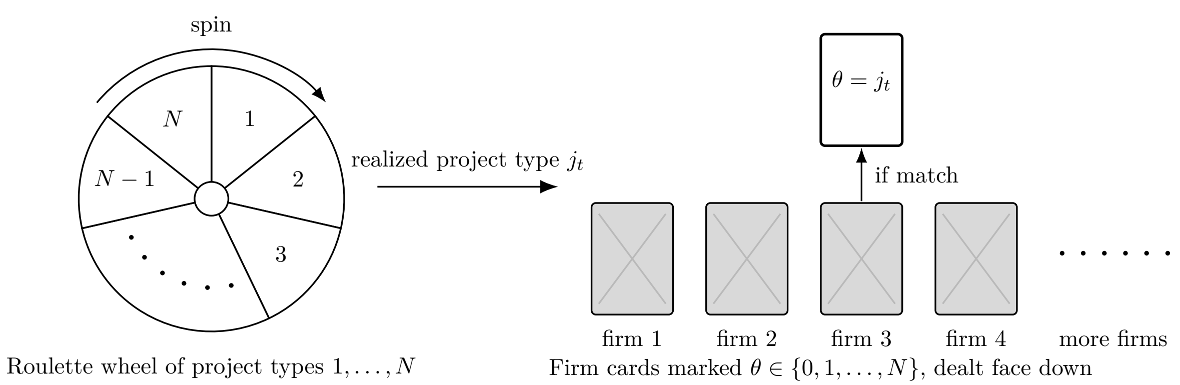


Figure 1: **The Roulette of New Technology**

**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

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

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

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

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

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

**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

$$V_t(k, m) = \max\left\{\Phi_t^{\text{wait}}(k, m),\ \Phi_t^{\text{inv}}(k, m)\right\}, \tag{3.8}$$

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

$$\Phi_t^{\text{wait}}(k,m) = \beta\, V_{t+1}\Big((1-\delta)k,\, m\Big). \tag{3.9}$$

*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,

$$\Phi_t^{\text{inv}}(k,m) = \max_{x\in[1,k]} \Big[\mu(m)\Big(Rx+\beta\, W_{t+1}\big((1-\delta)(k-x)\big)\Big)+\Big(1-\mu(m)\Big)\beta\, V_{t+1}\Big((1-\delta)(k-x),\, m-1\Big)\Big], \tag{3.10}$$

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

$$\mu(m) = \Pr\Big(\theta = j_t \,\Big|\, j_t \in M_t,\ |M_t| = m\Big) = \frac{\frac{p}{N}}{(1-p)+m\frac{p}{N}} \in (0,1). \tag{3.11}$$

**Proposition 3.1.** *In the searching regime* (3.8),

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

*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

$$\tilde{k}_1 := k_1, \qquad \tilde{k}_{t+1} := (1-\delta)\,(\tilde{k}_t - 1) \tag{3.12}$$

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

$$L \coloneqq \min\{\, T,\ \ell(k_1)\,\}, \tag{3.13}$$

where

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

**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 (3.11); 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,*

$$V_1(k_1, N) = R\,\frac{p}{N}\,k_1. \tag{3.14}$$

*(ii) For a patient firm,*

$$V_1(k_1, N) = R\,\frac{p}{N}\,\left[\sum_{t=1}^{L-1}\left(\beta^{t-1} + \beta^t \tilde{k}_{t+1}\right) + \beta^{L-1}\tilde{k}_L\right], \tag{3.15}$$

*where $(\tilde{k}_t)_{t\geq 1}$ is the feasible-testing capital path* (3.12)*.*

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 (3.14) and (3.15) 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.

**Investment Self-Rationing** Under standard investment logic, there is no reason to ration 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)). It is the patient firm that our mechanism sets apart: 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 commitment when that knowledge arrives. In practice, the distinction is one of degree: most firms confronting a new technology stage their commitments and then scale, as seen in the round-by-round financing of venture capital. 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. Thus the self-discovery motive leads to *investment self-rationing*: while searching, the firm withholds capital from itself.[4]

Self-rationing reverses the usual real-options intuition as well. There, the option is to wait and see, holding investment back until more is known (Dixit and Pindyck (1994)). Here, the option value of self-discovery derives instead from active search, with waiting made even less attractive by depreciating capital.

