Knowledge Cycles and Corporate Investment*

M. Cecilia Bustamante† Julien Cujean‡ Laurent Frésard§

December 21, 2021

Abstract

We examine how firms invest along their knowledge cycles. We show that if investment is only a means to accumulate capital, it declines over the cycle and moves exactly together with value. But in fact because investment also creates knowledge—serendipitously—it is high early *and* late in the cycle. Its relation with value spikes before new cycles start and declines thereafter. We uncover this pattern in the data, identifying new cycles using sharp changes in patents' citations to prior technologies. Cycles' length has tripled in recent years, coinciding with concurrent changes in the investment-value relation.

Keywords: Investment, Experimentation, Exploration, Knowledge Capital, Knowledge Cycles, Total Q, Intangible Capital.

---

*We thank Luke Taylor (EFA discussant) for suggesting the Intel example. We are grateful to Daniel Andrei, Serguey Braguinsky, Nicolas Crouzet (Cavalcade discussant), Howard Kung, Jerry Hoberg, Johan Hombert, Julien Hugonnier, Nobuhiro Kiyotaki, Mark Lowenstein, Albert Kyle, William Mann (ASU discussant), Rich Matthews, Erwan Morellec, Lukas Schmid (WFA discussant), Philip Valta, Laura Veldkamp, Alexei Zhdanov (FOM discussant) and Jake Zhao (MFA discussant) for their feedback, and to seminar participants at the 4Nations Cup at HEC Paris, the University of Southern Denmark, University of Bern, the USC Finance Organizations and Markets (FOM) Conference 2019, the ASU Sonoran Winter Finance Conference 2020, the WFA Annual Meetings 2020, the Federal Reserve Board, Universite Paris Dauphine, Queen Mary University, the Econometric Society World Congress 2020, the European Finance Association Annual meeting 2020, the Women Assistant Professors in Finance Conference at Stern, the SFS Cavalcade conference 2021, the MFA Annual Meetings 2021, and the Society for Economic Dynamics annual meeting 2021. We thank Tina Oreski for excellent research assistance.

†University of Maryland, R. H. Smith School of Business. Contact: 4423 Van Munching Hall, College Park, MD 20742, USA; +1 (301) 405 7934; mcbustam@rhsmith.umd.edu

‡University of Bern, Institute for Financial Management, and CEPR. Contact: Office 214, Engehaldenstrasse 4, CH - 3012 Bern; +41 (0) 31 631 3853; julien.cujean@ifm.unibe.ch

§Universita della Svizzera italiana (USI Lugano), Swiss Finance Institute, and CEPR. Contact: Via Giuseppe Buffi 13, CH - 6900 Lugano; +41 (058) 666 4491; laurent.fresard@usi.ch

1 Introduction

The creation of technological knowledge within firms follows cycles. Due to workers' mobility, social interactions, imitation or competitive pressures, firms cannot fully appropriate the returns on their knowledge. Because knowledge "dissipates" the dollar benefits of knowledge on a technology are *bounded* and operating a same technology forever is not optimal (Young, 1993). Rather, exploring new technologies is necessary to survive and generate future revenues. To the extent that knowledge is partially transferable across technologies (Jovanovic and Nyarko, 1996) the need to eventually invest in new technologies leads to partial knowledge resets—firms' knowledge follows cycles.

The case of *Intel* is a classic example (Casadesus-Masanell, Yoffie, and Mattu, 2014). From the early 1970s to 1982 *Intel* developed RAM chips technology, creating unique knowledge and substantial value for its shareholders. *Intel*'s profits were eventually eroded by competition from Japan and the resulting dissipation of its knowledge, precipitating a new knowledge cycle of central processors technology that started in 1988 and declined in the early 2000s. The value of *Intel* closely followed its knowledge cycles, peaking early in each cycle when its superior knowledge conferred a leading market position, and declining afterwards when this knowledge dissipated.

This example suggests that knowledge cycles likely have important implications for firms, and raises several questions into which the literature provides limited insights. If firms' value varies over knowledge cycles, at which stage of the cycle should firms invest most? Do investment and value move together along cycles, as predicted by neoclassical models, or in opposite directions? What determines the exploration of new technologies and what is the resulting length of knowledge cycles? These questions appear particularly relevant in a time of rapid technological changes (e.g., in artificial intelligence or clean technologies) in which many companies must decide whether and how to adjust their knowledge.

In neoclassical models investment is only a means to accumulate capital (Hayashi, 1982). However, knowledge creation is not separate from capital accumulation but a by-product of economic activity (Arrow, 1962). Reminiscent of Grossman, Kihlstrom, and Mirman (1977), investment is a form of experimentation with a technology and a means to accumulate knowledge. In the 1980s *Intel* faced two architectures in designing microprocessors (CISC and RISC). The risk was that the better-performing but costlier RISC architecture would eventually displace CISC. *Intel* decided to invest in both and eventually resolved uncertainty by incorporating RISC techniques it experimented with to improve the Pentium processor. This example illustrates that the outcomes of experimentation are uncertain (e.g., Callander (2011)). That is, distinct from capital, the accumulation of knowledge is genuinely noisy.

We develop a model that retains the main themes of neoclassical models while capturing knowledge cycles and considering investment as a "knowledge channel" through which firms accumulate knowledge serendipitously. The model delivers four main insights. In contrast to neoclassical predictions, *(1)* firm's investment is high both early *and* late in the cycle, *(2)* the relation between investment and value shifts around knowledge resets, and *(3)* in a specific way—it spikes late in the cycle and declines early in the cycle. In addition, whereas knowledge resets often follow pre-determined schedules in the literature (e.g., Berk, Green, and Naik (2004) or Pastor and Veronesi (2009)), cycles in this model are endogenous: *(4)* they are entirely determined by firm value, short, and expand when knowledge is easier to protect. We provide empirical support for these predictions and show that knowledge cycles and the need to create knowledge are important to understand firms' investment decisions.

Two separate intuitions suggest that investment should be high at the beginning of a cycle. At early stages a new technology has highest option value. In addition, early stages are associated with high uncertainty and thus high learning benefits from investing. Perhaps it would seem that investment should be low late in the cycle when both value and uncertainty are low. Yet, this intuition omits that when investment creates knowledge the possibility to explore a new technology is a put option on experimentation. The firm develops an attitude towards the uncertainty of experiments, as their outcomes affect value only after the firm has made its investment decision. Early in the cycle firm value is concave in knowledge—experimentation is a risk the firm would prefer to avoid. Now, since learning benefits are bounded, which we model by assuming that knowledge accumulation erodes revenues, value eventually becomes convex. The firm changes attitude and starts gambling on new technologies by investing actively when the end of a cycle is in sight.

Because firm value does not incorporate attitude towards noisy experiments it is no longer a sufficient statistic for investment—their relation varies along the knowledge cycle. Early on, when uncertainty about the new technology is high, the firm prefers to avoid experimentation and invests prudently when in fact value is highest; this weakens their relation. A cycle ends when exploration is optimal, and this occurs only when the marginal value of knowledge on the current technology is eventually zero. Since more knowledge harms value late in the cycle due to the revenue erosion it implies, the value of capital and that of knowledge together go up, exactly when the firm gambles on new technologies by raising investment. Thus the relation between value and investment is stronger late in the cycle.

This mechanism is based on a simple model of investment. A (risk-neutral) firm operates a given technology whose quality is unobservable. "Knowledge" is the firm's confidence, in a statistical sense, about the technology's quality. The firm combines this technology with capital to produce an output, and accumulates knowledge by observing productivity realizations and by actively experimenting with its technology through investment. Experiments incur adjustment costs and their outcome is uncertain although informative on average. The firm can decide at any time to explore new, unknown technologies (Jovanovic and Nyarko, 1996). Switching technologies makes part of existing capital obsolete and resets knowledge (Bresnahan, Greenstein, Brownstone, and Flamm, 1996). The firm explores when confident that its technology is poor, but also when valuable as in this case more knowledge precipitates dissipation (Lucas, 1988). Knowledge accumulates through realizations of productivity and noisy experiments until the firm explores and starts a new knowledge cycle.

A challenge in assessing the relevance of knowledge cycles for firms' investment is to identify knowledge resets in the data. We propose to measure them using information from the patents of a large sample of U.S. publicly-listed firms between 1976 and 2017. All patents a firm cites in its own patents (in recent years) constitute its "knowledge base", and captures technological knowledge it has accumulated over time (Ma, 2021). Aggregating cited patents across NBER technology classes we identify major annual changes in a firm's knowledge base through cosine similarity between the firm's current and lagged distributions of its technological knowledge across classes. A "knowledge reset" corresponds to a year in which the distributions of cited technologies suddenly becomes dissimilar.

To establish the knowledge channel empirically we examine whether the relation between investment and value shifts around knowledge resets. To map their definition in the model we follow Peters and Taylor (2017) and consider "total investment" in physical and intangible capital and "total Q", the ratio of firm value to physical and intangible capital. As predicted, we find substantial variation in investment$-Q$ sensitivity, with a significant spike immediately prior to the reset (late in the cycle) and a steady decline (of about 40%) following it (early in the cycle). Additional tests indicate this pattern is unlikely explained by changes in product market competition, financing frictions or ownership around resets, nor by mechanical changes in patenting activity nor by acquisition of firms with distinct technologies. We also find that cycles are short, with a median of 5 to 7 years, and are shorter in more competitive industries in which knowledge likely dissipates faster.

Notably, knowledge resets have become less prevalent over time, with cycle length almost tripling between 1980 and 2017. Though not our focus, longer cycles could originate from the increased difficulty to generate new ideas (Bloom, Jones, van Reenen, and Webb, 2017) or declining competition (Gutierrez and Philippon, 2018). Since the investment$-Q$ relation is closely linked to knowledge cycles, it is interesting to contrast their lengthening to the evolution of this relation. Gutierrez and Philippon (2017) and Alexander and Eberly (2018) document a decline in investment in the last twenty years but no decline in Q, which our sample confirms. Investment$-Q$ sensitivity being strongest around resets, its weakening could be tied to less frequent resets and underlying changes in firms' knowledge creation.

This paper adds to a growing literature on the role of intangible capital in firms' investment and valuation. Recent models do not consider knowledge cycles and often incorporate intangible capital as a production factor that accumulates similarly to physical capital, and under perfect information (e.g., Eisfeldt and Papanikolau (2013), Peters and Taylor (2017), Crouzet and Eberly (2018), Crouzet and Eberly (2020)). Closer to this paper, Andrei, Mann, and Moyen (2018) develop a model in which technological innovation occurs randomly and firms learn about it passively. In contrast, we consider that investment is a means to create knowledge, exploration of technologies is a decision firms make, and the resulting knowledge cycles are important for investment.

We also contribute to the literature in finance that explicitly models the process of knowledge creation within firms (e.g., Berk et al. (2004), Bergemann and Hege (2005), Pastor and Veronesi (2009), Manso (2011), or Manso, Balsmeier, and Fleming (2019)). Whereas we share the view that experimentation and exploration dictate the evolution of firms' technological knowledge, we associate experimentation with investment and we focus on the implications for the relation between investment and value over knowledge cycles.

This paper also belongs to the literature on experimentation, and is closest technically to Moscarini and Smith (2001).[1] Similar to their decision maker, our firm decides how

[1] For instance, Grossman et al. (1977), Weitzman (1979), Jovanovic and Rob (1990), Rob (1991), Aghion, Bolton, Harris, and Jullien (1991), Jovanovic and Nyarko (1996), Keller and Rady (1999), Callander (2011).

much to experiment (to invest) and when to take a payoff-relevant action (permanent stop in their case, and exploration in ours). We also find that experimentation intensity is high when the action is about to be taken. Their decision maker minimizes the present value of experimentation costs which delays high-intensity experimentation. Our firm experiments more late in the cycle because exploration induces gambling on new technologies. Another difference is in the formulation and scope. They focus on sequential analysis (Wald, 1947), whereas we study investment in a neoclassical context.

2 Model setup

We argue knowledge is not a typical input in a production function. A key distinction with physical capital is that knowledge accumulation is uncertain. Second, knowledge is a by-product of economic activity: investment is not only a way to accumulate capital, it is also a way to create knowledge by experimenting with a technology. Third, because the dollar benefits of learning on a technology are bounded, exploration of new technologies is necessary to regenerate revenues—knowledge accumulates and depreciates in cycles. Our objective is to develop a model that retains the main themes of neoclassical investment models while simultaneously capturing these three dimensions of knowledge creation.

2.1 Knowledge creation within the firm

Consider a firm that combines capital, K_t, with some technology to produce a unique non-storable output, Y_t, at time t according to the production function:

$$Y_t = A_{n,t} K_t^{\alpha},$$

where $A_{n,t}$ denotes the productivity of technology indexed by $n \in \mathbb{N}$ and $\alpha \in (0, 1)$ denotes returns to scale on capital. We use the concept of technology in a broad sense as a "technology" may refer for instance to production techniques, but also management practices (e.g., Bloom, Sadun, and van Reenen (2016)), or organization designs (e.g., Prescott and Visscher (1980)). At every date t, the firm can choose to continue operating the current technology n about which it has accumulated knowledge (to be defined), or to abandon it and to explore instead by selecting another unknown technology, say $n + 1$.

Each technology n is characterized by the growth rate of its productivity, M_n, which reflects its quality.[2] The quality of each technology is unobservable, and the link between qualities across technologies is informational. Following Jovanovic and Nyarko (1996), a switch to another technology reduces knowledge temporarily. In the main analysis, we focus on the case in which each technology resets knowledge completely, or equivalently the firm "learns and forgets" (e.g., Benkard (2000)); in Section 4.4 we relax this assumption, allowing knowledge to be partially transferable across technologies, and examine its implications. Upon selecting technology $n + 1$, its quality is randomly drawn as per:

$$M_{n+1} \sim \mathcal{N}(0, \tau_M^{-1}), \quad \text{with } M_n \perp M_{n+1}, \quad \forall n \in \mathbb{N},$$

so that the technology index n can be discarded, and the parameter τ_M captures prior precision regarding the quality of the technology.[3] Furthermore, as in Jovanovic and Nyarko (1996), we assume that once a technology has been abandoned it can never be "recalled".

Learning about the quality of a technology is a noisy process. Although the firm does not observe the quality of its current technology, it observes realizations of its productivity. Each incremental realization, $\mathrm{d}A/A$, reveals noisy information about M according to:

$$\mathrm{d}A_t/A_t = \tau_A^{1/2} \frac{M}{\Omega_t^{1/2}} \mathrm{d}t + \mathrm{d}B_t, \tag{1}$$

where B is a Brownian shock, which will account in part for changes in the firm's expectations about quality; Ω_t denotes the conditional variance of the firm's estimate of quality at time t (defined below). To simplify interpretation, we scale quality, M, by the conditional standard error, $\Omega^{1/2}$, so that the conditional variance of productivity realizations is normalized to one. This normalization has two economic implications. First, uncertainty about productivity is fixed—the parameter τ_A measures the informativeness of the productivity signal (about M) in time-invariant units of uncertainty (e.g., Kyle, Obizhaeva, and Wang (2018)). Second, learning affects directly the growth rate of productivity, A, through the conditional variance,

[2] The literature (e.g., Acemoglu (2009)) typically assumes that technological innovation modifies the level of productivity (i.e., A), as opposed to its growth rate (M). Adding this feature to the model is possible but is strictly equivalent to reducing obsolescence costs, given the way we model the evolution of capital below. Hence, we characterize technologies only in terms of their productivity growth.

[3] The assumption that each technology resets knowledge completely, along with the specification of the signals in Eqs. (1)–(2) below, jointly imply that the parameter τ_M will play no role in the analysis.

Ω. Specifically, because the growth rate in absolute value is decreasing in Ω and learning reduces Ω, learning increases the future absolute growth rate of productivity.

Abandoning the current technology to explore is costly, as it implies that a fraction ω of capital K becomes obsolete upon exploration. As is customary in the literature the firm accumulates capital through investment. Let $I_t \geq 0$ represent investment in capital at time t and $i_t = I_t/K_t$ be the corresponding investment rate. Further assume that capital depreciates at rate, δ. It follows that the net change in the firm's stock of capital is:

$$\mathrm{d}K_t = (i_t - \delta)K_t\mathrm{d}t - \omega K_{t-}\mathbf{1}_{t=\nu},$$

where, throughout the paper, ν denotes dates at which the firm decides to abandon the current technology and to explore an unknown technology. Absent exploration capital accumulation is identical to that in neoclassical models.

In this model knowledge is a serendipitous by-product of economic activity. Investment is not only a means to accumulate capital, it is also a means to create knowledge by actively experimenting with a technology (e.g., Arrow (1962); Grossman et al. (1977); Farboodi, Mihet, Philippon, and Veldkamp (2019)). For instance, Braguinsky, Ohyama, Okazaki, and Syverson (2021) show that expanding the set of machines used for production enlarges the productivity of machines and thus may refine the firm's information about its current technology.[4] Experimentation takes the form of an expansion of capital on which the current technology is applied. Formally, in addition to productivity realizations the firm observes passively, by investing ($i_t > 0$) the firm receives an informative signal flow, $\mathrm{d}S$, about M:

$$\mathrm{d}S_t = \frac{\tau_S^{1/2} i_t^{1/2}}{\Omega_t^{1/2}} M\mathrm{d}t + \mathrm{d}B_{S,t}, \tag{2}$$

where τ_S is a parameter that, together with investment, $\tau_S i$, captures how informative experimentation is; setting it to zero also allows us to shut down experimentation (see Section 3). The more the firm invests relative to its stock of capital, the more it learns about M.

When the firm invests to create knowledge, the outcome of its experiments is uncertain (e.g., Callander (2011)). We capture this genuine feature of experimentation by introducing

[4] As noted by Arrow (1962): "Each new machine produced and put into use is capable of changing the environment in which production takes place so that learning is taking place with continually new stimuli."

an independent Brownian, $B_S \perp B$, in signal realizations—adding new machines may harm the firm's current prior about M, e.g., if the new machines provide "contrarian" results. The two Brownians B and B_S together account for changes in firm's expectation about quality. Similar to the passive productivity signal, informativeness of S is measured in time-invariant units of uncertainty. To simplify intuition, we rule out learning through capital liquidation by making investment irreversible. Since investment augments both capital and knowledge, it represents total investment in capital and knowledge (Peters and Taylor, 2017).

The firm's information about the quality of the technology it operates entirely results from observing the passive and active signals:

$$\mathscr{F}_t = \sigma\left((A_s, S_s) : s \leq t\right).$$

Based on this information, the firm computes an estimate, $\widehat{M}_t = \mathbb{E}[M|\mathscr{F}_t]$, of its quality. The conditional error variance of this estimate is $\Omega_t = \mathbb{V}[M|\mathscr{F}_t]$ (introduced above). Standard arguments imply that the firm gradually updates these statistics according to:

$$\begin{aligned} \mathrm{d}\widehat{M}_t &= \Omega_t^{1/2}\tau_A^{1/2}\mathrm{d}\widehat{B}_t + \Omega_t^{1/2}\tau_S^{1/2}i_t^{1/2}\mathrm{d}\widehat{B}_{S,t} - \widehat{M}_{t-}\mathbf{1}_{t=\nu} \\ \mathrm{d}\Omega_t &= -\Omega_t\left(\tau_A + \tau_S i_t\right)\mathrm{d}t + (\tau_M^{-1} - \Omega_{t-})\mathbf{1}_{t=\nu}, \end{aligned}$$

where $\widehat{B}$ and $\widehat{B}_S$ are two independent Brownians under the firm's probability measure.[5] The firm updates its estimates about the quality of its current technology continuously by observing realized productivity and by experimenting with it, and resets its estimate and precision to priors (0 and τ_M, respectively) whenever it decides to explore another technology.

We can now be specific about what we mean by "knowledge." The firm's stock of knowledge, Z_t, represents the information it has at date t about the technology it operates. In the model the firm's estimate, $\widehat{M}$, and its standard error, $\Omega^{1/2}$, do not matter separately; they only matter as a $t-$statistic ratio, $Z \equiv \widehat{M}/\Omega^{1/2}$, which summarizes this stock of knowledge:

$$\mathrm{d}Z_t = \underbrace{\tau_A/2Z_t\mathrm{d}t + \tau_A^{1/2}\mathrm{d}\widehat{B}_t}_{\text{passive learning}} + \underbrace{\tau_S i_t/2Z_t\mathrm{d}t + \tau_S^{1/2}i_t^{1/2}\mathrm{d}\widehat{B}_{S,t}}_{\text{learning by experimenting}} \underbrace{-Z_{t-}\mathbf{1}_{t=\nu}}_{\text{knowledge reset}}. \tag{3}$$

Positive (negative) values of Z indicate that the firm is confident, in a statistical sense, that

[5] See Theorem 12.7 in Lipster and Shiryaev (2001).

the quality of its current technology is high (low). On average knowledge moves away from prior beliefs, $Z_0 = 0$, accumulating in continuation of the current estimate, Z_t, at the speed with which the passive and active signals reveal information ($\tau_A/2$ and $\tau_S i_t/2$, respectively). Exploring a new technology resets the firm's knowledge to prior beliefs, $Z_\nu = 0$. The stock of knowledge is the key variable based on which the firm decides to experiment or explore.

