Knowledge Cycles and Corporate Investment*

M. Cecilia Bustamante† Julien Cujean‡ Laurent Frésard§

May 28, 2020

Abstract

Investment is a means to experiment with technologies, creating valuable knowledge for firms. We analyze the feedback interaction between investment and knowledge through endogenous cycles of experimentation and exploration. Because experimentation is uncertain, firms become averse to investing when knowledge is scarce and q is high, overturning neoclassical insights. Investment is increasing in knowledge and its relation with q varies over the knowledge cycle. Experimentation further shortens cycle length. Using a text-based proxy for firms' knowledge cycles, we provide evidence supporting these novel predictions. Shorter knowledge cycles could explain weak aggregate investment despite high valuations in recent years.

Keywords: Investment, Experimentation, Exploration, Knowledge, Experimentation Aversion, Intangibles, Knowledge Capital.

*We are grateful to Daniel Andrei, Serguey Braguinsky, Jerry Hoberg, Johan Hombert, Julien Hugonnier, Nobuhiro Kiyotaki, Mark Lowenstein, Albert Kyle, William Mann, Rich Matthews, Erwan Morellec, Philip Valta, Laura Veldkamp and Alexei Zhdanov 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, and the ASU Sonoran Winter Finance Conference 2020. We also 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

Knowledge capital—employees' expertise, proprietary technologies, internal data, or organizational design—has become increasingly central to firms' growth and success. Recent research indicates that knowledge plays an important role in describing firms' investment decisions (e.g., Peters and Taylor (2017), Crouzet and Eberly (2018), Andrei, Mann, and Moyen (2018), and Crouzet and Eberly (2020)). Yet, existing research has been surprisingly silent on the role that investment itself plays in the process of knowledge creation within firms. In this paper, we show that considering the feedback interaction between knowledge and investment delivers several novel insights into the economic determinants of investment in a knowledge-based economy.

Building on the long-standing idea that knowledge arises as a by-product of economic activity (e.g., Arrow (1962)), we argue that investment produces knowledge because it is a means to experiment with new technologies. Our starting point is to recognize that, unlike capital and labor, knowledge is not a typical factor of production that firms can control precisely. Rather, knowledge is intangible and is generated as firms learn about the quality of their technologies (e.g., Pastor and Veronesi (2009)).[1] Thus, knowledge accumulates slowly over time, and often serendipitously through an internal process of uncertain trial and error (e.g., Callander (2011)). Knowledge can also dissipate quickly due to the well-known difficulty to protect it (e.g., Romer (1986)). To capture these unique dimensions, we consider that firms' knowledge follows endogenous cycles, accumulating as firms actively experiment with their existing technologies, and contracting when firms decide to explore new unknown technologies.

To study the interactions between such knowledge cycles and firms' investment decisions, we consider a firm that operates a given technology whose intrinsic quality is unobserved. The firm combines this technology with capital to produce an output, and accumulates information about its quality by observing productivity realizations. We formally define "knowledge" as the firm's confidence (in a statistical sense) about the quality of its technology. More knowledge is valuable since it improves the firm's decision. Besides passive

[1] Our definition of "technology" is generic and could encompass various forms, such as product design and features, a production process, an organizational design (e.g., decentralized authority, see Prescott and Visscher (1980)), or a management structure (e.g., lean management, see Bloom, Sadun, and van Reenen (2016)).

learning, the firm can increase its knowledge by actively experimenting with its current technology through multiple rounds of trial and error until the experiments become conclusive (i.e., produces information). Similar to Arrow (1962) and Farboodi, Mihet, Philippon, and Veldkamp (2019), experimentation takes the form of capital investment (e.g., expanding a research center or equipment with similar technological features). Experiments are costly and their outcome is uncertain, although they are informative on average. In this setting, investment affects firm value by directly altering output (the traditional neoclassical channel), but also by accelerating knowledge accumulation through noisy experimentation. We label this second effect the "knowledge channel".[2]

At any time, the firm can decide to explore new technologies and abandon its current technology. When it does, the quality of the new selected technology is unknown (e.g., Jovanovic and Nyarko (1996)). Switching technologies renders part of the existing stock of capital obsolete (e.g., Bresnahan, Greenstein, Brownstone, and Flamm (1996)) and resets knowledge. The firm thus explores when it becomes sufficiently confident that the quality of its existing technology is poor, but also when it is likely high, since knowledge dissipation erodes future profits emanating from high quality technologies (e.g., Lucas (1988), and Jovanovic and MacDonald (1994)). The trade-off between exploration and experimentation creates endogenous knowledge cycles, within which the firm's stock of knowledge evolves through passive observation and active experimentation, until the firm decides to explore and start a new cycle.

Our central result is that, when a firm invests, its investment is positively related to knowledge. To put this result in perspective and understand its implications, we shut down the knowledge channel by forcing experiments to be uninformative. This benchmark model delivers the insight from modern $q-$theory (e.g., Hayashi (1982)), namely that investment is increasing in firm value (i.e., marginal q). In that benchmark, when the firm invests, its value is inversely related with knowledge, so that the firm invests more when it has limited knowledge (i.e., when the marginal return to knowledge is high). It follows that, when it occurs, investment is *decreasing* in knowledge in the benchmark case. The knowledge channel, however, turns this neoclassical insight on its head.

When investment is a means to experiment, it follows a specific pattern along the knowl-

[2] In the model, the economic function of the knowledge channel operates as an information acquisition problem. See, for instance, van Nieuwerburgh and Veldkamp (2010) or Hellwig, Kohls, and Veldkamp (2012).

edge cycle. In particular, experimentation takes place at early stages of the cycle and ceases altogether as the cycle matures. Experimentation arises as a way to reap the marginal value of capital and of knowledge, both of which are highest following exploration. However, following exploration the quality of a new technology is highly uncertain, and thus experimentation is particularly risky (e.g., Pastor and Veronesi (2009)). As a result, experimentation is prudent at early stages and becomes bolder as the firm accumulates knowledge.

The driving force behind this result is that, although the firm is risk-neutral when it maximizes its value by choosing investment (whether or not to experiment) and its technology (whether or not to explore), it can become endogenously "experimentation averse." Because experimentation is uncertain, it affects firm value only after the firm has made its investment decision. Therefore, depending on the shape of the relation between firm value and knowledge, experimentation is either a gamble or a risk the firm would prefer to avoid. Intuitively, following exploration the relation between firm value and knowledge is concave. Thus, the firm is experimentation averse and invests by "tatonnement." As knowledge accumulates along the cycle, firm value becomes convex and experimentation becomes bolder as a way to gamble on future exploration. It follows that the firm invests less *despite* high marginal q when knowledge is limited, and raises investment despite lower q as knowledge accumulates, leading to an increasing relation between investment and knowledge.

The knowledge channel thus alters the traditional, and widely investigated, relationship between investment and q. In the model, q is given by the marginal value of one additional unit of capital averaged across all possible experimentation outcomes, which makes it unobservable to empiricists interested in testing the knowledge channel. We show, however, that the relationship between investment and average Q (i.e., the average product of capital), a quantity observable to empiricists and justified theoretically in our context (e.g., Caballero and Leahy (1996)), varies predictably over the knowledge cycle. This implication offers a way of assessing the relevance of the knowledge channel in the data. Using simulated series of investment and Q (and a baseline calibration), we estimate a positive and strong relation between investment and Q at early stages of the knowledge cycle (when knowledge is still scarce), but an insignificant relation later in the cycle. The positive investment-Q relation also materializes early in the cycle when knowledge dissipation is weak (e.g., when intellectual property is better protected). In contrast, when experiments yield better information (i.e., are less noisy), the positive investment-Q relation concentrates later in the cycle.

To provide supporting evidence for these novel predictions, we follow Abernathy and Utterback (1978), Klepper (1996) and Callander (2011) and proxy for firms knowledge cycles using product life cycles. In particular, we use the text-based measure of product life cycle recently developed by Hoberg and Maksimovic (2019). This measure exploits the text from public firms' annual reports between 1996 and 2016 to assign firm-year observations into four distinct product stages (i.e., product innovation, process innovation, maturity and decline). We first document an opposite pattern between investment and Q over the product cycle, in line with the distorting role of the knowledge channel. Furthermore, consistent with the model's novel prediction, the sensitivity of investment to Q varies considerably over the cycle, and is strongest in the second and third stage, at times when accumulated knowledge pushes firms to experiment by investing.

Another novel aspect of the model is that it produces endogenous knowledge cycles, and thus speaks to the drivers of experimentation and exploration within the firm. In particular, the presence of the knowledge channel leads to shorter knowledge cycles that are distributed more asymmetrically (i.e., right skewed) compared to those generated absent this channel. Furthermore, we show that weak knowledge dissipation lengthens knowledge cycles, because better protection of intellectual property reduces the firm's incentive to explore new technologies. In contrast, more informative experimentation yields shorter knowledge cycles, since it accelerates learning and renders exploration more frequent. Remarkably, the model-implied distribution of knowledge cycle length closely maps (in magnitude and shape) that observed in the data, measured as the time it takes firms that move out of the first stage of their current product cycle and fall back into the first stage of a new cycle. The model also characterizes the number of successive trials in each experimentation round, and whether rounds occur intermittently along the knowledge cycle. Under our baseline calibration, it most frequently takes a single round with three successive trials until experimentation becomes conclusive and the firm eventually explores a new technology.

Finally, we use the model as a laboratory to examine whether the knowledge channel could explain the weak aggregate investment observed in the last twenty years, its increasing disconnection with Q (e.g., Gutierrez and Philippon (2017), Alexander and Eberly (2018), and Crouzet and Eberly (2018)), and the role of intangibles (e.g., Crouzet and Eberly (2020)). In light of our model, these trends could stem from an aggregate shortening of firms' knowledge cycles, implying an increasing wedge between observed investment and Q. Indeed, weak investment despite high Q could arise if firms explore more often, thereby resetting their knowledge more frequently. Looking at the aggregate evolution of firms' position in their product cycles, we observe a clear trend towards shorter cycles in the data, mirroring the results of Hoberg and Maksimovic (2019). Using the insights from the model, we speculate that such an aggregate shortening of firms' knowledge cycles could be the result of more informative experimentation (e.g., based on computer simulations or rapid prototyping) or heightened knowledge dissipation (e.g., weaker intellectual property rights, or greater workers' mobility).

The paper adds to a growing literature analyzing the role of intangible assets in firms' investment decisions. Most recent papers (e.g., Eisfeldt and Papanikolau (2013), Peters and Taylor (2017), Crouzet and Eberly (2018) or Crouzet and Eberly (2020)) study the (one-way) impact of intangibles on investment decisions by introducing intangible capital as an extra factor of production in neoclassical investment models. Closer to our paper, Andrei et al. (2018) develop an investment model in which the firm learns passively and technological innovation occurs randomly. Distinct from existing research, we consider that investment in itself can create knowledge, and show that accounting for the feedback interaction between investment and knowledge delivers important, and non-trivial, implications for the determinants of firms' investment decisions.

Our analysis adds to the literature in finance that explicitly models the process of knowledge creation within firms (e.g., Berk, Green, and Naik (2004), Bergemann and Hege (2005), Pastor and Veronesi (2009), Manso (2011), or Manso, Balsmeier, and Fleming (2019)).[3] While we share the premise that knowledge arises from an internal trade-off between experimentation and exploration, our model generates endogenous (as opposed to pre-determined) knowledge cycles. This feature enables us to fully characterize their start, intensity, and length, and generate novel predictions regarding the determinants of firms' experimentation and exploration strategies (some of which we confront to the data). Distinct from existing papers, we study the implications of firms' using investment as a means to experiment.

The paper also belongs to a large literature on experimentation, and is closest technically

[3]These papers focus on questions related to valuation (Berk et al. (2004)), financing (Bergemann and Hege (2005)), asset price dynamics (Pastor and Veronesi (2009)), contracting (Manso (2011)), and innovation strategies (Manso et al. (2019)).

to Moscarini and Smith (2001).[4] Similar to their decision maker, our firm decides how much to experiment (i.e., invest) and when to take a payoff-relevant action (permanent stop in their case, and exploration in ours). Similar to their experimentation intensity, investment is increasing in knowledge and is highest when the action is about to be taken. A relevant difference is in what drives this pattern in both papers. Because the cost of delaying action is increasing in value, in their case experimentation must grow in value to raise the surplus from information production. In contrast, such surplus is uncertain in our paper and creates information aversion, which dictates higher investment despite lower value. Another relevant difference relates to the formulation and scope. They focus on a problem of sequential analysis (Wald, 1947), whereas we study investment in a neoclassical context.

2 Model setup

2.1 Knowledge, exploration and experimentation

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 processes, but also management practices (e.g., Bloom et al. (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, or to abandon it and 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

[4] See, for instance, Grossman, Kihlstrom, and Mirman (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).

reflects its quality.[5] 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. For simplicity, we focus on the case in which each technology is completely disruptive and resets knowledge completely, or equivalently the firm "learns and forgets" (e.g., Benkard (2000)). Upon selecting technology $n+1$, its quality is randomly drawn according to:

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

so that the technology index n can be discarded, and the parameter τ_M captures prior precision regarding the quality of the technology.[6] For later use, we denote by ν dates at which the firm decides to abandon the current technology and explore an unknown technology. Furthermore, as in Jovanovic and Nyarko (1996), we assume that once a technology has been abandoned it can never be "recalled".

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 noise and 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 variance, Ω, 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,

[5] 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.

[6] 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 in the main text.

and more importantly, learning affects directly the growth rate of productivity, A, through the conditional variance, 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. This effect is what the literature commonly refers to as "learning by doing", namely "*the role of* [learning from] *experience in increasing productivity*" (Arrow, 1962).

In addition to observing realizations of productivity passively, the firm can also learn actively about M through experimentation (e.g., Grossman et al. (1977)). Following Arrow (1962), experimentation takes the form of an expansion of capital on which the current technology is applied (see also Farboodi et al. (2019)). Consistent with this idea of "learning by investing", Braguinsky, Ohyama, Okazaki, and Syverson (2019) show that expanding the set of machines used for production enlarges the overall productivity of the machines and thus may refine the firm's information about its current technology.[7] Formally, let $I_t \geq 0$ represent investment in capital at time t and $i_t = I_t/K_{t-}$ be the corresponding investment rate. By investing (i.e., $i_t > 0$) the firm receives an informative signal flow, $\mathrm{d}S$, about the quality of the current technology:

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

Investment improves the precision of this signal, so that $\tau_S i$ captures the informativeness of experimentation. The more the firm invests relative to its current stock of capital, the more it learns about M. However, when the firm experiments through capacity expansion, the outcome of its experiments is uncertain (e.g., Callander (2011)). We capture this genuine feature of experimentation by introducing noise, ϵ, in $\mathrm{d}S$. Indeed adding new machines may sometimes harm the firm's current prior about M (e.g., if the new machines provide "contrarian" results). Similar to the passive productivity signal, informativeness of $\mathrm{d}S$ is measured in time-invariant units of uncertainty. To simplify intuition, we rule out learning through capital liquidation by making investment strictly irreversible.

The firm's information about the quality of the technology it operates entirely results

[7]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".

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:

$$\mathrm{d}\widehat{M}_t = \Omega_t^{1/2}\tau_A^{1/2}\mathrm{d}\widehat{B}_t + \Omega_{t-}^{1/2}\left(\frac{\tau_S i_t}{1+\tau_S i_t}\right)^{1/2}\widehat{\epsilon}_t\mathbf{1}_{i_t>0} - \widehat{M}_{t-}\mathbf{1}_{t=\nu}$$
$$\mathrm{d}\Omega_t = -\Omega_t\tau_A\mathrm{d}t - \Omega_{t-}\frac{\tau_S i_t}{1+\tau_S i_t}\mathbf{1}_{i_t>0} + (\tau_M^{-1} - \Omega_{t-})\mathbf{1}_{t=\nu},$$

where $\widehat{B}$ is a Brownian motion and $\widehat{\epsilon}$ is Gaussian with mean zero and unit variance under the firm's probability measure.[8] In short, the firm updates its estimates about the quality of its current technology continuously by observing realized productivity, punctually by experimenting with it, and resets its estimate to 0 and its prior precision to τ_M (i.e., its prior beliefs) whenever it decides to explore another technology.

The firm's stock of knowledge, Z_t, represents the information it has at date t about the technology it operates. In the context of 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 the firm's stock of knowledge:

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

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

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

experiment or explore.