## 3.3 Boom–Bust Investment Dynamics

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 (3.12) and (3.13). Let $\pi_{\mathrm{imp}} \in [0,1]$ be the population share of impatient firms, and $\pi_{\mathrm{pat}} := 1 - \pi_{\mathrm{imp}}$ the complementary share of patient firms.

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

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

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

[4] This is the mirror image of credit rationing in the sense of Stiglitz and Weiss (1981): there, lenders withhold capital and firms would invest more if they could.

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

$$\mathrm{Inv}_{\mathrm{pat},t} = \begin{cases} \pi_{\mathrm{pat}}, & t = 1, \\ \pi_{\mathrm{pat}}\left(S_t + \dfrac{p}{N}\tilde{k}_t\right), & t = 2, \dots, L-1, \\ \pi_{\mathrm{pat}}\left(S_L + \dfrac{p}{N}\right)\tilde{k}_L = \pi_{\mathrm{pat}}\left(1 - \dfrac{p}{N}(L-2)\right)\tilde{k}_L, & t = L, \\ 0, & t > L. \end{cases} \tag{3.16}$$

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

Aggregate investment (3.16) 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.

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

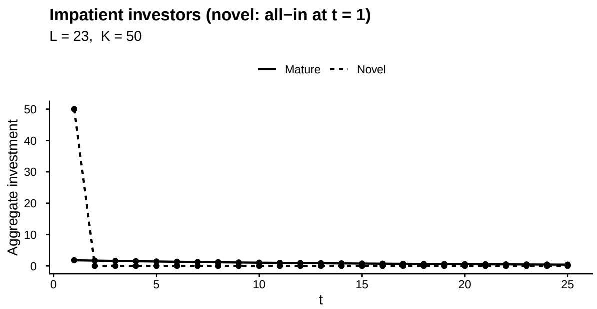


(a) Impatient firms.

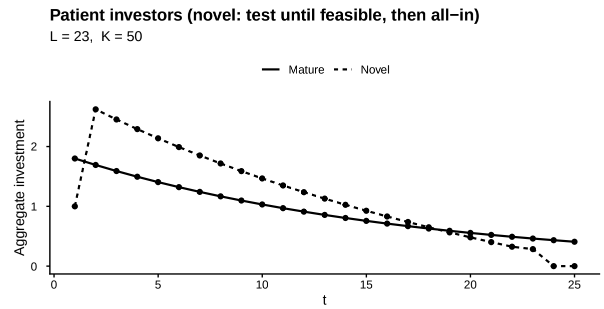


(b) Patient firms.

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.

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.

**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.

**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

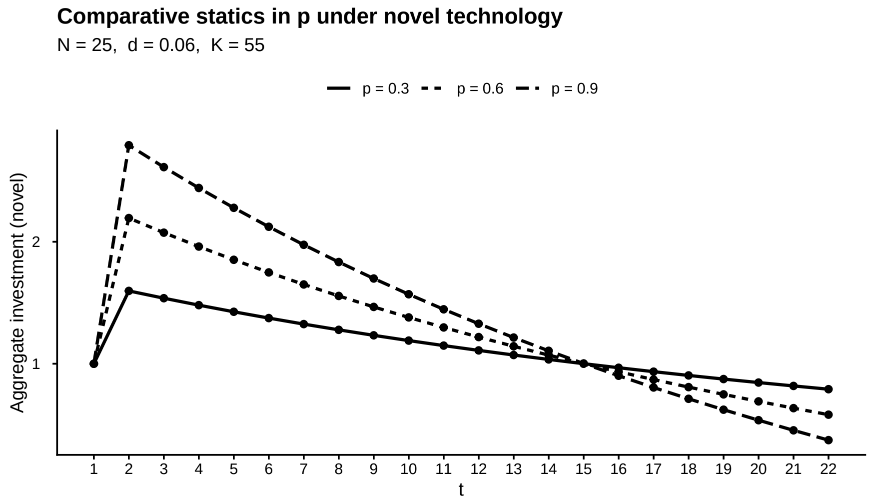


Figure 3: **Technology scope and boom–bust investment.** Each line plots patient-firm aggregate investment for a different scope $p = \Pr(\theta \neq 0)$. A larger scope $p$ raises the early boom because fit is more likely and more experiments turn into scale-up. Investment then declines as fewer firms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

