Self-Discovery in Real Investment Decisions

Christo A. Pirinsky Michael C. Tseng*

May 2026

Abstract

We analyze investment decisions when firms are uncertain about their ability 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 diagnostic learning. We show that the prospect of discovering each firm's comparative advantage incentivizes firms to invest more and earlier. In the aggregate, diagnostic learning gives rise to boom–bust investment patterns, with a surge in initial investment followed by high failure and exit rates. This pattern is driven by learning about idiosyncratic fit and contrasts with the smooth adjustment dynamics of Bayesian learning models in which uncertainty concerns a systematic payoff component. The boom–bust effect is stronger when the technology engages broader skill sets, the investor base is more impatient, and capital is more abundant.

JEL classification: D83; G31; O33.

Keywords: Introspection, diagnostic 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 Indrajit Mitra 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

Economic analysis has traditionally been extrospective. In most frameworks, even when agents have limited information about the outside world, they are assumed to possess complete knowledge about their own abilities and dispositions. As a result, learning in economic models is generally directed outwards, toward the environment, and not toward the decision maker. Yet an important aspect of life is introspective. A growing body of research across numerous disciplines suggests that through interaction with their environment, people and organizations learn about their own capabilities, strengths, and limitations (see Pirinsky (2026) and references therein). This research suggests that agents explore their environment not only to search for external opportunities, but also to discover their own capabilities. We refer to this process of self-discovery as *diagnostic learning* and explore its implications for investment under uncertainty.

Investment decisions are large-scale, long-term, and often irreversible. Not surprisingly, investors dedicate substantial resources to assess 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 to exploit a new technological opportunity. Often, that latent comparative advantage can be learned only through experimentation. As firms build and test projects, they may discover whether they can participate successfully in the emerging market and what their exact niche is. Investment is therefore not only capital allocation, but also the means by which firms discover where their capabilities apply—only in retrospect do the dots connect and reveal a firm's real edge.

We bring out this logic in a stylized model in which firms with different bundles of know-how confront a new technology whose commercial uses are only imperfectly understood. The technology requires a subset of these skills for successful development. Firms do not know ex ante whether their own capabilities are most productively deployed within that emerging opportunity set. In AI, for example, such project opportunities may include agentic service platforms (Dharani et al. (2025)), medical devices (FDA (2026)), drug-discovery platforms (Zitnik (2025)), and new outcome-based software businesses (Khanna et al. (2025)). As project opportunities arise, entrepreneurs can learn whether their capabilities are suited to a particular project by committing capital to it. If their capabilities match the project, the investment yields a high return that can be replicated in subsequent periods. If not, the entrepreneur loses the invested amount but learns that they are not well suited to that type of project.

We show that this diagnostic-learning mechanism alone is sufficient to generate boom–bust investment dynamics—front-loaded investment followed by high failure and exit rates. Each firm invests to resolve its own idiosyncratic match problem. Because success of the new technology ultimately draws on only a subset of the capabilities dispersed across firms, some firms discover a match and scale up, while others discover no match and exit. Early in the adoption wave, many firms invest to test where they fit within the opportunity set. As these experiments unfold, 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 erode their capital.

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. For example, Johnson (2007) studies a setting in which investors may optimally invest above the full-information level to learn the new technology's return to scale, but the resulting investment path remains one of gradual adjustment as information accumulates.

By contrast, the object of diagnostic 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 rather than simply refining beliefs about a shared environment. Aggregate investment is therefore no longer tied to the gradual adjustment logic of Bayesian learning. Instead, it exhibits the pattern often associated with new-technology episodes—a broad early participation boom followed by an endogenous shakeout as unsuccessful firms exit.

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).

We also find that boom–bust patterns intensify with investor impatience. More impatient investors want diagnostic 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., Bushee (1998); Brown et al. (1996); 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 shakeout sharper.

Our mechanism does not rely on cross-agent forces or departures from rationality emphasized in other accounts of boom–bust dynamics. In that literature, overinvestment is often attributed, for example, 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 need to learn their own fit with a new technology is sufficient to produce the same boom–bust pattern.

Because our diagnostic-learning mechanism does not rely on externalities, we deliberately strip out cross-agent effects. In our setting, experimentation reveals only a firm's own fit with the new technology; it neither changes other firms' payoffs nor transmits payoff-relevant information. Introducing externalities would leave this logic intact and may further reinforce it. For example, in a fixed market with a race for dominant share, strategic pressure to move early would amplify the front-loading we identify and sharpen the eventual shakeout.