2.2 The decision of the firm

Suppose the firm faces an isoelastic demand curve for the good (or service) it produces, $P_t = (Y_t N_t)^{-\eta}$, where $\eta \in (0,1)$ corresponds to the price elasticity of demand. It follows that the firm's revenues satisfy:

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

Revenues are increasing in the productivity of the current technology, A, but are decreasing in $(N_t)_{t\geq 0}$, an exogenous process capturing the dissipation of the firm's knowledge in reduced form. A large literature indicates that because knowledge is non-rival and largely embodied in (inalienable) workers, it is difficult for the firm to fully appropriate its returns. Proprietary knowledge thus typically dissipates over time (e.g., Romer (1986) or Lucas (1988)), so that larger values of N (i.e., the dissipation intensity) erodes revenues. 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 (i.e., corporate spying or thefts).[9]

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

$$\mathrm{d}N_t / N_t = \phi Z_t^2 \mathrm{d}t, \tag{4}$$

with $\phi > 0$.[10] We thus assume that the intensity with which the firm's knowledge dissipates increases as knowledge becomes stronger. The firm only has incentives to explore new technologies when the dollar benefits of knowledge are *bounded*. Therefore, it cannot be optimal

[9]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)).

[10]We assume that the growth rate of N is quadratic in Z, because the resulting affine-quadratic framework remains tractable (see Proposition 1). However, any *strictly* convex increase in N (even less convex than quadratic), will keep profits bounded and produce qualitatively similar results, except in terms of the asymmetry in the firm's choice between good and bad technologies.

for the firm to let knowledge grow unboundedly by sticking to the same technology forever, as the variable N would eventually erode its operating revenues to zero (e.g., Jovanovic and MacDonald (1994)). As a convention, we interpret N as a natural dissipation process that keeps knowledge stationary in the model.

The firm maximizes the expected value of its future profits (revenues net of costs) by optimally choosing its investment and its technology. Every time the firm decides to invest—times we denote by θ—it increases its stock of capital and knowledge. However, it also 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)):

$$\gamma_t \equiv \left(\kappa \mathbf{1}_{i>0} + \gamma/2i^2\right) \cdot \Pi\left(A_t, K_t, N_t\right) \equiv \widehat{\gamma}(i)\Pi\left(A_t, K_t, N_t\right), \tag{5}$$

where κ represents a fixed cost and the variable cost is convex (e.g., Abel and Eberly (1994)).[11] Similarly, abandoning the current technology to explore is costly, as it implies that a fraction ω of the firm's capital K becomes obsolete upon exploration. We further assume that capital depreciates at a fixed proportional rate, δ. It follows that the net change in the firm's stock of capital is:

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

Given the costs and benefits associated with exploration and experimentation, and assuming the firm is risk neutral and discounts its profits at a constant rate, r, its value, $V(\cdot)$, is given by:

$$\max_{\{\nu_n\}, i_{t+s}} \mathbb{E}\left[\int_0^\infty e^{-rs}\Pi\left(A_{t+s}, K_{t+s}, N_{t+s}\right)\mathrm{d}s - \sum_{n\geq 0}\gamma_{t+\theta_n}e^{-r\theta_n}\middle| \mathscr{F}_t\right]. \tag{6}$$

Our assumptions further imply that the firm value has the convenient representation:

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

where $v(\cdot)$ denotes the intensive value of the firm. This representation confirms that the

[11] Because the firm value (defined later) exhibits concave and convex regions, variable costs are needed to keep the investment rate finite in the convex regions.

firm's knowledge, Z, is the only state variable that determines the choice between exploration and experimentation.

3 A benchmark model without experimentation

To understand the key role played by experimentation in our model, we first construct a neoclassical benchmark in which experimentation through capital expansion is uninformative, $\tau_S \equiv 0$, which eliminates the effect of "learning-by-investing" on knowledge. When the firm does not learn actively from its investments, investment is determined by the traditional neoclassical trade-off (e.g., Hayashi (1982)).

To compare results between the benchmark and the model with experimentation, we assume that adjustment costs in the benchmark model are paid *per unit of time*. Specifically, investing over a period of time $[t, t + \mathrm{d}t)$ entails a cost $\gamma_t \mathrm{d}t$, where γ_t is defined in Eq. (5). Under this assumption, fixed investment costs create inaction regions but investment occurs continuously within these regions (e.g., Abel and Eberly (1994)). In contrast, in the full model with experimentation, we do not make this assumption so that investment occurs in lumps. In the model with experimentation, however, the effect of investment on firm value is uncertain due to experimentation noise, which makes it necessary for the firm to determine an optimal investment policy contingent on all possible experimentation outcomes. Hence, although investment is lumpy, the firm's "knowledge-contingent" investment plan is effectively continuous across all experimentation outcomes.

In the benchmark model, the problem of the firm simplifies to:

$$\begin{aligned} &\max_{\{\nu_n\}, i_{t+s}} \mathbb{E}\left[\int_0^\infty e^{-rs}\left(1-\widehat{\gamma}_{t+s}\right)\Pi\left(A_{t+s}, K_{t+s}, N_{t+s}\right)\mathrm{d}s\right] \\ \text{s.t. } &\mathrm{d}K_t = K_t((i_t-\delta)\mathrm{d}t - \omega \mathbf{1}_{t=\nu}) \\ &\mathrm{d}Z_t = \tau_A/2 Z_t \mathrm{d}t + \tau_A^{1/2}\mathrm{d}\widehat{B}_t - Z_{t-}\mathbf{1}_{t=\nu}. \end{aligned} \tag{8}$$

Because adjustment costs are paid per unit of time, investment (when it occurs) takes place continuously at rate i_t per unit of time $\mathrm{d}t$. The evolution of capital, K, and net revenues now reflect this aspect. Furthermore, because experiments are uninformative, investment does not affect knowledge; the dynamics of Z are modified accordingly. The dynamics of productivity, A, and knowledge dissipation, N, are unchanged relative to Eqs. (1) and (4).

The main insight from the neoclassical literature is that the marginal product of capital, hereafter marginal q, is a sufficient statistic for investment (e.g., Hayashi (1982)). We introduce the definition of marginal q that is relevant in our framework (benchmark and model with experimentation), and which is based on Abel and Blanchard (1986).

Definition 1. ***(Marginal Product of Capital)*** *Marginal q is the expected present value of a marginal unit of investment (Abel and Blanchard, 1986) scaled by relative profits, Π/K:*

$$q(Z) \equiv \int_{\mathbb{R}} \frac{V_K(\cdot, K+I, x)}{\Pi(\cdot)/K} \varphi(x; i(Z), Z)\, dx,$$

where $\varphi(\cdot; i, Z)$ denotes the distribution of experimentation outcomes for a given investment rate i and given a current level of knowledge, Z; it is defined explicitly in Section 4.

Note that, in the context of the benchmark model, this definition is more general than needed. Because we make experimentation uninformative, the effect of investment on firm value is known ex-ante. Thus, the distribution of experimentation outcomes is degenerate, $\varphi(x; i, Z) \equiv \delta_{x=Z}$, and we can drop the expectation in Definition 1. Furthermore, because investment over a period $\mathrm{d}t$ is infinitesimal, marginal q simplifies to:

$$q(Z) \equiv \frac{V_K(\cdot, K, x)}{\Pi(\cdot)/K} = \alpha(1-\eta)v(Z) \propto v(Z). \tag{9}$$

In the model with experimentation, this definition can be expressed in terms of intensive firm value as:

$$q(Z) = \alpha(1-\eta)(1+i(Z))^{\alpha(1-\eta)-1} \int_{\mathbb{R}} v(x)\varphi(x; i(Z), Z)\mathrm{d}x.$$

A (cosmetic) difference of Eq. 9 relative to Abel and Blanchard (1986) is on the relation between intensive firm value and the marginal product of capital in Definition 1. In their model and, more broadly, in the neoclassical investment literature, firm value is often assumed to be homogeneous in K so that the intensive firm value represents marginal q. Our specification of adjustment costs, however, considers that firm value is homogeneous in profits, Π, as opposed to K. We thus scale firm value by relative profits, Π/K. Moreover, because the effect of investment on knowledge is uncertain, marginal q in Definition 1 is the value of an additional unit of capital *averaged* over all possible experimentation outcomes.

The firm's problem in Eq. (8) is to choose its investment policy, as well as when to switch technology and explore. The decision to explore is uncommon in the neoclassical literature, and yet a relevant feature in our benchmark model. As we explain below, knowledge accumulates in cycles and depreciates at once when the firm explores. In the absence of experimentation through investment, knowledge accumulation within cycles is not materially different from capital accumulation in traditional neoclassical models; what makes its dynamics distinct is the necessity for the firm to explore. When a new technology is adopted, the firm learns passively about its productive potential. After some time, however, the learning benefits on this technology will decline, and so will the firm's profits. Exploration is necessary for the firm to regenerate future profits.

To determine the 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. (8) satisfies the Hamilton-Jacobi-Bellman (hereafter, HJB) equation:

$$\begin{aligned} &v(Z)\left((1-\eta)\left(\alpha\delta+\eta/2\right)-(1-\eta)\sqrt{\tau_A}Z+r+\eta\phi Z^2\right) \\ &\quad =1+\left((1-\eta)\sqrt{\tau_A}+\tau_A/2Z\right)v'(Z)+\tau_A/2v''(Z)+\psi(q(Z)), \end{aligned} \tag{10}$$

where we have defined ψ to be the maximized value of the maximand:

$$\psi(q)\equiv\max_{i\geq 0}\left\{qi-\gamma/2i^2-\kappa\mathbf{1}_{i>0}\right\}. \tag{11}$$

The function ψ is convex and increasing in q, and it is only optimal for the firm to invest when $\psi(q)\geq 0$. It follows that the firm's optimal investment policy is given by:

$$i_t\equiv i(q)=\frac{1}{\gamma}q\mathbf{1}_{q\geq\sqrt{2\gamma\kappa}},$$

thus delivering the traditional insight that q is a sufficient statistic for investment.

The solution assumes so far that it is optimal for the firm not to switch technology when $q\geq\sqrt{2\gamma\kappa}$. We thus need to determine the condition under which this assumption is verified. When the firm decides to explore by selecting a new technology, it resets its knowledge to priors and incurs a lump-sum cost of obsolescence, ωK, on its existing capital stock.

Considering that q in Eq. (9) is proportional to firm value, marginal q upon exploration is:

$$q^{\star} \equiv (1-\omega)^{\alpha(1-\eta)} q\,(0)\,, \tag{12}$$

so that the firm switches technology whenever $q \leq q^{\star}$; this also means that our assumption is verified if $q^{\star} < \sqrt{2\gamma\kappa}$, or equivalently if:

$$q\,(0) < \sqrt{2\gamma\kappa}/(1-\omega)^{\alpha(1-\eta)}. \tag{13}$$

Unfortunately, this is only an implicit condition, because exploration resets knowledge and q is itself a function of knowledge. We solve for firm value in Eq. (10) to determine whether the condition in Eq. (13) is satisfied.

Solving Eq. (10) requires some conjecture on the dependence of firm value on knowledge. Intuitively, as Z becomes increasingly negative the firm becomes increasingly confident that its technology is poor, and its value must decrease as Z further decreases; instead, as the firm becomes increasingly confident that the quality of its technology is high, knowledge dissipation intensifies, and its value must also decrease as Z further increases. We conjecture that firm value has a single peak, and that this peak is located in a neighborhood of the knowledge reset ($Z \equiv 0$) following exploration. Intuitively, firm value is high when knowledge dissipation is low, and this occurs following exploration. We later use parameter values under which this conjecture can be verified numerically.

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, and another two trigger levels, say $\underline{b} < \overline{b}$, at which $q(Z) = \sqrt{2\gamma\kappa}$ is satisfied and the firm decides to invest. Thus, $\mathscr{A} = (\underline{a}, \overline{a})$ denotes the region in which the firm does not explore, and $\mathscr{B} = (\underline{b}, \overline{b})$ the region in which the firm invests. Furthermore, if the condition in Eq. (13) is satisfied, then $\mathscr{B} \subset \mathscr{A}$. Our conjecture that firm value is largest following exploration also implies that $\{0\} \in \mathscr{B}$. In words, it is optimal for the firm to invest immediately following exploration, as this is when marginal q is highest. The times $\{\nu_k\}_{k=0}^{\infty}$ in Eq. (8) at which the firm explores correspond to hitting times at which knowledge exits $\mathscr{A}$:

$$\nu_n = \inf\{t \geq \nu_{n-1} : Z_t \notin \mathscr{A}\}, \quad \forall n \in \mathbb{N}. \tag{14}$$

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

Definition 2. ***(Knowledge Cycle)*** *A knowledge cycle starts when the firm resets its knowledge to 0 by exploring a new technology (i.e.,* $Z_0 = 0$*), and ends after a period of (random) length, as defined in Eq.* (14)*, when the firm abandons the technology to explore another, new technology (i.e.,* $Z \notin \mathscr{A}$*).*

Having defined the exploration and investment triggers, 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. In this case, the maximand in (11) is $\psi(q) \equiv 0$, and, as a result, Eq. (10) has an explicit solution. 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 A).

Proposition 1. *The intensive firm value in inaction regions,* $Z \in \mathscr{A} \setminus \mathscr{B}$*, satisfies:*

$$g\left(Z; C_1, C_2\right) := v_P\left(Z\right) + e^{g_1 Z + g_2 Z^2}\left(C_1 H_n\left(h_0 + h_1 Z\right) + C_2 M\left(-\frac{1}{2}n, \frac{1}{2}, m_0 + m_1 Z + m_2 Z^2\right)\right), \tag{15}$$

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

$$v_P(Z) = \int_0^\infty \exp\left(a_0(s) + a_1(s) Z + a_2(s) Z^2\right) ds, \tag{16}$$

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

$$n = \frac{1}{2}\left(\frac{2(\eta-1)^2\tau_A - \xi^2(-2(\eta-1)(2\alpha\delta+1) + 4r + \tau_A)}{\xi^3\sqrt{\tau_A}} - 1\right), \tag{17}$$

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* A*.*

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 1 when defined over the lower inaction region, and $\overline{C}_1$ and $\overline{C}_2$ over the upper inaction region.[12] We further determine the value of

[12]Because the inequality in Eq. (13) is strict, the boundaries of $\mathscr{A}$ and $\mathscr{B}$ never coincide.

the firm over the region $\mathscr{B}$ in which it invests. In this case, $\psi(q) \equiv \frac{1}{2\gamma}q^2 - \kappa$, and thus Eq. (10) no longer has an explicit solution. We proceed numerically, imposing the two boundary conditions, $v(\underline{b}) = v(\overline{b}) = \sqrt{2\gamma\kappa}/(\alpha(1-\eta))$. 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.[13] Figure 1 illustrates firm value as a function of the firm's stock of knowledge. This illustration relies on parameter values, which we discuss in Appendix B.

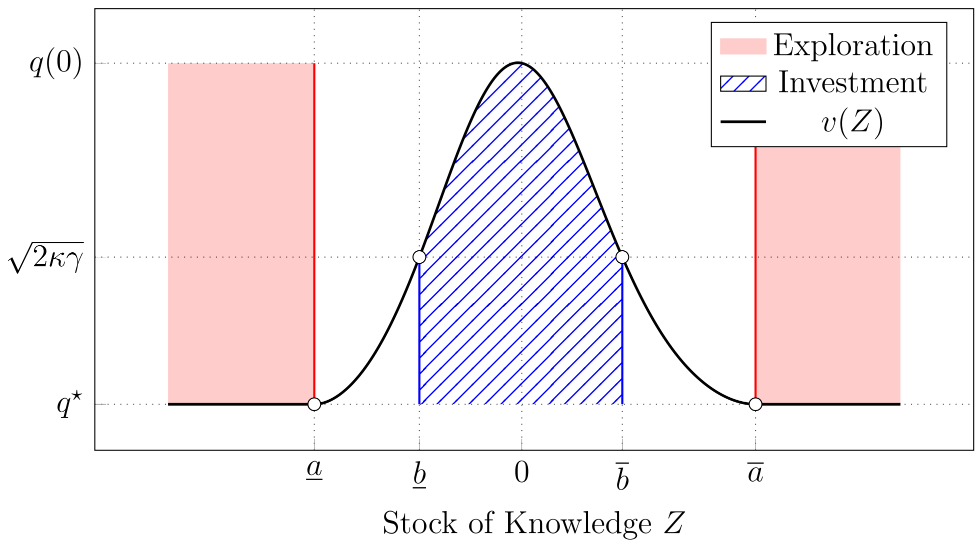

Figure 1: Firm value as a function of knowledge in the benchmark model. The figure plots intensive firm value as a function of knowledge, Z. Smooth-pasting for exploration requires that firm value lands flat at the edges of the domain defined by the exploration thresholds $\underline{a}$ and $\overline{a}$. Baseline parameter values are defined in Appendix B.

