

Growth Option Based Model for R&D-Data Asset Pricing

Hang Zeng^{1,2,†} and Liyan Han^{3,4,*}

¹Academy of Mathematics and Systems Science, Chinese Academy of Sciences,
Emails: 572579935@qq.com

²Beijing Institute of Mathematical Sciences and Applications

³Lab of Digital Economy, Beijing Institute of Mathematical Sciences and Applications

⁴School of Economics and Management, Beihang University, Email: hanly@buaa.edu.cn

Abstract

In response to the growing demand for valuing data alongside physical capital and technology, this paper develops a dynamic growth-option framework for pricing firm-level data and R&D assets in a market setting. We build a dynamic investment model with three capital stocks—physical capital, data assets, and firm-specific technology—embedded in a risk-neutral valuation measure with an innovation kernel and internal financing constraints. Data is modeled as an option-like intangible: the incentive to accumulate it is generated by continuation value, while the marginal payoff to data investment vanishes in a myopic static benchmark. We derive shadow prices of data and technology as marginal values of the dynamic program and obtain threshold-type investment rules via KKT conditions. The resulting switching frontiers characterize when the firm invests in data or R&D, or waits, and how the exercise regions shift with firm-specific technology and the shadow cost of internal funds. Endogenizing technology yields a coupled shadow-price system: under complementarity, stronger technology shifts the data-investment frontier outward, while richer data raises the shadow value of technology by strengthening innovation intensity. Numerical value iteration and frontier extraction illustrate these mechanisms and show that the main comparative statics are robust. Overall, data becomes economically significant primarily through its synergy with R&D and technology capital, rather than by operating in isolation.

Keywords: data assets; real options; R&D investment; shadow prices; dynamic programming; technology capital.

1 Introduction

In the digital economy era, data serves as the cornerstone of business intelligence. As a new type of production factor, data possesses the potential to be converted into corporate assets, thereby creating new profit streams and driving sustainable growth. From a broader academic perspective, Veldkamp and Chung (2024) examine how digitized information (“data”) can be modeled within an economic framework. They distinguish data from technology by emphasizing its core characteristics—non-rivalry, by-product nature, and limited spillover—and further embed data into analytical frameworks of growth, information frictions, and intangible capital. Their research illustrates how data reshapes productivity, firm dynamics, the measurement of GDP and welfare, as well as mechanisms of trade and monetary transmission, while highlighting unresolved issues regarding data valuation, ownership allocation, and the modeling of data-driven growth.

Pricing constitutes a core analytical focus and fundamental concept in economics. From a firm-level perspective, when integrating data into value-generating processes, the first and foremost requirement is to conduct economically reasonable pricing or valuation of data within specific business scenarios. However, our field investigations into firms’ specific business practices reveal that the isolated deployment of data assets rarely yields tangible outcomes. Instead, the true value of data is realized only when it is integrated with other assets—whether tangible or intangible. Abis and Veldkamp (2024) explore the impact of big data technology on knowledge production mechanisms through the lens of data-human labor combinations. Their findings indicate that the adoption of big data technology may reduce the labor share in knowledge work by 5%, a magnitude comparable to the labor structure changes induced by the Industrial Revolution. Furthermore, by leveraging publicly available recruitment and wage data, they have developed a novel data valuation tool to more accurately measure corporate value, market competition, and overall output and investment levels.

Complementing this firm-level evidence, Jones and Tonetti (2020) model data as a non-rivalrous production factor whose accumulation can sustain long-term economic growth. Cong et al. (2021) investigate how privacy regulations interact with data usage and knowledge accumulation in a data-driven economy. At the macroeconomic level, Veldkamp and Chung (2024) systematically review how data accumulation feeds into aggregate output, investment, and valuation. Notably, even though data may be technologically reusable within a single firm, the information advantage derived from it becomes rivalrous across competing firms and tends to dissipate once exploited due to market competition. While Cong et al. (2021) have developed a macroeconomic framework where data accumulation fuels aggregate knowledge growth, this paper complements their perspective by focusing on firm-level investment and valuation issues. We treat data as a proprietary, rivalrous input, whose value is determined by its option-like role in the innovation process under capital constraints.

Guided by these microeconomic, macroeconomic, and asset-pricing perspectives, this study focuses on the firm-level challenge of data valuation under conditions of capital constraint and intertemporal irreversibility of investment. Given that pricing is central to microeconomics and corporate decision-making, the valuation of data assets is pivotal to business digitization. Its significance permeates all aspects of firms’ digitization efforts, whether for internally assess-

ing the marginal contribution of data in business processes or for external trading activities. However, direct market prices for firm-level data assets are rarely observable, which necessitates a theoretical study employing a structural valuation approach. Veldkamp (2023) develops an explicit asset-pricing framework that values data as an intangible capital stock. Aligned with intangible asset valuation practices, the real options approach—with its robust theoretical development—emerges as a more scientific methodology (Ai et al. (2013) Ai et al. (2018)).

Pakes (1986) pioneers the application of real options theory to patent valuation and provides empirical illustrations using patent data from major European listed companies. Building on this foundation, Berk et al. (1999) take a groundbreaking step by integrating growth options into corporate valuation, proposing that firm value comprises two core components: the present value of cash flows generated by existing operations, and the value of growth options embedded in future investment opportunities. Within a dynamic asset-pricing framework, they demonstrate that growth options arise endogenously from uncertainties in future profitability and strategic flexibility. Their model shows that firms with superior growth opportunities (e.g., those in R&D-intensive industries) command higher valuations, as these options represent claims on uncertain future payoffs. This work establishes a direct link between growth option theory and firm value, highlighting that traditional discounted cash flow (DCF) methods tend to underestimate value by ignoring optionality, thereby laying the foundation for incorporating real options into corporate finance and asset pricing models. Aguerrevere (2009) extends framework of Berk et al. (1999) by integrating growth option theory into competitive market analysis, illustrating how firm valuations reflect both existing cash flows and strategic expansion options amid market competition. In a dynamic competitive setting, this research formalizes the interaction between growth options and market competition in shaping firm value. Kogan and Papanikolaou (2012) further demonstrate that growth options not only capture firms’ sensitivity to technological shocks but also serve as a valuable tool for corporate valuation. Their findings indicate that firms with significant growth options exhibit higher exposure to technological innovation, as these options embody future scalability and adaptability to disruptive technological advancements. Ai et al. (2013) expand growth option framework of Berk et al. (1999) by embedding growth options as intangible assets within a general equilibrium model. A key innovation of their work is the introduction of a “storage capacity” for growth options, enabling firms to retain unexercised options for future use—distinct from framework of Berk et al. (1999) where unexercised growth options expire immediately. This approach explicitly models growth options as intangible assets, linking them to balance sheet considerations.

Regarding the relationship between data and conventional intangible assets, Veldkamp (2023) emphasizes that “data may contribute to each of these intangibles, but it is not equivalent to any of them.” Distinguishing the value of data from that of other intangible assets presents a non-trivial challenge; it cannot be simply derived from market participants’ assessments and requires further theoretical and practical exploration. Building on this line of inquiry, we extend the real options framework by incorporating data assets as a central component of valuation analysis.

Leveraging data’s unique characteristics—such as scalability, non-rivalry as a production input, and dynamic adaptability—this extended framework models strategic decision-making

under uncertainty. This extension not only aligns with emerging empirical evidence regarding the growing economic impact of data-driven investments but also formalizes how intangible data repositories function as “optionable” resources capable of generating future optionality through iterative learning and technological evolution. Methodologically, such option-based pricing should be derived from an equilibrium state and embedded within an integrated framework that combines short-term and long-term perspectives. Consequently, the real options valuation of data assets must be incorporated into investment decision frameworks alongside tangible and other intangible assets. Ai et al. (2013) and Ai et al. (2018) further extend framework of Berk et al. (1999) by developing a general equilibrium model where intangible assets are conceptualized as growth options. In their model, growth options can be stored for future exercise—contrasting with Berk et al. (1999) setting where options must be exercised immediately or forfeited. Moreover, their model endogenizes growth options through agents’ optimal choices, allowing these options to dynamically adapt to productivity shocks within the equilibrium process.

Building on the analytical framework of endogenous market power dynamics, we propose a synergistic interaction among physical capital, technology capital, and data assets in driving firm-level investment returns. While physical asset investments generate direct output growth through capital deepening (consistent with Romer (1990)), technology capital and data assets exhibit time-dependent convexity in productivity spillovers—manifesting as delayed yet non-linear returns that amplify over the technology adoption lifecycle (in line with Arthur (1989)). Crucially, firm-level empirical evidence reveals that data assets cannot independently generate economic surplus; their value realization depends on symbiotic integration with complementary intangible assets (e.g., patents, organizational know-how) and physical infrastructure—consistent with the dynamic framework proposed by Acemoglu and Jensen (2023). Their work develops an economic growth model incorporating multiple behavioral preferences; addressing the equilibrium problem of unilateral growth behavior, they introduce the concept of price-adjusted equilibrium response, establishing an interactive mechanism between short-term direct responses and long-term equilibrium responses. This interdependence gives rise to a two-phase dominance mechanism: initially, data accumulation enhances the market power of technological assets through data-network complementarities, thereby increasing rent-capturing capacity (Gawer and Henderson (2008)); subsequently, the resulting monopolistic dominance enables firms to internalize cross-asset spillovers, effectively transforming speculative expectations of future cash flows into optionally exercisable economic rents (Dixit and Pindyck (1994)). Therefore, we aggregate data assets with technological intangible assets into physical investment to generate a new comprehensive productivity measure, while defining the marginal value of data assets at equilibrium as the growth option for data asset pricing.

Main results and contributions. We study how a financially constrained firm values and accumulates proprietary data when data is productive primarily through complementarities with firm-specific technology and when investment is partially irreversible. The key object is the *shadow price of data*—the marginal value of the firm’s continuation value with respect to the data stock—which naturally admits a real-options interpretation.

This paper makes three contributions. First, we develop a dynamic investment model (dy-

namic program) in which the firm allocates internal resources across physical capital, data assets, and technology capital. A central modeling choice is a Variable Elasticity of Substitution (VES) specification for the interaction between data and technology in the innovation channel, allowing complementarity strength to vary endogenously across states. Within this framework, data behaves as an option-like intangible: the incentive to accumulate it is primarily driven by continuation value in the dynamic problem, whereas the marginal payoff to data investment vanishes in the one-step myopic benchmark, where J_t is treated as a pure cost. We provide conditions under which the marginal innovation return to data is positive and economically sizable, so that data functions as a productive input in the innovation transformation rather than as an accounting artifact.

Second, we characterize threshold-type optimal policies that determine *when to invest in data, when to invest in R&D, and when to wait*. The optimality conditions generate state-dependent switching frontiers that partition the state space into inaction, data-investment, R&D-investment, and joint-investment regions. Under complementarity and monotonicity conditions, the data-investment region expands with the firm’s technology capital A_t : when A_t is low, data accumulation is typically deferred; when A_t is high, data investment becomes optimal over a substantially larger set of states. Under standard feasibility and regularity conditions (detailed in the Appendix), the resulting stationary policy is unique and yields a managerially interpretable capital-allocation rule. In addition, we use a continuous-time benchmark as a robustness check: it provides a free-boundary interpretation (HJB/variational-inequality) for the switching frontiers, and the shadow price admits a diffusion-based characterization under the valuation measure.

Third, we derive and interpret a coupled shadow-price system for data and technology capital. The shadow price of data internalizes how technology capital shifts innovation productivity, while the shadow price of technology internalizes how data shifts innovation productivity. Relative to the benchmark with exogenously fixed firm-specific technology, endogenizing technology through forward-looking R&D tends to raise the continuation value of data and shift the data-investment frontier outward; conversely, richer data stocks tend to increase the shadow value of technology by strengthening innovation intensity. Together, the model provides a tractable framework for valuing data and R&D as interacting real options and for generating comparative statics and empirically testable implications for long-run productivity and investment dynamics.

Organization of the paper. The remainder of the paper proceeds as follows. Section 2 introduces the baseline model and assumptions. Section 3 studies the benchmark with fixed firm-specific technology and characterizes optimal allocations between physical capital and data assets. Section 4 interprets data as a real option and derives the associated pricing and investment thresholds. Section 5 endogenizes firm-specific technology via R&D and develops the joint policy for physical capital, data, and technology capital. Section 6 contrasts regimes with exogenous versus endogenous technology. Section 7 provides numerical simulations and comparative statics. Section 8 concludes. The Appendix contains proofs and additional derivations.