Beyond aggregate investment, our framework also 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); Koellinger et al. (2007); Bernardo and Welch (2001)). We highlight a complementary rational margin—self-discovery. Starting a venture generates diagnostic information about whether an entrepreneur is capable of developing and commercializing that opportunity, or whether their effort is better directed elsewhere. Therefore, even failure can be valuable through the self-knowledge it produces—high entry and subsequent exit need not reflect miscalibration alone.

Finally, our learning channel is 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, diagnostic learning generates self-discovery.

The rest of the paper is organized as follows. Section 2 discusses the broad notion of exploration as a form of diagnostic self-discovery and how it naturally applies to investment in new technologies. Section 3 develops a stylized model in which investment serves as a diagnostic learning device and derives its implications for boom–bust investment dynamics. Section 4 studies the model's implications for cross-sectional sorting over time into winners and losers. Section 5 concludes. Proofs are provided in the Appendix.

2 Self-Discovery and Investment Behavior

Recent research across numerous disciplines suggests 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 new options to gain information (exploration) or choosing the best option among familiar outcomes (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.[1] 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. Critically, the mapping from "what we're 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

[1]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.

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 example. Its current AI advantage—massively parallel computation, programmable GPUs, and a deep developer ecosystem—was not 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 CEO Jensen Huang, "if we didn't build it, we'd never know" (Salian (2024)).

The late-1990s internet boom is another case. During the expansion, firms were not only learning whether the internet would become economically important but also whether they themselves could translate that new technology into viable business models. Many could not. The crash was therefore not just a change in market sentiment. It was also 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 booms 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 (Campbell and Turner (2012); Campbell (2013); Odlyzko (2010)).

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 Diagnostic Learning and Aggregate Investment

3.1 Model Setup

We now present a stylized 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. We refer to projects broadly as lines of experimentation rather than as standalone projects 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 outside type, meaning the firm has no fit.

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 $p \in (0, 1)$. Thus, with probability $1 - p$ a firm is an outside type, and conditional on having any fit ($\theta \neq 0$), that fit is uniformly distributed across the N project types. Here, p captures the *economic scope* of the new technology: higher p means the technology can be leveraged by a broader share of firms. N, in turn, reflects the technology's *granularity*: higher N means the emerging opportunity set is partitioned into a larger number of distinct skills.

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 β may vary across firms. Because learning about fit requires meaningful engagement with the technology, investment cannot be arbitrarily small. Any positive investment must therefore be at least a minimum threshold, normalized to one. Each firm is endowed with an initial capital stock k_1. Capital then evolves according to[2]

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

[2] We assume that any positive cash flow from the project is immediately paid out as a dividend.

where

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

is the investment level at date t.

Matching and Cash Flow If a type-θ 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.

Thus, a positive cash flow reveals that the firm's capabilities complement the new technology through that project type. A zero cash flow only shows that the current project type did not reveal a fit. Once matched, the firm can keep investing in that project type in subsequent periods, exploiting the fit it has uncovered through the experiment.

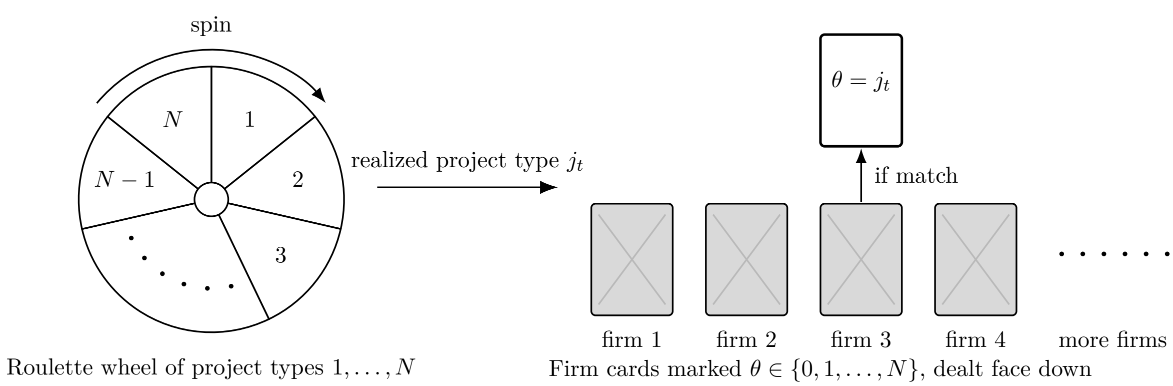

Figure 1: **The Roulette of New Technology**

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 $1, \dots, N$. A firm's type θ, 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.

This economy-wide opportunity set can be visualized as a roulette wheel, whose slots correspond to project types $1, \dots, N$, 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 that their capabilities fit the new technology. A card marked $\theta = 0$ indicates that no such fit exists within the opportunity set.