The main takeaway is that, absent experimentation, knowledge is irrelevant for corporate

[13] Based on Eq. (12) and piecing functions together, value-matching conditions satisfy:

$$\lim_{Z \searrow \underline{b}} v(Z) = \lim_{Z \nearrow \underline{b}} g\left(Z; \underline{C}_1, \underline{C}_2\right) = \lim_{Z \nearrow \overline{b}} v(Z) = \lim_{Z \searrow \overline{b}} g\left(Z; \overline{C}_1, \overline{C}_2\right), \ \lim_{Z \searrow \underline{a}} g\left(Z; \underline{C}_1, \underline{C}_2\right) = \lim_{Z \nearrow \overline{a}} g\left(Z; \overline{C}_1, \overline{C}_2\right) = \frac{q^\star}{\alpha(1-\eta)}.$$

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:

$$\lim_{Z \searrow \underline{a}} g'\left(Z; \underline{C}_1, \underline{C}_2\right) = \lim_{Z \nearrow \overline{a}} g'\left(Z; \overline{C}_1, \overline{C}_2\right) = 0, \ \lim_{Z \searrow \underline{b}} v'(Z) = \lim_{Z \nearrow \underline{b}} g'\left(Z; \underline{C}_1, \underline{C}_2\right), \ \lim_{Z \nearrow \overline{b}} v'(Z) = \lim_{Z \searrow \overline{b}} g'\left(Z; \overline{C}_1, \overline{C}_2\right). \tag{18}$$

investment—only q matters. Now, since q itself depends on knowledge, there exists an indirect relation between investment and knowledge. Inspecting Figure 1 shows that firm value is decreasing in knowledge, and thus so is investment. In particular, investment is highest following recent exploration (i.e., $Z \approx 0$), and weakest when exploration is imminent (i.e., $Z \approx \underline{a}$ or $Z \approx \overline{a}$). This mechanism arises because periods immediately following exploration are associated with high q, due to lower knowledge dissipation and high learning potential.

4 Trading off exploration against experimentation

We now study the model with experimentation ($\tau_S > 0$). Throughout the analysis, we define this "knowledge channel" as follows.

Definition 3. ***(Knowledge channel)*** *By "knowledge channel" we refer to the channel through which investment affects firm value (on average) by marginally improving knowledge, ignoring the marginal contribution of capital. Formally, the knowledge channel satisfies:*

$$c(Z) \equiv \int_{\mathbb{R}} \frac{V(\cdot, K+I, x)}{\Pi(\cdot)} \frac{\mathrm{d}}{\mathrm{d}i} \varphi(x; i(Z), Z)\, dx, \tag{19}$$

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)).

Eq. (19) can be equivalently expressed in terms of intensive firm value as:

$$c(Z) = (1+i(Z))^{\alpha(1-\eta)} \int_{\mathbb{R}} v(x) \frac{d}{di} \varphi(x; i(Z), Z) \mathrm{d}x,$$

so that the "knowledge channel" represents the sensitivity to investment of what the firm expects to be worth upon investing, while keeping the marginal effect of capital q fixed.

Building on the intuition gathered from the benchmark model, we conjecture that there exists a region $\mathscr{A}$, outside of which the firm decides to explore, and another region $\mathscr{B} \subset \mathcal{A}$, inside which the firm decides to invest. Unlike before, however, investment is now a means to experiment. 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 $Z \notin \mathscr{B}$). Furthermore, because the firm does not know ex-ante the outcome of its experiments, it must decide on a "knowledge-contingent investment plan", $i(Z)$, for every possible outcome that falls within the experimentation range.

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 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:

$$Z^{(n)} = Z^{(0)} \prod_{m=0}^{n-1} (1 + \tau_S i(Z^m))^{1/2} + \sum_{k=1}^{n-1} \sqrt{\tau_S i(Z^{(k-1)})}\widehat{\epsilon}^{(k)} \prod_{m=k}^{n-1} (1 + \tau_S i(Z^{(m)}))^{1/2} + \sqrt{\tau_S i(Z^{(n-1)})}\widehat{\epsilon}^{(n)}. \tag{20}$$

Eq. (20) 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. (20).

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 $Z \in \mathscr{B}$:

$$V(Z) = \max_{i \in \mathscr{D}} \underbrace{(1+i)^{\alpha(1-\eta)}}_{\text{returns on capital}} \underbrace{\int_{\mathbb{R}} v(x)\varphi(x; i(Z), Z)\mathrm{d}x}_{\text{expected returns on knowledge}} - \kappa - \gamma/2i^2, \tag{21}$$

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

$$h(Z) \equiv g(Z; \underline{C}_1, \underline{C}_2)\mathbf{1}_{Z \in (\underline{a}, \underline{b})} + (1-\omega)^{\alpha(1-\eta)} V(0)\mathbf{1}_{Z \notin \mathscr{A}} + g(Z; \overline{C}_1, \overline{C}_2)\mathbf{1}_{Z \in (\bar{b}, \bar{a})}. \tag{22}$$

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. (21) delivers the desired Bellman equation, which we highlight in the proposition below (see Appendix C).

Proposition 2. *The value function, $V(Z)$, defined in Eq. (21), associated with optimal experimentation satisfies the Bellman equation:*

$$V(Z) = \max_{i \in \mathscr{D}} (1+i)^{\alpha(1-\eta)} \int_{\mathbb{R}} h(y)\varphi(y; i(Z), Z)\,dy - \widehat{\gamma}(i) + (1+i)^{\alpha(1-\eta)} \int_{\mathscr{B}} V(y)\varphi(y; i(Z), Z)\,dy, \tag{23}$$

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 (23) 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.

The rest of the solution method is standard and mirrors that of the benchmark model. We impose continuity across the five smoothness regions, which gives us four equations for the four coefficients, $\underline{C}_1$, $\underline{C}_2$, $\overline{C}_1$, and $\overline{C}_2$. A technical difficulty at this stage is that the Bellman equation (23) does not have an explicit solution. In Appendix D, we explain how to proceed numerically. We then choose each threshold optimally by imposing that the firm value be continuously differentiable across regions.[14]

Figure 2 illustrates firm value as a function of knowledge based on the baseline parameter values discussed in Appendix B. Notably, firm value is similar to that in the benchmark model (see Figure 1). In particular, firm value is highest following recent exploration, thus confirming that the benefits of investing are highest at early stages; early investment is associated with high marginal revenue of knowledge and of capital. An important difference, however, is in the interpretation of the investment region (the hashed blue area). Because we assume adjustment costs are paid per unit of time in the benchmark case, investment occurs continuously within this region. In contrast, in the model with experimentation, investment occurs in lumps. Furthermore, and most importantly, in the model with experimentation, investment is a means to experiment and thus, upon entering the blue hashed area, the firm faces uncertainty as to the outcome.

We illustrate in Figure 3 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

[14] As in the benchmark model, the firm chooses optimally when to explore according to Eq. (18). Similarly, the firm chooses optimally when to experiment (i.e., the two experimentation triggers, $\underline{b}$ and $\bar{b}$) according to:

$$\lim_{Z\nearrow\underline{b}} g'(Z;\underline{C}_1,\underline{C}_2) = \lim_{Z\searrow\underline{b}} V'(Z) \quad \text{and} \quad \lim_{Z\nearrow\bar{b}} V'(Z) = \lim_{Z\searrow\bar{b}} g'(Z;\overline{C}_1,\overline{C}_2). \tag{24}$$

firm goes through possibly multiple rounds of experimentation, until experiments become conclusive (the blue dots).

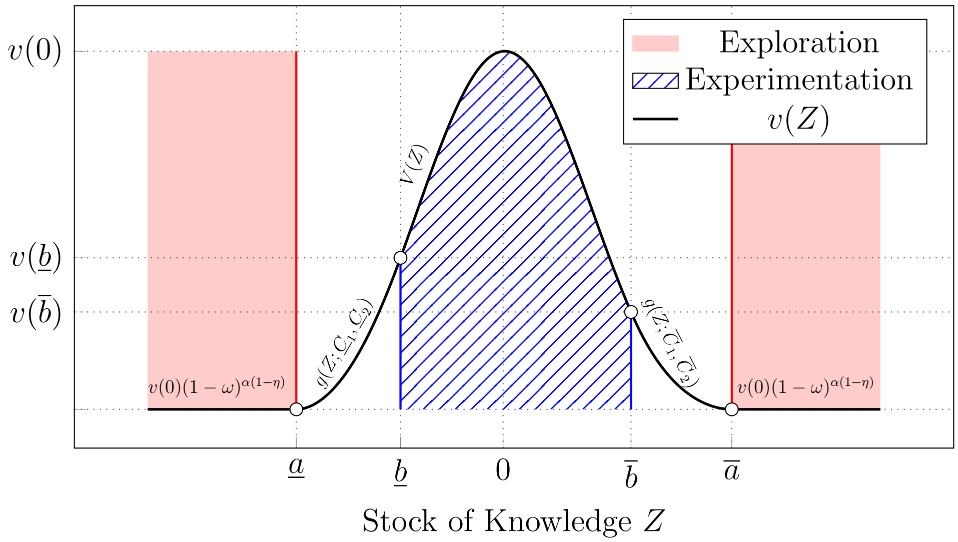

Figure 2: Firm value as a function of knowledge. The figure plots the firm value (solid black line) as a function of knowledge, Z, about the technology it operates. The blue hased area represents the experimentation regions, defined by the two thresholds, $\underline{b}$ and $\overline{b}$. The red shaded 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 in Eq. (33) is reported below the curve within each region. Baseline parameter values are defined in Appendix B.

To gain additional insights into how key parameters affect firm value, Figure 4 depicts the relation between firm value and knowledge for alternative values of the intensity of knowledge dissipation, ϕ, and of the informativeness of experimentation, τ_S. We denote by $\underline{b}'$ and $\overline{b}'$ the experimentation thresholds under these alternative parameters values. The left panel shows that the experimentation band narrows with low ϕ (in blue) relative to the baseline calibration (in black), and shifts towards more positive values of Z. Intuitively, when knowledge dissipation is low (i.e., intellectual property is better protected), the firm has less incentives to experiment unless it is confident that its current technology is of high quality. The right panel shows that, when experiments are more informative (i.e., less noisy), firm value (in red) rises relative to the baseline case (in black). Intuitively, as investment produces more valuable data, the experimentation band widens, the firm learns faster, and its value increases.[15]

[15]An increase in knowledge dissipation, ϕ, has a similar effect on the experimentation region, $\mathscr{B}$.

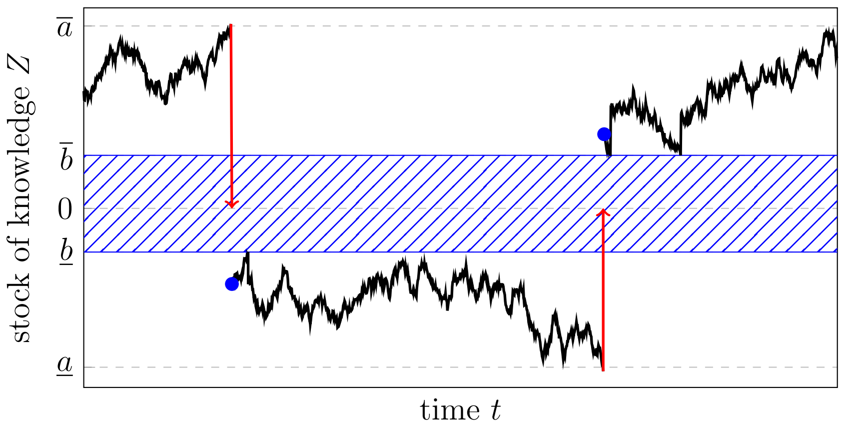

Figure 3: Sample path of Knowledge. The figure plots a sample path of the firm stock of knowledge (solid black line) over time. Blue dots correspond to times at which the firm chooses to experiment; red arrows correspond to times when the firm chooses to explore.

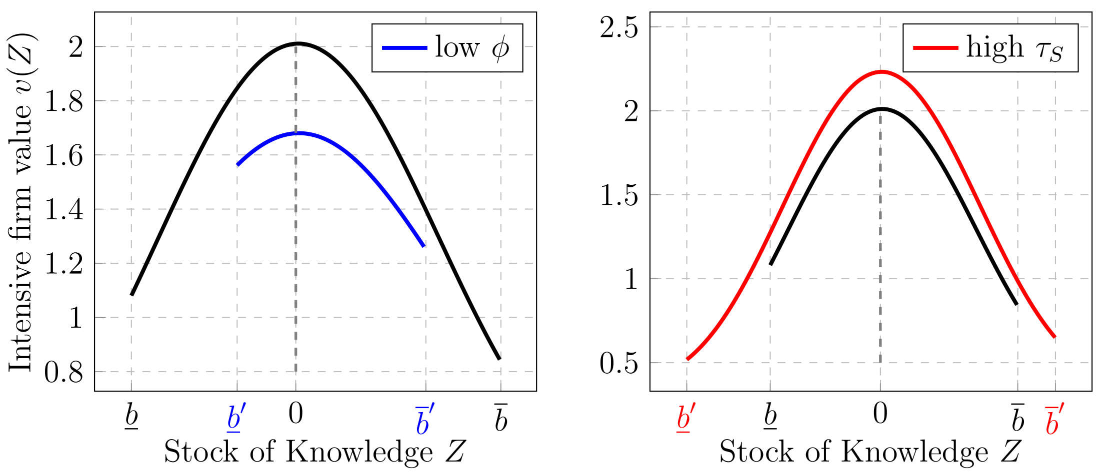

Figure 4: Comparative statics of the intensive firm value as a function of knowledge, Z, in the domain of the experimentation region, with respect to ϕ and τ_S. The black curve represents value under baseline parameter values; the blue curve represents the case of low knowledge dissipation ($\phi = 1.4 < 2$); the red line corresponds to low experimentation noise ($\tau_S = 0.3 > 0.25$). Baseline parameter values are defined in Appendix B.

5 Knowledge, Q, and Investment

The first set of novel insights from the model relates to its implications regarding the neoclassical relation between investment and marginal q. We start by presenting how the model with experimentation overturns neoclassical insights through endogenous experimentation aversion. Then, we discuss the resulting implications for the empirical relation between investment and q, and we present illustrative empirical evidence.

5.1 Endogenous Experimentation Aversion

To see the role played by the knowledge channel in the firm's investment decision, we use the first-order condition for optimal investment, i, from the Bellman Eq. (23). Substituting the definition of marginal q (Definition 1) and of the knowledge channel (Definition 3) yields:

$$0 = \underbrace{q(Z) - \gamma i(Z)}_{\text{neoclassical trade-off}} + \underbrace{c(Z)}_{\substack{\text{knowledge} \\ \text{channel}}} . \tag{25}$$

This equation determines a knowledge-contingent investment plan, $i(Z)$, for all values of knowledge within the experimentation band, $Z \in \mathscr{B}$. The first channel driving investment in Eq. (25) is the traditional neoclassical channel, as present in the benchmark model. When experimentation is uninformative, the marginal product of capital equals marginal adjustment costs (Hayashi, 1982). The second and novel channel driving investment is the knowledge channel, which captures how investment affects firm value by improving its stock of knowledge.

To illustrate this second channel, Figure 5 contrasts the relation between investment and the stock of knowledge obtained in the model with experimentation in the right panel (as determined from Eq. (25)) to that obtained in the benchmark model in the left panel (a scaled version of Figure 1). Strikingly, the knowledge channel turns the neoclassical relation between investment and knowledge on its head. In the benchmark model, because the neoclassical relation holds, investment is proportional to q. Investment is highest when marginal q is highest following recent exploration, and weakest when marginal q is low as exploration becomes imminent. Thus the benchmark model implies a *decreasing* relation between investment and knowledge, $|Z|$. When the knowledge channel operates (the right panel), investment also occurs following recent exploration when investment benefits are high (i.e., in the experimentation band). However, within the experiment band, investment is weakest following exploration, and strengthens as the firm accumulates knowledge. The knowledge channel thus implies an *increasing* relation between investment and the firm's stock of knowledge.[16]

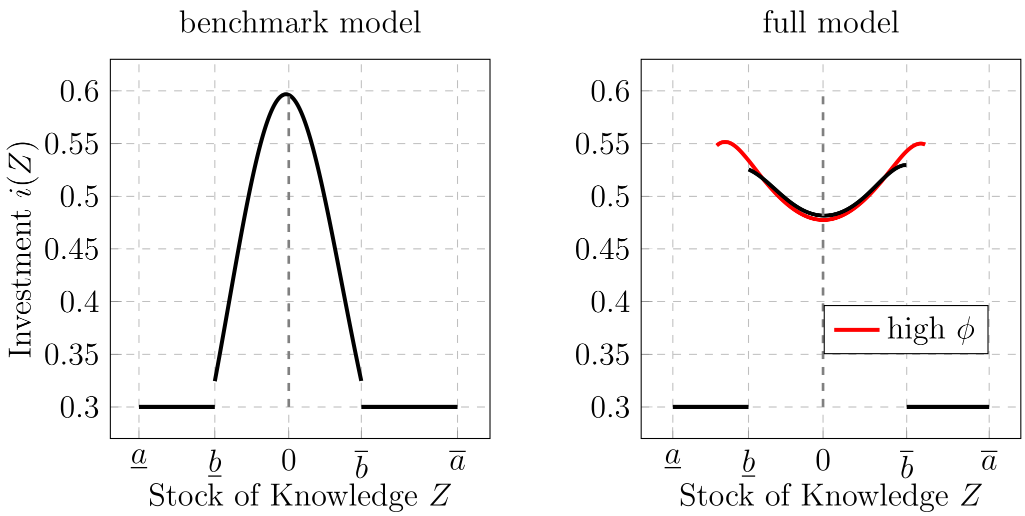

Figure 5: Knowledge-contingent investment plan. This figure plots investment as a function of knowledge, Z, in the benchmark model of Section 3 (left panel) and in the model with experimentation of Section 4 (right panel). In the right panel, the black curve represents investment under baseline parameter values; the red curve represents investment under high knowledge dissipation ($\phi = 3 > 2$). Baseline parameter values are defined in Appendix B.