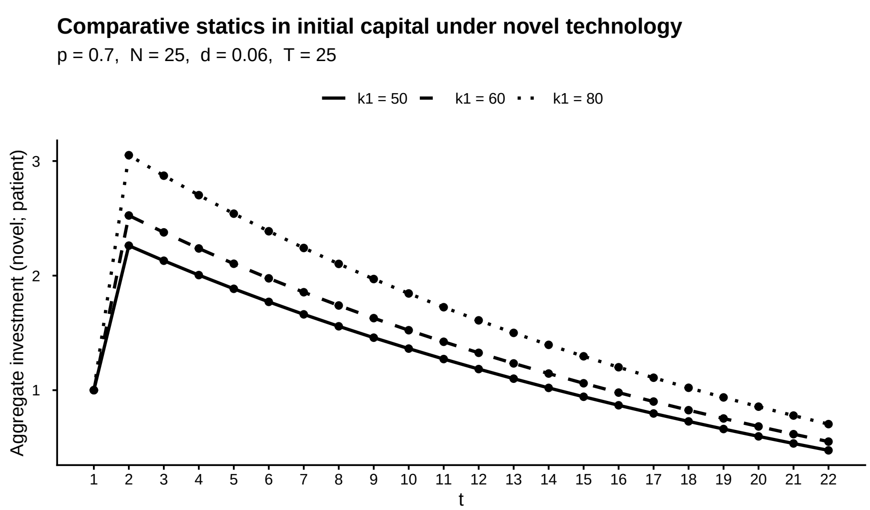


Figure 4: **Access to capital and boom–bust investment.** Each line plots patient-firm aggregate investment for a different initial capital stock $k_1$. Higher $k_1$ raises and prolongs the boom because firms can keep searching longer and successful firms scale with more capital left to deploy.

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.[5]

**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 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 Entrepreneurial 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*

[5] 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)).

*probability that the experiment does not reveal fit is*

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

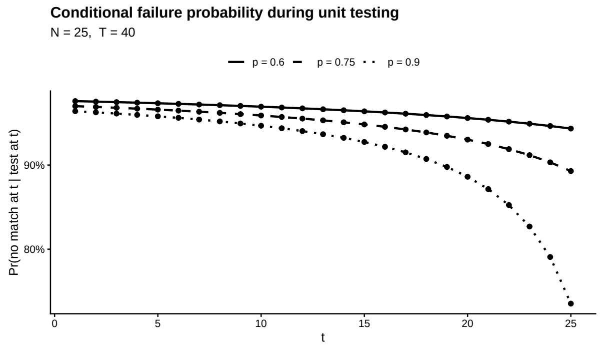


(a) Technology scope $p$. Larger scope lowers failure rates.

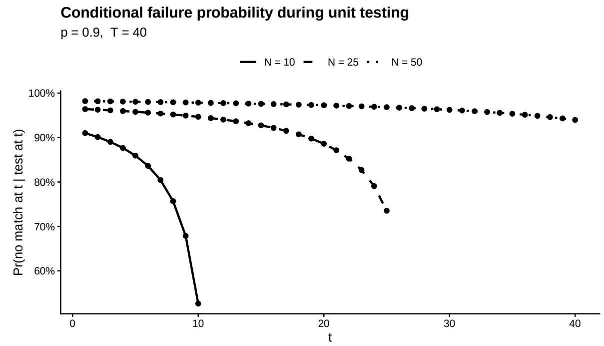


(b) Technology granularity $N$. Higher $N$ keeps failure rates elevated.

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.

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 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* (3.12) *and* (3.13)*, 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 Evidence from VC Investment during the Dot-Com Bubble

*"The future can be discovered by means of iterative, venture-backed experiments. It cannot be predicted."*
*—Sebastian Mallaby,* The Power Law*, 2022*

In this section, we present evidence on the role of firm self-discovery in investment decisions during the dot-com bubble. We focus on venture capital (VC) investments for two reasons. First, VC finances firms early in the life cycle, when uncertainty about their fit with a new technology is greatest and introspective learning is likely to be particularly important. Second, several institutional features of venture capital—including a sophisticated investor base, active involvement in portfolio companies, and the illiquidity of investments—help rule out alternative explanations for investment behavior during financial bubbles. The following predictions of our model bear directly on this episode.