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, that is, by investing at least one unit. A positive payoff reveals that a firm's capabilities fit the new technology through the realized project type. A zero payoff reveals that a firm's capabilities did not find a fit through that project. In this case, the firm's type remains unrevealed, and its card stays face down.

Only informative arrivals matter for the learning problem. If the wheel lands on a project type that unmatched firms have already seen, the draw carries no new diagnostic information and induces no new action. We therefore suppress repeated realizations and work in event time, indexed by arrivals of previously unseen project types.[3]

Thus, for the unmatched firms, the relevant uncertainty at a given date concerns only the set of project types not yet ruled out. At the start of period t, let

$$M_t \subseteq \{1, \ldots, N\}$$

denote that set of surviving project types. 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—is uniform over M_t:

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

Assumption 1. *We assume that the horizon is no longer than the number of distinct project types the firm may test, i.e., $T \leq N$.*

Going beyond N rounds is unnecessary. Each failed experiment rules out one distinct candidate project type, so after at most N such failures the firm has exhausted the relevant opportunity set and inferred that it has no fit.

Firm's Problem Each firm's problem is to choose how much capital to commit while still searching for its fit and, once fit is revealed, how much capital to commit to exploiting the now-identified direction. Formally, each firm chooses a history-dependent investment policy

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

[3] Equivalently, one could write the model in calendar time and keep track of uninformative repeats. The results are unchanged; event time simply avoids this bookkeeping.

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 solving (3.7), in which investment serves not only as capital deployment but also as 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 remaining uncertainty about where the technology is valuable for the firm, so the problem reduces to deploying capital in a known profitable direction. 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

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

That is, once fit is known, there is no option value to waiting. Delaying investment only exposes capital to depreciation and yields no offsetting informational benefit.

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 diagnostic value because it rules out one remaining project type and narrows the firm's future search.

At date t, suppose the firm is still searching for a fit, 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

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

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. This Bellman recursion closes with terminal and boundary conditions

$$V_{T+1}(\,\cdot\,,\,\cdot\,) = 0, \qquad V_t(k,m) = 0 \text{ if } k < 1 \text{ or } m = 0.$$

Waiting. If the firm waits, capital simply depreciates and no diagnostic information arrives, 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)*,*

$$\Phi_t^{\text{inv}}(k,m) > \Phi_t^{\text{wait}}(k,m).$$

That is, whenever a firm still has enough capital to undertake an experiment ($k \geq 1$), 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, whereas waiting generates no diagnostic information. Since fit can be learned only through experimentation, the value of information pulls investment forward in time.

As a corollary, when the firm is in the searching regime and positive investment is feasible, the only remaining choice is how aggressively to experiment, so the Bellman equation (3.8) for its continuation value $V_t(k,m)$ reduces to

$$V_t(k,m) = \max_{x\in[1,k]} \Big[\mu(m)\Big(Rx+\beta\, W_{t+1}\Big((1-\delta)(k-x)\Big)\Big)+\Big(1-\mu(m)\Big)\beta\, V_{t+1}\Big((1-\delta)(k-x),\, m-1\Big)\Big]. \tag{3.12}$$

In words, the firm trades off a larger experiment today against preserving capital for future experimentation if today's project fails. Solving (3.12) yields the characterization of the optimal investment policy.

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.[4] Along the path, the last date at which a one-unit experiment remains feasible is

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

This is the firm's experimentation runway. Capping this runway by the time horizon T defines the firm's *maximal experimentation horizon*

$$L := \min\{\, T,\ \ell(k_1)\,\}. \tag{3.15}$$

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

[4] Explicitly, this feasible-experimenting capital path is

$$\tilde{k}_t = (1-\delta)^{t-1}k_1 - \sum_{u=1}^{t-1}(1-\delta)^{t-u} = (1-\delta)^{t-1}k_1 - \frac{1-\delta}{\delta}\Big(1-(1-\delta)^{t-1}\Big). \tag{3.14}$$

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

Patient vs. Impatient Firms Theorem 3.2 shows that, under diagnostic learning, time preference does not merely tilt investment intensity, as it would in standard Bayesian learning. It determines how a firm learns—through an all-in diagnostic gamble or through sequential experimentation. While we abstract away from risk aversion, lower-β firms nonetheless behave as if they are more willing to gamble, because diagnostic information is produced only by putting capital at risk. They invest all at once, because delay mainly destroys capital and the value of preserving a learning option is small. Firms with higher β, by contrast, behave more cautiously, financing a sequence of small diagnostic experiments in order to learn where their fit lies before scaling up. 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*.