To better understand how the knowledge channel overturns neoclassical insights, we use Stein's lemma to decompose the knowledge channel in Eq. (19) as:[17]

$$c(Z) = (1+i)^{\alpha(1-\eta)} \frac{\tau_S}{2} \Big(\underbrace{\frac{Z}{\sqrt{1+\tau_S i}} \int_{\mathbb{R}} v'(x)\varphi(x; i(Z), Z) \mathrm{d}x}_{\text{marginal product of knowledge}} + \underbrace{\int_{\mathbb{R}} v''(x)\varphi(x; i(Z), Z) \mathrm{d}x}_{\text{attitude towards experimentation risk}} \Big) \quad (26)$$

The first term represents the marginal revenue product of knowledge, which captures the expected value of a marginal unit of knowledge. Because firm value is nonlinear, the firm's attitude towards noisy experimentation affects knowledge accumulation, captured by the second term in Eq. (26). That the firm's attitude towards noisy experimentation matters

[16]Experimentation also renders the variation of investment smoother. Absent experimentation, the effect of investment on firm value is known ex-ante, generating significant variation in investment. With experimentation, the effect of investment on knowledge is uncertain, and hence marginal q represents the value of an additional unit of capital *averaged* over all possible experimentation outcomes, dampening the variability of marginal q and reducing variation of investment.

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

for its investment decision is perhaps surprising if one considers that the firm is assumed to be risk-neutral. In particular, the firm is risk-neutral towards profit risk, an assumption that underlies the *ex-ante* representation of firm value in Eq. (6). However, noisy experimentation affects firm value *ex-post*. As a result, depending on the curvature of firm value (as measured by v'' in Eq. (26)), experimentation is either a gamble or a risk towards which the firm is averse. Since experimentation occurs over a region of knowledge where firm value is concave, the firm is endogenously "experimentation averse".

In light of the decomposition in Eq. (26), what Figure 5 actually shows is that the firm's attitude towards noisy experimentation dictates its investment decision, as opposed to q or the marginal product of knowledge. Note first that marginal product of knowledge in Eq. (26) works hand in hand with marginal q, thus reinforcing the neoclassical channel. Following recent exploration, the marginal product of knowledge is high; conversely, when the technology is likely of high (or low) quality, and knowledge dissipation makes exploration imminent, the marginal product of knowledge is low. However, following exploration the firm value is concave, the firm thus dislikes noisy experimentation, and invests by "tatonnement" despite high benefits to investment. As knowledge dissipation strengthens and exploration is imminent, the firm value becomes convex, and the firm is increasingly willing to take bets, so that experimentation becomes a "gamble" on future exploration. A key result from the model is thus that the firm's attitude towards noisy experiments acts as a dominating force that overturns the neoclassical channel.

5.2 Investment and Q over the knowledge cycle

We now study how the overturning of the neoclassical insights generated by the knowledge channel affects the relation between investment and q. We specifically focus on the investment-q relationship because, unlike knowledge and experimentation, both quantities can be reasonably constructed from the data by an empiricist interested in testing specific implications of the model.

A key feature of the model is that noisy experimentation makes both the knowledge channel and marginal q unobservable to the empiricist. Indeed, because the effect of investment on knowledge is uncertain, marginal q is the value of an additional unit of capital *averaged* over all possible experimentation outcomes.[18] As it is not possible to control for all experimentation outcomes that are possible *ex-ante*, an empiricist cannot infer marginal q so defined. She can only control for experimentation outcomes *ex-post*, introducing a mismeasurement in q from her point of view, which renders q a poor statistic to explain investment.[19] Moreover, investment is lumpy within the experimentation region, and firms need not experiment at all times within the same knowledge cycle. Investment and q are thus either unrelated (outside the experimentation region) or negatively related (within the experimentation region).

We argue however that the empiricist can use the observable *average* product of capital, or average Q, to test some novel implications of the model. Given our assumption that adjustment costs are proportional to revenue, average Q corresponds to the price-revenue ratio:

$$Q(Z) \equiv \frac{V(\cdot, K+I, Z)}{\Pi(\cdot, K)} = (1+i(Z))^{\alpha(1-\eta)} v(Z). \tag{27}$$

Suppose that the empiricist could estimate the typical investment equation:

$$i(Z) = a + bQ(Z) + \epsilon(Z). \tag{28}$$

The estimated coefficient b is given by:

$$\widehat{b} = \frac{1}{\gamma} \underbrace{\frac{\text{cov}\left(q(Z)\mathbf{1}_{Z\in\mathscr{B}}, Q(Z)\right)}{\mathbb{V}(Q(Z))}}_{\equiv \widehat{\beta}_q} + \frac{1}{\gamma} \underbrace{\frac{\text{cov}\left(c(Z)\mathbf{1}_{Z\in\mathscr{B}}, Q(Z)\right)}{\mathbb{V}(Q(Z))}}_{\equiv \widehat{\beta}_c}. \tag{29}$$

The quality of her estimate (i.e., whether Q is an adequate statistic for investment) thus depends on how observed average Q co-varies with the unobserved marginal q (the first term of (29)), and with the unobserved knowledge channel (the second term of (29)). Because marginal q is a poor statistic for investment, we expect $\widehat{\beta}_q$ to be weak, and $\widehat{\beta}_c$ to be substantially weaker than the true β_c (i.e., the correlation between $q(Z)$ and $c(Z)$). Numerical illustrations reported in Appendix E confirm that average Q performs well in explaining in-

[18] A similar definition applies when investment is productive with delay (e.g., Caballero and Leahy, 1996)

[19] The role of experimentation in investment decisions could thus be one reason underlying the well-known mismeasurement problem in q documented by Erickson and Whited (2000).

vestment in the model. This result echoes the conclusion of Caballero and Leahy (1996), who show that in the presence of fixed costs of investment (and absent our knowledge channel) average Q better explains investment than marginal q.

To understand the relation between investment and Q and its variation over the knowledge cycle, we use the model to simulate time series of *ex-post* average Q (as defined in Eq. (27)) and investment, which we illustrate in Figure 6. Due to fixed adjustment costs, investment occurs infrequently and in lumps, the size of which varies with the firm's stock of knowledge. Furthermore, because investment occurs within the experimentation range, it also clusters in time. Investment clustering is a feature that is specific to our premise that knowledge is a by-product of economic activity. Because average Q measures firm value even when the firm does not invest, it moves all the time.

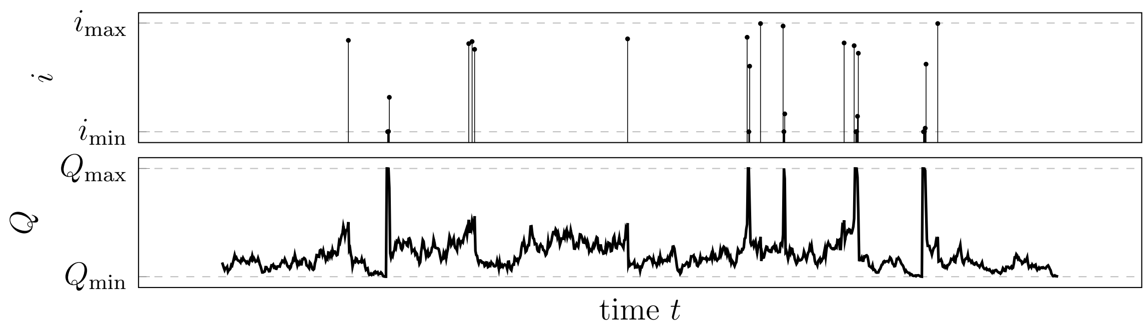

Figure 6: Sample path of investment and average Q. The upper and lower panel plot a sample path of the investment rate and the associated average Q, respectively, over time. The upper panel shows that investment occurs in lumps and clusters.

Figure 7 reports coefficients (b_k) from estimating Eq.(28) using the simulated series within deciles k of knowledge $|Z|$. Notably, the knowledge channel implies significant variation in the strength of the investment-Q relation over the knowledge cycle. We observe a positive relation between investment and average Q at early stages of the cycles when knowledge is limited. This is due to the presence of fixed adjustment costs that create inaction regions over which investment is mechanically unrelated to Q. This happens when knowledge is abundant (i.e., whenever $Z \notin \mathscr{B}$). Remarkably, the stage of the cycle at which the statistical power of average Q to explain investment is strongest depends on the extent of experimentation noise and knowledge dissipation. When experiments yield better information investment is more aggressive, the experimentation band widens and the sensitivity of investment to

Q concentrates in higher deciles of knowledge. In contrast, when knowledge dissipation is low, a significant investment-Q relation arises early in the knowledge cycle. Taken together, these novel predictions underline the importance of the knowledge channel in governing the relationship between investment and Q.

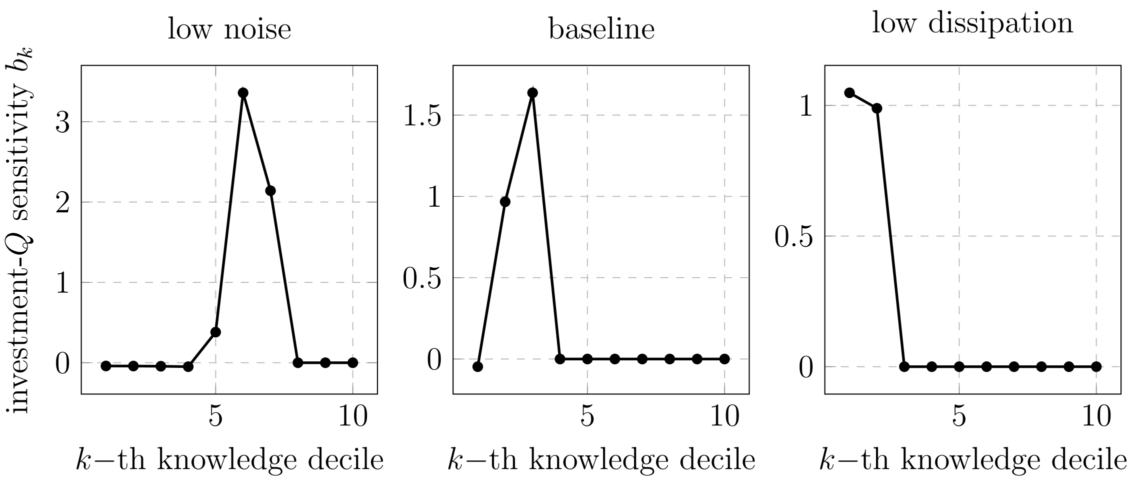

Figure 7: Investment$-q$ sensitivity as a function of knowledge quintiles. Each panel plots the sensitivity of investment to average Q within quintiles of the (absolute value of the) stock of knowledge, $|Z|$. The middle panel corresponds to the baseline calibration (Appendix B), the left panel considers low experimentation noise ($\tau_S = 0.3 > 0.25$), and the right panel considers low knowledge dissipation ($\phi = 1.4 < 2$).

5.3 Illustrative evidence

To provide qualitative evidence for the knowledge channel, we rely on a text-based measure of product life cycles as an empirical proxy for Z. Following the characterization of the life cycle of firms' products introduced by Abernathy and Utterback (1978), Hoberg and Maksimovic (2019) develop a textual analysis of firms' 10-K reports to identify the relative position of each firm (and year) in this cycle. We strictly follow their methodology and consider that product life cycle follows four distinct stages: (1) product innovation, (2) process innovation, (3) maturity, and (4) decline. We consider all firm-years present in the universe of Compustat between 1996 and 2016 with adequate 10-K data from EDGAR. We determine the intensity with which a given firm is in each stage in a given year using the number of paragraphs in its 10-K with words contained in four distinct (pre-determined) lists associated with each stages. We then define the four life cycles variables by dividing each of the four individual paragraph counts by the total paragraphs counts for the four. We refer to these four variables as Life1, Life2, Life3, and Life4, respectively.[20] We exclude observations with missing data on capital expenditures, sales and assets below $1 million. Our sample includes 63,199 firm-year observations.

We recognize that product and knowledge cycles may not perfectly overlap in reality. Thus, our proxy is arguably imperfect. Yet, following Abernathy and Utterback (1978), Klepper (1996) and Callander (2011), we posit that both cycles are likely to be positively correlated for most firms. We thus conjecture that firms' stock of knowledge is likely low (i.e., Z is close to zero) at the beginning of a new product cycle (i.e., when uncertainty is high), and accumulates (i.e., Z moves away from zero) in later stages as a result of (passive and active) learning.

Absent the knowledge channel, both investment and Q are *decreasing* in knowledge. In the presence of the knowledge channel, however, investment is *increasing* in knowledge (due to endogenous experimentation aversion), but Q continues to be *decreasing* in knowledge (see Figure 5). To assess which of these two possible patterns hold in the data, we replicate the empirical analysis of Hoberg and Maksimovic (2019) and regress investment and Q on the four product life cycle variables (with and without year fixed effects). We define investment as capital expenditures scaled by beginning of the period assets, and Q as the market value of assets divided by their book value.[21]

Figure 8 displays the estimated coefficients (together with their 95% confidence bounds). In the left panel, firms' investment is about twice lower at the beginning of the cycle (i.e., when knowledge is low) than in the second stage, when more knowledge has accumulated. However, the middle panel indicates that Q is more than three times larger in the first stage of the cycle than in the second stage.[22] These opposite patterns between investment and Q over the initial stages of the cycle are consistent with the predicted role of the knowledge

[20]For example, a firm scores highly on Life1 if it intensively discusses issues related to product or service innovation or development in its 10-K. By construction, Life1+Life2+Life3+Life4=1 for each firm-year.

[21]We standardize the four life cycles variables by their sample standard deviation to facilitate the economic interpretation of the estimated coefficients, and we cluster standard errors at the firm level.

[22]These results are consistent with the overall findings reported in Hoberg and Maksimovic (2019).

channel in the model.

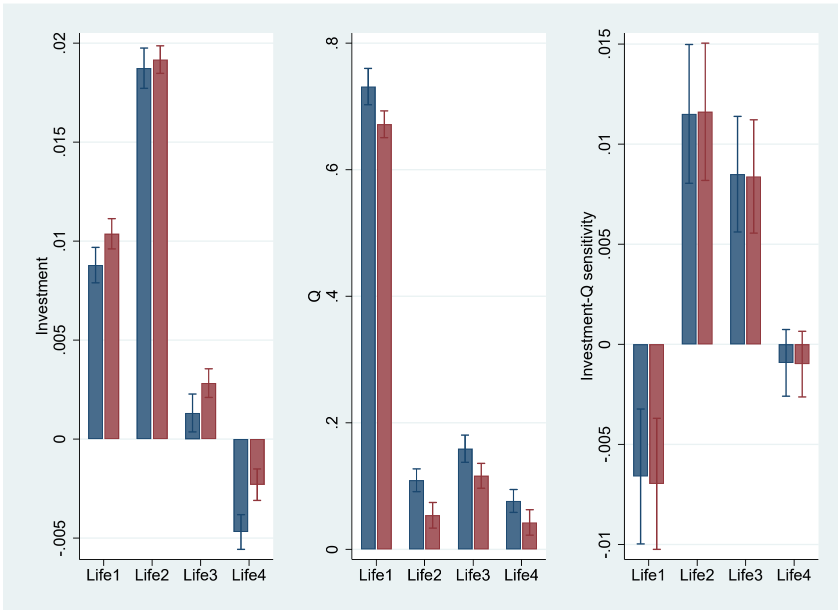

Figure 8: Investment and Q over the product life cycle. The left and middle panels plot investment and Q as a function of the four stages of the product cycle defined in Hoberg and Maksimovic (2019). The right panel plots estimates of the investment-Q sensitivity across the four product stages. The blue bars correspond to estimations without year fixed effects; the red bars represent the estimations with year fixed effects.