**Hypothesis 1.** *Aggregate investment in a new technology exhibits a pronounced boom followed by a bust.*

**Hypothesis 2.** *Firms in a new technology stage their investment across a larger number of financing rounds than firms in old technologies.*

**Hypothesis 3.** *Upon discovering a favorable match, firms in a new technology scale up their investment more aggressively than firms in old technologies.*

Hypothesis 1 corresponds to the boom–bust pattern that emerges from the model, illustrated in Figure 2b. Hypothesis 2 follows from the investment self-rationing established in Theorem 3.2: exploration demand prompts firms to spread their investment over a larger number of rounds. In testing this hypothesis, we examine the number of rounds and remain agnostic about the invested amounts: self-rationing is independent of firm size, and size could differ systematically across industries. Hypothesis 3 likewise follows from Theorem 3.2: upon discovering a match, firms switch from rationing to scaling up their investment aggressively.

## 5.1 Empirical Results

Our data come from the Thomson Reuters VentureXpert database and contain VC investments in U.S. companies. The database has been validated by previous researchers against known financing rounds (see, e.g., Kaplan et al. (2002)). The data contain detailed information on the dates of venture financing rounds, the investors and portfolio companies involved, and the estimated amounts invested by each firm. Given our interest in the dot-com bubble, we restrict the sample to investments by VC firms between 1995 and 2004.

We define industries at the Minor Industry Group level, based on the industry classifications provided by VentureXpert. Table 1 lists these industries for the sample of 39,390 VC investments. Two of the eleven industry groups are labeled "Internet Specific"; they define our treatment group.

Table 2 summarizes VC market activity year by year: the number of investments, the average amount invested, the share of investments directed toward Internet firms, and the average number of investment rounds, both overall and separately for Internet and non-Internet firms.[6] The rapid expansion of Internet investment throughout the second half of the 1990s can be seen clearly in the table: from 6 percent of VC investments in 1995 to nearly 30 percent by 1999. After 2000 the share fell as quickly as it had risen, and by 2003 it had returned to its mid-1990s level. This boom–bust pattern corroborates Hypothesis 1. Also, the number of investment rounds varies substantially across industries and over time, variation we will exploit in testing the remaining hypotheses.

To isolate differences in investment patterns between Internet and non-Internet firms over time, we estimate the following regression model:

$$Y_{i,t} = \alpha + \beta_1 \, Internet_i \times RUNUP_t + \gamma_j + \lambda_t + \epsilon_{i,t}, \tag{5.1}$$

where the dependent variables $Y_{i,t}$ measure various characteristics of VC investment in firm $i$ at time $t$, such as the number of investment rounds and the growth rate of investment across rounds, and $\gamma_j$ and $\lambda_t$ denote industry and year fixed effects. $RUNUP_t$ is an indicator variable equal to one during the 1995–1999 expansion period, and zero otherwise. $Internet_i$ is an indicator variable equal to one for Internet firms, and zero for firms in other industries. We include only the difference-in-differences term $Internet \times RUNUP$ because the main effects *Internet* and *RUNUP* are subsumed by the fixed effects.

In Table 3, we report estimates of (5.1) with the number of financing rounds as the dependent variable: column (1) takes the (log) number of rounds in a firm, and columns (2)–(4) indicators for firms that received more than one, two, and three rounds, respectively. In this regression, year fixed effects are dated by the time of the initial investment in a firm, absorbing differences across entry cohorts. Across all four columns, the estimates show that Internet firms were financed over more rounds during the expansion period, relative to

[6]To adjust for inflation, all dollar amounts are expressed in 2024 U.S. dollars using the GDP implicit price deflator from the Federal Reserve Bank of St. Louis (`https://fred.stlouisfed.org/series/GDPDEF`).

both non-Internet firms and Internet firms in the post-expansion period.[7] The difference is highly statistically significant and economically large: for example, column (2) implies that an Internet firm was about 20 percentage points more likely to raise a second round in the late 1990s. This staging pattern, stronger while the technology was new, is consistent with Hypothesis 2.

In Table 4, we report estimates of (5.1) with the growth rate of investment across rounds as the dependent variable: the change in log investment from one round to the next, so that each observation is a pair of consecutive rounds within a firm. Column (1) estimates the conditional mean by OLS; columns (2)–(4) replace it with the 25th, 50th, and 75th conditional quantiles. The results show that Internet firms scaled up more aggressively during the late-1990s expansion than firms in the control group. The quantile estimates are positive throughout and close to the OLS estimate: the whole distribution of round-to-round growth shifts, rather than a few firms in the right tail pulling the average.[8] This pattern corroborates Hypothesis 3.