2 Model Framework

We consider a discrete-time, infinite-horizon economy in which a financially constrained firm optimizes its investment strategy across multiple asset classes. The model integrates tangible physical capital with two distinct intangible assets—proprietary data and firm-specific technology—into a unified dynamic programming framework. The core mechanism is an innovation kernel that transforms investment inputs into productive capacity, subject to the availability of internal funds. The remainder of this section details the production technology, capital accumulation dynamics, and the optimization problem.

2.1 Environment, States, and Controls

We model a forward-looking firm that allocates internal resources across three interrelated capital stocks: physical capital K_t , data assets D_t , and firm-specific technology capital A_t .¹

States. The endogenous state vector at time t is

$$S_t = (K_t, D_t, A_t) \in \mathbb{R}_{++}^3,$$

augmented by an exogenous productivity shock Z_t (detailed in Section 2.2). Physical capital K_t represents installed productive capacity. Technology capital A_t captures accumulated capabilities—such as proprietary algorithms, patents, and organizational know-how—that enhance output and complement data in innovation. Data assets D_t consist of proprietary datasets and organizational knowledge that are deployed and partially depleted during innovation. We interpret D_t as the firm’s *effective proprietary information advantage* embodied in data (rather than raw storage), i.e., the subset of data value that can be converted into innovation before being eroded by diffusion, learning, and arbitrage. In our analytical characterization, we focus on the interior region where admissibility constraints are slack (as constructed in Appendix A), ensuring that first-order conditions are well-defined.

Controls and information. The control vector is $U_t = (I_t, J_t, R_t) \in \mathbb{R}_+^3$, where I_t denotes implementation expenditures (engineering integration, computational resources, and product testing), J_t represents data acquisition and governance, and R_t corresponds to R&D spending that replenishes technology capital. Controls are after observing the full state (S_t, Z_t) and before the realization of next-period shocks. Formally, U_t is $\mathcal{F}_t$ -measurable, where $(\mathcal{F}_t)$ is the filtration generated by the history of exogenous shocks and initial conditions. See Table 1 for a summary of all symbols and financial analogues.

Aggregate technology. Operating output is scaled by an aggregate technology factor Z_t according to (1). For clarity, Sections 3–5 adopt the normalization $Z_t \equiv 1$ and focus on the state vector $S_t = (K_t, D_t, A_t)$. Section 4.1 reinstates stochastic Z_t and utilizes the augmented state (K_t, D_t, A_t, Z_t) .

¹We use “assets” as the overarching concept, while “capital” specifically refers to the subset of assets deployed for value creation. Since equity issuance and payout policy design are outside our scope, we treat data and technology as firm-owned intangible assets that enter value creation and investment decisions through their economic usefulness and control rights.

Table 1: Model symbols and economic interpretation

Symbol	Definition	Economic role
K_t	Physical capital	Servers, equipment, facilities
D_t	Data assets	Proprietary databases, logs, behavioral datasets
A_t	Technology capital	Proprietary algorithms, patents, organizational know-how
I_t	Implementation effort	Deployment, integration, compute, experimentation
J_t	Data investment	Acquisition, cleaning, storage, governance
R_t	R&D expenditure	Frontier extension and capability building
G_t	Realized innovation flow	Implemented innovation produced from (I_t, D_t, A_t)
Z_t	Aggregate technology index	Economy-wide technology factor (common across firms)
w_t	Implementation efficiency shock	Efficiency converting G_t into K_{t+1} , execution risk
τ	Reinvestment rate	Fraction of output reinvested each period
p_t^D	Shadow price of D_t	Marginal value of data in the dynamic program
p_t^A	Shadow price of A_t	Marginal value of technology capital in the dynamic program

2.2 Production and Innovation

Operating output. Output is produced via a Cobb–Douglas technology:

$$Y_t = (A_t Z_t)^{1-\xi} K_t^\xi, \quad \xi \in (0, 1). \quad (1)$$

When $Z_t \equiv 1$ we write $Y_t = Y(K_t, A_t)$.

Innovation (implemented innovation flow). Innovation is generated by combining implementation effort with intangible stocks (technology and data). Define the intangible composite

$$H_t := (A_t^\alpha + \theta D_t^\alpha)^{1/\alpha},$$

where $\alpha \in (0, 1)$ and $\theta \in (0, 1]$ govern curvature and the relative weight of data in the composite. We adopt a Revankar-inspired specification in which the elasticity of substitution between implementation effort and the effective intangible capacity varies with the ratio I_t/H_t :²

$$G_t = \beta \exp\left(\frac{a I_t}{H_t}\right) I_t^{1-b} A_t^b D_t^b, \quad (2)$$

where $a, \beta > 0$ scale innovation and $b \in (0, 1)$ governs intangible intensity in the scale term. The factor $A_t^b D_t^b$ captures the scale contribution of technology and data, while $\exp(a I_t/H_t)$ captures a state-dependent implementation margin that can reflect deployment and integration frictions. As H_t increases for given I_t , the ratio I_t/H_t falls and the exponential term is dampened; this

²A standard Cobb–Douglas (or CES) innovation function imposes a constant elasticity of substitution between implementation effort and intangible capacity. This restriction rules out an empirically salient feature in data-intensive production: when data and technology are abundant but engineering capacity is scarce, scaling implementation becomes bottlenecked by integration, cleaning, governance, and compliance, so the marginal value of additional data can flatten and, locally, become non-monotone (and potentially weak at the margin) in effective terms. To capture this *scale–congestion trade-off* parsimoniously, we adopt a Revankar-inspired variable-elasticity specification in which the effective implementability of innovation depends on the intensity ratio I_t/H_t (“data gravity”): holding I_t fixed, a larger H_t dampens the implementation margin, reflecting deployment frictions that are not well represented by constant-elasticity kernels. This congestion channel is consistent with documented scaling frictions in data governance and compliance that rise with data scope and heterogeneity (The Royal Society and The British Academy, 2017).

feature allows data- and technology-rich firms to face integration and governance complexity when scaling implementation (e.g., cleaning, alignment, and compliance constraints).³

Lemma 1 (Positive marginal productivity threshold for data). *For $I_t > 0$ and interior states with $A_t > 0$ and $D_t > 0$, the marginal innovation return to data satisfies*

$$\frac{\partial G_t}{\partial D_t} > 0 \quad \text{if and only if} \quad b(A_t^\alpha + \theta D_t^\alpha)^{\frac{\alpha+1}{\alpha}} > \theta a I_t D_t^\alpha. \quad (3)$$

Lemma 1 follows from direct differentiation of (2): G_D is the sum of a positive scale term (from D_t^b) and a negative congestion term induced by the dependence of $\exp(aI_t/H_t)$ on D_t through H_t . Condition (3) therefore provides a nonlinear threshold for when additional data increases *effective* realized innovation. The left-hand side scales with the joint magnitude of (A_t, D_t) through the composite, while the right-hand side increases with implementation intensity I_t and the size of the data stock. Economically, (3) defines an *implementability threshold*: when implementation intensity I_t is high relative to effective intangible capacity H_t , additional data can raise governance/integration burden faster than it raises scalable innovation, weakening the incentive to expand D_t at the margin.

Remark 1 (Congestion and off-equilibrium states). Data enters innovation through two channels: it raises the multiplicative scale D_t^b but also increases H_t , which dampens the implementation term $\exp(aI_t/H_t)$. Economically, very large datasets may increase integration and governance complexity (cleaning, alignment, privacy/compliance, and overfitting risk), so that when implementation effort is intensive relative to effective intangible capacity (measured by H_t), additional data may reduce the *effective* conversion rate of implementation into realized innovation. In the dynamic program, such states correspond to regions where the shadow value of accumulating data is weak (or even locally non-positive), and optimal policies may feature corner solutions (e.g., $J_t = 0$ and waiting) rather than actively expanding the data stock. This non-monotonicity is central for our subsequent policy characterization: it generates regions in which the shadow value of data is weak and optimal policies exhibit corner behavior and switching frontiers rather than smooth interior adjustments.

2.3 State Transitions and Internal Budget Constraint

Capital dynamics. The three capital stocks evolve according to

$$K_{t+1} = (1 - \delta) K_t + w_t G_t, \quad (4a)$$

$$D_{t+1} = (1 - \delta_D) (D_t - G_t) + J_t, \quad (4b)$$

$$A_{t+1} = (1 - \delta_A) (A_t - \gamma G_t) + R_t, \quad (4c)$$

where $\delta, \delta_D, \delta_A \in (0, 1)$ denote depreciation and obsolescence rates. The implementation shock $w_t > 0$ is i.i.d. with unit mean, capturing uncertainty in the conversion of realized innovation

³Revankar's variable-elasticity form $G(I, H) = \beta e^{aI/H} I^{1-b} H^b$ yields an elasticity between I and H that depends on I/H (Revankar, 1971). Our kernel preserves the Revankar-style state dependence in $\exp(aI/H_t)$ while allowing technology and data to enter separately through $A_t^b D_t^b$ rather than a single aggregate H_t^b . This separation is central for the cross-capital comparative statics in Section 6.

into installed productive capital (e.g., implementation quality or market reception). The parameter $\gamma \in (0, 1]$ governs *opportunity depletion*: realized innovation exhausts a portion of the unexploited technology opportunities.

Admissibility and effective usage. To ensure non-negativity of the post-innovation stocks and preserve their interpretation as depletable opportunities, we define the *effective utilized innovation* as

$$\tilde{G}_t := \min \{G(I_t, D_t, A_t), D_t, A_t/\gamma\}, \quad \gamma \in (0, 1].$$

In the laws of motion (4), the innovation term G_t is understood as $\tilde{G}_t$ (with a slight abuse of notation). Throughout our analytical characterization, we focus on the interior (cap-slack) region where $\tilde{G}_t = G(I_t, D_t, A_t)$, so standard differentiable first-order conditions apply. When the cap binds, the same optimality conditions are interpreted in terms of directional derivatives/subgradients; numerically, admissibility is enforced by the cap by construction.

Economic interpretation. Equation (4a) states that installed productive capacity increases with implemented innovation $w_t G_t$. Implementation effort I_t does not enter K_{t+1} one-for-one; instead, it is a costly input required to generate G_t through (2). Equation (4b) models data as a *depletable information advantage*. Here, G_t should be interpreted as the amount of *proprietary informational advantage extracted and deployed* in period t , rather than the physical consumption of raw bits. Accordingly, $D_t - G_t$ captures the residual advantage that remains appropriate for future innovation after diffusion, imitation, learning, or arbitrage erode the exercised component. While raw data may be non-rival in a physical sense, the *economic value* of firm-specific data typically derives from proprietary signals, labeling, pipelines, and tacit know-how that confer temporary informational rents. When the firm converts data into realized innovation G_t and deploys it to the market, part of this advantage is revealed, learned, imitated, or arbitrated away (e.g., through competitor learning, user adaptation, reverse engineering, or market feedback), so the remaining stock $D_t - G_t$ represents the residual advantage available for future innovation. In this sense, (4b) does not assume that bits disappear; it assumes that *appropriate informational rents* embedded in data are partially exhausted upon exercise. Equation (4c) captures a fishing-out mechanism for technology capital: realized innovation consumes part of the current opportunity set, and R&D replenishes and extends it. Such “ideas-get-harder-to-find” (fishing-out) mechanisms are standard in R&D-based growth models and innovation search frameworks; see, e.g., Jones (1995); Kortum (1997); Segerstrom (1998).⁴

Internal budget constraint. The firm reinvests a fixed share $\tau \in (0, 1]$ of current output each period:

$$I_t + J_t + R_t = \tau Y_t. \tag{5}$$

The remaining share $(1 - \tau)Y_t$ is interpreted as payout. The internal budget makes implementation effort, data investment, and R&D explicitly compete for limited internal resources, which is central for the threshold policies and shadow prices derived later.

⁴Unlike macro frameworks where data can be treated as non-rival at the aggregate level, we model firm-level data as proprietary advantage that can be dissipated when exercised in competitive markets.