We also explore how the investment-Q relation varies over the knowledge cycle. Mimicking the approach used on the simulated data presented in Figure 7, we regress investment on the four life cycles stages variables and their interaction with Q (with and without year fixed effects).[23] The estimated coefficients on the four interaction variables track the empirical correlation between investment and Q across the four-cycle stages. Consistent with the knowledge channel, the right panel of Figure 8 indicates that the sensitivity of investment to Q varies considerably over the knowledge cycle. In particular, investment appears to be negatively related to Q in the first stage of the knowledge cycle. In the model, when firms' explore new technologies and possess little knowledge about them, Q is high but investment is low since firms do not experiment yet. In later stages, however, the accumulated knowledge pushes firms to experiment by investing, generating stronger investment-Q sensitivity.

[23]We again standardize the interacted variables by their sample standard deviation to facilitate the economic interpretation of the estimated coefficients, and we cluster standard errors at the firm level.

This prediction is supported by data, since we observe a positive and strong correlation between investment and Q in the second and third stage. Compared to the simulated results displayed in Figure 7, the patterns of Figure 8 suggest that the data appear consistent with calibrations of the model in which the noise associated with experimentation is low.

6 Knowledge Cycles and Experimentation Rounds

Another novel insight from the model with experimentation is that it generates endogenous knowledge cycles. This feature contrasts with models in which exploration and experimentation follow pre-determined schedules (e.g., Berk et al. (2004) or Pastor and Veronesi (2009)), and thus generates implications for the determinants of exploration and experimentation within the firm that we explore in this section.

6.1 The length of knowledge cycles

We first focus on the typical length of a knowledge cycle (i.e., the time spent by the firm operating a given technology), and display its distribution in Figure 9 under our baseline calibration. The left panel compares this distribution in the benchmark model (dashed line) to that with experimentation (solid line).[24] Because experimentation allows the firm to learn faster about the quality of its technology, it shortens considerably the expected length of the knowledge cycle. Note that we have calibrated the benchmark model so that the exploration thresholds coincide with those of the model with experimentation. Hence, the shortening of the cycle is a unique consequence of the learning benefits of active experimentation.[25]

The right panel characterizes how the model's key parameters affects the distribution of the length of knowledge cycles. Notably, consistent with the literature emphasizing the potential downsides of strong intellectual property rights (e.g., Williams (2013)), weak knowledge dissipation (the blue line) raises the length of the knowledge cycle, as it reduces the firm's incentive to explore a new technology. Similarly, the cycle is shortened when ex-

[24] See Appendix F for the derivations.

[25] In particular, we calibrate γ and κ in the benchmark model so that $\underline{a}$ and $\overline{a}$ coincide with those of the baseline calibration of the full model, and that the product $\gamma\kappa$ is the same in both models. All other parameters remain as reported in Table 1 of Appendix B.

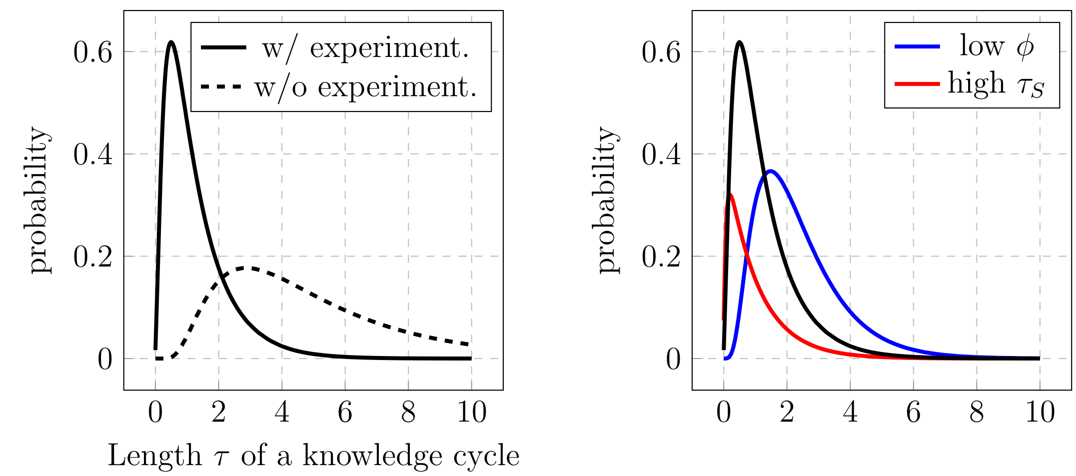

Figure 9: Distribution of the length of a knowledge cycle. The left-hand panel plots the distribution of the length of the cycle with (solid line) and without (dashed line) experimentation; derivations are reported in Appendix F. The right-hand panel does comparative statics, with the black curve denoting the baseline case, the blue curve corresponding to low knowledge dissipation ($\phi = 1.4 < 2$), and the red line corresponding to low experimentation noise ($\tau_S = 0.3 > 0.25$). Baseline parameters are reported in Appendix B.

periments produce better information (the red line).[26] Better information makes inaction less likely, the firm learns faster, and exploration is thus more frequent, consistent with the evidence reported by Braguinsky et al. (2019).

To contrast the length of a knowledge cycle in the model and in the data, we rely again on product cycles and assign each firm-year into four binary variables based on the largest value of the four stages.[27] As firms move in and out of each stage regularly (see Hoberg and Maksimovic (2019)), we approximate cycles' length using the number of years it takes a firm that leaves the initial stage to return to that stage. We thus implicitly assume that firms are in the first stage when they start a new cycle by exploring new technologies. Restricting our attention to firms that complete at least one entire cycle, we identify 1,716 unique cycles across 1,152 distinct firms.[28] Figure 10 displays the distribution of cycle length across firms. Mirroring closely the distribution implied by the model with experimentation, the empirical

[26] The expected number of rounds is distributed similarly as in the baseline case, and hence not shown.

[27] For instance $D_{Life1} = 1$ for a given firm in a given year if the value of Life1 is strictly larger than that of the three other Life variables in that year.

[28] Conservatively, we discard cycles when we cannot identify the beginning or the end, because of sample attrition (e.g., starting before 1995 and ending after 2015).

distribution is asymmetric, right skewed, and implies that firms experience short cycles on average. Consistent with the role of the knowledge channel, the majority of cycles last one year, and very few firms exhibit cycles that are longer than five years. This stands in sharp contrast with the distribution implied by the model without experimentation, which is much less asymmetric and implies significantly longer average cycles (see the dashed distribution in Figure 9).

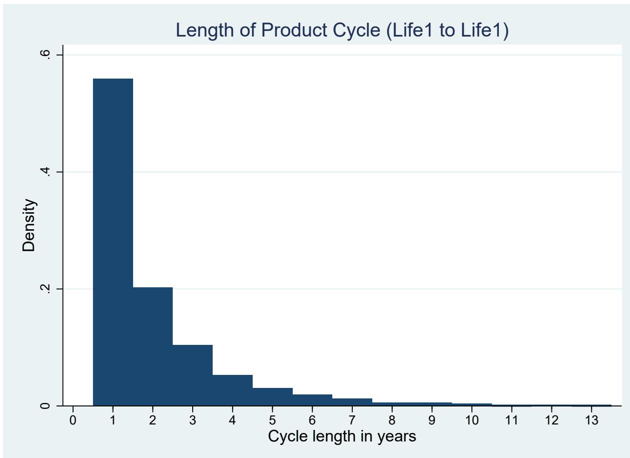

Figure 10: Empirical length of product cycles. This figure plots the distribution of the length of product cycle in our sample, where a cycle is the number of years it take firms to go from Life1 back into Life1.

6.2 Rounds of Experimentation

The model also has implications for the number of rounds of experimentation within each cycle. The left panel of Figure 11 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 rounds of three 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 instantaneously.

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.

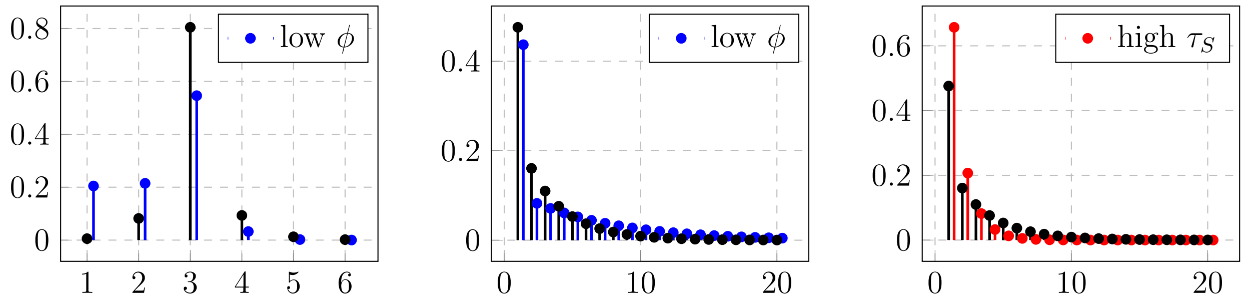

Figure 11: Distribution of experimentation trials. This figure plots the distribution of the number of instantaneous trials in an experimentation round as per Definition 4 (right panel) and "in-and-out" trials (middle and right panels). "In-and-out" trials correspond to the number of times a firm currently 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 ($\tau_S = 0.3 > 0.25$). Baseline parameters are provided in Appendix B.

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:

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

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 the sequence of knowledge stopped at these dates:

$$\left\{Z_{\theta_{0+}}, ..., Z_{\theta_{n+}}\right\}. \tag{30}$$

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 11. We relegate derivation details to Appendix G.

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 (Figure 9).

7 Conclusion and Speculations

To conclude, we examine whether potential aggregate changes in firms' knowledge cycles could explain the aggregate decline in investment in the last twenty years and its increasing disconnection with Q (e.g., Gutierrez and Philippon (2017), Alexander and Eberly (2018), and Crouzet and Eberly (2018)). Confirming these trends, the top panels of Figure 12 plots the average evolution of investment and Q (the left panel), as well the sensitivity of investment to Q (the right panel) in our working sample of public firms.[29]

Interpreted in light of the model, the aggregate shifts in investment and Q could stem from an aggregate shortening of the average knowledge cycle. Specifically, weaker physical investment *despite* high Q could result from firms exploring new technologies more frequently, thereby staying longer in the low knowledge region. Such a rise in exploration intensity could happen for instance if high quality technologies are getting harder to find as suggested by Bloom, Jones, van Reenen, and Webb (2017), pushing firms to explore more often. Alternatively, an improvement in the informativeness of experimentation (e.g., through the use computer simulations, mass screening, or rapid prototyping; see Thomke (1998)) or an intensification of knowledge dissipation (e.g., weaker intellectual property rights or greater inventors' mobility) could also induce a shortening of the knowledge cycle.

To assess whether weaker investment despite high Q in the recent period may be due to

[29]We obtain it by estimating annual investment regressions, and plotting the estimated coefficients.

Figure 12: Aggregate Trends. The upper left panel plots the average evolution of investment and average Q over the period 1996 to 2016. The upper right panel plots the evolution of sensitivity of investment to Q. The bottom left panel plots the evolution of the first two stages of firms' product cycles (Life1 and Life2), and the bottom right panel plots the evolution of the corresponding last two stages (Life3 and Life4).

a shortening of the knowledge cycle inducing more exploration, we plot in the bottom panels of Figure 12 the evolution of the four life-cycle stages over our sample period. Consistent with the rise in "business dynamism" noticed by Hoberg and Maksimovic (2019), the left panel unambiguously indicates a shortening of firms' knowledge cycles. The average value of Life1 is trending upwards, increasing from 27% in 1996 to 33% in 2016. In contrast, the average value of Life2 is decreasing from 43% to 36% over the last twenty years. The right panel suggests no significant change in the average value of Life3, and a marginal increase in Life4. Overall, Figure 12 suggests that firms are spending relatively more time in the exploration stage, where investment is predicted to be weak, and Q to be high.

While we acknowledge that our empirical analysis remains qualitative and cursory in nature, it broadly supports the predictions of the model. Moreover, the results are in line with the recent work of Crouzet and Eberly (2020) who estimate that the low level of aggregate investment depite high Q (which they call the "investment gap") stems majoritarily from the rise of intangibles. Our analysis complements theirs, and suggests that the internal process of knowledge creation within firms is an important intangible element that can explain the dynamics of corporate investment.

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.

Abernathy, W. and J. Utterback (1978). Patterns of industrial innovation. *The RAND Journal of Economics 80*, 41–47.

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.

Basak, G. and K.-W. Ho (2004). Level-crossing probabilities and first-passage times for linear processes. *Advances in Applied Probability 36*(2), 643–666.

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 (2019). 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.

Caballero, R. and J. Leahy (1996). Fixed costs: The demise of marginal q. *Unpublished Manuscript*.

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

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.

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.

Hellwig, C., S. Kohls, and L. Veldkamp (2012, May). Information choice technologies. *American Economic Review 102*(3), 35–40.

Hoberg, G. and V. Maksimovic (2019). Product life cycles in corporate finance. *Working Paper*.

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.

Klepper, S. (1996). Wentry, exit, growth, and innovation over the product life cycle. *American Economic Review* (86), 562–583.

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.

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.

Stokey, N. and R. Lucas (1996). *Recursive methods in economic dynamics* (4. print ed.). Cambridge, Mass. [u.a.]: Harvard Univ. Press.

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

Thomke, S. (1998). Managing experimentation in the design of new products. *Management Science 44*(6), 743–762.

van Nieuwerburgh, S. and L. Veldkamp (2010). Information acquisition and under-diversification. *The Review of Economic Studies 77*(2), 779–805.

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.

Williams, H. (2013). Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy* (121), 1–27.

A General solution in the inaction region

In this appendix we solve for the intensive firm value in the inaction region. The derivation applies both to the benchmark model (Section 3) and the model with experimentation (Section 4). In the inaction region, the processes for knowledge and physical capital move of their own accord. As a result, the firm value in Eq. (6) 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:

$$\left(r + \eta Z^2 \phi - (\eta - 1)\left(\alpha\delta + \frac{\eta}{2} - \sqrt{\tau_A} Z\right)\right) v = 1 + \sqrt{\tau_A} v' \left(\frac{\sqrt{\tau_A}}{2} Z + 1 - \eta\right) + \frac{\tau_A}{2} v'', \; Z \in \mathscr{I} \quad (31)$$

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

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

Introducing the change of probability measure,

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

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

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

accordingly, we define the function

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

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

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

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:

$$a_2(s) = \frac{1}{4}\left(-\frac{\xi \tanh\left(\frac{1}{2}\xi s\sqrt{\tau_A} - \coth^{-1}\left(\frac{\xi}{\sqrt{\tau_A}}\right)\right)}{\sqrt{\tau_A}} - 1\right), \tag{32}$$

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

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

where $\xi \equiv \sqrt{\tau_A + 8\eta\phi}$. We conclude that a particular solution to Eq. (31) is as in Eq. (16). Furthermore, the two solutions associated with the homogeneous version of Eq. (31) are:

$$G_1(Z) = e^{g_1 Z + g_2 Z^2} H_n\left(h_0 + h_1 Z\right),$$

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. (17)) 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. (31) takes the form in Eq. (15). Note that the existence of this solution depends on several conditions on the parameter values.

B Baseline calibration

In this appendix we define parameter values for our baseline calibration, and summarize them in Table 1. Most of the parameters are standard in the literature. Following Nikolov and Whited (2014), we set the discount rate to $r \equiv 5\%$, the capital depreciation rate to $\delta \equiv 13\%$, firms' variable adjustment costs to $\gamma \equiv 0.5$, and the production function is such that $\alpha \equiv 0.8$. Similarly, we assume that fixed adjustment costs represent $\kappa \equiv 9.5\%$ of profits, so that firms' reinvestment rate aligns with empirical evidence. 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).

Table 1: 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
γ	variable adjustment costs	0.5
κ	fixed adjustment costs	9.5%
η	inverse of price elasticity of demand	0.5
τ_A	informativeness of aggregate productivity	0.7
2. other parameters		
τ_S	informativeness of experiments	0.25
ω	obsolescence costs	98%
ϕ	knowledge dissipation	2

Other parameters determining the behavior of exploration and experimentation in our model are τ_S, ω and ϕ. We set the informativeness of experimentation to $\tau_S \equiv 0.25$, as well as the informativeness of the productivity signal to $\tau_A \equiv 0.7$, so that the expected growth rate of productivity aligns with the average historical growth rate of total factor productivity.[30] Similarly, we ensure

[30]Using estimates of the National Income and Product Accounts (NIPAs), private non-farm total factor productivity growth in the United States has fluctuated between 1% and 1.5% annually since 1950.

that exploration is infrequent by setting obsolescence costs ω sufficiently high. The stock of capital becomes obsolete so that firm value decreases by $1-(1-\omega)^{\alpha(1-\eta)} \approx 80\%$ 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 $\phi \equiv 2$, a value that triggers approximate symmetry across technologies.