## 5.2 Discussion

Market bubbles have been the subject of a long line of research (see Brunnermeier (2008)). Economists broadly classify this literature into three frameworks: *rational* (full information), *asymmetric information*, and *heterogeneous beliefs* (including biased beliefs).

In an influential paper, Tirole (1982) demonstrates that if traders have rational expectations and the same information sets, bubbles will not occur. Our framework does not fall under this category. Although investors in our model are rational, in the sense that they are self-interested and process information correctly, they do not possess complete knowledge about their own idiosyncratic abilities.[9]

When market participants have different information sets, market equilibria can allow for "excessive" participation rates. For example, Bikhchandani et al. (1992) show that individuals, having observed the actions of those ahead of them, may optimally choose to follow the

[7]Table 2 shows Internet firms with fewer rounds on average from 1999 through 2003, but those annual averages pool firms at different stages of their financing—the same composition underlying the boom–bust investment path predicted in Proposition 3.5. Hypotheses 2 and 3 concern a firm's own financing history, so our regressions synchronize firms by entry.

[8]The estimated shift is positive at each quantile we examine, as would be the case if the distribution of round-to-round growth for Internet firms during the expansion first-order stochastically dominates that of the control group.

[9]More recent models suggest that investment externalities can lead firms to overinvest relative to the socially optimal level (DeMarzo et al. (2007)). These distortions arise because firms fail to internalize the costs their actions impose on others. Our model deliberately abstracts from such externalities (e.g., competition for market share); incorporating them would likely amplify the boom–bust dynamics we identify.

behavior of others without regard to their own information. Even when every investor knows that the price of an asset exceeds its fundamental value, this information might not be common knowledge. Allen et al. (1993) show that, under certain conditions, this lack of higher-order mutual knowledge can also lead to the formation of bubbles.

In many asymmetric-information models, investors enter overheated markets to acquire or exploit superior information. The ability to ride a bubble, however, depends critically on investors' capacity to exit their positions quickly when expectations change. For example, Brunnermeier and Nagel (2004) show that hedge funds held substantial positions in highly valued technology stocks during the late 1990s and rapidly reduced their exposure as market conditions deteriorated. Private equity investors have no comparable exit. Because their investments are highly illiquid and distributions are typically realized only through portfolio exits or fund liquidation, they have little scope to capitalize on mispricing through dynamic trading. As a result, the bubble-riding mechanism emphasized by asymmetric-information models is unlikely to operate in private equity markets.

Bubbles can also arise when investors hold heterogeneous beliefs and face short-sale constraints. In this setting, optimists bid prices upward, while pessimists are unable to offset this pressure because they cannot short the asset (Miller (1977); Hong et al. (2006)). Ofek and Richardson (2003) further argue that prices may even exceed the valuation of the most optimistic investor. The reason is that the current owner derives value not only from the asset's fundamentals but also from the option to resell it in the future at an even higher price. As investors' beliefs evolve over time, other market participants may become more optimistic and therefore be willing to purchase the asset at a premium. This resale option creates an additional speculative component of value, allowing prices to exceed even the highest contemporaneous fundamental valuation (Harrison and Kreps (1978)).

Belief heterogeneity also underlies the famous beauty-contest analogy of Keynes (1936). Rather than valuing stocks according to their own assessment of fundamentals, investors attempt to anticipate the beliefs and actions of other market participants. When prices are driven by expectations about others' expectations, rather than by fundamentals alone, speculative dynamics and bubbles can emerge. Such speculative dynamics, however, are less likely to occur in the private equity market, given its high illiquidity.

The prevailing view in the venture capital literature is that staging primarily serves to mitigate agency problems. Gompers (1995) finds that stage financing is more prevalent in industries characterized by severe agency problems, such as those with high intangible assets and intensive research and development.[10] Similarly, Tian (2011) shows that venture

[10] See also Kaplan and Strömberg (2003, 2004).

capitalists located farther from entrepreneurial firms rely more heavily on stage financing.