2.4 Sequence of Events

The timing of events within each period t is summarized in Figure 1. The decision process proceeds as follows:

1. *Observation.* At the beginning of the period, the firm observes the state vector $S_t = (K_t, D_t, A_t)$ and the exogenous shock Z_t .
2. *Production.* Based on the current states, the firm generates operating output Y_t and determines the available reinvestment budget τY_t .
3. *Allocation.* The manager allocates this budget across implementation effort I_t , data acquisition J_t , and R&D R_t to maximize the firm's value.
4. *Innovation and Transition.* The innovation kernel $G(\cdot)$ determines the creation of new physical capital. Simultaneously, data and technology stocks are replenished or depreciated. The state then evolves to S_{t+1} according to the laws of motion (4).




Figure 1: Within-period timing and coupled state dynamics.

2.5 Static vs. Dynamic Decision Frameworks

To highlight the option-like nature of data and technology capital, we compare a myopic benchmark with the forward-looking dynamic program. Both face the budget constraint (5), but differ in their treatment of continuation value.

Static (one-step myopic) benchmark. Given $S_t = (K_t, D_t, A_t)$, the firm chooses (I_t, J_t, R_t) to maximize next-period output:

$$\max_{I_t, J_t, R_t \geq 0} Y_{t+1} \quad \text{s.t.} \quad I_t + J_t + R_t = \tau Y_t,$$

where Y_{t+1} is computed from (4) and (1). This benchmark ignores continuation value and serves as a counterfactual for quantifying the dynamic option value of data and technology. To isolate strategic exercise of growth options from aggregate risk, we first establish deterministic baselines with $Z_t \equiv 1$ and $w_t \equiv w$; without loss of generality, we normalize $w = 1$ when reporting benchmark policies.

Dynamic program. A forward-looking firm solves

$$V(S_t) = \max_{I_t, J_t, R_t \geq 0} \{(1 - \tau)Y_t + \rho \mathbb{E}[V(S_{t+1}) \mid S_t, U_t]\} \quad \text{s.t.} \quad (4), (5), \quad (6)$$

where $\rho \in (0, 1)$ is the discount factor, $U_t = (I_t, J_t, R_t)$, and the expectation is taken over w_t (and over Z_t once stochastic). Equation (6) highlights the trade-off between current payout $(1 - \tau)Y_t$ and the continuation value from accumulating capital. The value function $V(\cdot)$ is the basis for defining shadow prices (e.g., $p_t^D := \partial V / \partial D$ and $p_t^A := \partial V / \partial A$) and for deriving threshold-type investment rules in later sections. Interpreted as internal valuations, (p_t^D, p_t^A) summarize the firm's marginal willingness to pay for proprietary data advantage and technology capability under optimal continuation policies.

Comparing these two frameworks allows us to explicitly price the 'strategic patience' required for data accumulation—a value component ignored by myopic ROI-based heuristics often used in practice.

2.6 Uncertainty, Normalizations, and Continuous-Time Diagnostics

Unless stated otherwise, expectations in (6) are taken under the firm's baseline (physical) probability measure. In Section 4, we introduce a valuation measure $\mathbb{Q}$ as a diagnostic device to connect the discrete-time recursion to standard real-options representations.

We adopt normalizations to focus the exposition while ensuring comparability across regimes:

- *Deterministic benchmarks.* For the static benchmark and invariant-technology baseline, we normalize $Z_t \equiv 1$ and treat implementation efficiency as deterministic ($w_t \equiv 1$). Conditional expectations then simplify to evaluations at the next-period state.
- *Stochastic aggregate technology.* For the dynamic analysis and real-options diagnostics, aggregate technology risk follows a geometric random walk that converges to a geometric Brownian motion (GBM) in the continuous limit:

$$\frac{dZ_t}{Z_t} = \mu_Z dt + \sigma_Z dW_t^{\mathbb{Q}}.$$

In the discrete-time implementation, we employ the consistent log-normal update $Z_{t+1} = Z_t \exp((\mu_Z - \frac{1}{2}\sigma_Z^2)\Delta t + \sigma_Z\sqrt{\Delta t}\varepsilon_{t+1})$ with $\varepsilon \sim \mathcal{N}(0, 1)$. The deterministic case is recovered when $\sigma_Z = 0$.

- *Implementation risk (optional).* The idiosyncratic shock w_t can be retained without changing the structure of the Bellman equation or the shadow-price characterizations; it primarily affects the diffusion-limit representation in Appendix H.

Remark 2 (Discrete time and continuous-time diagnostics). We use a discrete-time Bellman framework to mirror quarterly decision cycles in capital budgeting, data acquisition, and R&D planning and to facilitate tractable numerics (Section 7). While the discrete-time model aligns with operational periodicity, the continuous-time limit allows us to leverage standard real-options theory to obtain analytical characterizations of the investment boundaries.

2.7 Assumptions

Throughout Sections 3–5, we impose assumptions that ensure a well-posed dynamic program, regularity of the value function, and monotone comparative statics.

A1 *Discounting and depreciation.* $\rho \in (0, 1)$ and $\delta, \delta_D, \delta_A \in (0, 1)$.

ρ represents the firm’s effective cost of capital. The depreciation rates capture physical wear and the obsolescence of data and technology.

A2 *Innovation-kernel parameters.* $\alpha, b \in (0, 1)$ and $a, \beta > 0$, with $\theta, \gamma \in (0, 1]$.

The pair (α, θ) governs curvature and the relative weight of data in the intangible composite. The parameter γ controls opportunity depletion induced by realized innovation.

A3 *Production and reinvestment.* $\tau \in (0, 1]$ and $\xi \in (0, 1)$.

τ is the internal reinvestment rate (retention ratio), while $\xi < 1$ ensures diminishing marginal returns to physical capital.

A4 *Effective domain and compactness.* We assume initial conditions and parameters (specifically, finite budget τY and positive depreciation) are such that the state process S_t remains within a forward-invariant compact domain $\mathcal{S} \subset \mathbb{R}_+^3$ almost surely. (See Appendix for the explicit construction of the bounding box.)

A5 *Innovation regularity and concavity.* On the admissible domain, the kernel G is C^2 in (I, D, A) for $I > 0$ and continuous at $I = 0$. Furthermore, we assume parameters satisfy the sufficiency conditions (derived in Appendix A) such that the map $I \mapsto G(I, D, A)$ is strictly concave on the feasible budget set.⁵

A6 *Complementarity.* On the admissible domain, technology and data are strategic complements in innovation: $G_{AD}(I, D, A) \geq 0$.

We do *not* require the marginal product of data G_D to be globally positive, consistent with the congestion threshold in Lemma 1. Assumption **A6** ensures that stronger algorithms increase the marginal usefulness of data and, conversely, richer data increase the marginal productivity of algorithmic capabilities.

Assumptions **A4–A5** ensure a well-posed dynamic program and the sufficiency of first-order conditions, while Assumption **A6** enables the monotone comparative statics developed later.

3 Benchmark Analysis: Invariant Firm-Specific Technology Regime

We begin with a benchmark regime where firm-specific technology capital is exogenous and time-invariant, i.e., $A_t \equiv A$ for all t . This setting corresponds to a mature “cash-cow” phase where the firm operates on a fixed technology frontier and does not actively expand its capabilities through R&D. Formally, relative to Section 2, we fix $A_t \equiv A$ and drop the A – law of motion (hence $R_t \equiv 0$).

⁵Since G contains the term $\exp(aI/H)$, which is convex in I , strict concavity requires that the parameter a be sufficiently small relative to the curvature parameter b . This ensures that the diminishing returns from I^{1-b} dominate the convexity of the exponential term within the relevant range of investment intensities.

In this regime, data does not enter the production function directly. Instead, it affects long-run performance solely through the innovation channel G_t and the subsequent accumulation of productive capacity. This benchmark is structurally important because it isolates the asset-accumulation value of data—arising from intertemporal smoothing and irreversibility—from the strategic option value associated with endogenous technology expansion.

Throughout this section, we normalize the aggregate technology index to $Z_t \equiv 1$. For the static benchmark, we assume a deterministic implementation efficiency $w_t \equiv w$. In the dynamic formulation, we allow w_t to be i.i.d. with $\mathbb{E}[w_t] = 1$; when w_t is deterministic, the expectation operator serves purely as notation to facilitate comparison with the stochastic model.

3.1 Static Optimization Problem

Fix the technology level $A_t \equiv A$ and a pre-decision state $S_t = (K_t, D_t) \in \mathbb{R}_{++}^2$. A myopic firm allocates its reinvestment budget to maximize next-period output:

$$Y_{t+1} = A^{1-\xi} \left[(1-\delta)K_t + w G_t \right]^\xi, \quad (7)$$

where the innovation flow is given by

$$G_t = \beta \exp \left[a I_t (A^\alpha + \theta D_t^\alpha)^{-1/\alpha} \right] I_t^{1-b} A^b D_t^b.$$

The optimization is subject to the budget and transition constraints:

$$I_t + J_t = \tau Y_t, \quad D_{t+1} = (1 - \delta_D)(D_t - G_t) + J_t, \quad (8)$$

and the admissibility constraint $G_t \leq D_t$. The static problem reduces to:

$$\max_{I_t, J_t \geq 0} Y_{t+1}(K_t, D_t; I_t, J_t) \quad \text{s.t.} \quad I_t + J_t = \tau Y_t.$$

Crucially, Y_{t+1} depends on the decision variables only through K_{t+1} , which in turn is driven by I_t via G_t . Since J_t does not enter the innovation function or the capital accumulation equation directly, data investment carries no marginal payoff in the one-step objective. Consequently, a myopic firm allocates the entire reinvestment budget to implementation effort.

Proposition 1 (Myopic preference for implementation effort). *Under Assumptions A1–A4, the unique optimal allocation for the static problem is*

$$(I_t^*, J_t^*) = (\tau Y_t, 0).$$

Proof sketch. Substitute the budget constraint $J_t = \tau Y_t - I_t$ into the system. The objective Y_{t+1} is strictly increasing in K_{t+1} , which is strictly increasing in G_t . Since G_t is strictly increasing in implementation effort I_t on the admissible domain (see Appendix B), the objective is maximized by setting I_t as large as possible. Thus, the corner solution $I_t = \tau Y_t$ (implying $J_t = 0$) is optimal.

Managerially, Proposition 1 formalizes the implications of short-termism: when continuation values are ignored, data acquisition is treated purely as a cost center rather than an investment

in future growth capacity. In the dynamic program, however, data acquires value through its potential to relax future innovation constraints.

3.2 Dynamic Program with Invariant Technology

We now consider a forward-looking firm that internalizes continuation values. The Bellman equation is:

$$V(K_t, D_t) = \max_{I_t, J_t \geq 0} \left\{ (1 - \tau)Y_t + \rho \mathbb{E}[V(K_{t+1}, D_{t+1})] \right\} \quad \text{s.t.} \quad I_t + J_t = \tau Y_t, \quad (9)$$

where the transitions follow (4) (with A fixed). Using the budget constraint to eliminate J_t , the problem simplifies to a single-control optimization over $I_t \in [0, \tau Y_t]$. We assume w_t has bounded support (or is truncated on $[\underline{w}, \bar{w}]$), consistent with the forward-invariant compact set in Assumption A4.

Theorem 1 (Existence and uniqueness). *Under Assumptions A1–A4, the Bellman operator associated with (9) is a ρ -contraction on $C_b(\mathcal{S}_0^{KD})$, the space of bounded continuous functions on a compact forward-invariant state space $\mathcal{S}_0^{KD} \subset \mathbb{R}_{++}^2$. Consequently, there exists a unique fixed point $V^* \in C_b(\mathcal{S}_0^{KD})$, and value iteration converges to V^* at a geometric rate.*

The proof follows standard arguments relying on bounded payoffs, compact feasible sets, and discounting. Detailed proofs are provided in Appendix B (specific model) and Appendix C (general contraction result).