2.2 Knowledge and bounded revenues of the firm

The firm has incentives to explore new technologies only if the benefits of knowledge are *bounded* (Young, 1993). A large literature indicates that because technological knowledge is non-rival and largely embodied in (inalienable) workers, firms cannot fully appropriate its returns. Proprietary knowledge thus typically dissipates over time (e.g., Romer (1986); Lucas (1988)). Knowledge dissipation could result, for instance, from workers moving across firms (e.g., Stoyanov and Zubanov (2012)), social interactions (e.g., Glaeser, Kallal, Scheinkman, and Shleifer (1992)), imitation (e.g., Lieberman and Asaba (2006)), or illegal actions (corporate spying or thefts).[6] Knowledge dissipation erodes revenues, making exploration necessary to regenerate future revenues.

Suppose the firm faces an isoelastic demand curve for the good (or service) it produces, $P_t = (Y_t N_t)^{-\eta}$. Importantly, $(N_t)_{t\geq 0}$ is an exogenous process that captures the dissipation of the firm's knowledge in reduced form and $\eta \in (0,1)$ corresponds to the price elasticity of demand. The firm's revenues are increasing in the productivity of the current technology, A, but are decreasing in the intensity of knowledge dissipation, N:

$$\Pi\left(A_t, K_t, N_t\right) \equiv P_t Y_t = A_t^{1-\eta} K_t^{\alpha(1-\eta)} N_t^{-\eta}.$$

Since (the logarithm of) revenues are increasing linearly in knowledge, Z, but decreasing linearly in $\log(N)$, we specify the evolution of the dissipation intensity according to:

$$\mathrm{d}N_t/N_t = \phi Z_t^2 \mathbf{1}_{Z_t \geq 0} \mathrm{d}t, \tag{4}$$

[6] The firm may limit the dissipation of its knowledge by using patents (or trademarks). Intellectual property tools are however limited in their scope because certain types of knowledge cannot be precisely codified, are costly to give away through public filings (e.g., trade secrets), or because patents may fail to adequately protect property rights (e.g., Lanjouw and Schankerman (2001)).

with $\phi > 0$.[7] In words, we assume that knowledge dissipates faster as it accumulates, and this dissipation only occurs for "good" technologies, $Z > 0$. When a technology is likely poor, the firm has strong incentives to explore. Yet, when the firm discovers a good technology exploration occurs only when the dollar benefits of knowledge are bounded. Thus the firm does not let knowledge grow unboundedly by sticking to the same technology forever, as dissipation eventually erodes its revenues to zero (e.g., Jovanovic and MacDonald (1994)).

The rest of the model contains usual elements of the neoclassical framework. In particular, investment incurs adjustment costs, $\gamma_t(i)$, for purchasing and installing new capital, which we assume proportional to revenues (e.g., Cooper (2006); Hackbarth and Johnson (2015)):

$$\gamma_t \equiv \underbrace{\left(i + \gamma/2(i-\delta)^2\right)}_{\equiv \widehat{\gamma}(i)} \cdot \Pi\left(A_t, K_t, N_t\right), \tag{5}$$

where the price of capital is taken to be proportional to $\Pi(\cdot)/K$ and where adjustment cost is convex (e.g., Abel and Eberly (1994)). In Appendix B we consider an extension with fixed costs, which causes investment and thus experimentation to take place in rounds.

The firm maximizes the expected value of its future profits (revenues net of costs) by optimally choosing its technology and its investment. Assuming the firm is risk neutral and discounts its profits at a constant rate, r, its value, $V(\cdot)$, is given by:

$$\max_{\{\nu_n\}, i_{t+s}} \mathbb{E}\left[\int_0^\infty e^{-rs}\left(\Pi\left(A_{t+s}, K_{t+s}, N_{t+s}\right) - \gamma_{t+s}\right) \mathrm{d}s \middle| \mathscr{F}_t\right]. \tag{6}$$

Our assumption on adjustment costs guarantees that firm value is homogeneous in revenues:

$$V\left(N_t, A_t, K_t, \widehat{M}_t, \Omega_t\right) \equiv \Pi\left(A_t, K_t, N_t\right) \cdot v\left(\widehat{M}_t/\Omega_t^{1/2}\right) \equiv \Pi\left(A_t, K_t, N_t\right) \cdot v\left(Z_t\right), \tag{7}$$

where $v(\cdot)$ denotes the intensive value of the firm. Hence, the firm's knowledge, Z, is the only state variable that determines the choice between exploration and experimentation.

[7] We assume that the growth rate of $\log(N)$ is quadratic in Z, because the resulting affine-quadratic framework remains tractable. However, any *strictly* convex increase in N will keep profits bounded and produce qualitatively similar results, except in terms of the asymmetry between good and bad technologies.

3 A benchmark model without experimentation

The choice of a technology and the knowledge resets it implies are uncommon in the neoclassical literature but a central feature in this model. The relevance of knowledge cycles raises at least three questions: what determines the exploration of a technology, at which stage of the cycle is it optimal for the firm to invest, and how do cycles affect neoclassical insights? This section examines this matter by removing the effect of investment on knowledge, $\tau_S \equiv 0$. We first define marginal product of capital (hereafter q) with or without experimentation.

Definition 1. *The* ***marginal product of capital*** q *is the ratio of value of an additional unit of capital to profits, ignoring the effect of investment on knowledge:*

$$q(Z) \equiv \frac{1}{\Pi} \cdot K \cdot V_K. \tag{8}$$

A (cosmetic) difference of Eq. (8) relative to the standard definition of q is on the relation between firm value and q. In neoclassical models firm value is often assumed to be homogeneous in K so that q is proportional to V/K (average q). Our specification of adjustment costs, however, considers that firm value is homogeneous in profits, Π, (see Eq. (7)) so that q is proportional to V/Π (intensive firm value v):

$$q(Z) = \alpha(1-\eta)v(Z) \propto v(Z).$$

The main insight from the neoclassical literature is that q is a sufficient statistic for investment (e.g., Hayashi (1982)). To determine the firm's optimal investment policy, suppose for the moment that it is optimal for the firm to stick to its current technology. In this case the (intensive) value of the firm, v, associated with Eq. (6) satisfies the HJB equation:

$$\begin{aligned} &v(Z)\left(r + \eta Z^2 \phi \mathbf{1}_{Z>0} - (\eta - 1)\left(\alpha\delta + \frac{\eta}{2} - \sqrt{\tau_A} Z\right)\right) \\ &= \psi(q(Z)) + \frac{1}{2}\tau_A v''(Z) + \sqrt{\tau_A} v'(Z)\left(-2\eta + \frac{\sqrt{\tau_A} Z}{2} + 1\right), \end{aligned} \tag{9}$$

where we have defined ψ to be the maximand:

$$\psi(q) \equiv \max_{i \geq 0} \{qi - \widehat{\gamma}(i)\} \equiv (q-1)(2\gamma\delta + q - 1)/2\gamma. \tag{10}$$

The firm's optimal investment policy is given by:

$$i_t \equiv i(q) = \delta + (q-1)/\gamma,$$

thus delivering the traditional insight that q is a sufficient statistic for investment. In the benchmark experiments are uninformative and Z evolves through passive learning and resets:

$$\mathrm{d}Z_t = \tau_A/2Z_t\mathrm{d}t + \tau_A^{1/2}\mathrm{d}\widehat{B}_t - Z_{t-}\mathbf{1}_{t=\nu}.$$

Since investment only affects value through capital accumulation, the marginal product of capital satisfies $q(Z) \equiv \frac{1}{\Pi}\frac{\mathrm{d}}{\mathrm{d}i}\frac{\mathrm{d}V}{\mathrm{d}t}$, and thus summarizes entirely how investment affects value.

Not only does q determine investment, but it is also a sufficient statistic for when to explore a new technology. The solution assumes so far that it is optimal for the firm not to switch technology. When the firm decides to explore a new technology, it resets its knowledge to priors and incurs a lump-sum cost of obsolescence, ωK, on its existing capital stock. Since q in Eq. (8) is proportional to (intensive) firm value, q upon exploration is:

$$q^{\star} \equiv (1-\omega)^{\alpha(1-\eta)} q\,(0)\,, \tag{11}$$

so that the firm switches technology whenever $q \leq q^{\star}$. Because exploration resets knowledge, of which q is a function, determining when exploration is optimal implies solving for q in Eq. (9); this requires some conjecture on the dependence of firm value on knowledge.

Intuitively, q peaks around knowledge resets, $Z \approx 0$, following exploration. As Z grows negative the firm is confident that its technology is poor, and its value must decrease as Z further declines; as the firm instead becomes confident that its technology is of high quality, knowledge dissipation intensifies and its value must decline as Z further rises. We conjecture that q has a single peak, and later use parameter values under which this conjecture is verified numerically. Thus, q is high when a technology has high option value and knowledge dissipation is low ($Z \approx 0$), and this occurs following exploration. Conversely, q is low when option value is low and dissipation is high, and this triggers exploration.

Under this conjecture, there must exist two "trigger levels" of knowledge Z, say $\underline{a} < \overline{a}$, at which $q(Z) = q^{\star}$ is satisfied and the firm decides to explore. Thus, $\mathscr{A} = (\underline{a}, \overline{a})$ denotes the region in which the firm does not explore. The times $\{\nu_k\}_{k=0}^{\infty}$ in Eq. (6) at which the firm explores correspond to hitting times at which knowledge exits $\mathscr{A}$:

$$\nu_n = \inf\{t \geq \nu_{n-1} : Z_t \notin \mathscr{A}\}, \ \ \forall n \in \mathbb{N}. \tag{12}$$

Every time the firm triggers exploration, it ends a cycle and begins a new one. We refer to the period in between exploration triggers as a "knowledge cycle".

Definition 2. *A* ***knowledge cycle*** *starts when the firm resets its knowledge to 0 by exploring a new technology ($Z_0 = 0$), and ends after a period of (random) length as per Eq. (12) when the firm abandons the technology to explore a new one ($Z \notin \mathscr{A}$).*

We compute firm value piecewise over the two regions $[\underline{a}, 0]$ and $(0, \overline{a}]$. Eq. (9) does not have an explicit solution; we proceed numerically, imposing the two boundary conditions implied by Eq. (11), $v(\underline{a}) = v(\overline{a}) = (1-\omega)^{\alpha(1-\eta)} v(0)$. We obtain another boundary condition by "piecing together" firm value across the two regions, $v(0_-) = v(0_+)$. Customary "smooth-pasting" conditions then provide the remaining conditions that determine the location of the optimal exploration thresholds.[8] Figure 1 illustrates q as a function of the stock of knowledge. This illustration relies on parameter values, which we discuss in Appendix A.2.

Recall that absent experimentation knowledge is irrelevant for investment—only q matters. Now, since q itself depends on knowledge, there exists an indirect relation between investment and knowledge. Figure 1 confirms that q is high (low) early (late) in the cycle. Thus, investment is highest following recent exploration ($Z \approx 0$), and weakest when q is low and triggers exploration ($Z \approx \underline{a}$ or $Z \approx \overline{a}$). However, because investment and q move exactly together their relation—investment$-q$ sensitivity—is flat (at $1/\gamma$) throughout the cycle.

4 The knowledge channel

Unlike in the benchmark model (Section 3), investment is now a means to create knowledge. We define this "knowledge channel" as follows.

[8] Smooth-pasting conditions satisfy:

$$\lim_{Z \downarrow \underline{a}} v'(Z) = \lim_{Z \uparrow \overline{a}} v'(Z) = 0, \ \ \text{and} \ \ \lim_{Z \uparrow 0} v'(Z) = \lim_{Z \downarrow 0} v'(Z). \tag{13}$$

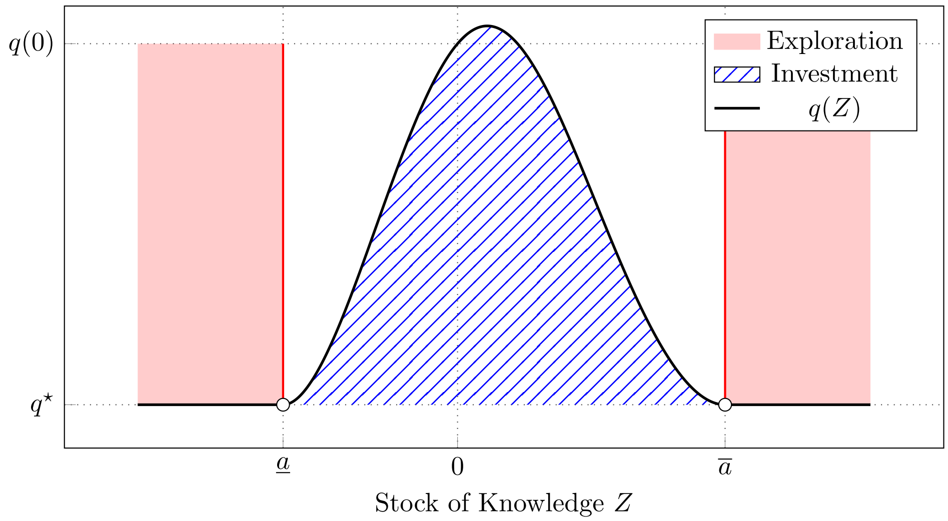

Figure 1: Marginal product of capital as a function of knowledge in the benchmark model. The figure plots q as a function of knowledge, Z. Smooth-pasting for exploration requires that q lands flat at the edges of the domain defined by the exploration thresholds $\underline{a}$ and $\overline{a}$. Parameter values are defined in Appendix A.2.

Definition 3. *The **knowledge channel** is the ratio of expected value of an additional unit of knowledge to profits, ignoring the (traditional) effect of investment on capital:*

$$c(Z) \equiv \frac{1}{\Pi} \cdot \tau_S \cdot \Omega \cdot \left(-V_\Omega + \frac{1}{2} \cdot V_{\widehat{M}\widehat{M}} \right). \tag{14}$$

The knowledge channel is the *expected* marginal benefit of knowledge, keeping the marginal effect of capital on value fixed. The effect of investment on knowledge, unlike traditional capital, is uncertain. Formally, we can rewrite c (over a period of length $\mathrm{d}t$) as:

$$c(Z)\mathrm{d}t = \frac{1}{\Pi}\frac{\mathrm{d}}{\mathrm{d}i}\mathbb{E}\left[V_M \mathrm{d}\widehat{M} + V_\Omega \mathrm{d}\Omega + \frac{1}{2}V_{\widehat{M}\widehat{M}}\mathrm{d}\langle\widehat{M}\rangle\right] = \frac{1}{\Pi}\frac{\mathrm{d}}{\mathrm{d}i}\left(V_\Omega \mathrm{d}\Omega + \frac{1}{2}V_{\widehat{M}\widehat{M}}\mathrm{d}\langle\widehat{M}\rangle\right).$$

The first equality averages the value of an additional unit of knowledge over the firm's revisions in quality estimates, $\mathrm{d}\widehat{M}$, associated with each possible experimentation outcome. From the firm's perspective $\widehat{M}$ is a martingale (because quality is fixed for a given technology) and these revisions are pure noise and thus have mean zero, the second equality. However, this noise matters for the knowledge channel through its instantaneous variance, $\mathrm{d}\langle\widehat{M}\rangle$.

The knowledge channel is a bundle of two forces that may work in opposite directions depending on the stage of the knowledge cycle. Eq. (7) implies $V_\Omega = -Z/2v'\Pi\Omega^{-1}$ and $V_{MM} = v''\Pi\Omega^{-1}$ and we can rewrite the knowledge channel in Eq. (14) as:

$$c(Z) \equiv \underbrace{\frac{\tau_S}{2}v'(Z)Z}_{\text{knowledge } q} + \underbrace{\frac{\tau_S}{2}v''(Z)}_{\substack{\text{attitude towards}\\ \text{noisy experimentation}}} . \tag{15}$$

Just as q measures the marginal value of capital, the first component measures maginal value of precision, Ω^{-1}, and can be thought of as "knowledge q." In a Gaussian setup the effect of experimentation on precision is known, much as that on capital accumulation. Yet a distinction with capital is that knowledge accumulation is noisy. The second component in Eq. (15), which is central to our argument, captures the firm's attitude towards this noise.

The solution method with experimentation is similar to the benchmark. We conjecture that there exists a region $\mathscr{A}$, outside of which the firm decides to explore, and inside which it experiments; we solve Eq. (9) imposing identical boundary conditions. The only difference is that the maximand in Eq. (10) not only involves q but the knowledge channel, c, too:

$$\psi(q, c) = \max_{i \geq 0} \left\{(c + q)i - \widehat{\gamma}(i)\right\}. \tag{16}$$

This section shows that the knowledge channel affects, to a large extent, investment and its relation to value. From Eq. (16) the first-order condition for optimal investment, i, is:

$$\underbrace{\widehat{\gamma}'(i(Z)) = q(Z)}_{\text{neoclassical trade-off}} + \underbrace{c(Z)}_{\substack{\text{knowledge}\\ \text{channel}}} . \tag{17}$$

The neoclassical trade-off, as present in the benchmark, equalizes the marginal value of capital, q, to marginal costs. When investment also creates knowledge, marginal costs equal the sum of q and the knowledge channel. Hence, q is no longer a sufficient statistic for investment and the firm now looks for a knowledge-contingent investment plan, $i(Z)$.

4.1 Investment and firm's attitude towards noisy experimentation

Our main argument is that when investment creates knowledge it is high at the beginning *and* at the end of the knowledge cycle. Intuitively, early stages of a cycle are associated with high q, the intuition the benchmark delivers. And if, in addition, investment is a means to learn, early stages are associated with high uncertainty and thus high learning benefits. It would seem that investment should be low late in the cycle when both q and uncertainty are low. But this conjecture ignores that when investment creates knowledge the possibility to explore new technologies is a put option on experimentation. Exploration induces firms to gamble on new technologies by investing actively when the end of a cycle is in sight.

The mechanism for our argument operates as follows. Focus first on knowledge q, the first component of the knowledge channel in Eq. (15), and use it to define "total q" as:

$$q_{\text{tot}}(Z) \equiv q(Z) + \frac{\tau_S}{2} v'(Z) Z, \tag{18}$$

that is, traditional q plus knowledge q. Total q has the notable advantage that it can be reasonably constructed from the data (Peters and Taylor, 2017). Figure 2 shows how traditional q (solid line) and total q (dashed red line) differ over the knowledge cycle, separately for good ($Z > 0$, left panel) and bad ($Z < 0$, right panel) technologies. Note that in the benchmark they are identical; with experimentation they may only coincide at the beginning of a knowledge cycle (because $Z \equiv 0$) and at the end of it (due to smooth-pasting, $v' \equiv 0$).

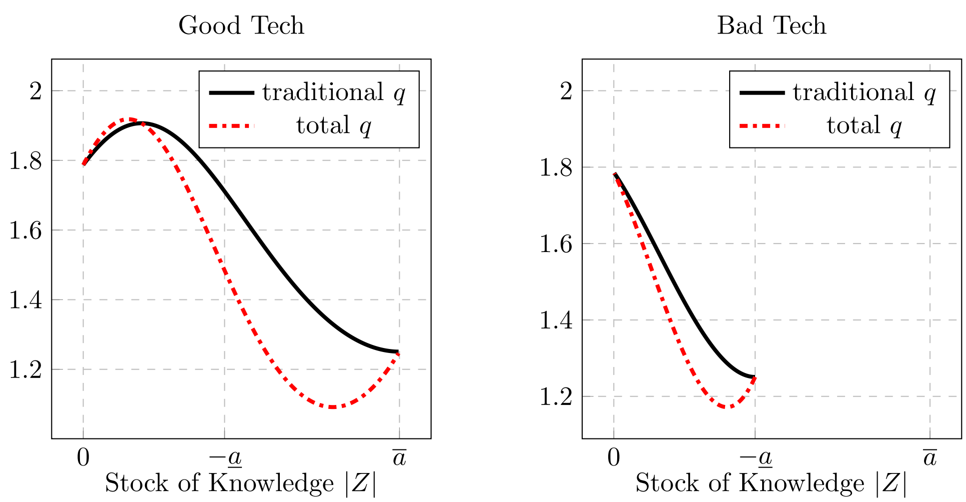

Figure 2: Traditional q and total q. This figure plots traditional q (solid black line) and total q (dashed red line) as functions of knowledge, Z, separately for good (left panel) and bad (right panel) technologies. Parameter values correspond to the baseline calibration, defined in Appendix A.2.

Knowledge q works hand in hand with traditional q (defined earlier), except late in the knowledge cycle. Note first that the knowledge channel does not affect the shape of traditional q (relative to that in Figure 1 in the benchmark): with or without experimentation, q is high early in the cycle and low late in the cycle. Since learning benefits are high early in the cycle and decline thereafter, knowledge q pushes total q above traditional q early in the cycle and causes it to decline faster after, thus reinforcing the neoclassical channel. However, for exploration to be optimal improvement in knowledge on the current technology must be eventually valueless (knowledge q is 0), causing total q to rise late in the cycle. Nevertheless, both total and traditional q may only imply low investment late in the cycle.