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 r is below the ex ante return from experimentation. Specifically, when $r < \frac{p}{N}R$, 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 diagnostic learning before the firm knows whether it has a fit, and therefore captures the private 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,

$$V_1(k_1, N) = R\,\frac{p}{N}\,k_1. \tag{3.16}$$

(ii) For a patient firm,

$$V_1(k_1, N) = R\,\frac{p}{N}\,\left[\sum_{t=1}^{L-1}\left(\beta^{t-1} + \beta^t \tilde{k}_{t+1}\right) + \beta^{L-1}\tilde{k}_L\right], \tag{3.17}$$

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 (3.16) and (3.17) is the probability that a given experiment uncovers fit. When the new technology has broader scope p, ex ante value rises because more firms have capabilities that may travel to the new domain. A 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.

For patient firms, the bracketed term in (3.17) is the option value of sequential experimentation. A higher granularity N lowers this option value by making the path to fit more diffuse: the matching direction is found later on average, failed tests consume capital along the way, and the firm reaches scale-up later with less capital to deploy.

Since fit is learned only by engaging with the technology, the option value here derives from actively testing while preserving enough capital to scale if a match is found. The capital-feasibility constraint reinforces this effect by adding a "deadline" incentive for learning, further amplifying front-loading. With a limited learning window, uncertainty creates value from experimenting sooner. In this sense, diagnostic learning reverses the usual real-options intuition that uncertainty creates value by allowing firms to wait and see (Dixit and Pindyck, 1994). Waiting does not preserve the option to learn. It narrows the window in which that option can be exercised.

3.3 Boom-Bust Investment Dynamics

We now turn from individual investment policy to aggregate investment dynamics. The key question is how diagnostic 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, who invest aggressively up front, and patient firms, who experiment before scaling.

We continue to write $(\tilde{k}_t)_{t=1}^{L}$ for the feasible-testing capital path and L for the associated maximal experimentation horizon, as defined in (3.13) and (3.15). Let $\pi_{\text{imp}} \in [0,1]$ be the population share of impatient agents, and $\pi_{\text{pat}} \coloneqq 1 - \pi_{\text{imp}}$ the complementary share of patient agents.

Proposition 3.5. *Fix initial capital $k_1 \geq 1$. For each period $t \in \mathbb{N}$, 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 path is:*

(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

$$\text{Inv}_{\text{pat},t} = \begin{cases} \pi_{\text{pat}}, & t = 1, \\ \pi_{\text{pat}}\left(q_t + \dfrac{p}{N}\tilde{k}_t\right), & t = 2, \ldots, L-1, \\ \pi_{\text{pat}}\left(q_L + \dfrac{p}{N}\right)\tilde{k}_L = \pi_{\text{pat}}\left(1 - \dfrac{p}{N}(L-2)\right)\tilde{k}_L, & t = L, \\ 0, & t > L. \end{cases} \tag{3.18}$$

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

Among patient firms, aggregate investment comes from two sources. In (3.18), the q_t term is the unit test by firms still searching for fit. The $\frac{p}{N}\tilde{k}_t$ term captures scale-up by firms that found fit one period earlier and now scale up with their remaining capital.

Aggregate Investment Path Figure 2 visualizes the aggregate investment time path characterized in Proposition 3.5. For comparison, it also plots the corresponding investment path under a mature technology—one whose relevant applications are already well understood, so firms know from the outset where their capabilities fit.[5] Comparing the mature-technology benchmark with the new-technology path isolates the self-discovery incentive 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 even in the absence of aggregate shocks or financial frictions. Under a novel technology, investment is not merely a scale choice but also a diagnostic experiment. Firms put capital at risk 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

[5] When the technology is already mature, firms know their type. Then there is no informational reason to invest early: firms with $\theta = 0$ never invest, while firms with $\theta \neq 0$ simply wait until the current draw matches their known type and then invest all available capital. Aggregate investment in period t is therefore

$$\text{Inv}_t^{\text{mat}} = \mathbf{1}_{\{t \leq N\}} \cdot \frac{p}{N}\,(1-\delta)^{t-1} k_1.$$

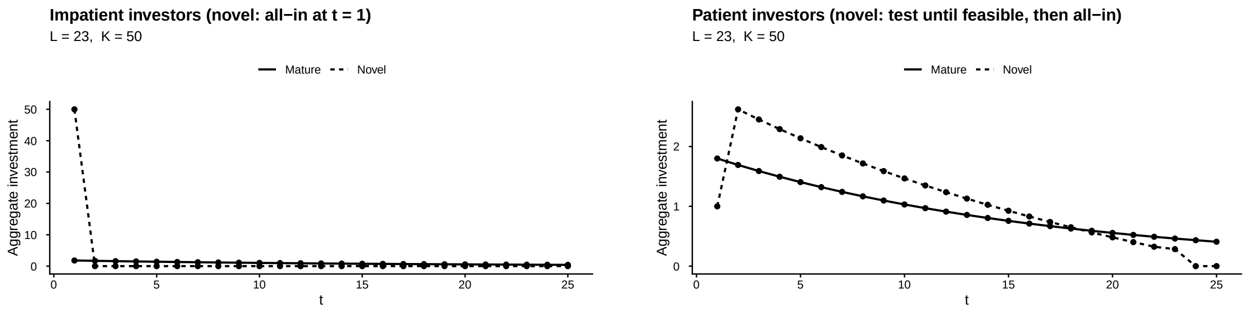

(a) Impatient firms.