3.3 Dynamic Investment Strategy and Threshold Behavior

Policy characterization. Consider interior states where the non-negativity constraints on controls do not bind. Let $M_t := \tau Y_t$ denote the budget capacity. Differentiating the continuation value in (9) with respect to I_t (keeping $J_t = M_t - I_t$) yields the first-order condition:

$$0 = \rho \mathbb{E} \left[V_K(S_{t+1}) w_t G_I(I_t, D_t) + V_D(S_{t+1}) (-(1 - \delta_D)G_I(I_t, D_t) - 1) \mid S_t \right], \quad (10)$$

where $S_{t+1} = (K_{t+1}, D_{t+1})$ and G_I denotes the partial derivative with respect to I_t . Rearranging terms yields the key trade-off:

$$\mathbb{E} \left[V_K(S_{t+1}) w_t G_I(I_t, D_t) \mid S_t \right] = \mathbb{E} \left[V_D(S_{t+1}) (1 + (1 - \delta_D)G_I(I_t, D_t)) \mid S_t \right]. \quad (11)$$

Economically, condition (11) equates the *marginal return on implementation* with the *full opportunity cost of data utilization*, both expressed in shadow-value units. In particular, $V_K(S_{t+1})$ is the shadow value (an *internal marginal q*) of installed productive capacity, while $V_D(S_{t+1})$ is the firm’s internal transfer price of proprietary data advantage. The right-hand side therefore summarizes the “shadow bill” that implementation must pay for drawing down the information stock: a mechanical crowd-out of data investment via the budget constraint (the unit term) and an operational cannibalization component, $(1 - \delta_D)G_I$, which reflects accelerated depletion/obsolescence of proprietary advantage as innovation is scaled.

Scale–obsolescence trade-off. Unlike physical capital, utilizing data to generate innovation endogenously accelerates the depletion of proprietary information advantage in the next-period state. Thus, high-growth implementation (high I_t) can erode the firm’s competitive moat even when reinvestment capacity is high, creating a natural brake on expansion through the shadow cost of data.

The left-hand side represents the *expected marginal benefit* of implementation: increasing I_t raises realized innovation G_t , which is converted into productive capital K_{t+1} at rate w_t and valued at $V_K(S_{t+1})$.

The right-hand side represents the *expected marginal opportunity cost* in terms of data. Increasing I_t reduces the next-period data advantage through two channels: (i) mechanically, because $J_t = M_t - I_t$ implies one unit less data investment per additional unit of implementation; and (ii) operationally, because greater implementation accelerates the “exercise” of proprietary signals, captured by the depletion term $(1 - \delta_D)G_I(I_t, D_t)$.

The second channel is the *cannibalization of information advantage*: by scaling up innovation G_t , the firm more aggressively converts proprietary signals into deployed products and decisions. Once exercised, part of this advantage is revealed, learned, imitated, or arbitrated away, making it less useful for future innovation. Consequently, rapid scaling (high I_t) can lower the continuation value by hastening the obsolescence of the very advantage that enables data-driven growth. Both channels are priced by the shadow value of data, $V_D(S_{t+1})$.

Switching frontier. Define the marginal value of relaxation at the upper budget corner $I_t = M_t$ (where $J_t = 0$):

$$\Delta Q(M_t; S_t) := \left. \frac{\partial}{\partial I_t} \rho \mathbb{E}[V(S_{t+1}) \mid S_t] \right|_{I_t=M_t^-}. \quad (12)$$

A corner allocation with $J_t^* = 0$ is optimal whenever $\Delta Q(M_t; S_t) \geq 0$. The condition $\Delta Q(M_t; S_t) = 0$ defines an implicit *switching frontier* Γ in the (K, D) -plane. This frontier separates the “inaction” region ($J_t^* = 0$, where data is abundant relative to capital) from the “investment” region ($J_t^* > 0$, where data is scarce).

Comparative statics. Under the monotonicity differences conditions in Appendix B.4, optimal implementation effort is weakly increasing in the data stock. Intuitively, a larger data stock raises the marginal productivity of implementation, incentivizing the firm to allocate more budget to I_t .

Proposition 2 (Unique interior steady state). *Under the additional regularity conditions in Appendix B.5 (Assumptions A8–A9), there exists a unique interior steady state $(\bar{K}, \bar{D})$ for the expected dynamics induced by the optimal policy. Starting from any admissible initial condition, the system converges to $(\bar{K}, \bar{D})$ in expectation at a geometric rate.*

This result implies a unique long-run target mix of physical capital and data assets. As the system converges to $(\bar{K}, \bar{D})$, the shadow price of data p_t^D stabilizes, providing a robust baseline for the real-options valuation analysis in Section 4. Furthermore, this invariant benchmark serves as a structural reference for clarifying the additional value derived from endogenous technology accumulation in Section 5. For related perspectives on shadow pricing and real

options, see, for example, Kuosmanen and Zhou (2021); Bellalah (2016); Lambrecht (2017); Lu et al. (2025).

4 Shadow Price as Real Options

We interpret the shadow value of data as the value of a perpetual real option. This links the dynamic program to the real-options tradition and yields state-dependent investment triggers. From this section on, we work in the invariant-technology benchmark with $A_t \equiv A > 0$ fixed (as in Section 3). The pre-decision state is

$$S_t = (K_t, D_t, Z_t),$$

where Z_t is an aggregate technology index that captures economy-wide technological conditions. The shadow price of data is

$$p_t^D := \partial_D V^*(S_t),$$

whenever V^* is differentiable; at kinks, we interpret p_t^D as a supergradient (an element of the superdifferential) of V^* with respect to D . Throughout this section, we formulate valuation under a pricing kernel and denote the induced measure by $\mathbb{Q}$. Accordingly, the Bellman recursion is written under $\mathbb{Q}$, which aligns the discrete-time recursion with the continuous-time real-options benchmark.

Remark 3 (Measure convention). In the sequel we evaluate continuation values under a valuation measure $\mathbb{Q}$ with one-period discount factor $\rho = e^{-r\Delta t}$, and write conditional expectations as $\mathbb{E}_t^{\mathbb{Q}}[\cdot | S_t]$. We adopt $\mathbb{Q}$ for both the discrete- and continuous-time representations. Relative to a physical measure $\mathbb{P}$, $\mathbb{Q}$ is an equivalent measure that reweights states according to the pricing kernel; in particular, the drift of Z under $\mathbb{Q}$ embeds risk premia. We state the Z -dynamics directly under $\mathbb{Q}$ to streamline the link between (14) and the continuous-time benchmark.

4.1 Stochastic Aggregate Technology Environment

Output is $Y_t = (AZ_t)^{1-\xi} K_t^\xi$. Under the valuation measure $\mathbb{Q}$, we model the aggregate technology index Z_t as a lognormal process with the exact transition (13),

$$Z_{t+\Delta t} = Z_t \exp\left((\mu_Z - \frac{1}{2}\sigma_Z^2)\Delta t + \sigma_Z\sqrt{\Delta t}\varepsilon_{t+1}\right), \quad \varepsilon_{t+1} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \quad (13)$$

which converges to a GBM as $\Delta t \downarrow 0$. This lognormal specification yields a tractable Markov state for valuation and supports closed-form continuous-time diagnostics as $\Delta t \downarrow 0$.

To focus on aggregate technology risk, we set the implementation factor $w_t \equiv 1$ in this section. We keep the internal resource constraint $I_t + J_t = \tau Y_t$ and, for expositional convenience, introduce a reduced-form convex adjustment cost for data investment,

$$\Pi_t = (1 - \tau)Y_t - \frac{\kappa}{2}J_t^2, \quad J_t \in [0, \tau Y_t],$$

which captures increasing marginal frictions in scaling data spending (e.g., vendor search, in-

tegration, governance, and privacy/compliance). The trigger logic derived below relies on the presence of convex scaling frictions (and the shadow cost of internal funds), rather than on the specific quadratic form. Accordingly, we write $I_t = \tau Y_t - J_t$ throughout this section.

The Bellman problem is

$$V(S_t) = \max_{0 \leq J_t \leq \tau Y_t} \left\{ \Pi_t + \rho \mathbb{E}_t^{\mathbb{Q}}[V(S_{t+1}) \mid S_t, J_t] \right\}, \quad S_t = (K_t, D_t, Z_t), \quad (14)$$

where $Z_{t+\Delta t}$ follows (13) and $(K_{t+\Delta t}, D_{t+\Delta t})$ follow (4) with $I_t = \tau Y_t - J_t$.

Proposition 3 (Well-posedness and convergence). *Under standard continuity and discounting conditions for the stochastic dynamic program with state $S_t = (K_t, D_t, Z_t)$, the Bellman operator in (14) is a contraction and admits a unique fixed point V^* . Moreover, value iteration converges to V^* at a geometric rate.*

A complete set of sufficient conditions (including unbounded-state extensions under weighted norms) and the proof are provided in Appendix C.

4.2 Shadow Price Recursion and Dynamic Value

Shadow price recursion. Applying the envelope theorem to (14) gives

$$p_t^D = \rho(1 - \delta_D) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] + \rho \mathbb{E}_t^{\mathbb{Q}}[(V_K(S_{t+1}) - (1 - \delta_D)p_{t+1}^D) G_D(I_t, D_t) \mid S_t]. \quad (15)$$

This recursion mirrors an Euler equation for intangible capital: the shadow price today balances the discounted continuation value of retained data with its indirect effect on future capital productivity.

Interior assumption. If $0 < J_t^* < \tau Y_t$, the first-order condition for J_t yields

$$\kappa J_t^* = \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D - (V_K(S_{t+1}) - (1 - \delta_D)p_{t+1}^D) G_I(I_t, D_t) \mid S_t]. \quad (16)$$

Substituting (16) into (15) yields

$$p_t^D = (1 - \delta_D) \kappa J_t^* + \rho \mathbb{E}_t^{\mathbb{Q}}[(V_K - (1 - \delta_D)p_{t+1}^D)(G_D + (1 - \delta_D)G_I) \mid S_t], \quad (17)$$

making the link between p_t^D and J_t^* explicit. Uniqueness of V^* follows from Proposition 3; whenever V^* is differentiable in D , the induced shadow price p_t^D is therefore uniquely pinned down.

Economically, this mirrors the no-arbitrage logic in option pricing: the marginal value of data must balance contemporaneous scaling frictions against the forward-looking option payoff generated through future productivity.

4.3 Option-Based Interpretation and Investment Triggers

We interpret the shadow price of data, $p_t^D := \partial_D V^*(S_t)$, through a real-options lens that links the KKT conditions to exercise logic and clarifies how per-unit wedges and the shadow value of internal funds shape state-dependent triggers. Unlike standard real-options models in which

the exercise price is exogenous, the effective cost of *internally* accumulating data is endogenous: it is pinned down by the shadow value of internal funds, i.e., the Lagrange multiplier on the reinvestment budget. This wedge is state-dependent and is typically highest when the firm is cash-constrained, shifting the investment trigger even when the marginal value of data is high.

Present value of marginal returns. Define the expected present value (EPV) of future marginal returns

$$\Omega_t^D := \mathbb{E}_t^{\mathbb{Q}} \left[\sum_{k=0}^{\infty} \rho^k p_{t+k}^D \mid S_t \right], \quad \Omega_t^D = p_t^D + \rho \mathbb{E}_t^{\mathbb{Q}} [\Omega_{t+1}^D \mid S_t],$$

which converges whenever $\sum_{k \geq 0} \rho^k \mathbb{E}_t^{\mathbb{Q}} [|p_{t+k}^D|] < \infty$; a sufficient condition is $\sup_t \mathbb{E}_t^{\mathbb{Q}} [|p_t^D|] < \infty$.

Two economically distinct “options.” There are two natural exercise margins. Both yield state-dependent triggers, but they differ in timing and in the relevant cost of exercise:

- (a) *Internal accumulation option (make-to-order) for $t+1$ (baseline).* Holding I_t fixed, one unit of J_t raises D_{t+1} one-for-one; under the budget $I_t + J_t = \tau Y_t$, this direct accumulation benefit trades off against a lower implementation intensity I_t and hence a lower G_t , which is exactly priced by λ_t . It is convenient to view (14) as a two-control problem in (I_t, J_t) subject to the linear budget $I_t + J_t = \tau Y_t$, and to let λ_t denote the associated Lagrange multiplier (the shadow value of relaxing the budget by one unit). A standard envelope argument yields

$$\lambda_t = \rho \mathbb{E}_t^{\mathbb{Q}} \left[(V_K(S_{t+1}) - (1 - \delta_D) p_{t+1}^D) G_I(I_t, D_t) \mid S_t \right].$$

Let $c_t^{\text{adj}} \geq 0$ collect any additional per-unit wedge beyond λ_t (e.g., an internal compliance fee).⁶ The relevant internal hurdle is therefore the *effective exercise cost* $\lambda_t + c_t^{\text{adj}}$, which is endogenous and can be countercyclical: when internal funds are scarce (high λ_t), the firm optimally delays data accumulation even if the fundamental marginal value $\rho \mathbb{E}_t^{\mathbb{Q}} [p_{t+1}^D \mid S_t]$ is high. This financing wedge is absent in standard real-options settings with a fixed exercise price.