This intuition is incomplete as noisy experimentation affects firm value *ex-post*, that is only after the firm has made its investment decision. Thus, the curvature of firm value, as measured by the second component, v'', in Eq. (15), matters for the firm's decision of ramping investment up or down. In particular, Eqs. (17), (15) and (18) together imply that total q does not constitute a sufficient statistic for investment:

$$i(Z) = \delta + \frac{1}{\gamma}\big(q_{\text{tot}}(Z) + \underbrace{\tau_S/2v''(Z)}_{\substack{\text{attitude towards} \\ \text{noisy experimentation}}} - 1\big), \tag{19}$$

as it ignores the firm's attitude towards noisy experimentation, as we argue next.

Through the knowledge channel the firm develops an attitude towards noisy experiments, which we illustrate in Figure 3. Early in the cycle firm value is concave, $v'' < 0$, and experimentation is a risk the firm would prefer to avoid—keeping capital fixed expected value post investment is below its known, pre-investment value. Halfway through the cycle, however, knowledge dissipation accelerates and value becomes convex, $v'' > 0$. The firm develops a gambling attitude towards noisy experiments and ramps up on investment. The possibility of ramping knowledge down through exploration is a put option on experimentation, inducing the firm to gamble on new technologies by raising investment late in the cycle.

To examine how attitude towards noisy experiments affects investment over the cycle, Figure 4 contrasts the relation between investment and knowledge in the benchmark (solid line, a scaled version of Figure 1) to that with experimentation (dashed red line), separately for good (left panel) and bad (right panel) technologies. Focus first on the benchmark (solid black line). Since neoclassical theory works in this case, investment is proportional to q: it is highest following recent exploration, and lowest late in the cycle. Furthermore, when the

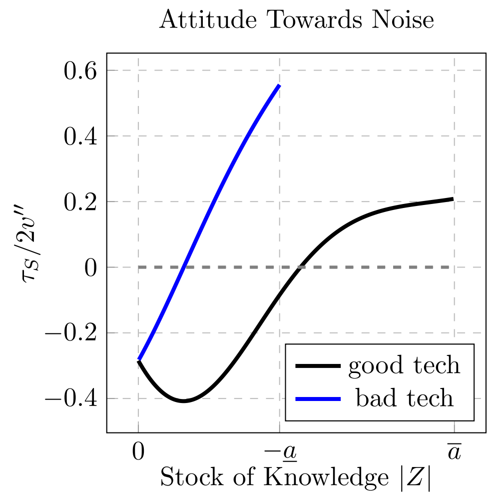

Figure 3: Attitude towards noisy experimentation over the knowledge cycle. This figure plots the second component of the knowledge channel in Eq. (15), $\tau_S/2v''$, as a function of knowledge, separately for good (black line) and bad (blue line) technologies. Parameter values correspond to the baseline calibration, defined in Appendix A.2.

technology is likely good investment rises early in the cycle (q goes up), and then declines.

Activating the knowledge channel (dashed red line) causes investment to rise late in the cycle. Early in the cycle, though total q dictates the shape of investment, the firm's reluctance towards experimentation noise flattens considerably the relation between investment and knowledge. As knowledge accumulates along the cycle the firm changes attitude and starts experimenting actively as a way to gamble on new technologies despite low total q. Figure 4 shows that this gambling attitude dominates late in the cycle and causes investment to rise. This effect is particularly important for bad technologies, and turns the neoclassical relation on its head (investment is decreasing in knowledge in the neoclassical case but increasing in knowledge when investment is a means to experiment); for good technologies instead it implies that investment is high early in—as in the benchmark model—and late in the cycle.

Finally, Figure 5 illustrates how key parameters of the model affect the relation between investment and knowledge. Making experiments more precise (a higher τ_S) may give full force to the firm's attitude towards noisy experiments, exacerbating the late rise in investment and reducing early investment (left panel). Another important parameter of the model is knowledge dissipation. When knowledge is easier to protect (lower ϕ), the investment pattern over the cycle is qualitatively similar, but the overall investment level is substantially higher.

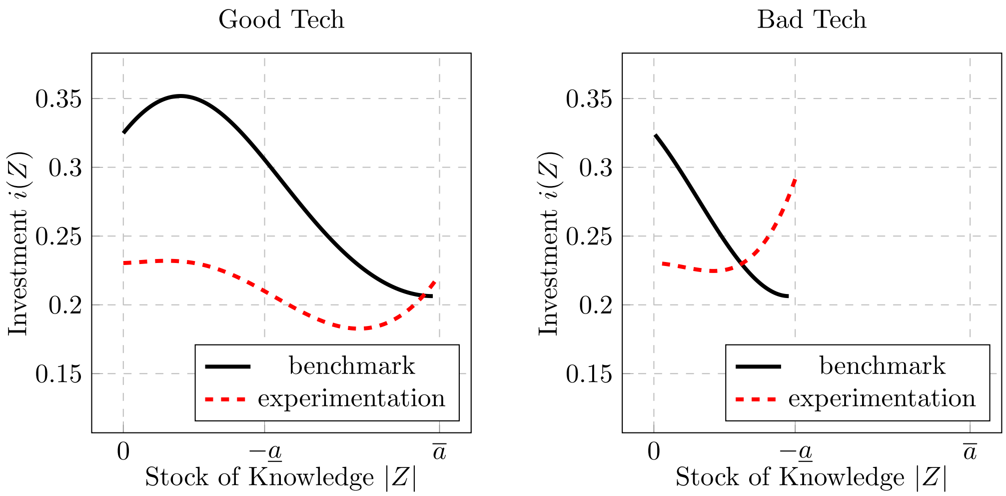

Figure 4: Knowledge-contingent investment with and without experimentation. This figure plots investment as a function of knowledge, Z, in the benchmark model of Section 3 (solid black) and in the model with experimentation of Section 4.4 (red dashed). The panel plots investment under good technologies; the right panel plots investment under bad technologies. Parameter values correspond to the baseline calibration, defined in Appendix A.2.

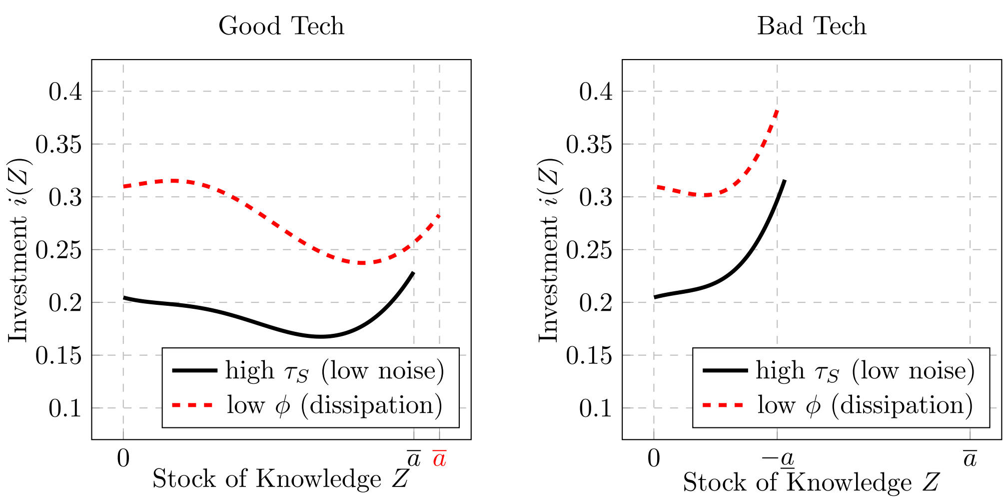

Figure 5: Investment and knowledge frictions. This figure plots investment as a function of knowledge, Z, with experimentation for alternative parameter values, separately for good (left panel) and bad (right panel) technologies. The blue line corresponds to the case of low experimentation noise ($\tau_S = 0.45 > 0.3$). The red line corresponds to the case of low knowledge dissipation ($\phi = 0.3 < 0.35$). Parameter values are defined in Appendix A.2.

4.2 Relation between investment and q over the knowledge cycle

One discriminating implication of the benchmark is that absent experimentation the relation between investment and q is flat along the knowledge cycle (Section 3). But when investment creates knowledge total q is no longer a sufficient statistic for investment and the relation between the two varies along the cycle. Experimentation causes total q and investment to rise late in the cycle but weakens their relation early on. Knowledge resets identify a shift in investment$-q$ sensitivity, with a spike before the reset and a drop upon and following it.

We approach the relation between investment and total q in a way that can be replicated in empirical tests. From the firm's perspective, the relevant statistic for investment is the sum of traditional q and the knowledge channel (or, total q and attitude towards noise, see Eq. (19)). The problem is that, unlike traditional q, the knowledge channel cannot be directly constructed from the data. More precisely, empirical proxies for total q exist (Peters and Taylor, 2017), but measuring attitude towards experimentation seems difficult, e.g., this requires averaging value over all possible experimentation outcomes before the firm invests. We circumvent this difficulty by exploiting the discriminating implication of the benchmark.

The main idea is to use knowledge resets as a way to control for a firm's attitude towards noisy experimentation. We know that absent the firm's attitude towards experiments the relation between investment and total q is flat around knowledge resets. Thus, if the firm's attitude towards noise is to matter empirically the data must reveal variation in this relation around resets. In addition, this variation must follow a predictable pattern, which we describe below. How to measure knowledge resets is itself an empirical challenge, which we tackle in Section 5. For the present section suppose these resets are available to the empiricist. We can then estimate investment$-q$ regressions, which are common in empirical studies, but examining specifically how the relation evolves around knowledge resets.

We simulate daily time series of total q and investment, which are long enough ($8,000$ years) to cover many different knowledge cycles ($\approx 2,460$ cycles). Because we use patent data in empirical tests, and that it is unlikely that bad technologies are patented, we remove all cycles that revealed a technology to be poor, which amount to about half of the sample. Every time the firm explores a new technology ($Z \notin \mathscr{A}$), we flag the exploration date (ν).

To map the data frequency in later tests we aggregate investment and total q annually:

$$I_t = \int_t^{t+1} i_s \mathrm{d}s \quad \text{and} \quad Q_t = \int_t^{t+1} q_{\mathrm{tot},s} \mathrm{d}s.$$

We keep track of whether a knowledge reset occurred over the year using the dummy variable:

$$D_t^{Reset} = \mathbf{1}_{\nu \in [t,t+1)}.$$

To examine how the relation between investment and total q evolves around resets, we track this relation 5 years preceding and following the reset by estimating the specification:

$$I_t = \alpha + \beta Q_t + \sum_{k=-5}^{+5} \gamma_k D_{t+k}^{Reset} Q_t + \sum_{k=-5}^{+5} \delta_k D_{t+k}^{Reset} + \epsilon_t, \tag{20}$$

where $k \in \{-5, ..., 5\}$ denotes the number of years preceding or following the reset year, t. The coefficient β captures the average level of investment$-q$ sensitivity; each coefficient γ_k captures variation in sensitivity relative to this level k years before or after the reset. Figure 6 reports the sensitivity, $\beta + \gamma_k$, in event time k around the reset ($k = 0$).

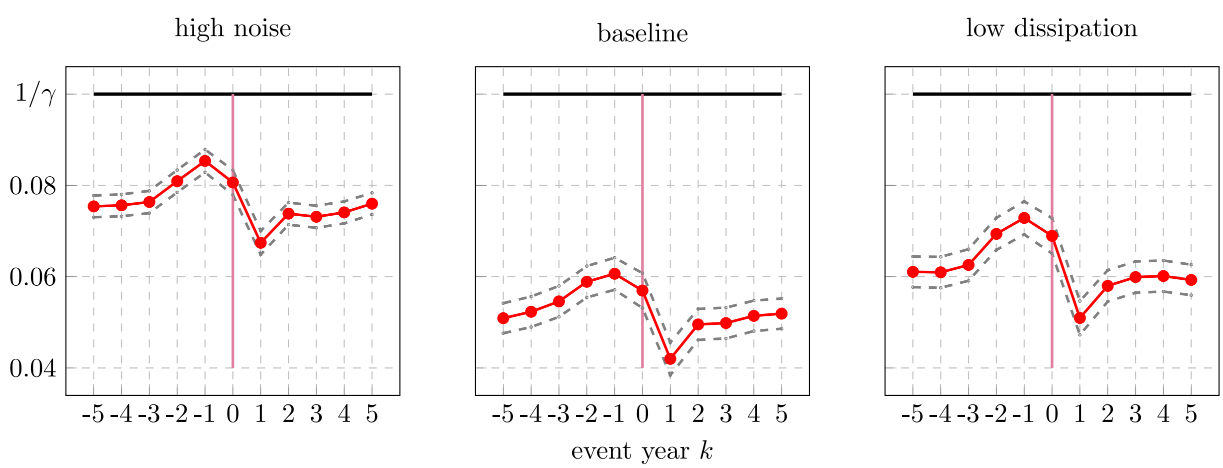

Figure 6: Investment$-q$ sensitivity in event time around a knowledge reset. Each panel plots the sensitivity of investment to total Q within the five years before and after the reset event, $k = 0$. We plot the sensitivity in the benchmark (flat black line) and with experimentation (red line). The middle panel corresponds to the baseline calibration (Appendix A.2), the left panel considers high experimentation noise ($\tau_S = 0.25 < 0.3$), and the right panel considers low knowledge dissipation ($\phi = 0.3 < 0.35$). Dashed gray lines correspond to the confidence intervals at the 95% level.

Recall that benchmark investment$-q$ sensitivity is flat at $1/\gamma$ throughout the cycle (the flat black line in Figure 6). Remarkably, this prediction means that although knowledge resets affect firm value and investment, they do not affect their relation absent experimentation. Note further that this prediction based on total q would carry over to experimentation if the firm's attitude towards noisy experiments did not matter. However, when this attitude matters total q is not a sufficient statistic for investment and the relation between the two must exhibit variation around resets, as is clearly apparent from the red lines in Figure 6.

Variation in the strength of the investment$-q$ relation follows a specific pattern around resets—the relation is weaker early in the cycle and stronger late in the cycle. Figure 3 shows that the firm's attitude towards noise weakens the relation between investment and total q early in the cycle ($v'' < 0$). When uncertainty about the current technology is high, the firm prefers to avoid experimentation risk and invests prudently when in fact total q is highest. Late in the cycle the firm gambles on new technologies by raising investment, when simultaneously knowledge q causes total q to rise. This pattern is robust to alternative values of key model parameters: it is more pronounced either when experiments are less informative or when knowledge is easier to protect (see left and right panels in Figure 6).

These two implications are key in assessing the relevance of the knowledge channel empirically. Of course, the knowledge channel may not be the only determinant of this relation around resets in the data. For instance, resets could be associated with a shift in competition as a new technology may improve a firm's competitive edge and thus affects its investment and value upon reset. However, this possibility can be quickly discarded in the model, as ϕ can also be regarded as capturing competitive pressures; the right panel of Figure 6 shows that competition does not affect the pattern.

4.3 What is the length of a knowledge cycle?

One of the theoretical contributions of this paper is to endogenize knowledge cycles (Definition 2), and one empirical contribution is to measure their length. However, the literature offers little guidance regarding the length of knowledge cycles. For instance, exploration and experimentation typically follow pre-determined schedules (e.g., Berk et al. (2004) or Pastor and Veronesi (2009)). Perhaps surprisingly the model suggests—and the data later confirms—that cycles are short. The model also offers intuitive but empirically useful comparative statics: cycles are shorter when knowledge is more difficult to protect.

Let $\tau_x = \inf\{t > \nu_\cdot : Z_t = x\}$ be the first time knowledge hits level x after switching to a new technology. The probability that the next exploration time for good technologies occurs before or at time $t \geq 0$ immediately following exploration (i.e., $Z_0 = 0$) is:

$$\mathbb{P}\left(\tau_{\overline{a}} \leq t | Z_0 = 0, \tau_{\overline{a}} < \tau_{\underline{a}}\right) \equiv \underbrace{\mathbb{E}\left(\mathbf{1}_{\tau_{\overline{a}} \leq t}\mathbf{1}_{\tau_{\overline{a}} < \tau_{\underline{a}}} | Z_0 = 0\right)}_{\equiv G(t;0)} / \underbrace{\mathbb{E}\left(\mathbf{1}_{\tau_{\overline{a}} < \tau_{\underline{a}}} | Z_0 = 0\right)}_{\equiv g(0)}.$$

For an arbitrary starting point $Z \in \mathscr{A}$ within the experimentation region, standard arguments show that $G(t, Z)$ solves the Kolmogorov Backward Equation (KBE):

$$G_t = \frac{1}{2}(\tau_A + \tau_S i(Z))(G_Z Z + G_{ZZ}),$$

with initial condition $G(0, Z) = 0$ and boundary conditions $G(t, \underline{a}) = 0$ and $G(t, \overline{a}) = 1$.[9] The distribution of a cycle length is $\frac{\mathrm{d}}{\mathrm{d}t}G(t; 0)/g(0)$ and is plotted in Figure 7 with and without experimentation (left panel) and with experimentation only but when knowledge is more difficult to protect (right panel).

All curves in Figure 7 suggest that knowledge cycles are short, with a median of approximately 6 years. Casual intuition may suggest otherwise. For instance, the cycles associated with *Intel*'s knowledge about RAM chips or central processor technology lasted around ten years. Thus, a median of 6 years seems substantially "shorter." However, we will show empirically in Section 5 that knowledge cycles are consistently short in the data.

Experimentation leads to shorter knowledge cycles (the left panel of Figure 7). Note that we have chosen parameter values in the benchmark model so that the benchmark exploration thresholds coincide with those under experimentation. Hence, the shortening of the cycle is a unique consequence of the learning benefits of active experimentation. Because experimentation allows the firm to learn faster about the quality of its technology, it shortens the average and median length of the cycle. Considering the substantial effect of experimentation on investment relative to neoclassical predictions, however, this effect is relatively small; this is because experimentation significantly reduces the initial level of investment due to the firm's reluctance to take experimentation risk early in the cycle.

[9]Similarly, $g(Z)$ solves $0 = (\tau_A + \tau_S i(Z))(g_Z Z + g_{ZZ})$, with boundary conditions $g(\overline{a}) = 1$ and $g(\underline{a}) = 0$.

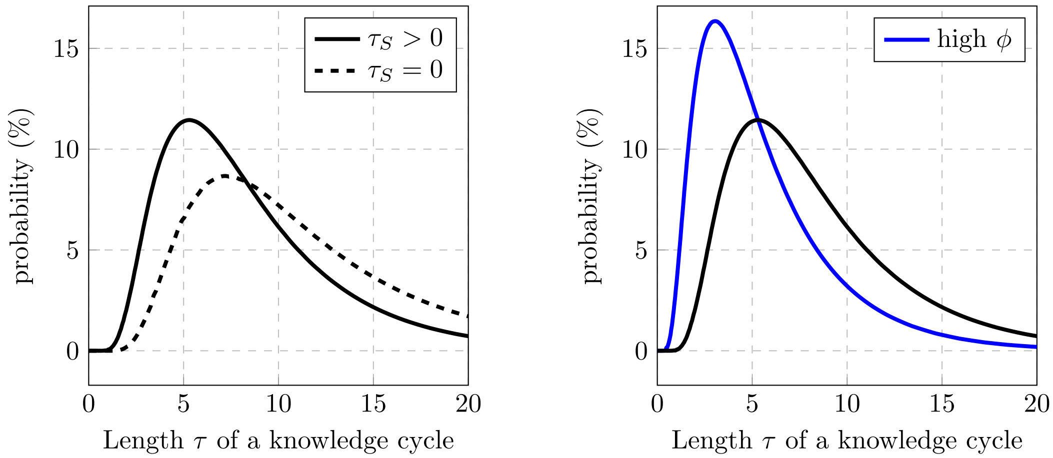

Figure 7: Distribution of the length of a knowledge cycle for good technologies. The left-hand panel plots the distribution of the length of the cycle with (solid line) and without (dashed line) experimentation. The right-hand panel does comparative statics, with the black curve denoting the case with experimentation, and the blue curve corresponding to the case with experimentation and high knowledge dissipation ($\phi = 0.8 > 0.35$). Parameters are reported in Appendix A.2

It is intuitive that knowledge dissipation, for instance in the form of competitive pressures, should reduce the length of knowledge cycles. In the case of RAM chips, competition from Japan precipitated to a large extent the end of *Intel*'s knowledge cycle. Alternatively, when property rights are harder to protect the firm's incentive to explore a new technology are higher. The right panel of Figure 7 shows that high knowledge dissipation (the blue line) reduces the length of the knowledge cycle, a prediction that we confirm later in the data.

4.4 Model extension: partially transferable knowledge

We now allow knowledge to be partially transferable across technologies, as opposed to being reset completely upon exploration. The paper's main insights carry over to this extension. It is likely that knowledge acquired about a technology is in fact partially transferable to another. Thus, we relax our baseline assumption that knowledge is completely specific to a technology. In the spirit of Jovanovic and Nyarko (1996), we assume that upon exploring a new technology (which occurs at time ν) knowledge is reset to:

$$Z_\nu = \lambda \cdot Z_{\nu-},$$

where $\lambda \in [0, 1)$. When $\lambda \equiv 0$, knowledge is specific to a technology and we recover the model of Section 2. When $\lambda > 0$, knowledge is partially transferable across technologies. The case $\lambda \equiv 1$ in which knowledge is freely transferable across technologies does not lead to a stationary solution and must be discarded.