C Expected firm value upon experimentation

To solve for the expected firm value upon experimentation of the model in Section 4, we start by formulating the dynamic optimization problem in Eq. (21) globally. We write the value function piecewise as:

$$v(Z) = \begin{cases} g(Z; \underline{C}_1, \underline{C}_2), & Z \in (\underline{a}, \underline{b}) \\ V(Z), & Z \in \mathscr{B} \\ g(Z; \overline{C}_1, \overline{C}_2), & Z \in (\bar{b}, \overline{a}) \\ (1-\omega)^{\alpha(1-\eta)} V(0), & Z \notin \mathscr{A} \end{cases} . \tag{33}$$

Using the function $h(Z)$ in Eq. (22) 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:

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

Furthermore, using the distribution $\varphi(\cdot; i(Z), Z)$ of incremental knowledge gained from experimentation (see Definition 3), define for $y \equiv Z^{(n+1)}$ and $z = Z^{(0)}$:

$$\phi^{(n)}(y, z) \equiv \int_{\mathscr{B}^n} \prod_{k=1}^{n} (1+i_k)^{\alpha(1-\eta)} \varphi(y; i_n, z^{(n)}) \varphi(z^{(k)}; i_{k-1}, z^{(k-1)}) \mathrm{d}z^{(k)}, \tag{34}$$

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:

$$V(Z) = \max_{\{i_k\}_{k=0}^{\infty}} (1+i_0)^{\alpha(1-\eta)} \sum_{n=0}^{\infty} \int_{\mathbb{R}} \left(h(y) - \widehat{\gamma}(i(y))\mathbf{1}_{y\in\mathscr{B}}\right) \phi^{(n)}(y; Z)\mathrm{d}y - \widehat{\gamma}(i_0). \tag{35}$$

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

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

Rewriting Eq. (34) 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:

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

and the second term in brackets as:

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

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

$$V(Z) \equiv \max_{i_0}(1+i_0)^{\alpha(1-\eta)} \int_{\mathbb{R}} h(y)\phi(y; Z, i_0)\mathrm{d}y - \kappa - \gamma/2i_0^2 + \max_{\{i_k\}_{k=0}^{\infty}} (1+i_0)^{\alpha(1-\eta)} \times \tag{36}$$
$$\int_{\mathscr{B}} \left(\begin{array}{c} (1+i(x))^{\alpha(1-\eta)} \sum_{n=0}^{\infty} \int_{\mathbb{R}} \left(h(y) - (\kappa + \gamma/2i(y)^2)\mathbf{1}_{y\in\mathscr{B}}\right) \phi^{(n)}(y; x, \{i_k\}_{k=0}^{n})\mathrm{d}y \\ -\kappa - \gamma/2i(x)^2 \end{array} \right) \phi(x; Z, i_0)\mathrm{d}x$$

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

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. (36), the optimal policy function must satisfy the first-order condition:

$$\begin{aligned} 0 = {} & \frac{\alpha(1-\eta)}{1+i(Z)}V(Z) + \frac{\alpha(1-\eta)(\kappa + \gamma/2i(Z)^2) - \gamma(1+i(Z))i(Z)}{1+i(Z)} \\ & + (1+i(Z))^{\alpha(1-\eta)} \left(\int_{\mathbb{R}} h(y) \frac{\partial}{\partial i}\phi(x; i, Z)\mathrm{d}y + \int_{\mathscr{B}} V(y)\frac{\partial}{\partial i}\phi(x; i, Z)\mathrm{d}y \right). \end{aligned} \tag{37}$$

We start by computing the first expectation in Eq. (36). 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

$$c_2 < \frac{1}{2\tau_S i(Z)}, \tag{38}$$

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. (15), we use the integral representation of the Hermite polynomial. Note that the order of the polynomial in Eq. (17) may be negative, in which case the following integral representation is appropriate:

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

We then obtain:

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

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. (36) 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. (38), 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:

$$i < \frac{2}{\tau_S(\xi/\sqrt{\tau_A}-1)} =: \bar{i}. \tag{39}$$

Furthermore, using Eq. (32) for a_2, since tanh : $\to [-1,1]$ it follows that the first equation is always satisfied if the condition in Eq. (39) 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

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

D Solution method

To solve the model with experimentation in Section 4 numerically, we start by imposing that the firm value be continuous over each region involved in its piecewise representation in Eq. (33). In particular, when the firm invests there can no jump in its value on average;

$$\begin{aligned} g(\underline{b}; \underline{C}_1, \underline{C}_2) &= V(\underline{b}), \\ g(\bar{b}; \overline{C}_1, \overline{C}_2) &= V(\bar{b}). \end{aligned} \tag{40}$$

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

$$\begin{aligned} g(\underline{a}; \underline{C}_1, \underline{C}_2) &= V(0)(1-\delta)^{\alpha(1-\eta)}, \\ g(\bar{a}; \overline{C}_1, \overline{C}_2) &= V(0)(1-\delta)^{\alpha(1-\eta)}. \end{aligned} \tag{41}$$

Eqs. (40)–(41) 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. (35). 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 (23) 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:

$$
\begin{aligned}
V(z_j) = (1+i(z_j))^{\alpha(1-\eta)} \int_{\mathbb{R}} (h(y) - \mathbf{1}_{y\in\mathscr{B}}(\kappa + \gamma/2 i(y)^2))\phi(y; i(z_j), z_j)\mathrm{d}y \\
- \kappa - \gamma/2 i(z_j)^2 + (1+i(z_j))^{\alpha(1-\eta)} \sum_{k=1}^{n} \omega_k \phi(x_k; i(z_j), z_j) V(x_k), \quad j = 1, \dots, n,
\end{aligned}
\tag{42}
$$

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

We then define the following (column) vectors:

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

and the following $n \times n$ matrix:

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

We can then rewrite the system in Eq. (42) 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:

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

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. (40)–(41) explicitly as:

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

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 knolwedge-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. (18) (the first two equation for optimal exploration choice) and Eq. (24), wich 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. (37). To find a solution, we use a Newton algorithm to solve for Eqs. (18)–(24), and at each Newton step we perform Policy Function Iteration to obtain the optimal investment policy.

E Average Q, marginal q and mismeasurement

In this appendix we discuss heuristically how average Q provides a better description of investment, relative to marginal q, in the presence of the knowledge channel. The analysis complements the discussion in Section 5.3. In particular, consider the regression in Eq. (28). Whether average Q is a good statistic for investment in this model depends on its relation to marginal q ($\widehat{\beta}_q$) and with the knowledge channel ($\widehat{\beta}_c$). In Figure 13, we plot the coefficients $\widehat{b}$, $\widehat{\beta}_q$ and $\widehat{\beta}_c$ against the values b, β_q, and β_c we would have obtained had we used marginal q as a right-hand side variable, instead of average Q. We do so for different values of knowledge dissipation intensity, ϕ.

We start by describing the coefficients β_q (left panel) and β_c (middle panel), and the resulting sensitivity estimate b (right panel) using marginal q as a right-hand variable. On the one hand, we expect the knowledge channel to covary negatively with marginal q within the experimentation region. On the other hand, the experimentation band coincides with high values of q on average, which should lead the knowledge channel to covary positively with marginal q unconditionally. The middle panel shows that the latter effect dominates for the lower range of knowledge dissipation intensity, whereas the former effect dominates as knowledge dissipation strengthens. Furthermore, absent the knowledge channel, and in the presence of fixed costs and investment irreversibility, investment and marginal q are negatively related (e.g., Caballero and Leahy (1996)). Intuitively, when investment occurs $q = 0$, and $q > 0$ when it does not. It follows that $[qI] = 0$, and thus $\text{cov}(q, I) = -[q][I]$, which is negative under irreversibility, $I \geq 0$. Adding a convex cost can make this covariance positive. The left panel shows that for low knowledge dissipation intensity the former effect dominates, whereas for high knowledge dissipation intensity the latter effect dominates. On balance, the net effect is that investment and marginal q are negatively related (the right panel), the more so the more intense knowledge dissipation is.

Using *ex-post* average Q as a right-hand variable produces a strong attenuation bias (dashed lines are of smaller magnitude relative to solid lines). Unlike covariances involving marginal q, the

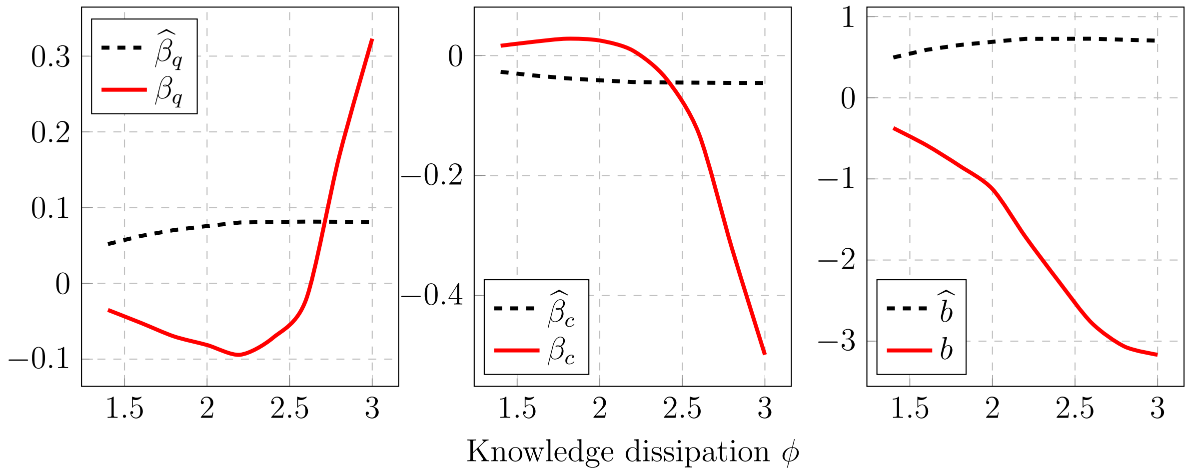

Figure 13: Investment$-q$ sensitivity as a function of knowledge dissipation. The figure plots the coefficients β_q and $\widehat{\beta}_q$ (left panel), β_c and $\widehat{\beta}_c$ (middle panel), and b and $\widehat{b}$ (right panel) as a function of competitive pressure, ϕ; these coefficients are defined in Eq. (29).

sign of covariances involving average Q is unambiguous. Average Q always covaries positively with marginal q, and negatively with the knowledge channel. On net the former covariance dominates leading to a positive measured sensitivity of investment to average Q, which strengthens with greater knowledge dissipation. In sum, *ex-post* average Q performs substantially better than marginal q as an explanatory variable for investment.

F Distribution of the Length of the Exploration Cycle

We derive the distribution of the length of the exploration cycle discussed in the main text based on a system of coupled Kolmogorov Backward Equations (KBEs). Formally, our goal is to compute the probability that the first exploration time θ (defined in Eq. (14)) occurs before or at time $t \geq 0$ given that Z_0 starts at 0 (i.e., immediately following an initial round of exploration):

$$G(t) \equiv \ (\theta \leq t | Z_0 = 0)\,.$$

To keep notations transparent, for this appendix only we change notations for stopping times and define the hitting time of Z at level k as:

$$\tau_k \equiv \inf\{t : Z_t = k\}.$$

To derive an expression for $G(\cdot)$, it is convenient to start by working piecewise over two regions, $(\underline{a}, \underline{b})$ and $(\bar{b}, \overline{a})$; accordingly, define the two functions:

$$f_1(t, Z) \equiv \left(\theta \leq t | Z_0 \in (\bar{b}, \overline{a})\right) = \left(\mathbf{1}_{\theta \leq t} | Z_0 \in (\bar{b}, \overline{a})\right),$$
$$f_2(t, Z) \equiv \left(\theta \leq t | Z_0 \in (\bar{b}, \overline{a})\right) = \left(\mathbf{1}_{\theta \leq t} | Z_0 \in (\bar{b}, \overline{a})\right).$$

These two functions satisfy a system of coupled equations. Suppose first that today Z_0 falls into the first region, $(\underline{a}, b)$. In this situation, there are two ways exploration occurs. Either we hit $\underline{a}$ first, or we hit $\underline{b}$ first and with probability:

$$\underline{\Theta} \equiv \int_{y \notin \mathscr{A}} \theta(y, \underline{b}) \mathrm{d}y,$$

experimentation leads us to explore. Note that $\underline{\Theta}$ takes into account all possible number of successive experimentation rounds that leads us to explore; it is thus based on the function θ defined in Eq. (53). Furthermore, upon hitting $\underline{b}$ first, multiple successive rounds of experimentation may either take us back to the first region $(\underline{a}, \underline{b})$ or to the second region $(\bar{b}, \overline{a})$, which precisely couples the system of equations for $f_1(\cdot)$ and $f_2(\cdot)$. Similarly, when starting in the second region $Z_0 \in (\bar{b}, \overline{a})$, exploration may either occur as we hit $\overline{a}$ first or as we hit $\bar{b}$ first and then with probability:

$$\overline{\Theta} \equiv \int_{y \notin \mathscr{A}} \theta(y, \bar{b}) \mathrm{d}y,$$

successive rounds of experimentation lead to exploration. Formally, they satisfy:

$$f_1(t, Z) = \left[\mathbf{1}_{\tau_{\underline{a}} \leq t} \mathbf{1}_{\tau_{\underline{a}} < \tau_{\underline{b}}} | Z \in (\underline{a}, \underline{b})\right] + \left[\mathbf{1}_{\tau_{\underline{b}} < \tau_{\underline{a}}} \left(\begin{array}{c} \underline{\Theta} \mathbf{1}_{\tau_{\underline{b}} \leq t} + \int_{\underline{a}}^{\underline{b}} f_1(t, x) \theta(x, \underline{b}) \mathrm{d}x \\ + \int_{\bar{b}}^{\overline{a}} f_2(t, x) \theta(x, \underline{b}) \mathrm{d}x \end{array} \right) \middle| Z \in (\underline{a}, \underline{b}) \right], \tag{43}$$
$$f_2(t, Z) = \left[\mathbf{1}_{\tau_{\overline{a}} \leq t} \mathbf{1}_{\tau_{\overline{a}} < \tau_{\bar{b}}} | Z \in (\bar{b}, \overline{a})\right] + \left[\mathbf{1}_{\tau_{\bar{b}} < \tau_{\overline{a}}} \left(\begin{array}{c} \overline{\Theta} \mathbf{1}_{\tau_{\bar{b}} \leq t} + \int_{\underline{a}}^{\underline{b}} f_1(t, x) \theta(x, \bar{b}) \mathrm{d}x \\ + \int_{\bar{b}}^{\overline{a}} f_2(t, x) \theta(x, \bar{b}) \mathrm{d}x \end{array} \right) \middle| Z \in (\bar{b}, \overline{a}) \right].$$

We further define the following functions:

$$\theta_1(Z) \equiv \left[\mathbf{1}_{\tau_{\underline{b}} < \tau_{\underline{a}}} \middle| Z \in (\underline{a}, \underline{b})\right],$$
$$\theta_2(Z) \equiv \left[\mathbf{1}_{\tau_{\bar{b}} < \tau_{\overline{a}}} \middle| Z \in (\bar{b}, \overline{a})\right].$$

These expressions satisfy the following ODE over their respective range:

$$0 = \theta_{k,Z} + \theta_{k,ZZ}, \quad k = 1, 2,$$

with boundary conditions $\theta_1(\underline{b}) = \theta_2(\bar{b}) = 1$ and $\theta_1(\underline{a}) = \theta_2(\overline{a}) = 0$, respectively. The solution is:

$$\theta_1(Z) = \frac{e^{\underline{b}}\left(e^{\underline{a} - Z} - 1\right)}{e^{\underline{a}} - e^{\underline{b}}} \quad \text{and} \quad \theta_2(Z) = \frac{e^{\bar{b}}\left(e^{\overline{a} - Z} - 1\right)}{e^{\overline{a}} - e^{\bar{b}}}. \tag{44}$$

Similarly, we define the functions:

$$\begin{aligned} \phi_1(t,Z) &\equiv \left[\mathbf{1}_{\tau_{\underline{a}}\le t}\mathbf{1}_{\tau_{\underline{a}}<\tau_{\underline{b}}}\middle| Z\in(\underline{a},\underline{b})\right], \quad & \phi_2(t,Z) &\equiv \left[\mathbf{1}_{\tau_{\overline{a}}\le t}\mathbf{1}_{\tau_{\overline{a}}<\tau_{\overline{b}}}\middle| Z\in(\overline{b},\overline{a})\right], \\ \psi_1(t,Z) &\equiv \left[\mathbf{1}_{\tau_{\underline{b}}\le t}\mathbf{1}_{\tau_{\tau_{\underline{b}}<\underline{a}}}\middle| Z\in(\underline{a},\underline{b})\right] & \psi_2(t,Z) &\equiv \left[\mathbf{1}_{\tau_{\overline{b}}\le t}\mathbf{1}_{\tau_{\tau_{\overline{b}}<\overline{a}}}\middle| Z\in(\overline{b},\overline{a})\right]. \end{aligned}$$