Define the discounted net marginal benefit

$$X_t^{\text{int}} := \rho \mathbb{E}_t^{\mathbb{Q}} [p_{t+1}^D \mid S_t] - \lambda_t - c_t^{\text{adj}}.$$

Then the interior FOC (16) implies $\kappa J_t^* = X_t^{\text{int}}$ whenever $0 < J_t^* < \tau Y_t$. More generally, with the box constraint $0 \leq J_t \leq \tau Y_t$, the KKT conditions yield the projected intensity rule

$$J_t^* = \min \left\{ \tau Y_t, (X_t^{\text{int}} / \kappa)_+ \right\}.$$

Thus $J_t^* > 0$ if and only if $X_t^{\text{int}} > 0$, while the upper corner $J_t^* = \tau Y_t$ occurs if and only if $X_t^{\text{int}} \geq \kappa \tau Y_t$. Economically, the internal option uses the discounted next-period marginal value $\rho \mathbb{E}_t^{\mathbb{Q}} [p_{t+1}^D \mid S_t]$ as the underlying and the state-dependent effective exercise cost $\lambda_t + c_t^{\text{adj}}$ as the hurdle. When the firm is cash-constrained (high λ_t), the exercise region shrinks even

⁶In the special case of a constant per-unit internal fee c_D^{int} attached to J_t , set $c_t^{\text{adj}} \equiv c_D^{\text{int}}$.

when p_t^D is large. Managerially, this financing wedge implies that data spending can be disproportionately procyclical: whenever internal funds are scarce (high λ_t), firms optimally cut J_t more aggressively than a static NPV rule based solely on p_t^D would suggest.

- (b) *Market acquisition option (spot purchase) at t (counterfactual benchmark)*. To contrast with the internal channel, we consider a theoretical benchmark distinct from the operational problem (14): suppose the firm can additionally acquire data at an exogenous spot price $c_D > 0$, instantaneously augmenting the pre-decision stock D_t . The net payoff is

$$X_t^{\text{ext}} := p_t^D - c_D.$$

The corresponding option value is the Snell envelope

$$S_t^{\text{ext}} = \max \left\{ X_t^{\text{ext}}, \rho \mathbb{E}_t^{\mathbb{Q}}[S_{t+1}^{\text{ext}} \mid S_t] \right\},$$

which yields an optimal-stopping trigger region characterized by $p_t^D \geq c_{\text{ext}}(S_t)$ under a Markov and monotonicity structure (Appendix E). In the myopic/zero-continuation limit, $c_{\text{ext}}(S_t) \downarrow c_D$, recovering the heuristic “invest when $p_t^D \geq c_D$ ” rule.

KKT frontier as trigger. Let

$$\Delta_J(S_t) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t - c_t^{\text{adj}}.$$

The switching surface

$$\Gamma_J := \{S : \Delta_J(S) = 0\}$$

is the KKT boundary between $J_t = 0$ and $J_t > 0$. The upper-corner boundary is given by $\Delta_J(S) = \kappa\tau Y(S)$. Under monotonicity of V and of the transition kernel, these boundaries are upper sets in relevant coordinates (Appendix B.4). Economically, the upper-set property implies a state-dependent “investment readiness” threshold: only when the firm’s productive state is sufficiently strong to justify the shadow cost of internal funds does it enter the region with $J_t > 0$.

Mapping between external and internal costs. The external trigger compares the *current* marginal value p_t^D to the market price c_D . The internal trigger compares the *discounted next-period* marginal value $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t]$ to the *state-dependent* internal hurdle $\lambda_t + c_t^{\text{adj}}$. When (i) the time step is small (so that p_t^D and $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t]$ are first-order close, ignoring $O(\Delta t)$ generator drift) and (ii) the firm faces a constant per-unit internal fee c_D^{int} , the two rules admit a local identification of effective costs,

$$c_{\text{eff}} \approx \lambda_t + c_t^{\text{adj}}.$$

In this sense, internal and external triggers can be locally aligned; in general, however, they represent distinct margins, reflecting a fundamental timing trade-off: spot purchases deliver immediate stock augmentation at price c_D , whereas internal accumulation shifts value intertemporally and is exposed to the state-dependent (and potentially volatile) hurdle λ_t .

Continuous-time limit. As $\Delta t \downarrow 0$, the internal trigger induces a free boundary in the

continuous-time HJB/variational-inequality representation with generator $\mathcal{L}$. The discrete frontier Γ_J can be viewed as the discrete-time counterpart of this free boundary, consistent with the diffusion-based interpretation in Section 4.4.

4.4 Continuous-Time Benchmark and Theoretical Robustness

Our baseline analysis is in discrete (quarterly) time. To check that results are not an artifact of time aggregation, we consider the continuous-time limit under $\mathbb{Q}$ with $dZ_t/Z_t = \mu_Z dt + \sigma_Z dW_t^{\mathbb{Q}}$, $w_t \equiv 1$, and $G = G(I, D)$. We focus on the economic interpretation in the main text; the rigorous justification of the limit (including viscosity solutions and convergence of the discrete scheme) is provided in Appendix H. To take $\Delta t \downarrow 0$, interpret $G(I, D)$ and J as flow rates so that per-period stock changes are $G(I, D)\Delta t$ and $J\Delta t$. Consistently, we scale depreciation as $\delta = \tilde{\delta}\Delta t$ and $\delta_D = \tilde{\delta}_D\Delta t$ when taking $\Delta t \downarrow 0$. The limiting laws of motion are

$$dK_t = (-\delta K_t + G(I_t, D_t)) dt, \quad dD_t = (-\delta_D D_t - G(I_t, D_t) + J_t) dt.$$

HJB and the shadow-price PDE. With continuous discount rate $r > 0$ (so $\rho = e^{-r\Delta t}$ in discrete time), the HJB is

$$rV = \max_{0 \leq J \leq \tau Y} \left\{ \Pi + V_K(-\delta K + G(I, D)) + V_D(-\delta_D D - G(I, D) + J) + \mathcal{L}_Z V \right\}, \quad (18)$$

where $I = \tau Y - J$ and $\mathcal{L}_Z V := \mu_Z Z V_Z + \frac{1}{2} \sigma_Z^2 Z^2 V_{ZZ}$. Here time arguments are omitted when no confusion arises. Let $p^D := V_D$. Differentiating (18) w.r.t. D and applying the envelope theorem (holding the optimal J fixed) yields

$$(r + \delta_D) p^D = (V_K - p^D) G_D(I, D) + (-\delta K + G) p_K^D + (-\delta_D D - G + J) p_D^D + \mathcal{L}_Z p^D. \quad (19)$$

If V (hence p^D) is not $C^{1,2}$, interpret (19) in the viscosity sense; the discrete scheme is monotone and consistent, so the limit solves the HJB equation in the viscosity sense.

Shadow-price diffusion. Applying Itô's formula to $p^D(K_t, D_t, Z_t)$ and using (19) gives

$$dp_t^D = \left[(r + \delta_D) p_t^D + (p_t^D - V_K(K_t, D_t, Z_t)) G_D(I_t, D_t) \right] dt + \sigma_Z Z_t p_Z^D dW_t^{\mathbb{Q}}. \quad (20)$$

In the drift, $(r + \delta_D)p^D$ is the discount-depreciation carry and $(p^D - V_K)G_D$ is the innovation channel; the diffusion term $\sigma_Z Z_t p_Z^D dW_t^{\mathbb{Q}}$ propagates common technology shocks.

Consistency with the discrete model. As $\Delta t \downarrow 0$, the discrete Bellman equation converges to the HJB (18), and the discrete shadow-price recursion converges to the PDE/SDE representations (19)–(20). The formal expansions and a monotone-scheme argument are provided in Appendix H.

5 Benchmark Analysis with Endogenous Firm-Specific Technology

Building on the real-options logic in Section 4, we now endogenize the firm-specific technology stock A_t . The firm allocates the reinvestment budget τY_t across implementation effort I_t , data acquisition J_t , and R&D expenditures R_t . This introduces a coupled real-options environment: the innovation kernel $G(I_t, D_t, A_t)$ links current choices to the joint evolution of physical capital, data, and technology, so that complementarities in G shape both contemporaneous marginal returns and the path of future productivity. Throughout, we take the implementation factor to be constant, $w_t \equiv w > 0$ (normalized to $w = 1$ when convenient). Unless explicitly stated, we suppress aggregate technology risk by setting $Z_t \equiv 1$.

To isolate the role of intertemporal foresight, we begin with a one-period (myopic) benchmark that maximizes next-period output only. Under the timing maintained throughout the paper, data spending J affects the *next* data stock but does not enter next-period output directly. Consequently, myopic optimization sets $J^* = 0$, underscoring why dynamic planning is essential to rationalize positive data investment.

5.1 Static Allocation without Intertemporal Foresight

Given $S_t = (K_t, D_t, A_t)$, the myopic firm chooses (I_t, J_t, R_t) at time t to maximize next-period output:

$$\max_{I_t, J_t, R_t \geq 0} Y_{t+1}(I_t, J_t, R_t; S_t) \quad \text{s.t.} \quad I_t + J_t + R_t = \tau Y_t, \quad (21)$$

where $Y_{t+1} = (A_{t+1} Z_{t+1})^{1-\xi} K_{t+1}^\xi$ with $Z_{t+1} \equiv 1$ and $w_t \equiv w$. The one-step-ahead stocks $(K_{t+1}, D_{t+1}, A_{t+1})$ are induced by the transition laws in (4).

Proposition 4 (One-Period Optimum with Endogenous Technology). *Suppose $G(\cdot, D_t, A_t)$ is concave in I_t on $[0, \tau Y_t]$ and the induced objective $Y_{t+1}(I_t, 0, R_t; S_t)$ is strictly concave in (I_t, R_t) on the budget simplex. Then problem (21) admits a unique solution (I_t^*, J_t^*, R_t^*) with:*

- (a) $J_t^* = 0$ (data acquisition affects D_{t+1} but does not enter Y_{t+1} under the one-period objective and timing);
- (b) $I_t^*, R_t^* \geq 0$ and $I_t^* + R_t^* = \tau Y_t$;
- (c) If the optimum is interior in (I_t, R_t) , then marginal contributions are equalized:

$$\partial_{I_t} Y_{t+1}(I_t^*, 0, R_t^*; S_t) = \partial_{R_t} Y_{t+1}(I_t^*, 0, R_t^*; S_t);$$

- (d) Equivalently,

$$w G_I(I_t^*, D_t, A_t) = \frac{1-\xi}{\xi} \frac{K_{t+1}}{A_{t+1}} \left[1 + \gamma(1 - \delta_A) G_I(I_t^*, D_t, A_t) \right], \quad (22)$$

where K_{t+1} and A_{t+1} are the one-step-ahead stocks induced by $(I_t^*, 0, R_t^*)$ via (4).

Proposition 4 mirrors Proposition 1 in the invariant-technology case. When the objective is purely one-period and output depends on (K_{t+1}, A_{t+1}) but not on D_{t+1} , the control J_t is

dominated: any $J_t > 0$ reduces $(I_t + R_t)$ without affecting Y_{t+1} , so $J_t^* = 0$. Conditional on $J_t^* = 0$, optimality equates the marginal effect of implementation spending I_t on next-period output (via K_{t+1} and the innovation kernel) with the marginal effect of R&D spending R_t on next-period output (via A_{t+1}). Equation (22) makes this trade-off explicit: implementation raises K_{t+1} by wG_I , while (under (4)) it also affects A_{t+1} through the innovation-induced technology erosion term proportional to $\gamma(1 - \delta_A)$.

The stark corner $J_t^* = 0$ under one-period objectives highlights that data is intrinsically *option-like*: its value is realized primarily through future innovation rather than contemporaneous output. Consequently, short-horizon evaluation criteria can mechanically bias capital budgeting away from data accumulation even when long-run value creation is high.