The solution method is identical to the baseline, except for the boundary conditions that apply when exploration is optimal. Value-matching conditions are now given by:

$$\lim_{Z \downarrow \underline{a}} v(Z) = (1-\omega)^{\alpha(1-\eta)} v(\lambda \cdot \underline{a}) \quad \text{and} \quad \lim_{Z \uparrow \overline{a}} v(Z) = (1-\omega)^{\alpha(1-\eta)} v(\lambda \cdot \overline{a}).$$

Smooth-pasting conditions are modified similarly.[10] Figure 8 compares investment under partially transferable knowledge ($\lambda > 0$, dashed lines) to that when knowledge is technology-specific ($\lambda = 0$, solid lines) separately for good and bad technologies.

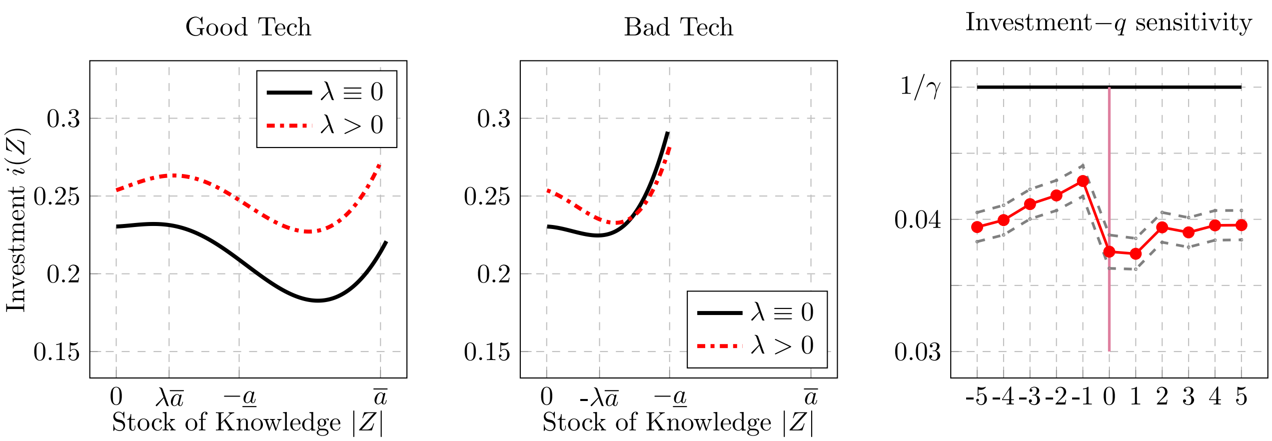

Figure 8: Investment and its sensitivity to total q when knowledge is partially transferable across technologies. The left-hand and center panels plot investment as a function of knowledge, Z, with partially transferable ($\lambda = 0.2$, dashed lines) and with technology-specific ($\lambda = 0$, solid lines) knowledge for good (left panel) and bad (center panel) technologies. The right-hand panel plots investment-q sensitivity according to the simulation procedure of Section 4.2 5 years before and after a reset event, $k = 0$, but with $\lambda = 0.2$. Parameter values are defined in Appendix A.2.

The main insights of the model remain, although the asymmetry between good and bad technologies becomes stronger. When the firm uncovers a good technology, it is now more likely to uncover another good one in the next cycle. Thus, the firm has greater incentives to invest when it is able to partially reuse its existing knowledge. The same argument implies

[10] $\lim_{Z \downarrow \underline{a}} v'(Z) = \lambda \cdot (1-\omega)^{\alpha(1-\eta)} v'(\lambda \cdot \underline{a})$ and $\lim_{Z \downarrow \overline{a}} v'(Z) = \lambda \cdot (1-\omega)^{\alpha(1-\eta)} v'(\lambda \cdot \overline{a})$.

that investment is lower for bad technologies relative to its level under technology-specific knowledge, except early in the cycle when partial transferability of knowledge raises total q. Around knowledge resets, the relation between knowledge and investment and the evolution of investment$-q$ sensitivity (right panel) remain qualitatively unchanged.

5 Empirical relevance of the knowledge channel

We develop a measure of knowledge resets to assess the empirical relevance of the knowledge channel. The main idea is to use patents' citations to capture firms' technological "knowledge base" and identify resets as major changes in this base. We find substantial variation in the investment$-Q$ around these resets, suggesting the knowledge channel is at play in the data.

5.1 Measuring knowledge resets

We obtain patent data from the United States Patent and Trademark Office (USPTO) PatentView platform. We consider all utility patents applications granted between 1976 and 2017 and link assignees to public firms in Compustat using the bridge file provided by NBER until 2006, and that extended by Stoffman, Woeppel, and Yavuz (2021) for later years. Central to our measurement of knowledge resets, the data contains all the citations each patent makes to prior patents, that we consider as the prior art that each patent is building upon. We only consider firms that were granted at least one patent over the 1976-2017 period, and require available data on total investment and total Q, as defined by Peters and Taylor (2017). Firms enter the sample when they apply for their first patent over that period and leave three years after their last patent grant. The sample comprises 1,442,813 distinct patent applications granted to 5,428 firms, or 54,250 firm-year observations. On average firms are granted 26 new patents per year (the median is 2), and each patent makes on average 15 citations to prior patents (with a median of 7).

We define firm f technological "knowledge base" at time t as the set of all the patents cited by the patents applied for by f between $t-5$ and t.[11] Following Ma (2021), this set captures the underlying technological knowledge that f has accumulated between $t-5$ and t.

[11] We consider patent application year instead of grant year to better reflect the true timing of firms' innovation. While we consider a period of 5 years to define the knowledge base, we obtain similar results if we consider longer periods of 10 or 15 years.

On average, firms' knowledge base is composed of 148 (median of 11) distinct patents cited from 133 patents (median of 11). To track the evolution of f's knowledge base, we aggregate the cited patents across the 38 distinct NBER patent technology sub-classes. Specifically, we define the vector v_t^f where each of its 38 elements represents the share of f's knowledge base in year t corresponding to each class. Intuitively, v_t^f represents the distribution of knowledge accumulated by f in year t across distinct technology areas. The innovation of some firms builds on patents from a few technology areas whereas other firms have a more dispersed knowledge base.

Then, we measure annual changes in f's' knowledge base from the cosine similarity between v_t^f and v_{t-1}^f according to:

$$\Delta v_t^f = \frac{v_t^f}{\|v_t^f\|} \cdot \frac{v_{t-1}^f}{\|v_{t-1}^f\|},$$

where the operator "$\cdot$" denotes the scalar product, and "$\|v\|$" denotes vector v's Euclidean norm. By definition, Δv_t^f is bounded in $[0, 1]$. Higher values of Δv_t^f reflect greater similarity between v_t^f and v_{t-1}^f, and indicate stability in the knowledge base. In contrast, lower values of Δv_t^f indicate instability. Knowledge base is instable when the technology developed by f in year t builds on patents from different technological classes compared to what was used previously. Hence, $\Delta v_t^f = 0$ indicates that new technology in year t builds on technologies that were not used until year $t-1$. Table 1 (Panel A) shows that, perhaps unsurprisingly, firms' knowledge base is remarkably stable since the average value of Δv_t^f is 0.96 and the median is 0.99. The sample standard deviation of Δv_t^f is 0.10, indicating that firms' innovations typically build on a stable distribution of technologies.

To identify knowledge "resets," defined as times when firms explore a new technology and Z declines sharply, we flag years in which Δv_t^f is abnormally low, using the following binary variable(s):

$$Reset_t^f(\theta) = \begin{cases} 1 & \text{if } \Delta v_t^f < \overline{\Delta v^f} - \theta \times \sigma(\Delta v^f) \\ 0 & \text{otherwise,} \end{cases}$$

where $\overline{\Delta v^f}$ and $\sigma(\Delta v^f)$ are the time-series average and standard deviation of Δv_t for firm f, and $\theta \in (1, 1.5, 2, 2.5)$. A reset occurs when Δv_t^f is θ standard deviations below its average.

Variable	mean	sd	p25	p50	p75	N
			Panel A: firm-year variables			
Δv	0.964	0.107	0.982	0.997	0.999	54,250
Reset(1.0)	0.108	0.310	0	0	0	54,250
Reset(1.5)	0.081	0.273	0	0	0	54,250
Reset(2.0)	0.058	0.234	0	0	0	54,250
Reset(2.5)	0.039	0.193	0	0	0	54,250
Δv|Reset(1.0)	0.779	0.229	0.678	0.859	0.961	5,827
Δv|Reset(1.5)	0.745	0.245	0.596	0.826	0.950	4,353
Δv|Reset(2.0)	0.708	0.258	0.535	0.780	0.930	3,107
Δv|Reset(2.5)	0.682	0.264	0.500	0.751	0.911	2,076
Total Q	1.194	1.949	0.223	0.615	1.345	54,250
Total Investment	0.236	0.185	0.124	0.183	0.278	54,250
Physical investment	0.066	0.080	0.022	0.042	0.079	54,250
Intangible Investment	0.169	0.143	0.076	0.130	0.213	54,250
			Panel B: Cycle variables			
$\tau_{data}(1.0)$	6.404	5.348	3	5	8	1,885
$\tau_{data}(1.5)$	7.526	6.429	3	6	9	1,095
$\tau_{data}(2.0)$	8.387	7.273	3	6	11	568
$\tau_{data}(2.5)$	9.674	8.298	3	7	13	227

Table 1: Summary Statistics. This table presents the summary statistics of all variables used in the empirical analysis. Panel A reports statistics for firm-year variables. Panel B reports statistics for cycle variables. The sample period is from 1976 to 2017.

Table 1 indicates that, across all firms and years, knowledge resets are infrequent as their occurence, the average of $Reset(\theta)$, ranges between 3.8% ($\theta = 2.5$) and 10.8% ($\theta = 1$). Although infrequent, the change in knowledge base during resets is sizable, as Δv_t^f ranges between 0.67 and 0.77, implying roughly a 30% change in the knowledge base.

To build intuition, we detail the knowledge resets of two firms in our sample.[12] The first is the business-to-business software company CA Technologies that experienced a knowlegde reset in 1997, and the second is the distributor of motor fuels Sunoco that experienced a reset in 1985. Table 2 displays the elements (larger than 5% for brevity) of both firms' vectors v for years t (the reset year), $t-1$, and $t-2$. Panel A reveals that CA's reset (Δv dropping from 0.99 in 1996 to 0.39 in 1997) is driven by an abrupt increase in its share of citations of patents in "computer hardware and software", "communication", and "peripherals", and "information storage", and a simulateneous decrease in citations from

[12]Selected among the 25 largest firms with a least 10 patents forming their knowledge base and experiencing large resets during the sample period ($\theta = 2.5$)

NBER Sub-category	v_t	v_{t-1}	v_{t-2}
Panel A: CA Technologies (reset in t=1997)			
Agriculture, Food, Textiles	0.01	0.04	0.05
Coating	0.02	0.12	0.07
Miscellaneous	0.02	0.08	0.07
Communications	**0.05**	**0.00**	**0.00**
Computer Hardware & Software	**0.43**	**0.10**	**0.11**
Computer Peripherials	**0.05**	**0.00**	**0.00**
Information Storage	**0.14**	**0.00**	**0.00**
Electronic business methods and software	0.09	0.00	0.00
Mat. Proc & Handling	0.11	0.53	0.55
Δv	0.39	0.99	
Panel B: Sunoco (reset in t=1985)			
Gas	0.05	0.09	0.08
Organic Compounds	**0.28**	**0.03**	**0.01**
Resins	0.02	0.01	0.06
Miscellaneous	0.38	0.53	0.48
Drugs	0.07	0.04	0.03
Measuring & Testing	0.00	0.06	0.10
Mat. Proc & Handling	0.05	0.01	0.01
Heating	0.07	0.13	0.11
Δv	0.73	0.98	

Table 2: Examples of knowledge resets.

"material processing and handling", "coating", and "agriculture, food, and textiles". These changes reflect the company's strategy in the mid nineties to improve compatibility with products from other vendors, e.g., Hewlett-Packard or Apple Computer.[13] Panel B indicates that the knowledge reset of Sunoco (Δv dropping from 0.98 in 1984 to 0.73 in 1985) stems from its increased reliance on technologies related to "organic compounds", and a decreased reliance on technologies related to "measuring and testing", and "heating", coinciding with its launch of the Sunoco ULTRA 94 in 1983, the market's highest octane unleaded gasoline.[14]

5.2 What is the length of knowledge cycles in the data?

We define the length of a cycle as the number of years between two resets ($\tau_{data(\theta)}$) for firms that have completed at least one cycle (two observable resets). This restricted sample

[13]See CA Technologies' history at: https://en.wikipedia.org/wiki/CA_Technologies.
[14]See Sunoco's history at: https://en.wikipedia.org/wiki/Sunoco.

comprises between 227 ($\theta = 2.5$) and 1,885 ($\theta = 1$) cycles for 190 and 1,242 distinct firms, respectively. Figure 9 displays the distributions of τ_{data} for the four values of the threshold θ. Mirroring the distribution in the model (see Figure 7), the empirical distributions are asymmetric and right skewed. Panel B of Table 1 indicates that the average cycle length varies between 6.40 and 9.67 years, and the median is between 5 and 7 years. Figure 10 reveals notable heterogeneity across broad (one-digit SIC) industries and confirms the prevalence of skewed cycles (averages are larger than medians). The longest cycles appear in mining with an average close to 10 years, and the shortest in retail with less than 6 years. Although existing research provides no point of comparison that we are aware of, knowledge cycles appear rather short in our sample.

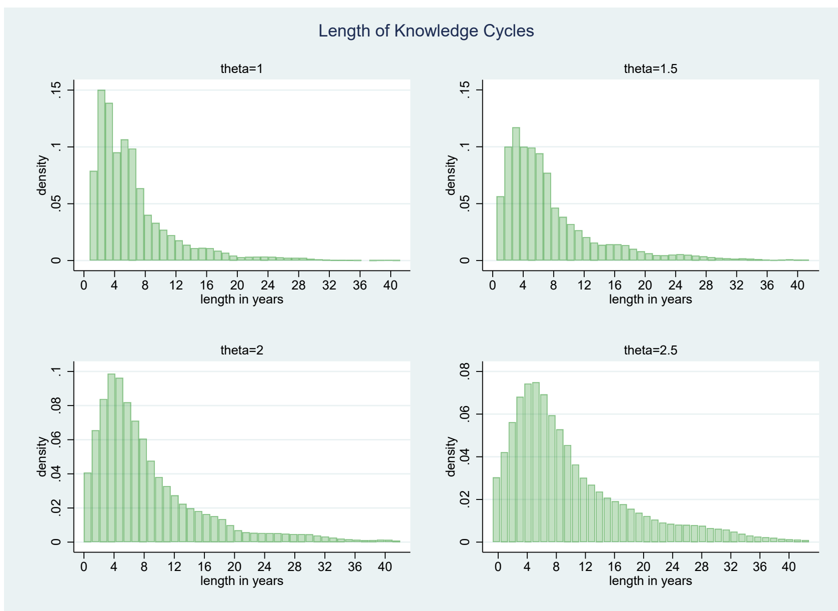

Figure 9: Empirical length of knowledge cycles. This figure plots the distribution of the length of knowledge cycles, where a cycle is given by the number of years between two knowledge resets. We consider resets based on four distinct values of θ.

Figure 11 displays length distributions separately for firms operating in competitive and concentrated industries, measured using the Herfindhal index (HHI) computed at the 3-digit SIC level below and above the sample median based on firms in Compustat. We use HHI as a proxy for the intensity of knowledge dissipation (ϕ), assuming lower dissipation in more concentrated industries (high HHI). In line with the model's prediction, Figure 11 reveals

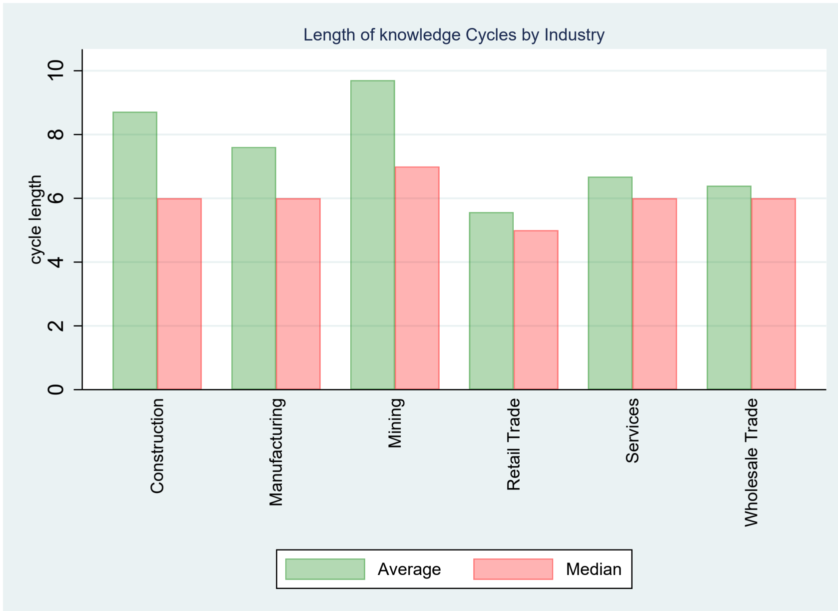

Figure 10: Empirical length of knowledge cycles across industries. This figure plots the average and median length of knowledge cycles across one-digit SIC industries, where a cycle is given by the number of years between two knowledge resets. We consider resets based on $\theta = 1.5$.

longer cycles for firms facing lower knowledge dissipation, with an average length of 8.37 years in concentrated industries compared to 6.69 years in competitive industries (the difference is statistically significant).[15] By highlighting that the intensity of knowledge dissipation is linked to the length of knowledge cycles, Figure 11 reinforces the economic plausibility of our measurement of knowledge resets and cycles.

5.3 Investment and total q around knowledge resets

We examine the relation between investment and total q around knowledge resets (see Figure 6) using "total" investment from Peters and Taylor (2017) to capture the optimal investment rate $i(Z)$. Similar to $i(Z)$, total investment is the sum of capital (capex) and intangible (R&D+$0.3\times$SG&A) expenditures divided by (lagged) total capital. Total capital is the sum of property, plant, and equipment, intangible assets on the balance sheet, research and organizational capital (accumulated R&D and SG&A spending). We use Peters and Taylor

[15] The model also predicts that the distribution of knowledge cycles should depend on the noise association with experimentation, τ_S. Unfortunately, we could not find a plausible empirical proxy for this parameter.

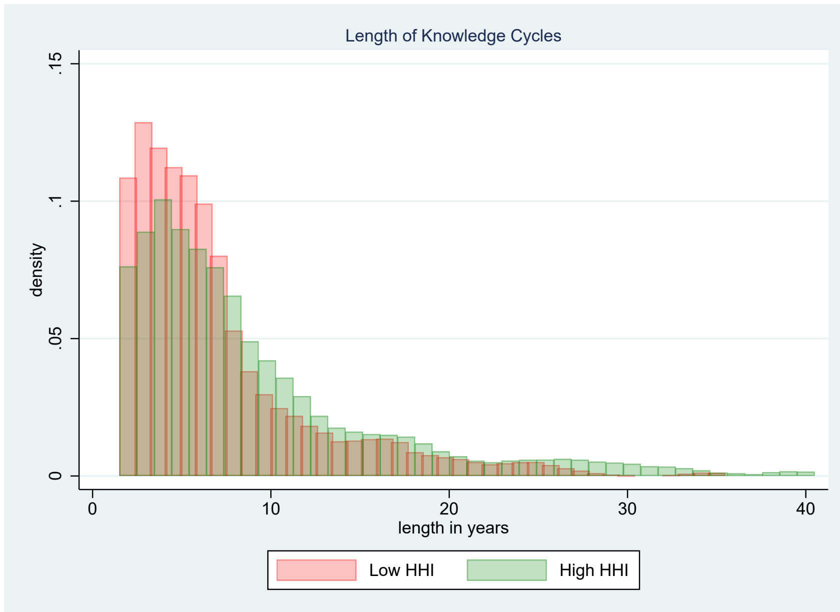

Figure 11: Empirical length of knowledge cycles by HHI. This figure plots the distribution of the length of knowledge cycles across firms in concentrated and competitive industries. Firms are assigned to groups based on whether the HHI of their industry is above or below the sample median. A cycle is given by the number of years between two knowledge resets. We consider resets based on $\theta = 1.5$.