(b) Patient firms.

Figure 2: **Boom–bust investment under a new technology.** The dotted 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.

reveals a direction worth scaling, while a failed one rules out a dead end. Investment is therefore pulled toward the front end of the technology's life cycle. The initial *boom* reflects broad participation by firms trying to discover their match. Capital is deployed before fit is known, because fit can be revealed only through engagement.

The subsequent *bust* follows from the endogenous exhaustion of the same learning motive. As experiments accumulate, uncertainty about fit is resolved. Some firms learn that their capabilities do not travel to the new domain and stop investing. Others discover a viable match and transition from experimentation to exploitation. Either way, the mass of firms still trying to learn where they fit shrinks over time. The decline in aggregate investment reflects this sorting process, as the economy has learned which firms can actually use the new technology and which cannot.

The Effect of Time Preference Proposition 3.5 also makes clear how time preference shapes the timing of the boom. Impatient firms are willing to put capital at risk immediately in order to learn whether they should scale or exit, rather than preserve capital for gradual experimentation. A larger share of impatient firms therefore raises initial aggregate investment. The initial boom becomes sharper, and the subsequent sorting process begins from a larger stock of capital already put at risk before fit is known.

Visually, a larger impatient share shifts 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. For any technology scope or capital level, it raises investment at the start of the boom. 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.

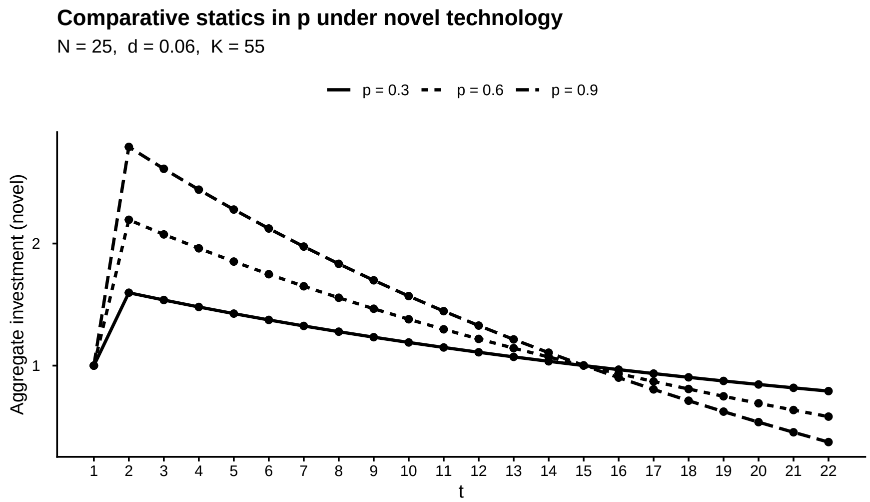

Figure 3: **Technology scope and boom–bust investment.** Each line plots patient-firm aggregate investment for a different scope $p = \Pr(\theta \neq 0)$. A larger scope p raises the early boom because fit is more likely and more experiments turn into scale-up. Investment then declines as fewer firms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

The Effect of Technology Scope Figure 3 illustrates the investment path of patient firms under the novel 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. In the aggregate investment path, this higher value is realized through the flow into scale-up. More firms therefore discover fit in each early round and begin deploying their remaining capital. This amplifies the initial boom, after which investment falls as the search pool thins.

This is the sense in which more inclusive technologies generate larger booms. When a new domain can be tried by many kinds of firms—for example, commercial real estate markets with relatively low entry barriers—many entrants have a plausible route to fit. Early participation is broad, more experiments reveal scalable uses, and the learning-driven boom is larger. A more specialized technology offers fewer such routes. Fewer experiments lead to scale-up, the expected benefits of learning are smaller, and the boom is correspondingly muted.