Fix $S = (K, D, A)$. Under standard smoothness and strong concavity, the one-period policy and objective are continuous (and locally Lipschitz) in primitive parameters such as $(\delta_A, \gamma, \alpha, \theta, b)$ and w . In particular, the corner $J_t^* = 0$ is stable under small perturbations in these primitives. Formal robustness statements are recorded in Propositions 9–10 and Lemma 7 (Appendix G.2).

5.2 Dynamic Optimization with Endogenous Technology

We now restore forward-looking behavior and embed the full three-asset system (4) into an infinite-horizon Bellman framework.

Per-period payoff. Cash flow in each period is

$$\Pi(K, D, A, Z; I, J, R) = (1 - \tau)Y(K, A, Z) - \frac{\kappa}{2}J^2 - \frac{\kappa'}{2}R^2, \quad Y(K, A, Z) := (ZA)^{1-\xi}K^\xi,$$

where $\kappa, \kappa' > 0$ capture increasing organizational frictions from rapidly scaling data teams and R&D labs (e.g., integration, coordination, and compliance overhead). Technically, these convex costs also ensure strict concavity of the intra-period allocation, yielding a unique interior best response whenever the simplex constraints are slack.

Managers observe the state $S = (K, D, A)$ (and Z when stochastic) and choose $(I, J, R) \in \mathbb{R}_+^3$ subject to the flow-budget constraint $I + J + R = \tau Y(K, A, Z)$.

Uncertainty. Continuation values are evaluated under the risk-neutral measure $\mathbb{Q}$ with one-period discount factor $\rho = e^{-r\Delta t}$. When aggregate technology risk is present, Z follows the lognormal update (13). Unless explicitly stated, we set $Z \equiv 1$ so that all uncertainty is shut down in this section.

Dynamic program. Under $\mathbb{Q}$, the Bellman operator on the augmented state $S = (K, D, A, Z)$ is

$$(\mathcal{TV})(K, D, A, Z) := \max_{\substack{I, J, R \geq 0 \\ I+J+R=\tau Y(K, A, Z)}} \left\{ (1 - \tau)Y(K, A, Z) - \frac{\kappa}{2}J^2 - \frac{\kappa'}{2}R^2 + \rho \mathbb{E}^{\mathbb{Q}}[V(K', D', A', Z') \mid K, D, A, Z] \right\}, \quad (23)$$

where (K', D', A') follow (4) and Z' follows (13) when stochastic.

Theorem 2 (Well-posedness and convergence). *Under Assumptions A1–A6 and $\rho \in (0, 1)$, the Bellman operator in (23) admits a unique fixed point V^* and an optimal stationary policy. Moreover, value iteration converges geometrically to V^* under standard continuity/compactness or weighted-drift regularity conditions stated in Appendix C.*

Theorem 2 is a direct specialization of Appendix C to the state $S = (K, D, A, Z)$ and the quadratic adjustment-cost specification in (23). Endogenizing A enlarges the state space but does not alter the contraction modulus: ρ in the bounded case and $\rho\zeta$ under the weighted norm.

5.3 Dynamic Policy and Shadow Price System

Define shadow prices $p_t^K := V_K(S_t)$, $p_t^D := V_D(S_t)$, and $p_t^A := V_A(S_t)$ (interpreted as super-gradients at kinks). Write $Y_t := (A_t Z_t)^{1-\xi} K_t^\xi$ and $Y_A := \partial Y / \partial A$. Let G_I, G_D, G_A denote first partial derivatives of $G(I, D, A)$ and $G_{AD} := \partial^2 G / \partial A \partial D$.

KKT system and the shadow value of funds. Let λ_t be the Lagrange multiplier on the flow budget $I_t + J_t + R_t = \tau Y_t$, i.e., the shadow value of internal funds. Using the per-period payoff $\Pi = (1 - \tau)Y - \frac{\kappa}{2}J_t^2 - \frac{\kappa'}{2}R_t^2$ and transitions (4), optimality implies the following KKT conditions (a.s.):

$$\rho \mathbb{E}_t^{\mathbb{Q}}[w p_{t+1}^K - (1 - \delta_D) p_{t+1}^D - \gamma(1 - \delta_A) p_{t+1}^A] \Big| S_t \Big] G_I(I_t, D_t, A_t) \leq \lambda_t, \quad I_t(\lambda_t - \mathcal{U}_t) = 0, \quad (24)$$

$$-\kappa J_t + \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t \leq 0, \quad J_t(\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \kappa J_t - \lambda_t) = 0, \quad (25)$$

$$-\kappa' R_t + \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A \mid S_t] - \lambda_t \leq 0, \quad R_t(\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A \mid S_t] - \kappa' R_t - \lambda_t) = 0, \quad (26)$$

$$I_t + J_t + R_t = \tau Y_t, \quad Y_t = (A_t Z_t)^{1-\xi} K_t^\xi, \quad (27)$$

where, for convenience, we define

$$\mathcal{U}_t := \rho \mathbb{E}_t^{\mathbb{Q}}[w p_{t+1}^K - (1 - \delta_D) p_{t+1}^D - \gamma(1 - \delta_A) p_{t+1}^A \mid S_t] G_I(I_t, D_t, A_t).$$

A convenient implication of (25)–(26) is the *marginal-intensity* representation: for given λ_t , the active margins satisfy

$$J_t = \left(\frac{\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t}{\kappa} \right)_+, \quad R_t = \left(\frac{\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A \mid S_t] - \lambda_t}{\kappa'} \right)_+,$$

and (27) pins down λ_t so that the implied (I_t, J_t, R_t) lie on the budget simplex with $I_t \geq 0$. In particular, whenever $I_t > 0$, (24) holds with equality:

$$\lambda_t = \rho \mathbb{E}_t^{\mathbb{Q}}[w p_{t+1}^K - (1 - \delta_D) p_{t+1}^D - \gamma(1 - \delta_A) p_{t+1}^A \mid S_t] G_I(I_t, D_t, A_t).$$

We maintain $G_I \geq 0$ on the admissible set; on sets where $G_I = 0$ the I -margin may be locally degenerate, in which case any measurable selector is admissible without affecting the comparative-statics results.⁷

⁷Expectations are under the risk-neutral measure $\mathbb{Q}$, consistent with discounting by $\rho = e^{-r\Delta t}$. Setting $w \equiv 1$ recovers the baseline without implementation risk.

Optional unit-cost extension (for reconciliation). If per-unit costs $c_D, c_A \geq 0$ are added to the payoff as $-c_D J_t - c_A R_t$, then (25)–(26) become

$$-\kappa J_t - c_D + \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] - \lambda_t \leq 0, \quad -\kappa' R_t - c_A + \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] - \lambda_t \leq 0,$$

so the excess-return gaps shift to $\Delta_J^c := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] - \lambda_t - c_D$, $\Delta_R^c := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] - \lambda_t - c_A$. With quadratic costs, intensities become $J_t = (\Delta_J^c / \kappa)_+$, $R_t = (\Delta_R^c / \kappa')_+$, with λ_t again determined endogenously by the flow budget.

Switching (exercise) surfaces. Define the (baseline) excess-return gaps

$$\Delta_J(S_t) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] - \lambda_t, \quad \Delta_R(S_t) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] - \lambda_t,$$

and switching surfaces

$$\Gamma_J := \{S : \Delta_J(S) = 0\}, \quad \Gamma_R := \{S : \Delta_R(S) = 0\}.$$

From (25)–(26),

$$J_t = 0 \iff \Delta_J \leq 0, \quad J_t > 0 \Rightarrow \kappa J_t = \Delta_J, \quad R_t = 0 \iff \Delta_R \leq 0, \quad R_t > 0 \Rightarrow \kappa' R_t = \Delta_R.$$

Thus (Γ_J, Γ_R) partition the state space into four regimes (both active, data-only, R&D-only, neither). When both $J_t, R_t > 0$, the common shadow value of funds ties the two margins:

$$\lambda_t = \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] - \kappa J_t = \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] - \kappa' R_t.$$

Real-options interpretation with endogenous technology. The inequalities (25)–(26) are endogenous exercise rules. For data (resp. R&D), the *underlying* is the discounted next-period marginal value $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t]$ (resp. $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t]$); the *effective strike* is the opportunity cost of internal funds λ_t (shifted by (c_D, c_A) when present); and quadratic adjustment costs govern exercise *intensity*.

Monotonicity and concise comparative statics. Under primitive complementarity in the innovation kernel ($G_D, G_A \geq 0$, $G_{AD} \geq 0$) and a small-gain bound that disciplines feedback (Appendix H, Lemma 8), one has

$$\partial_A \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] \geq 0, \quad \partial_D \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] \geq 0,$$

so higher technology raises the marginal value of data and vice versa. These signs line up with the steady-state and sensitivity results in Section 7.

Remark 4 (Embedded option valuation). Our framework admits two complementary risk-neutral valuations (with $\rho = e^{-r\Delta t}$).

- (i) *External (spot) acquisition.* The perpetual option to buy one unit of data at unit price

c_D is the Snell envelope

$$S_t^{\text{ext}} = \text{ess sup}_{\tau \geq t} \mathbb{E}_t^{\mathbb{Q}}[\rho^{(\tau-t)}(p_\tau^D - c_D) \mid S_t],$$

with the standard backward recursion and, in continuous time, a variational-inequality representation (see Appendix E).

(ii) *Internal (make-to-order) accumulation.* The value of flexibility along the J -margin can be expressed as the *flexibility premium*

$$\text{OV}_t^J := V^{\text{flex}}(S_t) - V^{\text{no-}J}(S_t),$$

i.e., the difference between the Bellman problems with and without the J -control (and analogously for R). These are risk-neutral managerial values, not necessarily market-traded prices in incomplete markets.

Theoretical robustness and limit properties. In the continuous-time limit (Appendix H), the same shadow value of funds λ_t appears as a budget-relaxation term in the p^A equation, and the diffusion system for (p^A, p^D) is driven by the aggregate technology factor Z_t . Hence the KKT exercise logic and the option interpretation remain aligned across discrete and continuous time. This alignment ensures that the state-dependent exercise regions and switching surfaces derived in discrete time are robust approximations of the continuous-time free-boundary objects, rather than artifacts of time discretization. A rigorous justification based on the viscosity-solution framework and a monotone-scheme argument is provided in Appendix H.

Comparative statics in the budget share τ and data weight θ . We record two directional comparative-statics results that formalize the discussion of switching frontiers above. Full proofs are provided in Appendix G.

Proposition 5 (Budget expansion enlarges internal exercise regions). *Suppose the inner problem is strictly concave in (I, J, R) on the feasible simplex.⁸ Then $\lambda_t(\tau)$ is weakly decreasing in τ . Consequently, the excess-return gaps*

$$\Delta_J(S_t; \tau) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t(\tau), \quad \Delta_R(S_t; \tau) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A \mid S_t] - \lambda_t(\tau)$$

are weakly increasing in τ , and the switching surfaces Γ_J and Γ_R shift outward as τ rises; equivalently, the exercise regions $\{J_t > 0\}$ and $\{R_t > 0\}$ expand with τ .

Proposition 6 (Data weight raises the shadow value and spending on data). *Let $G \in C^2$ depend on a parameter θ (e.g., the data weight in the innovation kernel) and suppose on the admissible set that $G_D \geq 0$ and $G_{D\theta} \geq 0$. Assume the monotonicity/small-gain conditions in Appendix H (Lemma 8) so that the shadow-price system is well-posed and its resolvent is positivity preserving. Then the data shadow price p_t^D is weakly increasing in θ (pointwise in S_t). In particular, the interior KKT rule $\kappa J_t = \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t$ (see (25)) implies that the optimal data spending J_t^* and the exercise region $\{J_t > 0\}$ are weakly increasing in θ .*

⁸This holds under the convex adjustment costs in (23) together with the regularity conditions A1–A5.