(2017)'s total Q as a proxy for the model's total q. Total Q is defined as firm value divided by total capital, which we view as a reasonable proxy for the sum of the model's traditional and knowledge q.[16]

Table 1 presents summary statistics. Average total Q is 1.280, with a median of 0.636. Total investment is 0.235 on average, with a median of 0.183. The average (median) physical investment is 0.065 (0.042) and the average (median) intangible investment is 0.168 (0.130). Overall, these statistics are close to those reported by Peters and Taylor (2017) on a larger sample that also includes non-patenting firms (see their Table 1). Firms in our sample invest more in intangible than in physical assets, which is expected as we focus on firms with patents.

To track the relation between total investment and total Q around resets we estimate

[16]In untabulated tests, we have used simulated data to compute the correlation between the model's total q and the variable $\frac{vK}{K+Z}$, which in the model corresponds to Peters and Taylor (2017)'s definition of total Q up to a constant. The corresponding correlation is 0.96 in the baseline calibration, and the evolution of investment$-q$ sensitivity is virtually unchanged.

the empirical counterpart to Eq. (20) in simulations:

$$I_{f,t} = \alpha_f + \eta_t + \beta Q_{f,t-1} + \sum_{k=-5}^{+5} \pi_k Q_{f,t-1} \times D_{t+k}^{Reset} + \sum_{k=-5}^{+5} \delta_k D_{t+k}^{Reset} + \varepsilon_{f,t}, \tag{21}$$

where $I_{f,t}$ is total investment of firm f in year t, and $Q_{f,t-1}$ is its lagged total Q. To closely map the model, we include firm fixed effects α_f to isolate the within-firm variation of investment and Q over resets. In addition, year fixed effects η_t absorb common variations across firms (e.g., investment or valuation booms or busts common to all firms in specific years). We cluster the standard errors at the firm level. The sample includes firms that experience, or not, a reset at some point. The indicators D_{t+k}^{Reset} identify years around resets, from 5 years before to 5 years after.[17] Hence, the coefficients π_k's on the interactions between $Q_{f,t-1}$ and D_{t+k}^{Reset} track the within-firm evolution of the investment$-Q$ sensitivity around resets. Because β captures average sensitivity outside periods surrounding resets or for control firms not experiencing resets, investment-Q sensitivity around resets is given by $\hat{\beta} + \hat{\pi_k}$ for $k \in [-5, +5]$, which Figure 12 plots (together with their 95% confidence bounds) for resets corresponding to $\theta = 1.5$.

Figure 12 reveals that the relation between total investment and total Q is *not* flat around resets, as predicted in the absence of experimentation. In contrast and consistent with the key implication of the knowledge channel, it closely mirrors the relation obtained in the simulation (see Figure 6). Resets mark a clear break in this relation. Investment-Q sensitivity increases and spikes before the reset (at $k = -1$), and drops following the reset (for $k \geq 0$), suggesting that the economic force underlying the knowledge channel—firms' attitude towards noisy experiments—is at play in the data. The evolution of the investment-Q sensitivity is both economically and statistically significant. It rises to 0.05 just before resets, and reaches to 0.03 five years into the new cycle, implying a drop of 40% that is significant at 1%. Reassuringly, the magnitude of the investment-Q sensitivities is in line with Peters and Taylor (2017). In their sample, the sensitivity is 0.049 (see their Table 2), whereas it is 0.042 in ours.

Figure 13 indicates similar patterns for physical and intangible investment. This finding suggests that investment in capital and knowledge are entangled, and should be modeled

[17] Because firms could have several resets within short periods, this specification accounts for the fact that some years could be simultaneously before and after a reset.

Figure 12: Investment-Q sensitivity around knowledge resets. This figure plots the sensitivity of investment to Q ($\hat{\beta} + \hat{\gamma}_k$ for $k \in [-5, +5]$) in event-time around resets, with resets corresponding to $\theta = 1.5$.

together as we do, and not as separate decisions. Suppose investment in capital and in knowledge were distinct decisions (e.g., capex vs. R&D), so that investment in capital would not create knowledge. In our model investment in capital would be declining over the cycle, whereas investment in knowledge would be U-shaped, much as total investment. However, Figure 13 indicates that this is not the case in the data.

5.4 Are alternative channels at play?

The timing of knowledge resets could coincide with changes in firms' environment, unrelated to the knowledge channel, but related to their investment and Q. We consider two possibilities, one mechanical and one economical. First, our measure of resets could be affected by changes in acquisition or innovation strategies related to investment and Q. For instance, resets may arise following the purchase of firms with distinct patented technologies, or large discontinuities in patenting (e.g., the rising importance secrecy). We track investment-Q sensitivity around resets seperately for firms that (i) were acquirers in the three years preceeding resets (or not), or (ii) experienced increase (or decrease) in the size of their patent

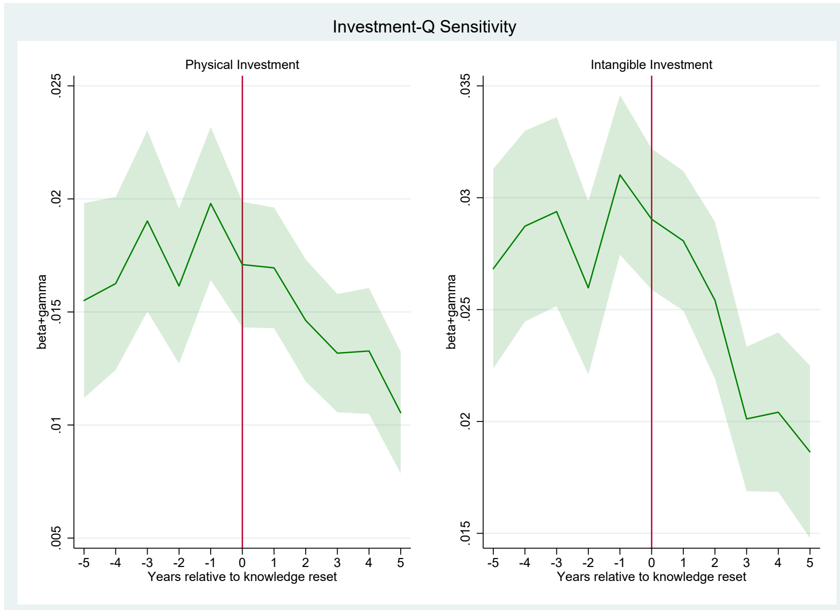

Figure 13: Investment-Q sensitivity around knowledge resets. This figure plots the sensitivity of physical (left panel) and intangible (right panel) investment to Q ($\hat{\beta} + \hat{\gamma}_k$ for $k \in [-5, +5]$) in event-time around resets, with resets corresponding to $\theta = 1.5$.

portfolios at the time of resets.[18] Figure 14 shows that the dynamics of investment-Q sensitivity around resets ($\hat{\beta} + \hat{\pi}_k$ for $\theta = 1.5$) are similar for both groups, labeled as "Low" and "High". Thus our conclusions are unlikely due to changes in firms' acquisition or patenting activities around resets.

Second, knowledge resets may correlate with firms' characteristics typically associated with the relation between investment and Q. Following Gutierrez and Philippon (2017) we consider three types of characteristics : (i) financial frictions, (ii) competition, and (iii) institutional ownership. We capture these characterisitcs using the ratio of cash and marketable securities to assets, industries' Herfindhal index computed at the 3-digit SIC level from Compustat, and the fraction of shares held by institutional investors. For each variable, we assign firms experiencing resets into two groups based on whether they are below or above the median taken across all resets, and track the investment-Q sensitivity seperately across groups. Figure 15 reveals no significant difference across groups, suggesting that the dynamics of investment-Q sensitivity around resets is unrelated to these channels.

[18] Specifically, we re-estimate specification (21) but interact all variables with two binary variables identifying the two groups in which firms are assigned at the time of resets.

Figure 14: Investment-Q sensitivity around knowledge resets: robustness I. This figure plots the sensitivity of investment to Q ($\hat{\beta} + \hat{\gamma}_k$ for $k \in [-5, +5]$) in event-time around resets across groups. Firms are assigned to groups based on whether they acquired other firms in the last three years (left panel), or whether the change in the size of their patent portfolio is above of below the sample median. Resets correspond to $\theta = 1.5$.

As an alternative robustness exercise, we include all five variables (whether a firm is an acquiror, the number of patent in its portfolio, cash, HHI and institutional ownership) in Eq. (21) as controls, both in levels and interacted with the indicators D_k^{Reset}. The lower-right panel of Figure 15 indicates that the inclusion of these variables does not affect the dynamics of the investment-Q sensitivity around resets (despite significant coefficients for some of them, not displayed for brevity).

We subject our results to other checks unreported for brevity. The dynamics of investment-Q sensitivity holds if we (i) focus only on firms experiencing resets at some point, (ii) consider other values for the threshold θ, (iii) eliminate years before 1980 to limit the potential truncation in reset measurement, (iv) include the interaction between year and industry (3-digit SIC) fixed effects to absorb time-varying industry shocks (e.g., unobserved common technological changes), and (v) add firms' size (log of assets) and cash-flows as additional controls, mimicking typical investment specifications (e.g., Peters and Taylor (2017)). The post-reset decline in investment-Q sensitivity also remains in a difference-in-differences specification ag-

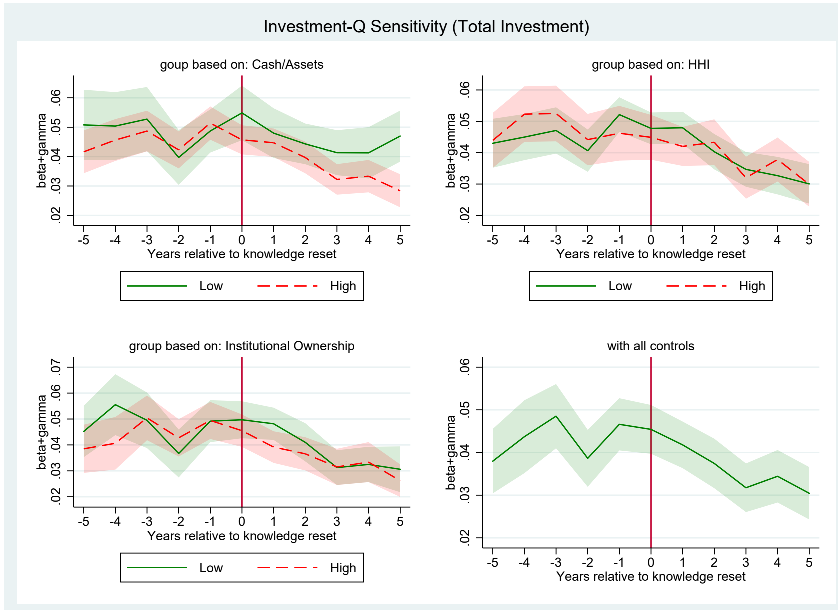

Figure 15: Investment-Q sensitivity around knowledge resets: robustness II. This figure plots the sensitivity of investment to Q ($\hat{\beta} + \hat{\gamma}_k$ for $k \in [-5, +5]$) in event-time around resets across groups. Firms are assigned to groups based on whether their cash-to-asset ratio (upper-left panel), industry HHI (upper-right panel), or institutional ownership (bottom-left panel) is above of below the sample median. In the bottom-right panel, we display the sensitivity of investment to Q after controlling for whether a firm is an acquiror, the number of patent in its portfolio, cash, HHI and institutional ownership). Resets correspond to $\theta = 1.5$.

gregating all D_k^{Reset} indicators into one capturing the post period $k > 0$ for firms experiencing resets, accounting for EIV in Q using the estimator of Erickson and Whited (2000).

6 Concluding Remarks

The relation between investment and Q drops when firms explore new technologies and reset their technological knowledge. This pattern is novel and consistent with a model in which investment is a means to accumulate capital, as is customary in the literature, but also a means to create knowledge. Investment as a "knowledge channel" is high early in the cycle (after adopting a new technology) but also late in the cycle (just before exploring new technologies). The relation between investment and value is weaker early in the cycle and stronger late in the cycle. The model generates endogeneous knowledge cycles that are short, and expand when knowledge is easier to protect. We find empirical support for these predictions, confirming that knowledge cycles are an important determinant of investment.

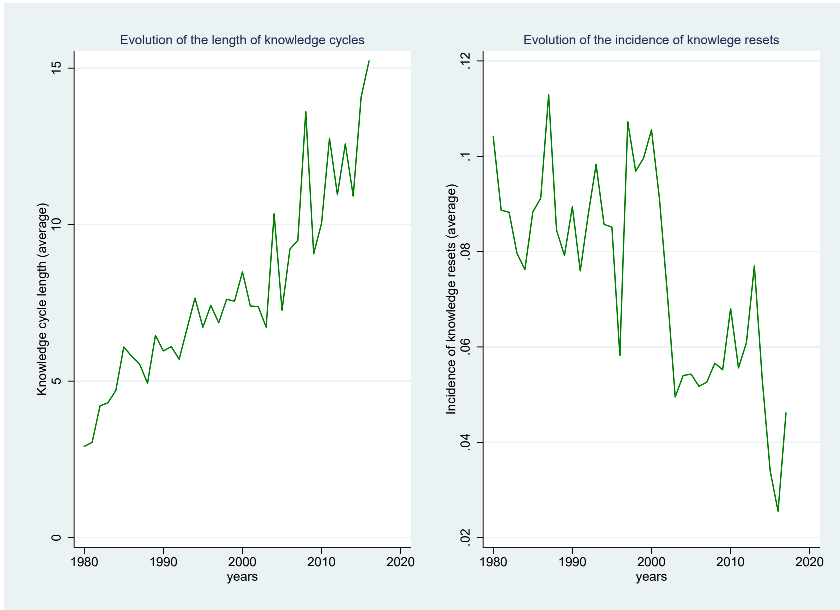

Figure 16: Evolution of cycles length and reset incidence. This figure plots the average length of knowledge cycles by year (left panel) and the average incidence of knowledge resets by year. Resets correspond to $\theta = 1.5$.

Knowledge resets have become less prevalent over time, as illustrated in Figure 16, resulting in cycles whose length almost tripled between 1980 and 2017.[19] The lengthening of cycles (which we believe we are first to document) could have various origins, for instance, the increased difficulty to generate new ideas (Bloom et al., 2017) or declining product market competition (Gutierrez and Philippon, 2018). Since investment$-Q$ sensitivity is linked to knowledge cycles, it is interesting to contrast its evolution to that of knowledge cycles. Gutierrez and Philippon (2017) and Alexander and Eberly (2018) document a decline in investment in the last twenty years, despite no decline in Q, which we confirm in our sample. Because the relation between investment and Q is strongest around resets (see Figure 12), its recent weakening could be connected to less frequent resets and the underlying changes in the process of knowledge creation within firms. We leave these questions for future research.

[19]We obtain similar evolutions if we consider other values for θ, and if we control for changes in the composition of the sample by removing firm fixed effects.

References

Abel, A. and O. Blanchard (1986). The present value of profits and cyclical movements in investment. *Econometrica 54*(2), 249–273.

Abel, A. and J. Eberly (1994). A unified model of investment under uncertainty. *American Economic Review 84*(5), 1369–84.

Abel, A. and J. Eberly (2011). How q and cash flow affect investment without frictions: An analytic explanation. *Review of Economic Studies 78*(4), 1179–1200.

Acemoglu, D. (2009). *Introduction to modern economic growth*. Princeton University Press.

Aghion, P., P. Bolton, C. Harris, and B. Jullien (1991). Optimal learning by experimentation. *The Review of Economic Studies 58*(4), 621–654.

Alexander, L. and J. Eberly (2018). Investment hollowing out. *IMF Economic Review* (66), 5–30.

Andrei, D., W. Mann, and N. Moyen (2018). Why did the q theory of investment start working? *The Journal of Financial Economics*.

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

Benkard, L. (2000, September). Learning and forgetting: The dynamics of aircraft production. *American Economic Review 90*(4), 1034–1054.

Bergemann, D. and U. Hege (2005). The financing of innovation: Learning and stopping. *The RAND Journal of Economics 36*(4), 719–752.

Berk, J., R. Green, and V. Naik (2004, 01). Valuation and return dynamics of new ventures. *The Review of Financial Studies 17*(1), 1–35.

Bloom, N., C. Jones, J. van Reenen, and M. Webb (2017). Are ideas getting harder to get? NBER Working Paper 23782.

Bloom, N., R. Sadun, and J. van Reenen (2016). Management as a technology? *Working Paper*.

Braguinsky, S., A. Ohyama, T. Okazaki, and C. Syverson (2021). Product innovation, product diversification, and firm growth: Evidence from japans early industrialization.

Bresnahan, T., S. Greenstein, D. Brownstone, and K. Flamm (1996). Technical progress and co-invention in computing and in the uses of computers. *Brookings Papers on Economic Activity* (1996), 1–83.

Callander, S. (2011, October). Searching and learning by trial and error. *American Economic Review 101*(6), 2277–2308.

Casadesus-Masanell, R., D. Yoffie, and S. Mattu (2014). Intel corporation: 1968-2003. Harvard Business School Case 703-427, November 2002. (Revised February 2010.).

Cooper, I. (2006). Asset pricing implications of nonconvex adjustment costs and irreversibility of investment. *The Journal of Finance 61*(1), 139–170.

Crouzet, N. and J. Eberly (2018, May). Intangibles, investment, and efficiency. *AEA Papers and Proceedings 108*, 426–31.

Crouzet, N. and J. Eberly (2020). Rents and intangible capital: A $q+$ framework.

De Loecker, J., J. Eeckhout, and G. Unger (2019). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, forthcoming.

Eisfeldt, A. and D. Papanikolau (2013). Organization capital and the cross-section of expected returns. *Journal of Finance 68*, 1365–1406.

Erickson, T. and T. Whited (2000). Measurement error and the relationship between investmentand ¡em¿q¡/em¿. *Journal of Political Economy 108*(5), 1027–1057.

Farboodi, M., R. Mihet, T. Philippon, and L. Veldkamp (2019, May). Big data and firm dynamics. *AEA Papers and Proceedings 109*, 38–42.

Foster, L., J. Haltinwanger, and C. Krizan (2001). *Aggregate Productivity Growth: Lessons from Microeconomic Evidence*, Volume 1. University of Chicago Press, Chicago IL.

Glaeser, E., H. Kallal, J. Scheinkman, and A. Shleifer (1992). Growth in cities. *Journal of Political Economy* (100), 1126–1152.

Grossman, S., R. Kihlstrom, and L. Mirman (1977). A bayesian approach to the production of information and learning by doing. *The Review of Economic Studies 44*(3), 533–547.

Gutierrez, G. and T. Philippon (2017). Investmentless growth: An empirical investigation. *Brookings Papers on Economic Activity*, 89–169.

Gutierrez, G. and T. Philippon (2018). Declining competition and investment in the us. *Working Paper*.

Hackbarth, D. and T. Johnson (2015). Real options and risk dynamics. *Review of Economic Studies 82*(4), 1449–1482.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica 50*(1), 213–224.

Jazwinski, A. (1970). *Stochastic Processes and Filtering Theory.* Academic Press, New York.

Jovanovic, B. and G. MacDonald (1994). Competitive diffusion. *Journal of Political Economy* (102), 24–52.

Jovanovic, B. and Y. Nyarko (1996, November). Learning by Doing and the Choice of Technology. *Econometrica 64*(6), 1299–1310.

Jovanovic, B. and R. Rob (1990). Long waves and short waves: Growth through intensive and extensive search. *Econometrica 58*(6), 1391–1409.

Judd, K. (1998, January). *Numerical Methods in Economics*, Volume 1 of *MIT Press Books*. The MIT Press.

Keller, G. and S. Rady (1999). Optimal experimentation in a changing environment. *Review of Economic Studies 66*(3), 475–507.

Kyle, A., A. Obizhaeva, and Y. Wang (2018). Smooth trading with overconfidence and market power. *Review of Economic Studies 85*(1), 611–662.

Lanjouw, J. and M. Schankerman (2001). Characteristics of patent litigation: A window on competition. *Rand Journal of Economics* (32), 129–151.

Lieberman, M. and S. Asaba (2006). Why do firms initate each other? *Academy of Management Review* (31), 366–385.

Lipster, R. and A. Shiryaev (2001). *Statistics of Random Processes II*. Springer Verlag, New York.

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

Ma, S. (2021). Obsolete Firms. *Working Paper*.

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

Manso, G., B. Balsmeier, and L. Fleming (2019). Heterogenous innovation and the antifragile economy. *Working Paper*.

Moscarini, G. and L. Smith (2001). The optimal level of experimentation. *Econometrica 69*(6), 1629–1644.

Nikolov, B. and T. Whited (2014). Agency conflicts and cash estimates from a structural model. *Journal of Finance 69*(5), 1883–1921.

Pastor, L. and P. Veronesi (2009). Technological revolutions and stock prices. *The American Economic Review 99*(4), 1451–1483.

Peters, R. and L. Taylor (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics 100*, 251–272.

Prescott, E. and M. Visscher (1980). Organization capital. *Journal of Political Economy* (88), 446–461.

Rob, R. (1991). Learning and capacity expansion under demand uncertainty. *The Review of Economic Studies 58*(4), 655–675.

Romer, P. (1986). Increasing returns and long run growth. *Journal of Political Economy* (94), 1002–1037.

Stoffman, N., M. Woeppel, and D. Yavuz (2021). Small innovators: No risk, no return. *Working Paper*.