These expressions satisfy the following PDE over their respective range:

$$\begin{aligned} \phi_{k,t} &= \tau_A/2(\phi_{k,Z}+\phi_{k,ZZ}) \quad \phi_k(0,Z)=0, \quad k=1,2, \\ \psi_{k,t} &= \tau_A/2(\psi_{k,Z}+\phi_{k,ZZ}) \quad \psi_k(0,Z)=0, \quad k=1,2, \end{aligned}$$

with boundary conditions:

$$\begin{aligned} \phi_1(t,\underline{a}) &= 1, \quad \phi_1(t,\underline{b})=0, \\ \phi_2(t,\overline{a}) &= 1, \quad \phi_2(t,\overline{b})=0, \\ \psi_1(t,\underline{b}) &= 1, \quad \psi_1(t,\underline{a})=0 \\ \psi_2(t,\overline{b}) &= 1, \quad \psi_2(t,\overline{a})=0. \end{aligned}$$

Although these expressions do not have analytical solutions, their Laplace transforms do. Denoting with a $\widehat{g}$ the Laplace transform of a function g and denoting by λ its frequency, we have:

$$\begin{aligned} \widehat{\phi}_1(\lambda,Z) &= \left(e^{\frac{(\underline{a}-Z)(\sqrt{8\lambda+\tau_A}+\sqrt{\tau_A})}{2\sqrt{\tau_A}}}\left(e^{\frac{Z\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{\underline{b}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}\right)\right)\Bigg/\left(e^{\frac{\underline{a}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{\underline{b}\sqrt{8\lambda_A+\tau}}{\sqrt{\tau_A}}}\right), \\ \widehat{\phi}_2(\lambda,Z) &= \left(e^{\frac{(\overline{a}-Z)(\sqrt{8\lambda+\tau_A}+\sqrt{\tau_A})}{2\sqrt{\tau_A}}}\left(e^{\frac{Z\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{\overline{b}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}\right)\right)\Bigg/\left(e^{\frac{\overline{a}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{\overline{b}\sqrt{8\lambda_A+\tau}}{\sqrt{\tau_A}}}\right), \\ \psi_1(\lambda,Z) &= \left(\left(e^{\frac{\underline{a}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{Z\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}\right)e^{\frac{(\underline{b}-Z)(\sqrt{8\lambda+\tau_A}+\sqrt{\tau_A})}{2\sqrt{\tau_A}}}\right)\Bigg/\left(e^{\frac{\underline{a}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{\underline{b}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}\right), \\ \psi_2(\lambda,Z) &= \left(\left(e^{\frac{\overline{a}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{Z\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}\right)e^{\frac{(\overline{b}-Z)(\sqrt{8\lambda+\tau_A}+\sqrt{\tau_A})}{2\sqrt{\tau_A}}}\right)\Bigg/\left(e^{\frac{\overline{a}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}-e^{\frac{\overline{b}\sqrt{8\lambda+\tau_A}}{\sqrt{\tau_A}}}\right). \end{aligned}$$

We can then recover each function through Laplace inversion. We finally rewrite the system in Eq. (43) as:

$$\begin{aligned} f_1(t,Z) &= \phi_1(t,Z)+\underline{\Theta}\psi_1(t,Z)+\theta_1(Z)\left(\int_{\underline{a}}^{\underline{b}} f_1(t,x)\theta(x,\underline{b})\mathrm{d}x+\int_{\overline{b}}^{\overline{a}} f_2(t,x)\theta(x,\underline{b})\mathrm{d}x\right) \\ f_2(t,Z) &= \phi_2(t,Z)+\overline{\Theta}\psi_2(t,Z)+\theta_2(Z)\left(\int_{\underline{a}}^{\overline{b}} f_1(t,x)\theta(x,\overline{b})\mathrm{d}x+\int_{\overline{b}}^{\overline{a}} f_2(t,x)\theta(x,\overline{b})\mathrm{d}x\right). \end{aligned}$$

Multiply first both sides of each equation by $\theta(x,\underline{b})$, integrate the first one over the first region, the second one over the second region, add these two equations together; then multiply both sides of each equation by $\theta(x,\bar{b})$ and repeat the same operations. The resulting system of equations yields:

$$\begin{pmatrix} \int_{\underline{a}}^{\underline{b}} f_1(t,x)\theta(x,\underline{b})\mathrm{d}x + \int_{\bar{b}}^{\bar{a}} f_2(t,x)\theta(x,\underline{b})\mathrm{d}x \\ \int_{\underline{a}}^{\underline{b}} f_1(t,x)\theta(x,\bar{b})\mathrm{d}x + \int_{\bar{b}}^{\bar{a}} f_2(t,x)\theta(x,\bar{b})\mathrm{d}x \end{pmatrix} = -\begin{pmatrix} \int_{\underline{a}}^{\underline{b}} \theta_1(x)\theta(x,\underline{b})\mathrm{d}x - 1 & \int_{\bar{b}}^{\bar{a}} \theta_2(x)\theta(x,\underline{b})\mathrm{d}x \\ \int_{\underline{a}}^{\underline{b}} \theta_1(x)\theta(x,\bar{b})\mathrm{d}x & \int_{\bar{b}}^{\bar{a}} \theta_2(x)\theta(x,\bar{b})\mathrm{d}x - 1 \end{pmatrix}^{-1}$$

$$\begin{pmatrix} \int_{\underline{a}}^{\underline{b}} (\phi_1(t,x) + \underline{\Theta}\psi_1(t,x))\theta(x;\underline{b})\mathrm{d}s + \int_{\bar{b}}^{\bar{a}} (\phi_2(t,x) + \overline{\Theta}\psi_2(t,x))\theta(x;\underline{b})\mathrm{d}s \\ \int_{\underline{a}}^{\underline{b}} (\phi_1(t,x) + \underline{\Theta}\psi_1(t,x))\theta(x;\bar{b})\mathrm{d}s + \int_{\bar{b}}^{\bar{a}} (\phi_2(t,x) + \overline{\Theta}\psi_2(t,x))\theta(x;\bar{b})\mathrm{d}s \end{pmatrix}.$$

Substituting back into Eq. (43) delivers an explicit solution for the two functions $f_1(\cdot)$ and $f_2(\cdot)$. Finally, using these functions we compute the cdf when starting in the experiment region, $Z \in \mathscr{B}$:

$$\begin{aligned} f_3(t,Z) \equiv & \int_{\underline{a}}^{\underline{b}} f_1(t,x)\varphi(x,i(Z),Z)\mathrm{d}x + \int_{\bar{b}}^{\bar{a}} f_2(t,x)\varphi(x,i(Z),Z)\mathrm{d}x \\ & + \int_{\mathscr{B}} f_3(t,x)\varphi(x,i(Z),Z)\mathrm{d}x + \int_{x\notin\mathscr{A}} \varphi(x,i(Z),Z)\mathrm{d}x. \end{aligned}$$

Iterating over this relation and evaluating it at $Z = 0$ we obtain the desired cdf:

$$G(t) \equiv f_3(t,0) = \int_{\underline{a}}^{\underline{b}} f_1(t,x)\theta(x,0)\mathrm{d}x + \int_{\bar{b}}^{\bar{a}} f_2(t,x)\theta(x,0)\mathrm{d}x + \int_{x\notin\mathscr{A}} \theta(x,0)\mathrm{d}x.$$

G Distribution of "In-and-Out" Experiments

In this appendix we derive the distribution of "in-and-out" experimentation trials discussed in Section 6.2. We begin by noting that, upon entering the experimentation band $\mathscr{B}$, knowledge Z becomes an AR(1) process. Hence, the distribution of number of experimentation rounds (as per Definition 4) is that of the first time exit of an AR(1) out of a range (e.g., Basak and Ho (2004)). Formally, the probability that experimentation ends at the $n-$th round is:

$$\left[Z^{(n)} \notin \mathscr{B}, Z^{(n-1)} \in \mathscr{B}, \cdots, Z^{(1)} \in \mathscr{B} \middle| Z^{(0)} = z \in \mathscr{B} \right] \equiv p(z,n).$$

This probability can be written as

$$p(z,n) = \begin{cases} 1 - q_1(z) & n = 1 \\ q_{n-1}(z) - q_n(z) & n > 1 \end{cases}, \tag{45}$$

where $q_n(z) = \left[\cap_{i=1}^n \{Z^{(i)} \in \mathscr{B}\} | Z^{(0)} = z\right]$.

Similarly, to define the density of the time at which the sequence in Eq. (30) ends, we can use the computations of Appendix F. In particular, we need the probability in Eq. (44), $\theta_1(z) = [\tau_{\underline{b}} < \tau_{\underline{a}} | z \in (\underline{a}, \underline{b})]$, that Z starting in the lower inaction region hits the experimentation band first.

Similarly, we need $\theta_2(z) = [\tau_{\bar{b}} < \tau_{\bar{a}} | z \in (\bar{b}, \bar{a})]$. We will also use the density of the stopped process $Z_{\theta+} = y$ just after it got of the experimentation band; this density is given by $\theta(y,z)$, which is defined in Eq. (53) below. We can then define the density g of $Z_{\theta_n +}$ after going in and out back into an inaction region $n-1$ times recursively as:

$$
\begin{aligned}
g_n(y,z) &\equiv \left[Z_{\theta_n +} = y \middle| \cap_{i=1}^{n} \{Z_{\theta_i +} \in \mathscr{A} \setminus \mathscr{B}\}, Z_0 = z \right], \\
&= \theta(y, \underline{b}) \int_{\underline{a}}^{\underline{b}} g_{n-1}(x,z)\theta_1(x)\mathrm{d}x + \theta(y, \bar{b}) \int_{\bar{b}}^{\bar{a}} g_{n-1}(x,z)\theta_2(x)\mathrm{d}x,
\end{aligned}
$$

condional on z drawn from an inaction region; accordingly, defining:

$$
q_n(z) \equiv \int_{\underline{a}}^{\underline{b}} g_n(x,z)\mathrm{d}x + \int_{\bar{b}}^{\bar{a}} g_n(x,z)\mathrm{d}x,
$$

we obtain the probability that the cycle of "in-and-out" rounds ends after the $n-th$ round using (45). Note that the density $p(z,n)$ holds for any z drawn from an inaction region. We finally need to average over these inaction regions. To do so, we use the stationary density of knowledge, which we compute next.

We derive the stationary distribution of knowledge based on its KBE. Define the cumulative probability that Z_T is below level y at time T given that it started at level x at time t as:

$$
\phi(y, x, T-t) := \left[\mathbf{1}_{Z_T \leq y} \middle| Z_t = x \right].
$$

This expression is a martingale, which implies that in the inaction regions this function satisfies the PDE:

$$
\phi_t = \tfrac{\tau_A}{2}(\phi_x x + \phi_{xx}), \quad \phi(y,x,0) = \mathbf{1}_{x \leq y}, \quad Z \in (\underline{a}, \underline{b}) \cup (\bar{b}, \bar{a}) \ . \tag{46}
$$

Furthermore, for this expression to be a martingale at $Z = \underline{a}$, $Z \equiv \bar{a}$, and over the experimentation region $Z \in \mathscr{B}$, we must further ensure that there is no jump in the function at these points and over the experimentation region:

$$
\phi(y, \bar{a}, T-t) = \phi(y, \underline{a}, T-t) = \phi(y, 0, T-t) \tag{47}
$$

and

$$
\phi(y, Z, T-t) = \int_{\underline{a}}^{\bar{a}} \phi(y,x,T-t)\varphi(x, i(Z), Z)\mathrm{d}x + \phi(y,0,T-t) \int_{x \notin (\underline{a}, \bar{a})} \varphi(x, i(Z), Z)\mathrm{d}x, \quad Z \in \mathscr{B}, \tag{48}
$$

where $\varphi(\cdot)$ denotes the normal density with mean $\sqrt{1 + \tau_S i(Z)}Z$ and variance $\tau_S i(Z)$.

Since we are interested in deriving the stationary distribution, we introduce the following time transform:

$$\psi(y,x;\lambda) \equiv \int_0^{+\infty} \exp(-\lambda\tau)\phi(y,x,\tau)\mathrm{d}\tau.$$

The Kolmogorov Backward Equation (46) becomes an ODE:

$$0 = \lambda(\mathbf{1}_{x\leq y} - \psi) + \frac{\tau_A}{2}(\psi_x x + \psi_{xx}), \quad x \in (\underline{a},\underline{b}) \cup (\overline{b},\overline{a}) \tag{49}$$

and Conditions (47)–(48) become:

$$\psi(y,\overline{a};\lambda) = \psi(y,\underline{a};\lambda) = \psi(y,0;\lambda) \tag{50}$$

and

$$\psi(y,Z;\lambda) = \int_{\underline{a}}^{\overline{a}} \psi(y,x;\lambda)\varphi(x,i(Z),Z)\mathrm{d}x + \psi(y,0;\lambda)\int_{x\notin(\underline{a},\overline{a})} \varphi(x,i(Z),Z)\mathrm{d}x, \quad Z \in \mathscr{B}. \tag{51}$$

We will further conjecture that:

$$\phi \in \mathscr{C}^1((\underline{a},\underline{b}) \cup (\overline{b},\overline{a})) \cap \mathscr{C}^0([\underline{a},\overline{a}]).$$

This in turn provides the remaining conditions needed to solve the problem:

$$\begin{aligned} \lim_{x\nearrow y} \phi(y,y;\lambda) &= \lim_{x\searrow y} \phi(y,y;\lambda), \quad \forall y \in [\underline{a},\overline{a}] \\ \lim_{x\nearrow y} \phi'(y,y;\lambda) &= \lim_{x\searrow y} \phi'(y,y;\lambda), \quad \forall y \in (\underline{a},\underline{b}) \cup (\overline{b},\overline{a}). \end{aligned} \tag{52}$$

The idea is to solve the ODE in Eq. (49) piecewise over the relevant smoothness regions of Z and the integral equation in Eq. (50) over $\mathscr{B}$. For $y < \underline{b}$, there are three smoothness regions: $x \in (\underline{a}, y)$, $x \in (y, \underline{b})$, and $x \in (\overline{b}, \overline{a})$; Importantly, for $y \in \mathscr{B}$, the cdf is flat everywhere, except at 0. At $y = 0$ the density exhibits an atom, which we denote by A, because every time firms experiment there is a finite probability that they end up exploring. Otherwise, i.e., $y \in \mathscr{B} \setminus \{0\}$, the density is flat as experimentation makes sure mass does not accumulate there; finally, for $y \in (\overline{b}, \overline{a})$, the smoothness regions are: $(\underline{a}, \underline{b})$, $(\overline{b}, y)$, and $(y, \overline{a})$. We thus solve the cdf for the cases $y < \underline{b}$ and $y > \overline{b}$, and then obtain the atom, A, by requiring that the cdf reaches 1. Starting with $y < \underline{b}$, let us index the regions $(\underline{a}, y)$, $(y, \underline{b})$, and $(\overline{b}, \overline{a})$ by (1), (2), and (3), respectively, and accordingly denote the solution to the ODE in Eq. (49) over these three regions by:

$$\psi^{(1)}(y,x;\lambda) = 1 + C_{1,l}(y,\lambda)e^{-x^2/2}H(-1-2\lambda/\tau_A, x/\sqrt{2}) + C_{2,l}(y,\lambda)e^{-x^2/2}M(-2\lambda/\tau_A, 1/2, -x^2/\sqrt{2}),$$

$$\psi^{(2)}(y,x;\lambda) = C_{1,m}(y,\lambda)e^{-x^2/2}H(-1-2\lambda/\tau_A, x/\sqrt{2}) + C_{2,m}(y,\lambda)e^{-x^2/2}M(-2\lambda/\tau_A, 1/2, -x^2/\sqrt{2}),$$

$$\psi^{(3)}(y,x;\lambda) = C_{1,h}(y,\lambda)e^{-x^2/2}H(-1-2\lambda/\tau_A, x/\sqrt{2}) + C_{2,h}(y,\lambda)e^{-x^2/2}M(-2\lambda/\tau_A, 1/2, -x^2/\sqrt{2}),$$

where $C_{1,l}$, $C_{2,l}$, $C_{1,m}$, $C_{2,m}$, $C_{1,h}$, and $C_{2,h}$ are six unknown integration coefficients. Furthermore, to write the solution to the integral equation in Eq. (51), we first define the density outside the experimentation region but inside the no-exploration region, i.e., within $\mathscr{A} \setminus \mathscr{B}$, as:

$$p(y,x;\lambda) \equiv \psi^{(1)}(y,x;\lambda)\mathbf{1}_{x\in(\underline{a},y]} + \psi^{(2)}(y,x;\lambda)\mathbf{1}_{x\in(y,\underline{b})} + \psi^{(3)}(y,x;\lambda)\mathbf{1}_{x\in[\bar{b},\bar{a})}.$$