5.4 The Dual-Option Mechanism: Complementarities Between Data and Technology

Shadow-price recursions and interpretation. Let $Y_t := Y(K_t, A_t, Z_t)$ and $Y_A := \partial Y / \partial A$. Applying the constrained envelope theorem to the Bellman equation (23) with transitions (4) yields

$$p_t^A = (1 - \tau) Y_A(K_t, A_t, Z_t) + \lambda_t \tau Y_A(K_t, A_t, Z_t) + \rho(1 - \delta_A) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] + \rho \mathbb{E}_t^{\mathbb{Q}}[(w p_{t+1}^K - (1 - \delta_D) p_{t+1}^D - \gamma(1 - \delta_A) p_{t+1}^A) G_A(I_t, D_t, A_t) | S_t]. \quad (28)$$

Analogously, the shadow value of data satisfies

$$p_t^D = \rho(1 - \delta_D) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] + \rho \mathbb{E}_t^{\mathbb{Q}}[(w p_{t+1}^K - (1 - \delta_D) p_{t+1}^D - \gamma(1 - \delta_A) p_{t+1}^A) G_D(I_t, D_t, A_t) | S_t]. \quad (29)$$

Equation (28) highlights the dual role of technology A_t : it raises contemporaneous output (via Y_A) and amplifies future productivity (via G_A). By contrast, (29) has no immediate Y_A term; under our production specification, data affects value purely through the innovation channel. Quadratic adjustment costs enter through the KKT system (25)–(26), not directly in (28)–(29). Together, (28)–(29) imply a cross-asset feedback channel: a higher A_t raises the marginal option value of data via $G_D(I_t, D_t, A_t)$, while a higher D_t raises the marginal option value of R&D via $G_A(I_t, D_t, A_t)$. This self-reinforcing mechanism can amplify heterogeneity in investment paths when complementarities are strong, helping rationalize divergent growth trajectories among data-intensive firms.

On interior sets where $I_t > 0$, the budget shadow value satisfies $\lambda_t = \mathcal{U}_t$, so (28) can be written in the compact form

$$p_t^A = (1 - \tau) Y_A + \rho(1 - \delta_A) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] + \rho \mathbb{E}_t^{\mathbb{Q}}[(w p_{t+1}^K - (1 - \delta_D) p_{t+1}^D - \gamma(1 - \delta_A) p_{t+1}^A) (G_A + \tau Y_A G_I) | S_t].$$

At corners, equalities are interpreted via the KKT inequalities. Appendix G.9 provides detailed derivations and the interior/corner characterization.

5.5 Continuous-Time Limit: Shadow Price Diffusions

The continuous-time counterpart provides a valuation benchmark for the discrete shadow-price recursions and makes the role of aggregate risk transparent. In particular, it shows that the switching surfaces obtained in discrete time are consistent with diffusion-based free-boundary logic in standard real-options theory. Under the regularity conditions in Appendix H, the discrete recursions converge to the diffusion system below.

Theorem 3 (Shadow-Price Diffusions under Endogenous Technology (general G)). *Let $V(K, D, A, Z)$ solve the continuous-time HJB under the optimal policy, with $p^A := V_A$ and $p^D := V_D$. With $Y(K, A, Z) = (AZ)^{1-\xi} K^\xi$, $0 < \xi < 1$, cash flow $\Pi_t = (1 - \tau)Y_t - \frac{\kappa}{2} J_t^2 - \frac{\kappa'}{2} R_t^2$, and*

$$\begin{aligned} dK_t &= (-\delta K_t + G(I_t, D_t, A_t)) dt, & dD_t &= (-\delta_D D_t - G(I_t, D_t, A_t) + J_t) dt, \\ dA_t &= (-\delta_A A_t - \gamma G(I_t, D_t, A_t) + R_t) dt, & \frac{dZ_t}{Z_t} &= \mu_Z dt + \sigma_Z dW_t^{\mathbb{Q}}, \end{aligned} \quad I_t + J_t + R_t = \tau Y_t,$$

define $V_K := V_K$ and $Y_A = (1 - \xi)Z^{1-\xi}K^\xi A^{-\xi}$. Let

$$\mathcal{L}f := (-\delta_K + G)f_K + (-\delta_D D - G + J)f_D + (-\delta_A A - \gamma G + R)f_A + \mu_Z Z f_Z + \frac{1}{2}\sigma_Z^2 Z^2 f_{ZZ}.$$

Then, along the optimal policy,

$$\begin{aligned} (r + \delta_A - \mathcal{L})p^A &= (V_K - p^D - \gamma p^A)G_A(I, D, A) + (1 - \tau + \lambda_t \tau)Y_A, \\ (r + \delta_D - \mathcal{L})p^D &= (V_K - p^D - \gamma p^A)G_D(I, D, A), \end{aligned} \tag{30}$$

where $\lambda_t = (V_K - p^D - \gamma p^A)G_I(I_t, D_t, A_t)$ at interior points.

The operators $(r + \delta_\bullet - \mathcal{L})$ have the standard risk-neutral “carry-drift” form: discounting plus economic depreciation must be offset by expected marginal payoffs generated through innovation. In particular, the wedge $(V_K - p^D - \gamma p^A)$ governs whether it is more valuable at the margin to transform intangible option value into installed productive capacity or to preserve the intangible base for future flexibility.

Equivalently,

$$dp_t^A = \left[(r + \delta_A)p_t^A + (p_t^D + \gamma p_t^A - V_K)G_A(I_t, D_t, A_t) - (1 - \tau + \lambda_t \tau)Y_A \right] dt + \sigma_Z Z_t \partial_Z p_t^A dW_t^\mathbb{Q}, \tag{31a}$$

$$dp_t^D = \left[(r + \delta_D)p_t^D + (p_t^D + \gamma p_t^A - V_K)G_D(I_t, D_t, A_t) \right] dt + \sigma_Z Z_t \partial_Z p_t^D dW_t^\mathbb{Q}. \tag{31b}$$

Risk exposure. The diffusion coefficients $\sigma_Z Z_t \partial_Z p_t^A$ and $\sigma_Z Z_t \partial_Z p_t^D$ are the shadow-price exposures (“betas”) to aggregate technology risk under $\mathbb{Q}$. These exposures feed directly into continuation values and therefore into the geometry of the state-dependent exercise regions.

Coupling and drift interpretation. Equations (30)–(31) are coupled through the *funds wedge*

$$\Delta_t := V_K - p_t^D - \gamma p_t^A,$$

and through $\lambda_t = \Delta_t G_I(I_t, D_t, A_t)$. We refer to Δ_t as the *innovation spread* (or *transformation spread*): it is the net marginal gain from converting intangible option value into tangible productive capacity. When $\Delta_t > 0$, implementation creates net value at the margin and the optimal policy tilts toward scaling innovation; when $\Delta_t < 0$, preserving the intangible base is more valuable, shrinking the effective exercise region. The terms $(p_t^D + \gamma p_t^A - V_K)G_A$ and $(p_t^D + \gamma p_t^A - V_K)G_D$ show how the innovation elasticities G_A and G_D load the shadow-price drifts on this spread.

Remark 5 (Implementation shocks and normalization). We focus on aggregate-technology risk and set $w_t \equiv 1$ in the baseline. Allowing for stochastic implementation efficiency w_t (measurable in S_t with mean one) leaves the discrete-time KKT system unchanged: w_t enters only through the conditional expectations (e.g., in (24)). In the continuous-time limit, normalizing $w_t = 1 + o(\sqrt{\Delta t})$ makes its contribution second order. If one wishes implementation risk to survive, one can augment the state by w_t (or, equivalently, introduce a diffusion in K_t); the generator $\mathcal{L}$ then gains the corresponding terms, while the well-posedness and KKT arguments carry over without changing the economic trigger logic.

6 Comparative Analysis of Technology Regimes

6.1 Structural Differences

We contrast the benchmark regime with *static* firm-specific technology ($A_t \equiv \bar{A}$) against the *dynamic* regime in which A_t evolves endogenously through R&D. Three structural differences are central.

(1) *Investment opportunity set.* With A_t fixed, the firm allocates the reinvestment budget across implementation effort and data investment, (I_t, J_t) (with $R_t \equiv 0$ by construction). The data margin is governed by the excess-value gap

$$\Delta_J(S_t) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] - \lambda_t,$$

(similar to (25) and (27)), and the problem features a single *intangible* shadow-price process p_t^D .⁹ With endogenous A_t , the choice set expands to (I_t, J_t, R_t) : R&D becomes an additional lever that raises future productivity and, through innovation, feeds back into the marginal value of data. The resulting exercise-region problem is genuinely multi-dimensional.

(2) *Coupled shadow-price system.* Under static A_t , the shadow price p_t^D follows a one-asset recursion driven by discounting, depreciation, and innovation intensity. With endogenous A_t , a second shadow price p_t^A appears. The two prices are *coupled* through the innovation kernel $G(I, D, A)$ and the common budget shadow value λ_t ; see the coupled recursions and the PDE/SDE system in Section 5.3 and Appendix H. Accordingly, shocks to the fundamentals of one intangible asset propagate to the valuation of the other through the coupled system.

(3) *Real-option interaction.* Static A_t yields a stand-alone perpetual real option for data with a single state-dependent trigger. Endogenous A_t introduces a *second* real option (technology), and the two exercise rules interact: exercising the R&D option (raising A_t) shifts both the payoff and the trigger for the data option (investing in D_t). When the innovation kernel exhibits complementarity, data and R&D act as *dynamic complements* rather than substitutes. The static- A benchmark is nested by *exogenizing* the technology stock (i.e., imposing $A_t \equiv \bar{A}$ so that the p^A block drops out), which collapses the coupled system back to a single-option problem for data.

6.2 Cross-Asset Feedback and Valuation Coupling

Endogenous technology induces structural coupling between the marginal valuations of data and R&D through the innovation kernel $G(I, D, A)$ and the shadow value of internal funds λ_t .

(1) *Data $\rightarrow$ technology (innovation channel).* Suppose $G_{DA}(I, D, A) \geq 0$ (data-technology complementarity) and the Bellman operator preserves monotonicity (see Appendix H, Lemma 8). Then, along optimal policies, a larger data stock weakly raises the continuation value of technology. In smooth regions, this can be expressed as

$$\frac{\partial}{\partial D_t} \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A | S_t] \geq 0,$$

⁹The shadow value of internal funds λ_t governs feasibility and hence the trigger; see Section 5.3.

which follows from the p^A recursion/PDE under $G_{DA} \geq 0$ together with the small-gain/positivity conditions ensuring a positivity-preserving resolvent (Appendix H, Proposition 11 and Lemma 8; see also (28) and (30)).

(2) *Technology $\rightarrow$ data (financial-slack channel)*. Under the same order-preserving and small-gain bounds,

$$\frac{\partial}{\partial A_t} \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] \geq 0.$$

Two mechanisms reinforce this inequality. First, a higher A_t increases the productivity of data in innovation when G_D is (weakly) increasing in A . Second, A_t raises current output $Y(K_t, A_t)$ and hence the budget level τY_t ; under concavity of the inner allocation problem, the associated multiplier λ_t is weakly decreasing in the budget level (the same comparative-statics logic as in Proposition 5). Thus, higher technology can both raise the discounted marginal value of data and relax the effective funding wedge. We refer to this mechanism as the *financial-slack effect*: stronger technology relaxes the internal funding wedge by raising cash-generation capacity, thereby lowering the effective hurdle λ_t for data accumulation.

(3) *Stability (dampened feedback loop)*. While data and technology reinforce each other through $G(\cdot)$, discounting and depreciation dampen this feedback so that the coupled valuations remain well behaved rather than explosive. Formally, Appendix H provides a sufficient “small-gain” condition ensuring that cross-asset complementarity is dominated by the effective cost of capital (equivalently, the coupled resolvent is a contraction in the relevant order). In the discrete-time implementation, Assumption A9 gives convenient sufficient inequalities (including the analogous bound for the A -block) ensuring that valuations are monotone in (D_t, A_t) yet remain uniformly bounded and well behaved.

6.3 Strategic and Practical Implications

Strategic synchronization of data and R&D. Data’s option value is realized most strongly when paired with sufficient innovation intensity. The KKT triggers imply data investment when $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] > \lambda_t$. With endogenous A_t , R&D raises the left-hand side through the coupled shadow-price system and can (via higher Y) lower the right-hand side, thereby enlarging the exercise region. Coordination of (J_t, R_t) is therefore value-critical: a firm that accumulates data without a corresponding R&D strategy behaves closer to the static- A benchmark, in which data investment is primarily episodic and less persistent. Conversely, a firm that accumulates data without sufficient R&D capacity may fall into a “data trap”: $J_t > 0$ raises measured data assets, but a weakly evolving technology frontier cannot fully monetize them, compressing marginal returns and triggering eventual retrenchment.