Stoyanov, A. and N. Zubanov (2012). Productivity spillovers across firms through worker mobility. *American Economic Journal: Applied Economics* (4), 168–98.

Wald, A. (1947). Foundations of a general theory of sequential decision functions. *Econometrica 15*(4), 279–313.

Weitzman, M. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641–54.

Young, A. (1993). Invention and bounded learning by doing. *Journal of Political Economy 101*(3), 443–472.

A Model Calibration

A.1 Baseline Calibration

We first define parameter values for our baseline calibration, and summarize them in Table 3. Many parameters are standard, and thus calibrated following common practices in the literature. Consistent with Nikolov and Whited (2014), we set the discount rate to $r \equiv 5\%$, the capital depreciation rate to $\delta \equiv 13\%$, and the production function is such that $\alpha \equiv 0.8$. The profit function exhibits significant decreasing returns to scale, consistent with the rise of market power documented in Gutierrez and Philippon (2017). We set the inverse of the elasticity of demand to $\eta \equiv 0.5$, to capture the average markup of 1.5 in recent years as estimated in De Loecker, Eeckhout, and Unger (2019). Importantly, conjecturing that firm value is concave when experimentation is optimal, the"expected returns on knowledge" is endogenously lower than 0.5 (by Jensen's inequality); accordingly, we choose returns on capital, $\alpha(1-\eta) = 0.4 < 0.5$. The curvature of profits with respect to capital is therefore on the lower bound as discussed in Abel and Eberly (2011). Last, because firms' variable adjustment costs in our setting are proportional to profits, we calibrate firms' variable adjustment costs to $\gamma \equiv 5$ so that the investment rate is consistent with the summary statistics observed in our sample (see Table 1) and in Table 1 of Peters and Taylor (2017).

Table 3: Baseline parameter values in our full model. The table summarizes standard parameters in neoclassical models (part 1), and others specific to our model (part 2).

Symbol	Definition	Value
1. standard parameters		
r	risk-free rate	5%
δ	depreciation rate	13%
α	returns to scale on capital	0.8
η	inverse of price elasticity of demand	0.5
γ	variable adjustment costs	5
2. other parameters		
τ_A	informativeness of aggregate productivity	0.6
τ_S	informativeness of experiments	0.3
ω	obsolescence costs	59%
ϕ	knowledge dissipation	0.4

Other parameters determining the behavior of exploration and experimentation in our model are τ_A, τ_S, ω and ϕ. In our baseline calibration, we set the informativeness of experimentation to $\tau_S \equiv 0.3$, and the informativeness of the productivity signal to $\tau_A \equiv 0.6$, to match the observed variation in firm-level growth rates of output (i.e., Ma (2021)). Similarly, we ensure that exploration is infrequent by setting sufficiently high obsolescence costs ω. The stock of capital becomes obsolete so that firm value decreases by $1-(1-\omega)^{\alpha(1-\eta)} \approx 30\%$ upon exploration. We thus assume significant creative destruction upon breakthroughs in exploration, consistent with the evidence in Foster, Haltinwanger, and Krizan (2001). Finally, the dissipation of knowledge occuring upon experimentation determines how desirable a technology is. We set $\phi \equiv 0.35$, a value that triggers sufficient asymmetry in the exploration thresholds, so that the firm is more likely to experiment with an existing technology when the firm is confident that the technology is good ($Z > 0$).

A.2 Alternative Parametrizations

In Figure 7, we illustrate the properties of the model with respect to the expected length of the knowledge cycle by considering the same parameters as in Table 3, with the exception of $\tau_A = 0.3$.

In the extension with partially transferable knowledge, we illustrate the properties of the model numerically by considering the same parameters as in Table 3, in addition to imposing $\lambda = 0.2$.

In the extension with fixed costs, we consider all standard paramaters as stated in the first panel of Table 3, as well as $\omega = 0.59$, which we also use in the baseline calibration. Considering that our assumptions on adjustment costs, knowledge dissipation and the signals extracted from experimentation in this extension are different from the baseline model, we modify the remaining parameters so that $\gamma = 0.5$, $\kappa = 0.095$, $\tau_A = 0.7$, $\tau_S = 0.2$, and $\phi = 2$.

B Fixed adjustment costs and experimentation rounds

Our baseline specification of adjustment costs in Eq. (5) implies that the firm experiments actively at all times. Yet investment usually involves both variable and fixed costs, so that firms may find it optimal not to invest, particularly so when the outcome of experimentation is uncertain. The main insigh of the extension with fixed costs is that experimentation takes place in endogenous rounds. All proofs are in Section B.4 of this Appendix.

To introduce the possibility that firms find it optimal not to experiment, we keep the model of Section 2 unchanged, except along the following dimensions. First we follow Abel and Eberly (1994) and allow adjustment costs to exhibit a quasi fixed part, $\kappa > 0$, proportional to revenue:

$$\gamma_t \equiv \left(\kappa \mathbf{1}_{i>0} + \gamma/2 i^2\right) \cdot \Pi\left(A_t, K_t, N_t\right) \equiv \widehat{\gamma}(i) \Pi\left(A_t, K_t, N_t\right),$$

where, for simplicity, we have removed the linear part in Eq. (5).[20] Introducing fixed costs in turn

[20] Because the firm value exhibits concave and convex regions, we keep variable quadratic costs as we need them to keep the investment rate finite in the convex regions.

makes investment lumpy; accordingly, we redefine the flow of active signals in Eq. (2) as:

$$\mathrm{d}S_t = \mathbf{1}_{i_t>0}\left(\frac{\tau_S^{1/2} i_t^{1/2}}{\Omega_{t-}^{1/2}} M + \epsilon_t\right), \quad \epsilon_t \sim \mathscr{N}(0,1) \text{ and i.i.d.},$$

where $i_t = I_t/K_{t-}$. Investment lumpiness also implies that the firm's stock of capital and stock of knowledge now satisfy:

$$\mathrm{d}K_t/K_t = i_t \mathbf{1}_{t=\theta} - \delta \mathrm{d}t - \omega \mathbf{1}_{t=\nu},$$

and

$$\mathrm{d}Z_t = \underbrace{\tau_A/2 Z_t \mathrm{d}t + \tau_A^{1/2}\mathrm{d}\widehat{B}_t}_{\text{passive learning}} + \underbrace{Z_{t-}\left((1+\tau_S i_t)^{1/2} - 1\right) + (\tau_S i_t)^{1/2}\,\widehat{\epsilon}_t}_{\text{learning by experimenting}} \underbrace{-Z_{t-}\mathbf{1}_{t=\nu}}_{\text{knowledge reset}},$$

where $\widehat{\epsilon}$ is Gaussian with mean zero and unit variance under the firm's probability measure.[21] A difference relative to the model of Section 2 is that experimentation now updates knowledge punctually. The firm's problem then becomes:

$$\max_{\{\nu_n\}, i_{t+s}} \mathbb{E}\left[\int_0^\infty e^{-rs}\Pi\left(A_{t+s}, K_{t+s}, N_{t+s}\right)\mathrm{d}s - \sum_{n\geq 0} \gamma_{t+\theta_n} e^{-r\theta_n} \,\middle|\, \mathscr{F}_t\right], \tag{22}$$

where θ_n denotes the $n-$th time at which the firm invests. Finally, for simplicity, we assume that knowledge dissipation in Eq. (4) is symmetric across good and bad technologies, i.e., it evolves as $\mathrm{d}N_t/N_t = \phi Z_t^2 \mathrm{d}t$.

B.1 Solution with endogenous experimentation rounds

Building on the baseline case without fixed costs, we first conjecture that there exists a region $\mathscr{A}$, outside of which the firm decides to explore. In the presence of fixed costs, we further conjecture that there is another region $\mathscr{B} = (\underline{b}, \overline{b}) \subset \mathscr{A}$, inside which the firm decides to invest. The baseline model also shows that firm value is largest following exploration, which also suggests that $\{0\} \in \mathscr{B}$. In words, it is optimal for the firm to invest immediately following exploration, as this is when q is highest. Because fixed costs imply inactions regions (i.e., $\mathscr{A} \setminus \mathscr{B}$, areas inside $\mathscr{A}$ but outside $\mathscr{B}$), experimentation optimally occurs in rounds. Whenever knowledge enters the range $\mathscr{B}$, the firm starts experimenting. Since the outcome of each experiment is uncertain, the firm may re-enter the experimentation region after the first trial. The firm will typically go through multiple rounds of trial and error until the experiments eventually becomes conclusive (so that $Z \notin \mathscr{B}$).

Formally, let $Z^{(0)}$ be the stock of knowledge upon first entry in the experimentation region, and $i(Z^{(0)})$ be the associated investment choice. Based on the knowledge dynamics in Eq. (3), this first

[21] See Theorem 12.7 in Lipster and Shiryaev (2001) for the continuous part, and Theorem 7.1 Jazwinski (1970) for the discontinuous part.

experiment will update knowledge to a new (random) level:

$$Z^{(1)} \equiv Z^{(0)}(1+\tau_S i(Z^{(0)}))^{1/2} + \sqrt{\tau_S i(Z^{(0)})}\widehat{\epsilon}^{(1)}.$$

Suppose further that the firm goes through n such experimentation rounds. Since each round of experimentation takes place instantaneously (each round has Lebesgue measure zero), knowledge, $Z^{(n)}$, after the $n-$th round of experimentation satisfies:

$$Z^{(n)} = Z^{(0)} \prod_{m=0}^{n-1}(1+\tau_S i(Z^m))^{1/2} + \sum_{k=1}^{n-1}\sqrt{\tau_S i(Z^{(k-1)})}\widehat{\epsilon}^{(k)} \prod_{m=k}^{n-1}(1+\tau_S i(Z^{(m)}))^{1/2} + \sqrt{\tau_S i(Z^{(n-1)})}\widehat{\epsilon}^{(n)}. \tag{23}$$

Eq. (23) shows that, across rounds of experimentation, knowledge evolves as an AR(1) process, the persistence and volatility of which are endogenously determined by investment. Based on this law of motion, we define an "experimentation round" as the number of successive trials it takes until first exit of the experimentation region, $\mathscr{B}$.

Definition 4. ***(Experimentation Round)*** *A round of experimentation starts upon first entry in the experimentation region, $Z^{(0)} \in \mathscr{B}$, and ends after the $n-$th trial:*

$$n = \inf\{k \in \mathbb{N} : Z^{(k)} \notin \mathscr{B}\}$$

upon first exit of the experimentation region, where $Z^{(k)}$ is defined in Eq. (23).

Because the firm anticipates that each experiment may lead to another round of experimentation, it must determine a knowledge-contingent policy for investment depending on *all* possible experimentation outcomes. In particular, the possibility of successive rounds of experimentation makes experimentation a dynamic optimization problem across rounds. Let $V(Z)$ denote the value associated with this problem (i.e., firm value in the experimentation range). Denoting by $\mathscr{D}$ the feasible investment set (to be defined below), firm value $V(Z)$ solves the following dynamic programming problem for $Z \in \mathscr{B}$:

$$V(Z) = \max_{i\in\mathscr{D}} \underbrace{(1+i)^{\alpha(1-\eta)}}_{\text{returns on capital}} \underbrace{\int_{\mathbb{R}} v(x)\varphi(x; i(Z), Z)\mathrm{d}x}_{\text{expected returns on knowledge}} -\kappa - \gamma/2i^2, \tag{24}$$

where $\varphi(\cdot; i(Z), Z)$ is the distribution of incremental knowledge built from experimenting (as defined in Definition 3). Because the outcome of experimentation is uncertain, the firm does not know *ex-ante* what information it will get from experimenting. It only knows is that experimentation through active investment improves knowledge on average.

We now reformulate the problem in Eq. (24) recursively in the form of a Bellman equation. We denote by $h(Z)$ firm value when not experimenting and write it piecewise as:

$$h(Z) \equiv g(Z; \underline{C}_1, \underline{C}_2)\mathbf{1}_{Z\in(\underline{a},\underline{b})} + (1-\omega)^{\alpha(1-\eta)}V(0)\mathbf{1}_{Z\notin\mathscr{A}} + g(Z; \overline{C}_1, \overline{C}_2)\mathbf{1}_{Z\in(\overline{b},\overline{a})}. \tag{25}$$

Accordingly, the value function satisfies, $v(Z) = \mathbf{1}_{Z\in\mathscr{B}}V(Z) + \mathbf{1}_{Z\notin\mathscr{B}}h(Z)$. Substituting this expression in Eq. (24) delivers the desired Bellman equation, which we highlight in the proposition below (see Appendix B.4.1).

Proposition 1. *The value function, $V(Z)$, defined in Eq. (24), associated with optimal experimentation satisfies the Bellman equation:*

$$V(Z) = \max_{i\in\mathscr{D}}(1+i)^{\alpha(1-\eta)}\int_{\mathbb{R}} h(y)\varphi(y;i(Z),Z)\,dy - \widehat{\gamma}(i) + (1+i)^{\alpha(1-\eta)}\int_{\mathscr{B}} V(y)\varphi(y;i(Z),Z)\,dy, \tag{26}$$

where the admissible investment set is given by:

$$\mathscr{D} = \left[0, \frac{2}{\tau_S(\sqrt{\tau_A + 8\eta\phi}/\sqrt{\tau_A} - 1)}\right).$$

The Bellman equation (26) formalizes the firm's dynamic investment problem across experimentation rounds. At every experimentation round, the firm trades off adjustment costs against the marginal revenue product of knowledge and capital, which the firm re-evaluates depending on the updated likelihood of each experimentation outcome. Through the first integral the firm anticipates that experimentation may lead to future inaction or immediate exploration; through the second integral the firm anticipates that experimentation may lead to yet another round of experimentation, which precisely makes the problem recursive across successive rounds of experimentation.

We now compute firm value piecewise over the two regions $\mathscr{A}$ and $\mathscr{B}$. We start in regions, $\mathscr{A}\setminus\mathscr{B}$, in which the firm is inactive—it neither explores nor invests. We denote the solution for the value of the firm when inactive by $g(\cdot)$, and define it in the proposition below (see also Appendix B.4.2).

Proposition 2. *The intensive firm value in inaction regions, $Z \in \mathscr{A}\setminus\mathscr{B}$, satisfies:*

$$g(Z; C_1, C_2) := v_P(Z) + e^{g_1 Z + g_2 Z^2}\left(C_1 H_n(h_0 + h_1 Z) + C_2 M\left(-\frac{1}{2}n, \frac{1}{2}, m_0 + m_1 Z + m_2 Z^2\right)\right) \tag{27}$$

where C_1 and C_2 are two integration constants, $v_P(\cdot)$, is a particular solution:

$$v_P(Z) = \int_0^\infty \exp\left(a_0(s) + a_1(s)Z + a_2(s)Z^2\right)ds, \tag{28}$$

the function, $H_n(\cdot)$, represents the Hermite polynomial of order:

$$n = \frac{1}{2}\left(\frac{2(\eta-1)^2\tau_A - \xi^2(-2(\eta-1)(2\alpha\delta+1) + 4r + \tau_A)}{\xi^3\sqrt{\tau_A}} - 1\right), \tag{29}$$

and $M(\cdot)$ denotes the the Kummer confluent hyper-geometric function. The definitions of the functions $a_0(\cdot)$, $a_1(\cdot)$ and $a_2(\cdot)$, and of other parameters are presented in Appendix B.4.2.

Because there are two inaction regions in the model, $(\underline{a}, \underline{b})$ and $(\overline{b}, \overline{a})$, we use the notation $\underline{C}_1$ and $\underline{C}_2$ for the two constants in Proposition 2 when defined over the lower inaction region, and $\overline{C}_1$

and $\overline{C}_2$ over the upper inaction region. We then determine the value of the firm over the region $\mathscr{B}$ in which it invests. A technical difficulty at this stage is that the Bellman equation (26) does not have an explicit solution. In Appendix B.4.3, we explain how to proceed numerically. Finally, we obtain $\underline{C}_1$, $\underline{C}_2$, $\overline{C}_1$ and $\overline{C}_2$ by "piecing together" firm value across the investment region, $\mathscr{B}$, and inaction regions, $\mathscr{A} \setminus \mathscr{B}$, and determine the location of the optimal thresholds through customary "smooth-pasting" conditions.[22]

Figure 17 illustrates firm value as a function of knowledge based on the parameter values discussed in Appendix A.2. Although of similar shape relative to the baseline model, an important difference is in the interpretation of the investment region (the hashed blue area). In the model with fixed costs, investment occurs in lumps. Furthermore, upon entering the blue hashed area, the firm faces uncertainty as to the outcome, as it may lead to another round of investment, to inaction or to explore a new technology. We illustrate in Figure 18 the corresponding exploration and experimentation strategy based on a representative simulated sample path of the stock of knowledge. As knowledge accumulates away from its reset point, it becomes optimal to explore a new technology (red arrows), either because the current technology is likely poor, or because knowledge dissipation intensifies. Upon entering the experimentation zone (the shaded blue area), the firm goes through possibly multiple rounds of experimentation, until experiments become conclusive (the blue dots).

B.2 Knowledge channel in the presence of fixed costs

We first examine the relation between investment and knowledge in the presence of fixed costs. In this case, the knowledge channel becomes:

$$c(Z) \equiv \int_{\mathbb{R}} \frac{V(\cdot, K+I, x)}{\Pi(\cdot)} \frac{d}{di}\varphi(x; i(Z), Z)\mathrm{d}x = (1+i(Z))^{\alpha(1-\eta)} \int_{\mathbb{R}} v(x) \frac{d}{di}\varphi(x; i(Z), Z)\mathrm{d}x, \quad (31)$$

where $\varphi(\cdot; i(Z), Z)$ denotes the normal density with mean $\sqrt{1+\tau_S i(Z)}Z$ and variance $\tau_S i(Z)$ (i.e., the distribution of incremental knowledge gathered from investing, as per Eq. (3)). Furthermore, q, which in this case is equivalent to that of Abel and Blanchard (1986) (except for scaling by

[22] Based on Eq. (11) and piecing functions together, value-matching conditions satisfy:

$$\lim_{Z \downarrow \underline{b}} V(Z) = \lim_{Z \uparrow \underline{b}} g\left(Z; \underline{C}_1, \underline{C}_2\right) = \lim_{Z \uparrow \bar{b}} V(Z) = \lim_{Z \downarrow \bar{b}} g\left(Z; \overline{C}_1, \overline{C}_2\right),$$
$$\lim_{Z \searrow \underline{a}} g\left(Z; \underline{C}_1, \underline{C}_2\right) = \lim_{Z \uparrow \bar{a}} g\left(Z; \overline{C}_1, \overline{C}_2\right) = V(0)(1-\omega)^{\alpha(1-\eta)}$$

and together serve to pin down the 4 constants, $\underline{C}_1$, $\underline{C}_2$, $\overline{C}_1$ and $\overline{C}_2$. Smooth-pasting conditions satisfy:

$$\lim_{Z \downarrow \underline{a}} g'\left(Z; \underline{C}_1, \underline{C}_2\right) = \lim_{Z \uparrow \bar{a}} g'\left(Z; \overline{C}_1, \overline{C}_2\right) = 0, \ \lim_{Z \uparrow \underline{b}} g'(Z; \underline{C}_1, \underline{C}_2) = \lim_{Z \downarrow \underline{b}} V'(Z), \ \lim_{Z \uparrow \bar{b}} V'(Z) = \lim_{Z \downarrow \bar{b}} g'(Z; \overline{C}_1, \overline{C}_2) \quad (30)$$

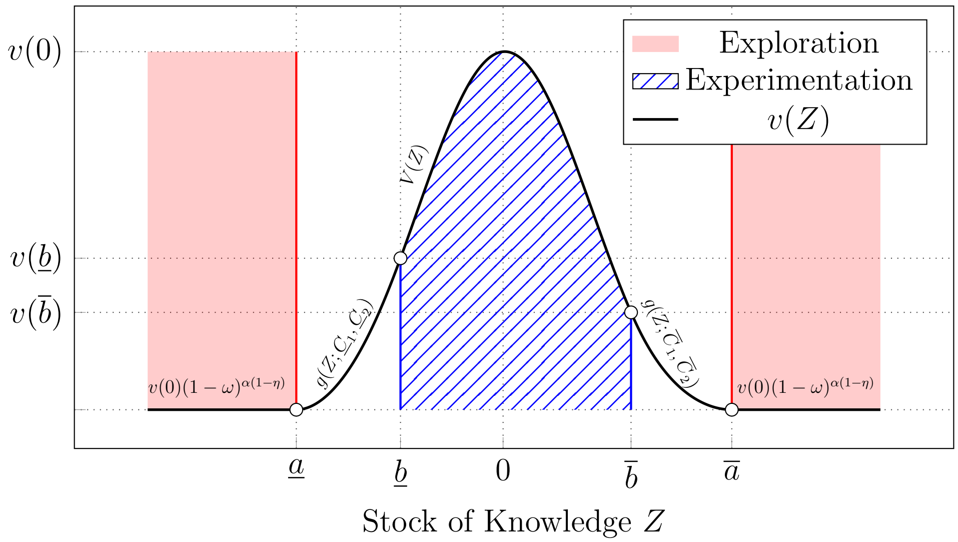

Figure 17: Firm value as a function of knowledge in the presence of fixed costs. The figure plots the firm value (solid black line) as a function of knowledge, Z, about the technology it operates. The blue hased area represents the experimentation regions, defined by the two thresholds, $\underline{b}$ and $\overline{b}$. The red shaded areas corresponds to the two exploration regions, each defined by the thresholds $\underline{a}$ and $\overline{a}$, respectively. The piecewise representation of the firm value is reported below the curve within each region. Parameter values are defined in Appendix A.2.

revenues), satisfies:

$$q(Z) \equiv \int_{\mathbb{R}} \frac{V_K(\cdot, K+I, x)}{\Pi(\cdot)/K} \varphi(x; i(Z), Z)\mathrm{d}x = \alpha(1-\eta)(1+i(Z))^{\alpha(1-\eta)-1} \int_{\mathbb{R}} v(x)\varphi(x; i(Z), Z)\mathrm{d}x.$$

Based on these redefinitions, the first-order condition for optimal investment is exactly as in Eq. (17) of the baseline.