Denoting the density over $\mathscr{B}$ by $\psi^{(4)}$ we can rewrite the integral equation in Eq. (51) as:

$$\begin{aligned}\psi^{(4)}(y,x;\lambda) = &\int_{\mathscr{A}\setminus\mathscr{B}} p(y,x;\lambda)\varphi(x,i(Z),Z)\mathrm{d}x + \int_{\mathscr{B}} \psi^{(4)}(y,x;\lambda)\varphi(x,i(Z),Z)\mathrm{d}x \\ &+ \psi^{(4)}(y,0;\lambda)\int_{x\notin\mathscr{A}} \varphi(x,i(Z),Z)\mathrm{d}x\end{aligned}$$

Iterating over this equation, define the following Liouville-Neumann series:

$$\theta(y,Z) \equiv \sum_{n=0}^{\infty} \theta^{(n)}(y,Z), \tag{53}$$

where $\theta^{(n)}$ satisfies for $n \in$:

$$\theta^{(n)}(y,Z) \equiv \int_{\mathscr{B}^n} \varphi(y,i(y^{(n)}),y^{(n)})\varphi(y^{(n)},i(y^{(n-1)}),y^{(n-1)})\cdots\varphi(y^{(1)},i(Z),Z)\mathrm{d}y^{(n)}\mathrm{d}y^{(n-1)}\cdots\mathrm{d}y^{(1)}.$$

Note that $\theta(y;Z)$ in Eq. (53) can be alternatively represented recursively as a Fredholm equation of the second kind:

$$\theta(y,Z) = \varphi(y,i(Z),Z) + \int_{\mathscr{B}} \theta(x,Z)\varphi(y,i(x),x)\mathrm{d}x.$$

The relevance of this alternative representation is that, given that φ is a probability density, the contraction mapping theorem (Theorem 3.2, Stokey and Lucas (1996)) guarantees that the representation as an infinite sum in Eq. (53) exists. We can then write the solution to Eq. (51)

as:

$$\psi^{(4)}(y,Z;\lambda)=\int_{\mathscr{A}\setminus\mathscr{B}}p(y,x;\lambda)\theta(x,Z)\mathrm{d}x+\psi^{(4)}(y,0;\lambda)\int_{x\notin\mathscr{A}}\theta(x,Z)\mathrm{d}x$$

Evaluating this expression at $x\equiv 0$ and solving for $\psi^{(4)}(y,0;\lambda)$ we get an explicit solution of the form:

$$\psi^{(4)}(y,x;\lambda)=\int_{\mathscr{A}\setminus\mathscr{B}}p(y,x;\lambda)\theta(x,Z)\mathrm{d}x+\frac{\int_{\mathscr{A}\setminus\mathscr{B}}p(y,x;\lambda)\theta(x,0)\mathrm{d}x}{1-\int_{x\notin\mathscr{A}}\theta(x,0)\mathrm{d}x}\int_{x\notin\mathscr{A}}\theta(x,Z)\mathrm{d}x. \tag{54}$$

It will be subsequently convenient to simplify notations:

$$H(x;\lambda)\equiv e^{-x^2/2}H(-1-2\lambda/\tau_A,x/\sqrt{2})\quad\text{and}\quad M(x;\lambda)\equiv e^{-x^2/2}M(-2\lambda/\tau_A,1/2,-x^2/\sqrt{2}),$$

and define accordingly the following functions:

$$h_1(Z,\underline{x},\overline{x},\lambda)\equiv\textstyle\int_{\underline{x}}^{\overline{x}}H(x;\lambda)\theta(x,Z)\mathrm{d}x\quad\text{and}\quad h_2(Z,\underline{x},\overline{x},\lambda)\equiv\int_{\underline{x}}^{\overline{x}}M(x;\lambda)\theta(x,Z)\mathrm{d}x,$$

the following probabilities:

$$\Theta(Z)\equiv\textstyle\int_{x\notin\mathscr{A}}\theta(x,Z)\mathrm{d}x\quad\text{and}\quad\Phi(y,Z)\equiv\int_{\underline{a}}^{y}\theta(x,Z)\mathrm{d}x.$$

We then use these functions to further define:

$$\begin{aligned}\phi(y,Z)&\equiv\Phi(y,Z)+\frac{\Theta(Z)}{1-\Theta(0)}\Phi(y,0)\\ g_k(Z,\underline{x},\overline{x},\lambda)&\equiv h_k(Z,\underline{x},\overline{x},\lambda)+\frac{\Theta(Z)}{1-\Theta(0)}h_k(0,\underline{x},\overline{x},\lambda),\quad k=1,2.\end{aligned}$$

Using this simplified notation we can rewrite the solution in Eq. (54) as:

$$\begin{aligned}\psi^{(4)}(y,Z;\lambda)=\phi(y,Z)&+\sum_{k=1,2}C_{l,k}(y,\lambda)g_k(Z,\underline{a},y,\lambda)\\&+\sum_{k=1,2}C_{m,k}(y,\lambda)g_k(Z,y,\underline{b},\lambda)+\sum_{k=1,2}C_{h,k}(y,\lambda)g_k(Z,\overline{b},\overline{a},\lambda).\end{aligned}$$

We finally obtain a piecewise representation of the density for $y<\underline{b}$:

$$\phi(y,x;\lambda)=\begin{cases}\psi^{(1)}(y,x;\lambda) & x\in(\underline{a},y]\\ \psi^{(2)}(y,x;\lambda) & x\in(y,\underline{b})\\ \psi^{(4)}(y,x;\lambda) & x\in\mathscr{B}\\ \psi^{(3)}(y,x;\lambda) & x\in(\overline{b},\overline{a})\end{cases},$$

Similarly, for $y > \overline{b}$, we can write:

$$\phi(y,x;\lambda) = \begin{cases} \psi^{(1)}(y,x;\lambda) & x \in (\underline{a},\underline{b}) \\ \psi^{(4)}(y,x;\lambda) & x \in \mathscr{B} \\ 1+\psi^{(2)}(y,x;\lambda) & x \in (\overline{b},y] \\ \psi^{(3)}(y,x;\lambda) & x \in (y,\overline{a}) \end{cases},$$

where $\psi^{(4)}(y,x;\lambda)$ for $y > \overline{b}$ satisfies:

$$\psi^{(4)}(y,Z;\lambda) = \phi(y,Z) + \sum_{k=1,2} C_{l,k}(y,\lambda) g_k(Z,\underline{a},\underline{b},\lambda)$$
$$+ \sum_{k=1,2} C_{m,k}(y,\lambda) g_k(Z,\overline{b},y,\lambda) + \sum_{k=1,2} C_{h,k}(y,\lambda) g_k(Z,y,\overline{a},\lambda).$$

Given this piecewise representation of the cdf, Eqs. (50) and (52) constitute a system of 6 equations that pins down the 6 coefficients of integration, $C_{1,l}$, $C_{2,l}$, $C_{1,m}$, $C_{2,m}$, $C_{1,h}$, and $C_{2,h}$. The plan is to solve for these 6 coefficients and then recover the stationary CDF, π, through the Final Value Theorem:

$$\pi(y) \equiv \lim_{T\to\infty} \psi(y,x;T-t) = \lim_{\lambda\to 0} \lambda\phi(y,x;\lambda). \tag{55}$$

Imposing Eqs. (50) and (51) for $y \in (\underline{a},\underline{b}) \cup (\overline{b},\overline{a})$ yields:

$$\begin{aligned} \psi^{(1)}(y,\underline{a};\lambda) &= \psi^{(4)}(y,0;\lambda), \\ \psi^{(3)}(y,\overline{a};\lambda) &= \psi^{(4)}(y,0;\lambda), \\ \psi^{(4)}(y,\underline{b};\lambda) &= \psi^{(2)}(y,\underline{b};\lambda)\mathbf{1}_{y<\underline{b}} + \psi^{(1)}(y,\underline{b};\lambda)\mathbf{1}_{y>\overline{b}}, \\ \psi^{(4)}(y,\overline{b};\lambda) &= \psi^{(3)}(y,\overline{b};\lambda)\mathbf{1}_{y<\underline{b}} + (1+\psi^{(2)}(y,\overline{b};\lambda))\mathbf{1}_{y>\overline{b}}, \\ \psi^{(2)}(y,y;\lambda) &= \psi^{(1)}(y,y;\lambda)\mathbf{1}_{y<\underline{b}} + (\psi^{(3)}(y,y;\lambda)-1)\mathbf{1}_{y>\overline{b}}, \\ \psi^{(2)'}(y,y;\lambda) &= \psi^{(1)'}(y,y;\lambda)\mathbf{1}_{y<\underline{b}} + \psi^{(3)'}(y,y;\lambda)\mathbf{1}_{y>\overline{b}}. \end{aligned} \tag{56}$$

It is convenient to rewrite this system in matrix form, and then solve it separately for $y \in (\underline{a},\underline{b})$ and $y \in (\overline{b},\overline{a})$; we gather the 6 unknown coefficients in a vector and index them with a superscript l for $y \in (\underline{a},\underline{b})$ and h for $y \in (\overline{b},\overline{a})$:

$$C_k(y;\lambda) \equiv \begin{pmatrix} C^k_{1l}(y,\lambda) & C^k_{2l}(y,\lambda) & C^k_{1m}(y,\lambda) & C^k_{2m}(y,\lambda) & C^k_{1h}(y,\lambda) & C^k_{2h}(y,\lambda) \end{pmatrix}', \quad k = \{l,h\}.$$

Using this notation, rewrite the system in Eq. (56) over the region $y \in (\underline{a},\underline{b})$ as:

$$B_l(y) = A_l(y;\lambda) C_l(y;\lambda),$$

where

$$A_l(y;\lambda) = \begin{pmatrix} H(\underline{a};\lambda) - g_1(0,\underline{a},y,\lambda) & M(\underline{a};\lambda) - g_2(0,\underline{a},y,\lambda) & -g_1(0,y,\underline{b},\lambda) & -g_2(0,y,\underline{b},\lambda) & -g_1(0,\overline{b},\overline{a},\lambda) & -g_2(0,\overline{b},\overline{a},\lambda) \\ -g_1(0,\underline{a},y,\lambda) & -g_2(0,\underline{a},y,\lambda) & -g_1(0,y,\underline{b},\lambda) & -g_2(0,y,\underline{b},\lambda) & H(\overline{a};\lambda) - g_1(0,\overline{b},\overline{a},\lambda) & M(\overline{a};\lambda) - g_2(0,\overline{b},\overline{a},\lambda) \\ g_1(\underline{b},\underline{a},y,\lambda) & g_2(\underline{b},\underline{a},y,\lambda) & g_1(\underline{b},y,\underline{b},\lambda) - H(\underline{b};\lambda) & g_2(\underline{b},y,\underline{b},\lambda) - M(\underline{b});\lambda & g_1(\underline{b},\overline{b},\overline{a},\lambda) & g_2(\underline{b},\overline{b},\overline{a},\lambda) \\ g_1(\overline{b},\underline{a},y,\lambda) & g_2(\overline{b},\underline{a},y,\lambda) & g_1(\overline{b},y,\underline{b},\lambda) & g_2(\overline{b},y,\underline{b},\lambda) & g_1(\overline{b},\overline{b},\overline{a},\lambda) - H(\overline{b};\lambda) & g_2(\overline{b},\overline{b},\overline{a},\lambda) - H(\overline{b};\lambda) \\ -H(y;\lambda) & -M(y;\lambda) & H(y;\lambda) & M(y;\lambda) & 0 & 0 \\ -H'(y;\lambda) & -M'(y;\lambda) & H'(y;\lambda) & M'(y;\lambda) & 0 & 0 \end{pmatrix},$$

and

$$B_l(y) = \begin{pmatrix} \phi(y,0)-1 & \phi(y,0) & -\phi(y,\underline{b}) & -\phi(y,\overline{b}) & 1 & 0 \end{pmatrix}'.$$

Similarly, in the region $y \in (\overline{b},\overline{a})$ we write:

$$B_h(y) = A_h(y;\lambda)C_h(y;\lambda),$$

where

$$A_h(y;\lambda) = \begin{pmatrix} H(\underline{a};\lambda) - g_1(0,\underline{a},\underline{b},\lambda) & M(\underline{a};\lambda) - g_2(0,\underline{a},\underline{b},\lambda) & -g_1(0,\overline{b},y,\lambda) & -g_2(0,\overline{b},y,\lambda) & -g_1(0,y,\overline{a},\lambda) & -g_2(0,y,\overline{a},\lambda) \\ -g_1(0,\underline{a},\underline{b},\lambda) & -g_2(0,\underline{a},\underline{b},\lambda) & -g_1(0,\overline{b},y,\lambda) & -g_2(0,\overline{b},y,\lambda) & H(\overline{a};\lambda) - g_1(0,y,\overline{a},\lambda) & M(\overline{a};\lambda) - g_2(0,y,\overline{a},\lambda) \\ g_1(\underline{b},\underline{a},\underline{b},\lambda) - H(\underline{b};\lambda) & g_2(\underline{b},\underline{a},\underline{b},\lambda) - M(\underline{b};\lambda) & g_1(\underline{b},\overline{b},y,\lambda) & g_2(\underline{b},\overline{b},y,\lambda) & g_1(\underline{b},y,\overline{a},\lambda) & g_2(\underline{b},y,\overline{a},\lambda) \\ g_1(\overline{b},\underline{a},\underline{b},\lambda) & g_2(\overline{b},\underline{a},\underline{b},\lambda) & g_1(\overline{b},\overline{b},y,\lambda) - H(\overline{b};\lambda) & g_2(\overline{b},\overline{b},y,\lambda) - M(\overline{b};\lambda) & g_1(\overline{b},y,\overline{a},\lambda) & g_2(\overline{b},y,\overline{a},\lambda) \\ 0 & 0 & H(y;\lambda) & M(y;\lambda) & -H(y;\lambda) & -M(y;\lambda) \\ 0 & 0 & H'(y;\lambda) & M'(y;\lambda) & -H'(y;\lambda) & -M'(y;\lambda) \end{pmatrix},$$

and

$$B_h(y) = \begin{pmatrix} \phi(y,0)-1 & \phi(y,0) & 1-\phi(y,\underline{b}) & 1-\phi(y,\overline{b}) & -1 & 0 \end{pmatrix}'.$$

We can then write the solution in the two regions as:

$$C_k(y,\lambda) = A_k(y;\lambda)^{-1}B_k(y), \quad k = \{l,h\}.$$

Since we are interested in the stationary distribution only, the last step is to apply the Final Value Theorem. This implies taking the limit in Eq. (55). Note that

$$\lim_{\lambda\to 0} H(x;\lambda) = \frac{\sqrt{\pi}}{2}(1-\mathrm{erf}(x/\sqrt{2})),$$
$$\lim_{\lambda\to 0} M(y;\lambda) = 1,$$

and that at a stationary state we lose the dependence on the initial state, which implies that (and which can be verified by direct computation):

$$\begin{aligned}
&\lim_{\lambda\to 0} C_{1,l}(y,\lambda) \equiv \lim_{\lambda\to 0} C_{1,m}(y,\lambda) \equiv \lim_{\lambda\to 0} C_{1,h}(y,\lambda) \equiv 0, \\
&\lim_{\lambda\to 0} C_{2,l}(y,\lambda) \equiv \lim_{\lambda\to 0} C_{2,m}(y,\lambda) \equiv \lim_{\lambda\to 0} C_{2,h}(y,\lambda).
\end{aligned}$$

These observations, along with Eq. (55), imply that we can write the stationary cdf as:

$$\pi(y) \equiv \lim_{\lambda\to 0} 1 + C_{2,l}(y,\lambda) \equiv \lim_{\lambda\to 0} C_{2,m}(y,\lambda) \equiv \lim_{\lambda\to 0} C_{2,h}(y,\lambda).$$

To compute these limits we rewrite the solution for coefficients C as:

$$C_k(y,\lambda) = \tfrac{1}{\det(A_k(y;\lambda))} \mathrm{adj}(A_k(y;\lambda)) B_k(y), \quad k = \{l,h\},$$

where $\mathrm{adj}(A)$ denotes the adjugate of A. Note that the dependence on λ is both in the denominator through the determinant, and in the numerator through the adjugate matrix; accordingly we apply L'Hospital Theorem to write this limit as:

$$\lim_{\lambda\to 0} C_k(y,\lambda) = \frac{\lim_{\lambda\to 0} \frac{\mathrm{d}}{\mathrm{d}\lambda} \mathrm{adj}(A_k(y;\lambda)) B_k(y)}{\lim_{\lambda\to 0} \frac{\mathrm{d}}{\mathrm{d}\lambda} \det(A_k(y;\lambda))}, \quad k = \{l,h\}.$$

Finally, we obtain the density that "in-and-out" experiments end after the $n-th$ round as:

$$\int_{\underline{a}}^{\underline{b}} p(z,n)\pi'(z)\mathrm{d}z + \int_{\overline{b}}^{\overline{a}} p(z,n)\pi'(z)\mathrm{d}z.$$