Although agency considerations provide a compelling rationale for stage financing, it is not clear that they explain why investment accelerates across successive financing rounds in an emerging technology. By contrast, our framework predicts both a greater number of financing rounds and increasingly intensive investment over time. As firms learn about their fit with a new technology, successful ventures optimally scale up investment while unsuccessful ones exit. Introspective learning therefore provides a complementary motive for stage financing, one that does not rely on agency frictions.

# 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 (3.10) 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, $x = k$ (all-in gamble). Accordingly, (3.10) 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 (3.10) 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

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

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 (3.11). 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 $s = 1, \dots, t$ 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 (3.12), which is decreasing over time. Hence early winners scale sooner and at larger size than late winners, proving part (iii).

# 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.

Allen, F., S. Morris, and A. Postlewaite (1993). Finite bubbles with short sale constraints and asymmetric information. *Journal of Economic Theory 61*(2), 206–229.

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

Bailey, W., A. Kumar, and D. Ng (2011). Behavioral biases of mutual fund investors. *Journal of Financial Economics 102*(1), 1–27.

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

Berridge, K. C. and M. L. Kringelbach (2015). Pleasure systems in the brain. *Neuron 86*(3), 646–664.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy 100*(5), 992–1026.

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.

Brunnermeier, M. K. (2008). Bubbles. In S. N. Durlauf and L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics* (2nd ed.). London: Palgrave Macmillan.

Brunnermeier, M. K. and S. Nagel (2004). Hedge funds and the technology bubble. *Journal of Finance 59*(5), 2013–2040.

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.

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.

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

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.

Harrison, J. M. and D. M. Kreps (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *Quarterly Journal of Economics 92*(2), 323–336.

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.

Hong, H., J. Scheinkman, and W. Xiong (2006). Asset float and speculative bubbles. *Journal of Finance 61*(3), 1073–1117.

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. 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 (2002). How well do venture capital databases reflect actual investments? Working paper, University of Chicago.

Kaplan, S. N. and P. Strömberg (2003). Financial contracting theory meets the real world: An empirical analysis of venture capital contracts. *Review of Economic Studies 70*(2), 281–315.

Kaplan, S. N. and P. Strömberg (2004). Characteristics, contracts, and actions: Evidence from venture capitalist analyses. *Journal of Finance 59*(5), 2177–2210.

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

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan.

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

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.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *Journal of Finance 32*(4), 1151–1168.

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.

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.

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

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.

Stiglitz, J. E. and A. Weiss (1981). Credit rationing in markets with imperfect information. *American Economic Review 71*(3), 393–410.

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.

Tian, X. (2011). The causes and consequences of venture capital stage financing. *Journal of Financial Economics 101*(1), 132–159.

Tirole, J. (1982). On the possibility of speculation under rational expectations. *Econometrica 50*(5), 1163–1181.

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

Wilson, T. D. and D. T. Gilbert (2005). Affective forecasting: Knowing what to want. *Current Directions in Psychological Science 14*(3), 131–134.

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

Table 1: Industry Classification of VC Investment Rounds

This table lists the industries represented in the sample of 39,390 VC investments between 1995 and 2004. Industries are defined at the Minor Industry Group level, based on the industry classifications provided by VentureXpert. The two minor groups labeled "Internet Specific" (in bold) define the treatment group.

| Industry class | Major group | Minor group | Obs. |
|---|---|---|---|
| Information Technology | Communications and Media | Communications and Media | 4,323 |
| | | **Internet Specific** | 2,694 |
| | Computer Related | Computer Hardware | 1,378 |
| | | Computer Software and Services | 9,525 |
| | | **Internet Specific** | 6,877 |
| | Semiconductors/Other Electronics | Semiconductors/Other Electronics | 2,475 |
| Medical/Health/Life Sci. | Biotechnology | Biotechnology | 2,406 |
| | Medical/Health/Life Sciences | Medical/Health | 4,138 |
| Non-High Technology | Non-High Technology | Consumer Related | 1,782 |
| | | Industrial/Energy | 1,163 |
| | | Other Products | 2,629 |

Table 2: VC Market Activity by Year

The table shows, by year, the number of investments, the average invested amount (in thousands of 2024 U.S. dollars), the fraction of investments in Internet firms, and the average number of investment rounds, both overall and separately for Internet and non-Internet firms. The sample is obtained from VentureXpert and contains 39,390 investments between 1995 and 2004.