Within the experimentation band, investment behaves very much like in the baseline model; it is weakest following exploration, and strengthens as the firm accumulates knowledge. In fact an application of Stein's lemma shows that the knowledge channel in Eq. (31) takes a form very similar to its counterpart in the baseline model (i.e., Eq. (15)):[23]

$$c(Z) = (1+i)^{\alpha(1-\eta)} \frac{\tau_S}{2} \Big(\underbrace{\frac{Z}{\sqrt{1+\tau_S i}} \int_{\mathbb{R}} v'(x)\varphi(x; i(Z), Z)\mathrm{d}x}_{\text{marginal product of knowledge}} + \underbrace{\int_{\mathbb{R}} v''(x)\varphi(x; i(Z), Z)\mathrm{d}x}_{\text{attitude towards experimentation risk}} \Big).$$

As a result, except for creating inaction the presence of fixed costs does not affect the main insight

[23]The resulting relation only makes sense if one ignores the points at which the firm value is not twice continuously differentiable.

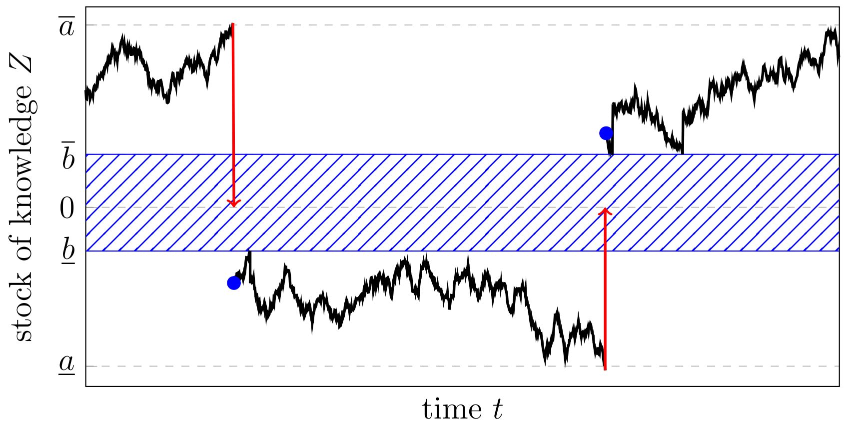

Figure 18: Sample path of Knowledge. The figure plots a sample path of the firm stock of knowledge (solid black line) over time. Blue dots correspond to times at which the firm chooses to experiment; red arrows correspond to times when the firm chooses to explore.

of the paper.

B.3 Knowledge cycles and experimentation rounds

The extension with fixed costs delivers novel implications regarding the number of rounds of experimentation within each cycle. The left panel of Figure 19 illustrates the distribution of experimentation rounds as defined in Definition 4 (i.e., the number of successive trials until first exit of $\mathscr{B}$). Under the baseline calibration, after exploring a new technology it most frequently takes three rounds of experimentation trials until the firm becomes either inactive or explores again. This panel further shows that a reduction in knowledge dissipation (the blue bars) typically reduces the average number of rounds by narrowing the experimentation band. The experimentation rounds in the left panel take place instantly.

Another meaningful way of characterizing the dynamics of experimentation is in terms of "in-and-out" trials, whereby rounds occur intermittently within a given knowledge cycle. Fix a given technology and suppose the firm is currently inactive. We can then compute the number of times the firm will go in and out of the experimentation region $\mathscr{B}$ and back into the inaction region, until it eventually explores a new technology. The cycle eventually ends because either the last round of experimentation leads to immediate exploration, or because it leads to an inaction region within which knowledge accumulates in a way to trigger subsequent exploration. Formally, consider the sequence of stopping times at which the firm experiments:

$$\theta_n = \inf\{t \geq \theta_{n-1} : Z_t \in \mathscr{B}\}, \;\; \forall n \in \mathbb{N}.$$

Let us further denote by θ_{n+} the moment the firm exits the experimentation range right after the $n-$th time it entered it, recalling that the firm goes in and out of $\mathscr{B}$ instantly. Finally, consider

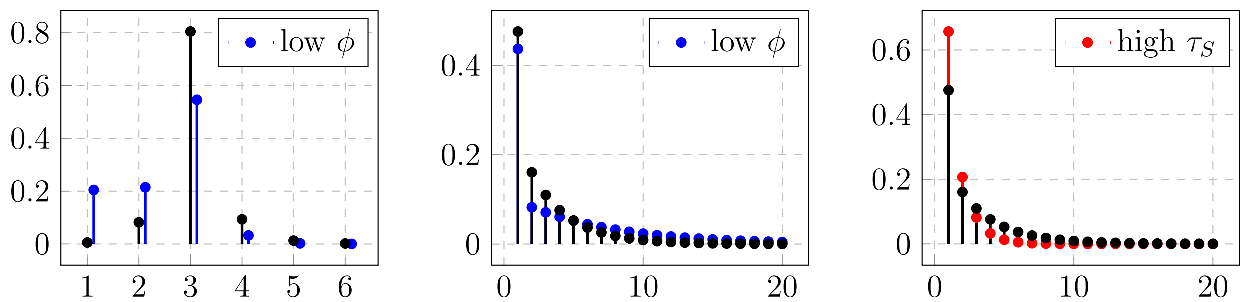

Figure 19: Distribution of experimentation trials. This figure plots the distribution of the number of instantaneous trials in an experimentation round as per Definition 4 (right panel) and "in-and-out" trials (middle and right panels). "In-and-out" trials correspond to the number of times a firm currently in an inaction region goes back and forth between the region $\mathscr{B}$ and inaction until the knowledge cycle ends. The black curve sets the baseline; the blue curve corresponds to knowledge dissipation ($\phi = 1.4 < 2$), and the red line to low experimentation noise ($\tau_S = 0.3 > 0.25$). Parameter values are provided in Appendix A.2.

the sequence of knowledge stopped at these dates:

$$\left\{Z_{\theta_{0+}}, ..., Z_{\theta_{n+}}\right\}.$$

We illustrate the distribution of the number n of in-and-out trials at which this sequence ends (i.e., $Z_{\theta_{n+}} \notin \mathscr{A}$) in the middle and right panels of Figure 19.

A firm will typically go through multiple "in-and-out" trials, with a single trial being the most frequent outcome but with a nontrivial probability of going through as many as ten trials. Considering both the left and middle panels, the firm may go in and out of the experimentation region back to inaction up to 10 times, each time typically experimenting in rounds of three instantaneous successive trials. The right panel further shows how more informative experiments (in red) affects these "in-and-out" trials. Informative experimentation makes longer sequences of experimentation less likely, as the inaction region shrinks. Reduced knowledge dissipation (the blue bars in the middle panel) has the opposite effect. In this case, experimentation is more likely to occur intermittently, as the firm's intellectual property is better protected. This kind of intermittent experimentation is associated with longer knowledge cycles.

B.4 Proofs

B.4.1 Proof of Proposition 1: expected firm value upon experimentation

To solve for the expected firm value upon experimentation, we start by formulating the dynamic optimization problem in Eq. (24) globally. We write the value function piecewise as:

$$v(Z) = \begin{cases} g(Z; \underline{C}_1, \underline{C}_2), & Z \in (\underline{a}, \underline{b}) \\ V(Z), & Z \in \mathscr{B} \\ g(Z; \overline{C}_1, \overline{C}_2), & Z \in (\overline{b}, \overline{a}) \\ (1-\omega)^{\alpha(1-\eta)} V(0), & Z \notin \mathscr{A} \end{cases}. \tag{32}$$

Using the function $h(Z)$ in Eq. (25) and letting the (yet unknown) optimal investment policy $i(Z)$ be a function of the current state Z, we can rewrite firm value as:

$$v(Z) \equiv h(Z) + \mathbf{1}_{Z\in\mathscr{B}} \left((1+i(Z))^{\alpha(1-\eta)} \int_{\mathbb{R}} v(y)\phi(y; Z, i(Z))\mathrm{d}y - (\kappa + \gamma/2 i(Z)^2) \right).$$

The firm will choose a different investment policy for each realization of knowledge, $Z^{(n)}$, at round n:

$$i_n \equiv i(Z^{(n)}), \quad n \in \mathbb{N}.$$

Furthermore, let $\varphi(\cdot; i(Z), Z)$ denote the distribution of incremental knowledge gathered from investing, as per Eq. (3); specifically, it is a normal density with mean $\sqrt{1+\tau_S i(Z)}Z$ and variance $\tau_S i(Z)$. of incremental knowledge gained from experimentation, define for $y \equiv Z^{(n+1)}$ and $z = Z^{(0)}$:

$$\phi^{(n)}(y,z) \equiv \int_{\mathscr{B}^n} \prod_{k=1}^{n} (1+i_k)^{\alpha(1-\eta)} \varphi(y; i_n, z^{(n)}) \varphi(z^{(k)}; i_{k-1}, z^{(k-1)}) \mathrm{d}z^{(k)}, \tag{33}$$

with $\phi^{(0)}(y;z) \equiv \varphi(y; i_0, z)$. This expression corresponds to the density of $Z^{(n+1)}$ after n consecutive rounds, scaled by a function of investment at each round. Using this notation, we can then express the expected firm value as

$$\int_{\mathbb{R}} v(y)\phi(y; Z, i(Z))\mathrm{d}y = \sum_{n=0}^{\infty} \int_{\mathbb{R}} \left(h(y) - (\kappa + \gamma/2 i(y)^2)\mathbf{1}_{y\in\mathscr{B}}\right) \phi^{(n)}(y; Z, \{i_k\}_{k=0}^{n})\mathrm{d}y.$$

It follows that for $Z \in \mathscr{B}$ the firm faces the dynamic programming problem:

$$V(Z) = \max_{\{i_k\}_{k=0}^{\infty}} (1+i_0)^{\alpha(1-\eta)} \sum_{n=0}^{\infty} \int_{\mathbb{R}} \left(h(y) - \widehat{\gamma}(i(y))\mathbf{1}_{y\in\mathscr{B}}\right) \phi^{(n)}(y; Z)\mathrm{d}y - \widehat{\gamma}(i_0). \tag{34}$$

As is customary we wish to formulate this program recursively in the form of a Bellman equation:

$$
\begin{aligned}
V(Z) \equiv \max_{i_0} & (1+i_0)^{\alpha(1-\eta)} \int_{\mathbb{R}} h(y)\phi(y;Z,i_0)\mathrm{d}y - (\kappa + \gamma/2 i_0^2) \\
& + \max_{\{i_k\}_{k=0}^{\infty}} (1+i_0)^{\alpha(1-\eta)} \left(\sum_{n=1}^{\infty} \int_{\mathbb{R}} h(y)\phi^{(n)}(y;Z,\{i_k\}_{k=0}^{n})\mathrm{d}y - \sum_{n=0}^{\infty} \int_{\mathscr{B}} (\kappa + \gamma/2 i(y)^2)\phi^{(n)}(y;Z,\{i_k\}_{k=0}^{n})\mathrm{d}y \right).
\end{aligned}
$$

Rewriting Eq. (33) recursively as:

$$
\phi^{(n+1)}(y;Z,\{i_k\}_{k=0}^{n+1}) = \int_{\mathscr{B}} (1+i(x))^{\alpha(1-\eta)} \phi^{(n)}(y;Z,\{i_k\}_{k=1}^{n})\phi(x;Z,i_0)\mathrm{d}x,
$$

we can rewrite the first term in brackets as:

$$
\begin{aligned}
\sum_{n=1}^{\infty} \int_{\mathbb{R}} h(y)\phi^{(n)}(y;Z,\{i_k\}_{k=0}^{n})\mathrm{d}y &= \sum_{n=0}^{\infty} \int_{\mathbb{R}} h(y)\phi^{(n+1)}(y;Z,\{i_k\}_{k=0}^{n+1})\mathrm{d}y \\
&= \int_{\mathscr{B}} (1+i(x))^{\alpha(1-\eta)} \sum_{n=0}^{\infty} \int_{\mathbb{R}} h(y)\phi^{(n)}(y;Z,\{i_k\}_{k=0}^{n})\mathrm{d}y\phi(x;Z,i_0)\mathrm{d}x.
\end{aligned}
$$

and the second term in brackets as:

$$
\begin{aligned}
& \sum_{n=0}^{\infty} \int_{\mathscr{B}} (\kappa + \gamma/2 i(y)^2)\phi^{(n)}(y;Z,\{i_k\}_{k=0}^{n})\mathrm{d}y \\
&= \int_{\mathscr{B}} \left(\kappa + \gamma/2 i(x)^2 + (1+i(x))^{\alpha(1-\eta)} \sum_{n=0}^{\infty} \int_{\mathscr{B}} (\kappa + \gamma/2 i(y)^2)\phi^{(n)}(y;x,\{i_k\}_{k=0}^{n})\mathrm{d}y \right) \phi(x;Z,i_0)\mathrm{d}x.
\end{aligned}
$$

We then substitute this expression in the value function to obtain:

$$
\begin{aligned}
V(Z) \equiv \max_{i_0} & (1+i_0)^{\alpha(1-\eta)} \int_{\mathbb{R}} h(y)\phi(y;Z,i_0)\mathrm{d}y - \kappa - \gamma/2 i_0^2 + \max_{\{i_k\}_{k=0}^{\infty}} (1+i_0)^{\alpha(1-\eta)} \times \qquad (35) \\
& \int_{\mathscr{B}} \left(\begin{array}{c} (1+i(x))^{\alpha(1-\eta)} \sum_{n=0}^{\infty} \int_{\mathbb{R}} \left(h(y) - (\kappa + \gamma/2 i(y)^2)\mathbf{1}_{y\in\mathscr{B}} \right) \phi^{(n)}(y;x,\{i_k\}_{k=0}^{n})\mathrm{d}y \\ -\kappa - \gamma/2 i(x)^2 \end{array} \right) \phi(x;Z,i_0)\mathrm{d}x \\
& = \max_{i} (1+i)^{\alpha(1-\eta)} \int_{\mathbb{R}} h(y)\phi(y;Z,i)\mathrm{d}y - \kappa - \gamma/2 i^2 + (1+i)^{\alpha(1-\eta)} \int_{\mathscr{B}} \phi(x;Z,i)V(x)\mathrm{d}x,
\end{aligned}
$$

where the last equation yields the Bellman equation we are looking for. We are further looking for an optimal policy function $i(\cdot)$ such that

$$V(z) = (1+i(Z))^{\alpha(1-\eta)} \int_{\mathbb{R}} h(y)\phi(y;Z,i)\mathrm{d}y - \kappa - \gamma/2i(Z)^2 + (1+i(Z))^{\alpha(1-\eta)} \int_{\mathscr{B}} \phi(x;Z,i)V(x)\mathrm{d}x.$$

Differentiating Eq. (35), the optimal policy function must satisfy the first-order condition:

$$\begin{aligned} 0 &= \frac{\alpha(1-\eta)}{1+i(Z)}V(Z) + \frac{\alpha(1-\eta)(\kappa+\gamma/2i(Z)^2) - \gamma(1+i(Z))i(Z)}{1+i(Z)} \\ &+ (1+i(Z))^{\alpha(1-\eta)} \left(\int_{\mathbb{R}} h(y)\frac{\partial}{\partial i}\phi(x;i,Z)\mathrm{d}y + \int_{\mathscr{B}} V(y)\frac{\partial}{\partial i}\phi(x;i,Z)\mathrm{d}y \right). \end{aligned} \tag{36}$$

We start by computing the first expectation in Eq. (35). It is convenient to define

$$\begin{aligned} \psi(Z;i,c_0,c_1,c_2,\underline{x},\overline{x}) &:= \int_{\underline{x}}^{\overline{x}} \exp\left(c_0 + c_1 y + c_2 y^2\right)\phi(y;i(Z),Z)\mathrm{d}y \\ &= \frac{e^{c_0 + \frac{c_1^2 i(Z)\tau_S + 2Z(c_1\sqrt{1+\tau_S i} + c_2 Z(1+\tau_S i(Z)))}{2i\tau_S(1-2c_2\tau_S i(Z))}}}{\sqrt{2(1-2c_2\tau_S i(Z))}} \left(\begin{array}{c} \operatorname{erf}\left(\frac{c_1\tau_S i(Z) + Z\sqrt{1+\tau_S i(Z)} - \overline{x}(1-2c_2\tau_s i(Z))}{\sqrt{2\tau_S i(Z)(1-2c_2\tau_S i(Z))}}\right) \\ -\operatorname{erf}\left(\frac{c_1\tau_S i(Z) + Z\sqrt{1+\tau_S i(Z)} - \underline{x}(1-2c_2\tau_s i(Z))}{\sqrt{2\tau_S i(Z)(1-2c_2\tau_S i(Z))}}\right) \end{array} \right), \end{aligned}$$

where c_0, c_1, c_2, and $\overline{x} \geq \underline{x}$ are constant coefficients. We note that this function exists provided that

$$c_2 < \frac{1}{2\tau_S i(Z)}, \tag{37}$$

which imposes an upper bound on the investment policy function $i(Z)$ when $c_2 > 0$. Using this function and applying Fubini's theorem it follows that

$$\int_{\underline{x}}^{\overline{x}} v_P(y)\,\phi(y;i(Z),Z)\mathrm{d}y = \int_0^{\infty} \psi(Z;i(Z),a_0(s),a_1(s),a_2(s),\underline{x},\overline{x})\mathrm{d}s =: \overline{v}_P(Z,\underline{x},\overline{x}).$$

To compute the expectation of the second term involved in Eq. (27), we use the integral representation of the Hermite polynomial. Note that the order of the polynomial in Eq. (29) may be negative, in which case the following integral representation is appropriate:

$$H_n(x) \equiv \frac{(-2\mathbb{i})^{-n}}{\sqrt{\pi}} \int_{\mathbb{R}} t^{-n}\exp(-(t-\mathbb{i}x)^2)\mathrm{d}t.$$

We then obtain:

$$\begin{aligned}\int_{\underline{x}}^{\overline{x}} G_1(y)\,\phi(y;i(Z),Z)\mathrm{d}y &= \int_{\underline{x}}^{\overline{x}} e^{g_1 y + g_2 y^2} H_n(h_0 + h_1 y)\,\phi(y;i(Z),Z)\mathrm{d}y \\ &= \int_{\underline{x}}^{\overline{x}} e^{g_1 y + g_2 y^2} \frac{(-2\mathbb{i})^{-n}}{\sqrt{\pi}} \int_{\mathbb{R}} t^{-n} \exp\left(-(t-\mathbb{i}(h_0+h_1 y))^2\right) \mathrm{d}t\phi(y;i(Z),Z)\mathrm{d}y \\ &= \frac{(-2\mathbb{i})^{-n}}{\sqrt{\pi}} \int_{\mathbb{R}} t^{-n} \psi(Z;i(Z),(h_0+\mathbb{i}t)^2, g_1 + 2h_1(h_0+\mathbb{i}t), g_2 + h_1^2, \underline{x}, \overline{x})\mathrm{d}t \\ &=: \overline{G}_1(Z,\underline{x},\overline{x}),\end{aligned}$$

where the last equality uses Fubini's theorem to slide the expectation inside the integral. Similarly, using the integral representation of the confluent hypergeometric function we further obtain:

$$\begin{aligned}\int_{\underline{x}}^{\overline{x}} G_2(y)\,\phi(y;i(Z),Z)\mathrm{d}y &= \int_{\underline{x}}^{\overline{x}} e^{g_1 y + g_2 y^2} M\left(\frac{1}{2}n, \frac{1}{2}, m_0 + m_1 y + m_2 y^2\right)\mathrm{d}y \\ &= \int_{\underline{x}}^{\overline{x}} e^{g_1 y + g_2 y^2} \int_0^1 \frac{t^{n/2-1}(1-t)^{-\frac{n+1}{2}}}{\Gamma(n/2)\Gamma((1-n)/2)} e^{(m_0+m_1 y + m_2 y^2)t} \mathrm{d}t\mathrm{d}y \\ &= \int_0^1 \frac{t^{n/2-1}(1-t)^{-\frac{n+1}{2}}}{\Gamma(n/2)\Gamma((1-n)/2)} \psi(Z;i(Z), m_0 t, g_1 + m_1 t, g_2 + m_2 t, \underline{x}, \overline{x})\mathrm{d}t \\ &=: \overline{G}_2(Z,\underline{x},\overline{x}),\end{aligned}$$

where the last equality uses Fubini's theorem. Using these expressions we can write the expectation in Eq. (35) as:

$$\begin{aligned}\int_{\mathbb{R}} h(y)\phi(y;Z,i)\mathrm{d}y = \overline{v}_P(Z,\underline{a},\underline{b}) + \underline{C}_1\overline{G}_1(Z,\underline{a},\underline{b}) + \underline{C}_2\overline{G}_2(Z,\underline{a},\underline{b}) + \overline{v}_P(Z,\overline{b},\overline{a}) + \overline{C}_1\overline{G}_1(Z,\overline{b},\overline{a}) \\ + \overline{C}_2\overline{G}_2(Z,\overline{b},\overline{a}) + (1-\delta)^{\alpha(1-\eta)}V(0)\int_{y\notin\mathscr{A}} \phi(y;i(Z),Z)\mathrm{d}y.\end{aligned}$$

Using the existence condition in Eq. (37), the existence of this expectation depends on the following conditions being satisfied:

$$\begin{aligned}&a_2(s)\tau_S i(Z) < 1/2, \quad \forall s > 0 \\ &\tau_S i(Z)\left(-\frac{1}{2} + \left(y - \frac{1}{2}\right)\frac{\xi}{\sqrt{\tau_A}}\right) < 1, \quad \forall y \in (0,1)\end{aligned}$$

These conditions place an upper bound on the investment rate. In particular, the last condition is always satisfied provided that the investment rate satisfies:

$$i < \frac{2}{\tau_S(\xi/\sqrt{\tau_A} - 1)} =: \bar{i}. \tag{38}$$

Furthermore, using Eq. (40) for a_2, since $\tanh : \mathbb{R} \to [-1, 1]$ it follows that the first equation is always satisfied if the condition in Eq. (38) is satisfied; thus it is a sufficient condition for the expectations above to be well-defined. We will thus take the admissible investment set to be

$$\mathscr{D} := [0, \bar{i}).$$

B.4.2 Proof of Proposition 2: general solution in the inaction region

In this appendix we solve for the intensive firm value in the inaction region. In the inaction region, the processes for knowledge and physical capital move of their own accord. As a result, the firm value in Eq. (22) satisfies the standard HJB equation, which we omit for the sake of brevity. Substituting the functional form in Eq. (7) into this equation, we obtain an ODE for the intensive value, $v(\cdot)$, that it must satisfy within the inaction region:

$$\left(r + \eta Z^2\phi - (\eta - 1)\left(\alpha\delta + \frac{\eta}{2} - \sqrt{\tau_A}Z\right)\right)v = 1 + \sqrt{\tau_A}v'\left(\frac{\sqrt{\tau_A}}{2}Z + 1 - \eta\right) + \frac{\tau_A}{2}v'',\ Z \in \mathscr{I} \tag{39}$$

which is a second-order inhomogeneous linear equation.