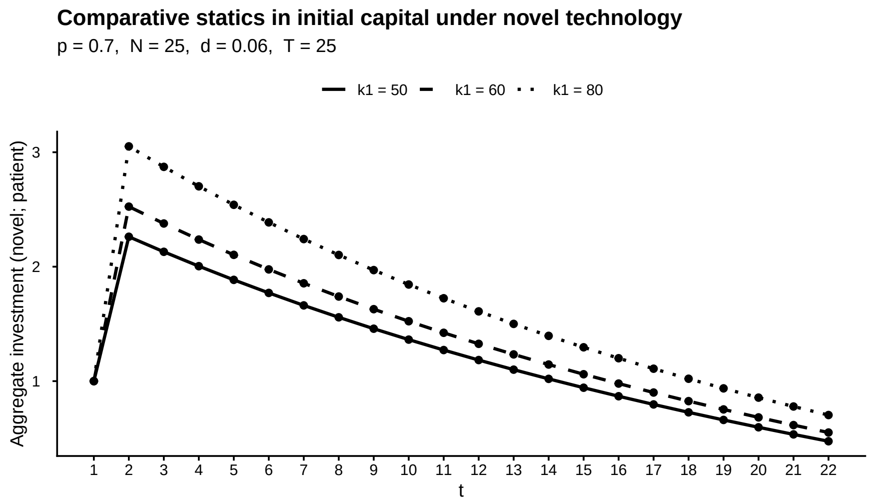

Figure 4: **Access to capital and boom–bust investment.** Each line plots patient-firm aggregate investment for a different initial capital stock k_1. Higher k_1 raises and prolongs the boom because firms can keep searching longer and successful firms scale with more capital left to deploy.

The Effect of Access to Capital More initial capital k_1 extends the boom phase because firms can keep searching, 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 paths lie above lower-k_1 paths and remain elevated for longer. Greater access to capital therefore raises and prolongs the boom–bust cycle by expanding the economy's capacity to fund diagnostic experimentation.

The traditional view of financial constraints is that they prevent firms from undertaking positive-NPV projects or force them to cut R&D when cash flow is low (see Fazzari et al. (1988); Hall and Lerner (2010); Kerr and Nanda (2015)). We highlight a different channel. 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 whether such projects exist for them. Access to finance lets firms absorb zero-payoff experiments long enough for fit to become visible.[6] When financing is abundant, the search phase lasts longer and successful firms reach scale-up with more capital left to deploy. When financing is scarce, experimentation is cut short before learning has fully played out. The bust then need not reflect deteriorating fundamentals. It can occur because the economy's capacity to fund

[6]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)).

diagnostic experimentation has been exhausted.

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 reinforce the diagnostic-learning channel, making the boom–bust investment dynamics sharper still.

As the incumbent alternative becomes less attractive, the opportunity cost of delaying increases and firms have stronger incentives to experiment early. Creative destruction therefore pulls investment further toward the early adoption phase and steepens the subsequent decline after firm-level self-discovery has run its course.

4 Sorting into Winners and Losers

The aggregate investment path characterized above reflects an underlying firm-level sorting process. As experimentation proceeds, some firms discover scalable fit, while others learn that their capabilities do not carry over to the new technology. At the disaggregate level, the same process appears as broad entry, failure, exit, and the eventual concentration of activity among a smaller set of scaled survivors. We now characterize these dynamics, beginning with failure rates during experimentation and then turning to the dynamic sorting of firms into winners and losers.

4.1 Entrepreneur 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 firms that remain active search over a narrower set of unresolved directions. Conditional on still experimenting, the probability of another no-match outcome therefore declines over time.

Proposition 4.1 records this conditional failure probability in closed form 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*

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

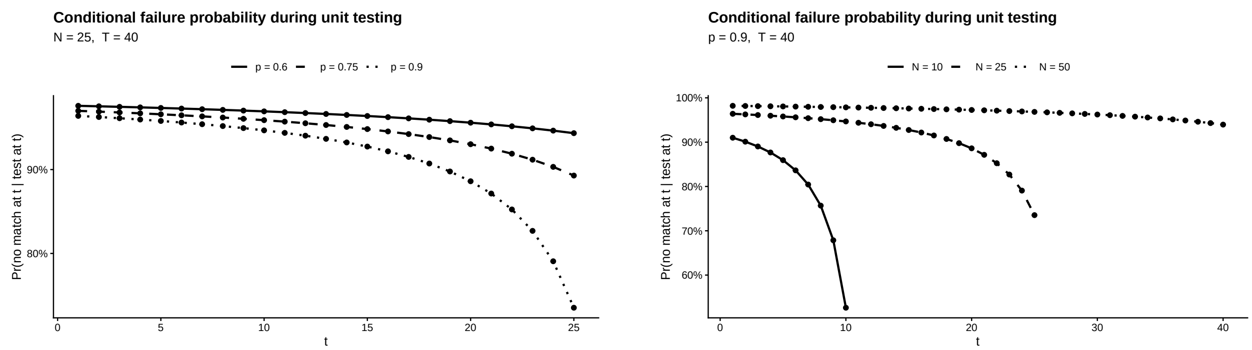

(a) Technology scope p. Larger scope lowers failure rates.