| | | Amount | Fraction of | # rounds | | |
|---|---|---|---|---|---|---|
| Year | # investments | (thousands, 2024 $) | Internet firms | All | Internet | Other |
| 1995 | 1,876 | 26,717 | 0.06 | 5.45 | 5.86 | 5.43 |
| 1996 | 2,756 | 30,523 | 0.09 | 4.57 | 5.25 | 4.49 |
| 1997 | 3,341 | 32,003 | 0.10 | 4.07 | 4.88 | 3.99 |
| 1998 | 3,871 | 46,517 | 0.14 | 3.95 | 4.06 | 3.94 |
| 1999 | 5,496 | 90,081 | 0.29 | 3.61 | 3.17 | 3.84 |
| 2000 | 7,902 | 131,886 | 0.28 | 3.21 | 2.40 | 3.53 |
| 2001 | 4,696 | 80,727 | 0.18 | 3.82 | 2.64 | 3.99 |
| 2002 | 3,246 | 65,179 | 0.13 | 4.40 | 3.59 | 4.47 |
| 2003 | 3,032 | 63,910 | 0.09 | 4.59 | 3.63 | 4.64 |
| 2004 | 3,174 | 68,152 | 0.10 | 4.28 | 4.30 | 4.28 |

Table 3: Number of Financing Rounds Per Firm

The table reports coefficient estimates ($t$-statistics in parentheses) from linear regression models examining the difference in the number of financing rounds between Internet and non-Internet firms in the late 1990s and early 2000s. The dependent variables are the (log) number of investment rounds in a firm and indicator variables for firms experiencing more than one, two, and three rounds, respectively. The variable *Internet* is an indicator for Internet firms, and *RUNUP* is an indicator for the 1995–1999 period. The sample covers all VC investments between 1995 and 2004. All models include industry and year fixed effects, and standard errors are clustered at the industry and year levels. The last two rows report the number of observations and the adjusted $R^2$ of each model. $^{***}$, $^{**}$, and $^{*}$ denote significance at the 0.01, 0.05, and 0.10 levels.

| | # Rounds (1) | >1 Rounds (2) | >2 Rounds (3) | >3 Rounds (4) |
|---|---|---|---|---|
| *Internet* × *RUNUP* | 0.33$^{***}$ | 0.19$^{***}$ | 0.23$^{***}$ | 0.17$^{***}$ |
| | (3.83) | (4.80) | (4.06) | (2.96) |
| Year fixed effects | Yes | Yes | Yes | Yes |
| Industry fixed effects | Yes | Yes | Yes | Yes |
| Observations | 15,356 | 15,356 | 15,356 | 15,356 |
| Adjusted $R^2$ (%) | 9.92 | 6.27 | 8.11 | 8.20 |

Table 4: Investment Rate Across Rounds

The table reports coefficient estimates ($t$-statistics in parentheses) from regression models examining the difference in the investment rate across consecutive financing rounds between Internet and non-Internet firms in the late 1990s and early 2000s. The dependent variable is the log growth rate of investment amounts across subsequent rounds of a firm. The first column fits an OLS regression model, while the next three columns fit quantile regression models estimating the conditional quantiles (25th, 50th, and 75th percentiles) of the dependent variable. The variable *Internet* is an indicator for Internet firms, and *RUNUP* is an indicator for the 1995–1999 period. The sample covers all VC investments between 1995 and 2004. All models include industry and year fixed effects. The last two rows report the number of observations and the adjusted (pseudo) $R^2$ of the OLS (quantile) models. $^{***}$, $^{**}$, and $^{*}$ denote significance at the 0.01, 0.05, and 0.10 levels.

| | OLS | Quantile regression | | |
|---|---|---|---|---|
| | Mean (1) | 0.25 (2) | 0.50 (3) | 0.75 (4) |
| *Internet* × *RUNUP* | 0.74$^{***}$ | 0.75$^{***}$ | 0.84$^{***}$ | 0.77$^{***}$ |
| | (5.42) | (7.65) | (10.99) | (9.78) |
| Year fixed effects | Yes | Yes | Yes | Yes |
| Industry fixed effects | Yes | Yes | Yes | Yes |
| Observations | 24,034 | 24,034 | 24,034 | 24,034 |
| Adj. / pseudo $R^2$ (%) | 3.84 | 1.65 | 3.16 | 3.88 |