Clearly, the firm could always choose never to invest nor to explore, and this outcome must thus be a particular solution, which we denote v_P, of this ODE. In particular, the value of the firm when it remains inactive forever is

$$\begin{aligned} V(N_t, A_t, K_t, \widehat{M}_t, \Omega_t) &\equiv \mathbb{E}\left[\int_0^\infty e^{-rs}\Pi(A_{t+s}, K_{t+s}, N_{t+s})\mathrm{d}s \middle| \mathscr{F}_t\right] \\ &= \Pi(A_t, K_t, N_t)\mathbb{E}\left[\int_0^\infty e^{\int_t^{t+s}\left((1-\eta)\left(\tau_A^{1/2}Z_u - \alpha\delta - 1/2\right) - \eta\phi Z_u^2 - r\right)\mathrm{d}u + (1-\eta)(\widehat{B}_{t+s} - \widehat{B}_t)}\mathrm{d}s \middle| \mathscr{F}_t\right]. \end{aligned}$$

Introducing the change of probability measure,

$$\left.\frac{\widetilde{\mathbb{P}}}{\widehat{\mathbb{P}}}\right|_{\mathscr{F}_t} = e^{-1/2(1-\eta)^2 t + (1-\eta)B_t},$$

and applying Fubini's theorem, we write this value as:

$$\begin{aligned} V(N_t, A_t, K_t, \widehat{M}_t, \Omega_t) &= \Pi(A_t, K_t, N_t)\int_0^\infty \widetilde{\mathbb{E}}\left[e^{\int_t^{t+s}\left((1-\eta)\left(\tau_A^{1/2}Z_u - \alpha\delta - \eta/2\right) - \eta\phi Z_u^2 - r\right)\mathrm{d}u} \middle| \mathscr{F}_t\right]\mathrm{d}s \\ &\equiv \Pi(A_t, K_t, N_t)v_P(Z_t); \end{aligned}$$

accordingly, we define the function

$$f(s,Z)=\widetilde{\mathbb{E}}\left[e^{\int_t^{t+s}\left((1-\eta)\left(\tau_A^{1/2}Z_u-\alpha\delta-\eta/2\right)-\eta\phi Z_u^2-r\right)\mathrm{du}}\middle|\mathscr{F}_t\right],$$

which, expressing the dynamics of Z under $\widetilde{\mathbb{P}}$, satisfies the following PDE:

$$f_s=\left((1-\eta)\left(\tau_A^{1/2}Z_u-\alpha\delta-\eta/2\right)-\eta\phi Z_u^2-r\right)f+(\tau_A^{1/2}(1-\eta)+\tau_A/2Z)f_Z+\tau_A/2f_{ZZ},$$

with boundary condition, $f(0,Z)\equiv 1$. Since Z is a Gaussian process, the solution takes the usual exponential affine quadratic form:

$$f(s,Z)=\exp\left(a_0(s)+a_1(s)Z+a_2(s)Z^2\right),$$

with coefficients satisfying the system of Riccati equations:

$$\begin{aligned}
a_0'&=(\eta-1)\left(\alpha\delta+\frac{\eta}{2}\right)+\frac{a_1^2\tau_A}{2}+a_1(1-\eta)\sqrt{\tau_A}+a_2\tau_A-r,\\
a_1'&=a_1(4a_2\tau_A+\tau_A)+(2a_2+1)(1-\eta)\sqrt{\tau_A}\\
a_2'&=-\eta\phi+a_2(2a_2+1)\tau_A,
\end{aligned}$$

with boundary conditions, $a_0(0)=a_1(0)=a_2(0)=0$. These equations have closed-form solutions of the form:

$$a_2(s)=\frac{1}{4}\left(-\frac{\xi\tanh\left(\frac{1}{2}\xi s\sqrt{\tau_A}-\coth^{-1}\left(\frac{\xi}{\sqrt{\tau_A}}\right)\right)}{\sqrt{\tau_A}}-1\right),\tag{40}$$

$$a_1(s)=\frac{(\eta-1)\left(\left(\xi+\sqrt{\tau_A}\right)^2-2\left(\xi^2+\tau_A\right)e^{\frac{1}{2}\xi s\sqrt{\tau_A}}+\left(\xi-\sqrt{\tau_A}\right)^2e^{\xi s\sqrt{\tau_A}}\right)}{\xi\sqrt{\tau_A}\left(\xi+\left(\xi-\sqrt{\tau_A}\right)e^{\xi s\sqrt{\tau_A}}+\sqrt{\tau_A}\right)},$$

$$\begin{aligned}
a_0(s)=&-\frac{4\eta^2+s\tau_A(4\alpha\delta+4r+\tau_A+2)-2\eta(s(2\alpha\delta\tau_A+\tau_A)+4)-\tau_A\log(4)+4}{4\tau_A}\\
&+\frac{(\eta-1)^2(s\tau_A+2)}{2\xi^2\tau_A}+\frac{1}{4}\xi s\sqrt{\tau_A}+\frac{1}{2}\left(\log(\xi)-\log\left(\xi+(\xi-\sqrt{\tau_A})\,e^{\xi s\sqrt{\tau_A}}+\sqrt{\tau_A}\right)\right)\\
&\frac{(\eta-1)^2\left(\left(\xi^4+\xi^2\tau_A+2\tau_A^2\right)\cosh\left(\frac{1}{2}\xi s\sqrt{\tau_A}\right)-2\tau_A\left(\xi^2+\tau_A\right)\right)}{\xi^3\tau_A\left(\xi\cosh\left(\frac{1}{2}\xi s\sqrt{\tau_A}\right)-\sqrt{\tau_A}\sinh\left(\frac{1}{2}\xi s\sqrt{\tau_A}\right)\right)},
\end{aligned}$$

where $\xi\equiv\sqrt{\tau_A+8\eta\phi}$. We conclude that a particular solution to Eq. (39) is as in Eq. (28). Furthermore, the two solutions associated with the homogeneous version of Eq. (39) are:

$$G_1(Z)=e^{g_1Z+g_2Z^2}H_n\left(h_0+h_1Z\right),$$

with $g_1=(\eta-1)\left(\frac{1}{\sqrt{\tau_A}}-\frac{1}{\xi}\right)$, $g_2=-\frac{\sqrt{\tau_A}+\xi}{4\sqrt{\tau_A}}$, $h_0=\frac{\sqrt{2}(\eta-1)\sqrt[4]{\tau_A}}{\xi^{3/2}}$, $h_1=\sqrt{\frac{\xi}{2\sqrt{\tau_A}}}$, and H_n denotes the

Hermite polynomial of order n (as defined in Eq. (29)) and

$$G_2(Z) = e^{g_1 Z + g_2 Z^2} M\left(-\frac{1}{2}n, \frac{1}{2}, m_0 + m_1 Z + m_2 Z^2\right),$$

with $m_0 = \frac{2(\eta-1)^2\sqrt{\tau_A}}{\xi^3}$, $m_1 = \frac{2(\eta-1)}{\xi}$ and $m_2 = \frac{\xi}{2\sqrt{\tau_A}}$, where M denotes the Kummer confluent hypergeometric function. We conclude that, in the inaction region, the general solution to Eq. (39) takes the form in Eq. (27). Note that the existence of this solution requires conditions on parameters.

B.4.3 Solution method

To solve the model numerically, we start by imposing that the firm value be continuous over each region involved in its piecewise representation in Eq. (32). In particular, when the firm invests there can no jump in its value on average;

$$\begin{aligned} g(\underline{b}; \underline{C}_1, \underline{C}_2) &= V(\underline{b}), \\ g(\bar{b}; \overline{C}_1, \overline{C}_2) &= V(\bar{b}). \end{aligned} \tag{41}$$

Similarly, there can be no jump in the firm value when it decides to explore:

$$\begin{aligned} g(\underline{a}; \underline{C}_1, \underline{C}_2) &= V(0)(1-\delta)^{\alpha(1-\eta)}, \\ g(\overline{a}; \overline{C}_1, \overline{C}_2) &= V(0)(1-\delta)^{\alpha(1-\eta)}. \end{aligned} \tag{42}$$

Eqs. (41)–(42) constitute a system of four equations that determines the four unknown constant of integration, $\underline{C}_1$, $\underline{C}_2$, $\overline{C}_1$, and $\overline{C}_2$.

A difficulty in solving this system of equations is that the functional form of the value function, $V(\cdot)$, is unknown. There are two alternative routes in handling this problem numerically. One is to use the representation of the value function as infinite sum in Eq. (34). Another, more appropriate way is to use Gauss-Legendre quadrature (see, e.g., Judd (1998)) to obtain an explicit representation of the value function. This latter approach is the one we adopt. In particular, we approximate the integral involved in the Bellman equation (26) as

$$\int_{\mathscr{B}} V(y)\phi(y; i(z), z)\mathrm{d}y \approx \sum_{k=1}^{n} \omega_k \phi(x_k; i(z), z) V(x_k),$$

where x_k are the quadrature nodes and ω_k are the quadrature weights over the interval $\mathscr{B}$. We then collect the value function evaluated at each quadrature node in a vector, which we denote by $\mathbf{V} \equiv \{V(x_j)\}_{j=1}^{n}$. The goal is now to solve for this vector explicitly. To do so, we plug the integral approximation back into the Bellman equation and evaluate it at the quadrature nodes:

$$\begin{aligned} V(z_j) = {} & (1+i(z_j))^{\alpha(1-\eta)} \int_{\mathbb{R}} (h(y) - \mathbf{1}_{y\in\mathscr{B}}(\kappa + \gamma/2 i(y)^2))\phi(y; i(z_j), z_j)\mathrm{d}y \\ & - \kappa - \gamma/2 i(z_j)^2 + (1+i(z_j))^{\alpha(1-\eta)} \sum_{k=1}^{n} \omega_k \phi(x_k; i(z_j), z_j) V(x_k), \quad j = 1,\dots,n, \end{aligned} \tag{43}$$

which yields a system of n equations for the n unknown elements composing the vector $\mathbf{V}$.

We then define the following (column) vectors:

$$\begin{aligned} \mathbf{\Phi} &\equiv \Big((1+i(z_j))^{\alpha(1-\eta)}(\overline{v}_P(z_j, \underline{a}, \underline{b}) + \overline{v}_P(z_j, \overline{b}, \overline{a})) - \kappa - \gamma/2 i(z_j)^2\Big)_{j=1,\dots,n}, \\ \underline{\mathbf{G}}_k &\equiv \Big((1+i(z_j))^{\alpha(1-\eta)}(\overline{G}_k(z_j, \underline{a}, \underline{b})\Big)_{j=1,\dots,n}, \quad k = 1,2, \\ \overline{\mathbf{G}}_k &\equiv \Big((1+i(z_j))^{\alpha(1-\eta)}(\overline{G}_k(z_j, \overline{b}, \overline{a})\Big)_{j=1,\dots,n}, \quad k = 1,2, \\ \mathbf{A} &\equiv \left((1-\omega)^{\alpha(1-\eta)} \int_{y\notin(\underline{a},\overline{a})} \phi(y; i(z_j), z_j)\mathrm{d}y\right)_{j=1,\dots,n}, \end{aligned}$$

and the following $n \times n$ matrix:

$$\mathbf{\Sigma} \equiv \Big((1+i(z_j))^{\alpha(1-\eta)}\omega_i \phi(x_i; i(z_j), z_j)\Big)_{i,j=1,\dots,n}.$$

We can then rewrite the system in Eq. (43) in vector form and express the vector $\mathbf{V}$ as:

$$\mathbf{V} = (\mathbf{I} - \mathbf{\Sigma})^{-1}\left(\mathbf{\Phi} + \sum_{k=1,2} \underline{C}_k \underline{\mathbf{G}}_k + \sum_{k=1,2} \overline{C}_k \overline{\mathbf{G}}_k + \mathbf{A}V(0)\right).$$

By analogy to the vectors we defined above, we further define the following functions:

$$\begin{aligned} \Phi(Z) &\equiv (1+i(Z))^{\alpha(1-\eta)}(\overline{v}_P(Z, \underline{a}, \underline{b}) + \overline{v}_P(Z, \overline{b}, \overline{a})) - \kappa - \gamma/2 i(Z)^2, \\ \underline{G}_k(Z) &\equiv (1+i(Z))^{\alpha(1-\eta)}(\overline{G}_k(Z, \underline{a}, \underline{b}), \quad k = 1,2, \\ \overline{G}_k(Z) &\equiv (1+i(Z))^{\alpha(1-\eta)}(\overline{G}_k(Z, \overline{b}, \overline{a}), \quad k = 1,2, \\ A(Z) &\equiv (1-\omega)^{\alpha(1-\eta)} \int_{y\notin(\underline{a},\overline{a})} \phi(y; i(Z), Z)\mathrm{d}y, \end{aligned}$$

and the following vector:

$$\Sigma(Z) \equiv \Big((1+i(Z))^{\alpha(1-\eta)}\omega_i \phi(x_i; i(Z), Z)\Big)_{i=1,\dots,n}.$$

We use these functions to write the Nystrom extension (see, e.g., Judd (1998)) to the value function as:

$$V(Z) \approx \Phi(Z) + \sum_{k=1,2} \underline{C}_k \underline{G}_k(Z) + \sum_{k=1,2} \overline{C}_k \overline{G}_k(Z) + A(Z)V(0) +$$

$$\Sigma(Z)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \left(\mathbf{\Phi} + \sum_{k=1,2} \underline{C}_k \underline{\mathbf{G}}_k + \sum_{k=1,2} \overline{C}_k \overline{\mathbf{G}}_k + \mathbf{A} V(0) \right).$$

Evaluating the Nystrom extension at 0 yields:

$$V(0) = \frac{1}{1 - A(0) - \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{A}} \left(\begin{array}{c} \Phi(0) + \sum_{k=1,2} \underline{C}_k \underline{G}_k(0) + \sum_{k=1,2} \overline{C}_k \overline{G}_k(0) \\ + \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \left(\mathbf{\Phi} + \sum_{k=1,2} \underline{C}_k \underline{\mathbf{G}}_k + \sum_{k=1,2} \overline{C}_k \overline{\mathbf{G}}_k \right) \end{array} \right).$$

Plugging back in the Nystrom extension and defining:

$$\phi_0(Z) \equiv \Phi(Z) + \Sigma(Z)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{\Phi} + \frac{A(Z) + \Sigma(Z)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{A}}{1 - A(0) - \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{A}} \left(\Phi(0) + \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{\Phi} \right),$$

$$\underline{g}_k(Z) \equiv \underline{G}_k(Z) + \Sigma(Z)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \underline{\mathbf{G}}_k + \frac{A(Z) + \Sigma(Z)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{A}}{1 - A(0) - \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{A}} \left(\underline{G}_k(0) + \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \underline{\mathbf{G}}_k \right),$$

$$\overline{g}_k(Z) \equiv \overline{G}_k(Z) + \Sigma(Z)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \overline{\mathbf{G}}_k + \frac{A(Z) + \Sigma(Z)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{A}}{1 - A(0) - \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \mathbf{A}} \left(\overline{G}_k(0) + \Sigma(0)' (\mathbf{I} - \mathbf{\Sigma})^{-1} \overline{\mathbf{G}}_k \right),$$

provides an approximate but explicit representation of the value function in terms of the four unknown integration constants:

$$V(Z) \approx \phi_0(Z) + \sum_{k=1,2} \underline{C}_k \underline{g}_k(Z) + \sum_{k=1,2} \overline{C}_k \overline{g}_k(Z).$$

Using this representation, and letting $\widehat{\omega} \equiv (1 - \omega)^{\alpha(1-\eta)}$ we define the following matrix:

$$B = \left(\begin{array}{cccc} G_1(\underline{a}) - \widehat{\omega} \underline{g}_1(0) & G_2(\underline{a}) - \widehat{\omega} \underline{g}_2(0) & -\widehat{\omega} \overline{g}_2(0) & -\widehat{\omega} \overline{g}_2(0) \\ -\widehat{\omega} \underline{g}_1(0) & -\widehat{\omega} \underline{g}_2(0) & G_1(\overline{a}) - \widehat{\omega} \overline{g}_2(0) & G_2(\overline{a}) - \widehat{\omega} \overline{g}_2(0) \\ G_1(\underline{b}) - \underline{g}_1(\underline{b}) & G_2(\underline{b}) - \underline{g}_2(\underline{b}) & -\overline{g}_2(\underline{b}) & -\overline{g}_2(\underline{b}) \\ -\underline{g}_1(\overline{b}) & -\underline{g}_2(\overline{b}) & G_1(\overline{b}) - \overline{g}_2(\overline{b}) & G_2(\overline{b}) - \overline{g}_2(\overline{b}) \end{array} \right).$$

we can in turn write the solution of the system in Eqs. (41)–(42) explicitly as:

$$\begin{pmatrix} \underline{C}_1 \\ \underline{C}_2 \\ \overline{C}_1 \\ \overline{C}_2 \end{pmatrix} = B^{-1} \begin{pmatrix} \phi_0(0)\widehat{\omega} - v_P(\underline{a}) \\ \phi_0(0)\widehat{\omega} - v_P(\overline{a}) \\ \phi_0(\underline{b}) - v_P(\underline{b}) \\ \phi_0(\overline{b}) - v_P(\overline{b}) \end{pmatrix}.$$

This solution is unique if and only if the determinant of $\det(B) \neq 0$.

The solution we constructed holds for given thresholds $\underline{a}$, $\overline{a}$, $\underline{b}$, and $\overline{b}$, and a given knowledge-contingent investment policy $i(Z)$. The last step is to pick these elements optimally. To solve for the optimal thresholds we use the system of smooth-pasting condition, Eq. (13) (the first two equation for optimal exploration choice) and Eq. (30), which provides four equations for the four unknowns, $\underline{a}$, $\overline{a}$, $\underline{b}$, and $\overline{b}$. The optimal investment policy further solves the first-order condition in Eq. (36). To find a solution, we use a Newton algorithm to solve for Eqs. (13)–(30), and at each Newton step we perform Policy Function Iteration to obtain the optimal investment policy.