(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.

In the expression (4.1), $\frac{p}{N}$ is the mass that discovers fit in the current round, and $1 - \frac{p}{N}(t-1)$ is the mass still searching at the start of t. Their ratio is the success probability for still-searching firms. The failure rate is the complement. Figure 5 visualizes this result. Each path plots the probability that a firm still searching for fit receives another no-match outcome in period t, for a given parameter value. The common downward slope reflects the selection induced by diagnostic 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 can be tried by a wider range of firms, so more bundles of know-how, routines, and complementary assets have a plausible route to fit. The same breadth that makes more early experiments turn into scale-up also makes a still-searching firm less likely to receive another no-match outcome. Larger scope therefore 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 shifts the failure path upward and makes it decline more slowly. 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 be the one that uncovers a firm's fit. A failed attempt is still informative, but it rules out only one direction at a time. Search therefore remains failure-intensive for longer. Greater granularity makes self-discovery slower and more failure-prone by spreading fit across a larger set of possible directions.

The same early participation that generates the investment boom also produces high early failure. Firms enter early because investment is the only way to learn whether their capabilities fit the new technology. Many of those early diagnostic attempts therefore return no match. Failure is not only evidence of mistaken entry. It is also the information produced by entry. Cumulative exit rises even as the conditional failure probability among continuing firms falls.

These failure dynamics also anticipate the sorting results in Section 4.2. They are not a natural prediction of standard Bayesian learning with a common payoff primitive. In that setting, early outcomes mainly update a shared posterior and move firms together as common beliefs change, rather than generating firm-level separation. Under diagnostic learning, by contrast, investment becomes the process through which firms are sorted by fit. Each experiment reveals something about a particular firm's capabilities, so broad early participation and high early failure are two sides of the same self-discovery mechanism.

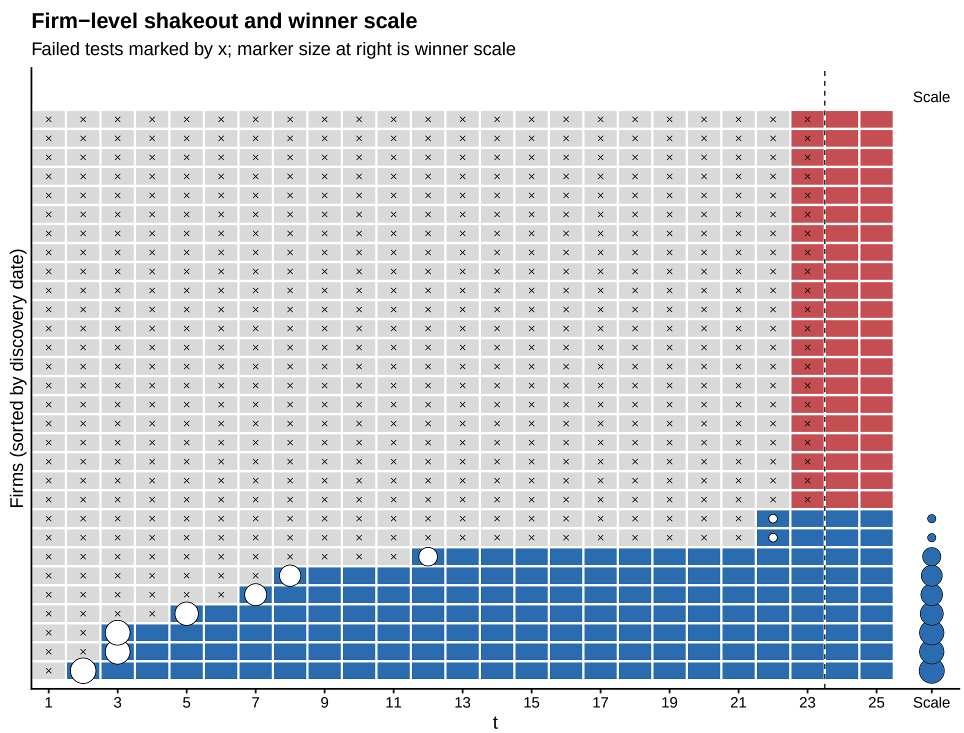

Figure 6: Firm-level shakeout and winner scale.

4.2 Winner-Loser Sorting over Time

The conditional failure rates above describe the experience of firms that remain in search. We now turn to the cross-sectional composition of firms, disaggregating the aggregate boom–bust dynamics characterized in Section 3.3. Diagnostic learning turns broad early participation into a growing winner pool, a shrinking search pool, and a residual loser pool. Within the winner pool, earlier discovery also translates into larger scale.

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 window, as in* (3.13) *and* (3.15)*, 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

$$W_t = \frac{tp}{N}.$$

(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 $s < r < L$, then $\tilde{k}_{s+1} > \tilde{k}_{r+1}$. Thus, earlier winners enter exploitation at larger scale.

For impatient firms, the sorting is immediate: the all-in experiment at $t = 1$ creates a winner mass p/N and a loser mass $1 - p/N$. Thus, the impatient firms add an initial jump in the winner and loser pools, while the gradual shakeout comes from patient firms.

Early in the wave, firms put capital at risk to learn whether the new technology fits their capabilities. Impatient firms compress this sorting into the initial diagnostic test, while patient firms generate a gradual shakeout: each round peels off a thin but constant cohort of winners, and the remaining search pool becomes smaller and increasingly tilted toward firms that were never a good fit.

There is also sorting within the winner pool. Firms that win early become the dominant adopters because they discover fit before depreciation and failed tests erode their capital, so they scale up sooner and at larger size. Late winners still find fit, but they do so with a more depleted capital stock. Hence sorting occurs on both the extensive margin—who becomes a winner—and the intensive margin—how large each winner becomes.

Taken together, diagnostic learning generates the familiar empirical pattern of broad early participation, high early failures and exits, and eventual concentration among a narrower set of scaled survivors.

Figure 6 visualizes this result.[7] Each row is a firm, sorted by the date at which it first discovers fit, and each column is a period. Light cells indicate firms still experimenting, blue cells indicate firms that have discovered fit, and red cells indicate firms that reach the end of the experimentation window without discovering fit. The symbol $\times$ marks failed diagnostic tests along the search path.

Over time, some firms peel off into the winner pool of blue cells, while others continue searching until the feasible experimentation window closes at L, marked by the dashed vertical line. The rightmost column records the intensive margin—marker size is proportional to winner scale, with larger dots corresponding to firms that discovered fit earlier and therefore reached exploitation with more remaining capital.

In classic industry life-cycle and shakeout models such as Jovanovic and MacDonald (1994), the industry learns which firms can implement the new technology, and market selection reallocates activity toward the more efficient producers. The learning is therefore external to the firm: the market discovers which firms are productive. Here, the central learning problem is internal to the firm. A firm invests in order to learn whether its own capabilities fit the new technology at all. The resulting shakeout is therefore not only market selection among known producers, but the aggregate trace of private self-discovery.

5 Conclusion

We develop a stylized model that isolates an underexplored motive for real investment in new technologies. Firms invest not only for expected cash flows, but also to learn where their own capabilities apply. The key idea is simple, and familiar from the experience of firms building at the frontier—comparative advantage is often understood only retrospectively, after firms try to build with the technology themselves. Investment is therefore not merely capital allocation under uncertainty, but also a diagnostic act through which firms learn whether their capabilities fit the new domain.

This diagnostic-learning channel generates several familiar empirical patterns in a trans-

[7]The figure is a finite-firm simulation with 30 firmes, used to make the sorting process visible at the firm level. Proposition 4.2 characterizes the corresponding continuum masses.

parent way. During periods of rapid technological change, our model predicts an unusually early investment surge, elevated early failure and exit, and eventual concentration among firms whose comparative advantages have been revealed through experimentation. These boom–bust dynamics are stronger when the technology is accessible to a broader range of firms, when the investor base is more impatient, and when firms have greater access to capital. That this full set of dynamics arises from diagnostic learning alone suggests that the self-discovery motive may be a first-order driver of new-technology 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 π that begins with $x_t = 0$.

If π 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 π does invest at some future date. Let $r \geq 1$ denote the number of initial waiting periods under π, 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 π would take one period later after the same sequence of success/failure outcomes.

The first shifted investment is feasible because under π,

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

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^{π}_{s+1} denotes the capital under π 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 π 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 π 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 π, 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 π succeeds with conditional probability $\mu(m) > 0$. Therefore $\hat{\pi}$ yields strictly higher expected payoff than π.

Since π 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 (3.12) 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, (3.12) 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

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

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 π_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 π_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 (3.12) 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-agent case is immediate. We consider the patient agents. Let S_t denote the event "still searching at the start of period t." Along the feasible experimentation path, define

$$m_t := N-t+1.$$

Then

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

Moreover, for $s\le L-1$ (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,

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

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

A.5 Proof of Proposition 4.2

Fix $t \le L$. 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. So 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).

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.

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.

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.

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.

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.

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.

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

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.

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.

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.

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.

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.

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.

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.