Heterogeneity across industries. The coupled mechanism generates qualitative cross-industry predictions. In knowledge-intensive sectors (e.g., AI, biotech, digital platforms) where technology responds strongly to R&D, the model predicts pronounced comovement between data and R&D spending and frequent operation near the joint-investment region. In capital-intensive sectors with comparatively stagnant firm-specific technology (approximating the static- A benchmark), data investment is expected to be more sporadic and front-loaded, driven more by initial conditions and weaker dynamic feedback.

Policy spillovers. Data-access rules affect both p_t^D and p_t^A . Regulations that restrict effective data access (e.g., privacy rules that reduce usable D_t) lower the shadow price of data and, through cross-asset coupling, can also depress the shadow value of technology and incentives for R&D. In this sense, privacy and competition policies influence long-run technology growth through the real-option value of data and its interaction with endogenous innovation. This highlights a potential unintended consequence of data regulation: by reducing effective usable data, stricter access constraints can indirectly weaken incentives for R&D through the complementarity channel, even when the policy target is narrowly privacy-oriented.

Under static A_t , data behaves as a stand-alone option with a single trigger and a tendency to depreciate absent continued investment. With endogenous A_t , the valuation landscape becomes *compound*: (p_t^D, p_t^A) are jointly determined, triggers interact through λ_t and G , and data acquires durable value primarily when paired with R&D. This coupled real-option structure underlies the quantitative patterns and robustness documented in Section 7.

7 Numerical Implementation and Results

We solve the Bellman equation on the four-dimensional state (K, D, A, Z) using value iteration with Gauss–Seidel updates and Gauss–Hermite (GH) quadrature for the aggregate-technology shock. Throughout, conditional expectations are evaluated under the valuation measure $\mathbb{Q}$ for consistency with the real-options representations (Remark 3); this measure convention is used solely for valuation and does not alter the policy characterization.

State grid and expectations. The value function is computed on a tensor-product grid with $(N_K, N_D, N_A, N_Z) = (60, 60, 60, 11)$. The (K, D, A) grids are log-spaced around the baseline normalization $(1, 1, 1)$; in the constant- A benchmark, the A dimension collapses to a singleton node at $A \equiv 1$. The Z grid is centered at $Z = 1$ and aligned with the GH nodes. Expectations over Z' use a probability-form GH rule with $n_{\text{GH}} = 5$ nodes in the baseline, mapping standard-normal nodes through the exact lognormal step

$$Z' = Z \exp\left((\mu_Z - \tfrac{1}{2}\sigma_Z^2)\Delta + \sigma_Z\sqrt{\Delta}\varepsilon\right), \quad \varepsilon \sim \mathcal{N}(0, 1),$$

so quadrature weights integrate the standard normal and sum to one. Increasing n_{GH} to 7 or 9 leaves the reported objects essentially unchanged (Appendix I). Off-grid continuation values are evaluated by multilinear interpolation; boundary handling and robustness checks are documented in Appendix I.

Intra-period allocation and policy computation. At each state, we compute the shadow value of internal funds λ so that the flow budget holds, $I + J + R = \tau Y$, and recover the adjustment-cost intensities from the KKT system:

$$J = \frac{(\rho \mathbb{E}^{\mathbb{Q}}[p_{t+1}^D | S] - \lambda)_+}{\kappa_J}, \quad R = \frac{(\rho \mathbb{E}^{\mathbb{Q}}[p_{t+1}^A | S] - \lambda)_+}{\kappa_R}, \quad I = \tau Y - J - R.$$

Shadow prices (p^D, p^A) are computed by finite differences on the grid and are numerically consistent with the envelope recursions in Appendix G.9. Value iteration terminates when




Figure 2: Baseline paths (constant A ; mid- Z slice). Upper panel: (K_t, D_t, Y_t) with $A_t \equiv \bar{A}$ (dashed). Lower panel: the data-investment gap $\Delta_J = \rho \mathbb{E}_t^Q[p_{t+1}^D] - \lambda_t$. Along this slice the gap is mildly positive, so the firm continues to invest in data. Net of depreciation and data usage through innovation, however, D_t erodes and output drifts down under the constant- A benchmark.

Bellman and feasibility diagnostics fall below tight tolerances; the full set of stopping rules and residual definitions is reported in Appendix I. Convergence follows from the standard contraction arguments for discounted dynamic programs (Theorem 2 and Appendix C).

7.1 Baseline calibration and diagnostics

Unless otherwise noted we set (quarterly) $\Delta = 0.25$, $\rho = e^{-0.05\Delta}$, $\xi = 0.35$, $\tau = 0.30$, and $(\delta, \delta_D, \delta_A) = (0.04, 0.06, 0.06)$. The innovation kernel $G(I, D, A; \theta)$ is Revankar-CES with parameters $(\alpha, \beta, a, b, \theta, \gamma) = (0.60, 1.0, 0.50, 0.30, 0.80, 0.50)$, and quadratic adjustment costs are imposed on data acquisition and R&D with $(\kappa_J, \kappa_R) = (6, 8)$.¹⁰ Aggregate technology follows lognormal growth with annual volatility 0.15 and zero drift.

Convergence is fast and numerically tight in both regimes. Representative diagnostics are reported in Table 3 (Appendix): Bellman residuals are on the order of 10^{-4} (relative) with feasibility and complementarity residuals typically below 10^{-7} . Grid refinement ($N_\bullet : 60 \rightarrow 80$) and quadrature refinement ($n_{\text{GH}} : 5 \rightarrow 7$) yield sup-norm differences below 10^{-3} and negligible frontier displacement, indicating that the baseline discretization is sufficient for the comparative statics reported below.

Figures 2–3 illustrate the economic contrast between regimes. With constant A , the firm can rationalize continued data spending along the baseline slice, but depreciation and data usage dominate on net, so the data stock and output drift down. With endogenous A , intermittent activation of the R&D margin raises the technology stock and sustains output growth even

¹⁰Notation mapping: theory uses (κ, κ') , while numerics use (κ_J, κ_R) for the data and R&D adjustment coefficients.




Figure 3: Baseline paths (dynamic A ; mid- Z slice). Upper panel: (K_t, D_t, A_t, Y_t) . Lower panel: internal gaps (Δ_J, Δ_R) . Along this slice $\Delta_J \leq 0$ so new data acquisition is not triggered, while Δ_R turns positive intermittently, activating R&D and driving technology accumulation. Endogenous technology sustains higher output relative to the constant- A benchmark.

when the baseline slice lies in the data-inaction region.

7.2 Frontier extraction and comparative statics in τ and θ

To summarize state-dependent exercise behavior, we extract switching frontiers as zero-level sets of the signed KKT gaps Δ_J and Δ_R on a mid- Z slice. Figure 4 varies the reinvestment share $\tau \in \{0.25, 0.30, 0.35\}$. Relaxing the internal budget constraint expands the data-investment exercise region, consistent with Proposition 5. Figure 5 varies the data weight $\theta \in \{0.7, 0.9, 1.1\}$ and shows that higher data effectiveness in the innovation kernel shifts the data frontier outward, consistent with Proposition 6. Implementation details for frontier extraction and interpolation are reported in Appendix I.

7.3 Robustness and sensitivity analysis

We complement the frontier comparisons with a global sensitivity analysis that perturbs $(\tau, \delta_A, \delta_D, \theta)$ over randomized parameter draws in a local neighborhood of the baseline calibration. The resulting tornado chart (Figure 6) confirms that τ is quantitatively dominant, confirming that internal financing capacity is the primary driver of value creation in this friction-constrained environment. δ_A has a moderate negative effect, and (θ, δ_D) are locally smaller. Frontier-shift histograms (Figure 7) show that switching surfaces move little under these local perturbations. Full sampling design and regression details are provided in Appendix I.

Finally, we verify time-discretization consistency by re-solving on a monthly grid ($\Delta = 1/12$) with the same annual primitives and comparing to the quarterly baseline. Value-function differences are on the order of 10^{-3} – 10^{-2} in sup norm and frontier displacement is negligible;




Figure 4: Frontier overlay in τ on (K, D) at mid- Z . Solid: Γ_J with endogenous A ; dashed: constant- A benchmark. Higher τ expands the internal data-investment exercise region; the right-axis profile reports the average vertical shift ΔD^* between the lowest and highest τ .




Figure 5: Frontier overlay in θ on (K, D) at mid- Z . Higher θ expands the data-investment exercise region; the right-axis profile reports the average vertical shift ΔD^* between the lowest and highest θ .

details and mappings are reported in Appendix I.

Replication package. The complete replication package (code, parameter files, and logs) is provided in the e-companion; Appendix I documents the numerical design and diagnostics supporting the reported results.

8 Conclusion

This paper develops a dynamic growth-option framework for valuing firm-level data and R&D under internal financing constraints. By embedding physical capital, proprietary data advantage, and firm-specific technology into a unified dynamic program, we connect intangible accumulation to state-dependent investment triggers and shadow valuations.

Our analysis yields three main insights. First, data behaves as an option-like intangible: the incentive to accumulate it is generated by continuation value, and optimal data spending is




Figure 6: Global sensitivity analysis for standardized effects on mean output $\bar{Y}$ ($N = 40$). τ is quantitatively dominant; δ_A is moderate and negative; θ and δ_D are locally smaller.




Figure 7: Distribution of D -frontier shifts across randomized parameter draws. Mass concentrates near zero, indicating small relocation of switching surfaces under local perturbations.

triggered when the discounted marginal value of data exceeds the shadow cost of internal funds. This endogenous funding wedge links data investment to the firm's reinvestment capacity and implies a procyclical tendency in data accumulation when internal funds become scarce.

Second, endogenizing firm-specific technology reveals a dual-option mechanism in which data and R&D act as dynamic complements. A stronger technology base tends to raise the option value of data and expand the data-investment region, while richer data stocks tend to increase the shadow value of technology by strengthening innovation intensity. Conversely, the model highlights a potential "data trap": accumulating data without a synchronized technology and

R&D strategy can leave data under-monetized when the innovation frontier is weak.

Third, our numerical implementation confirms that these mechanisms generate stable switching frontiers and economically meaningful comparative statics under baseline calibration and robustness checks. In particular, relaxing the reinvestment constraint and increasing data effectiveness in the innovation channel shift the data-investment frontier outward, consistent with the theory.

Managerially, the results suggest that optimal data strategy requires pricing the strategic patience of internal accumulation and coordinating data acquisition with R&D to avoid diminishing returns. From a policy perspective, restrictions that reduce effective data access may have unintended spillovers by weakening R&D incentives through the complementarity channel. A natural direction for future work is to incorporate strategic competition and endogenous data-market pricing, as well as richer financing frictions (e.g., costly external finance) to sharpen the interaction between financial structure and intangible investment.

References

- Abis S, Veldkamp L (2024) The changing economics of knowledge production. *The Review of Financial Studies* 37(1):89–118.
- Acemoglu D, Jensen MK (2023) Equilibrium analysis in behavioural one-sector growth models. *Review of Economic Studies* 91(2):599–640.
- Aguerrevere FL (2009) Real options, product market competition, and asset returns. *The Journal of Finance* 64(2):957–983.
- Ahn Y (2019) Intangible capital, volatility shock, and the value premium. *Financial Review* 54(4):739–762.
- Ai H, Croce MM, Diercks AM, Li K (2018) News shocks and the production-based term structure of equity returns. *The Review of Financial Studies* 31(7):2423–2467.
- Ai H, Croce MM, Li K (2013) Toward a quantitative general equilibrium asset pricing model with intangible capital. *The Review of Financial Studies* 26(2):491–530.
- Arthur WB (1989) Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal* 99(394):116–131.
- Bai H, Hou K, Kung H, Li E, Zhang L (2019) The CAPM strikes back? An equilibrium model with disasters. *Journal of Financial Economics* 131(2):269–298.
- Begenau J, Farboodi M, Veldkamp L (2018) Big data in finance and the growth of large firms. *Journal of Monetary Economics* 97:71–87.
- Bellalah M (2016) Issues in real options with shadow costs of incomplete information and short sales. *The Journal of Economic Asymmetries* 13:45–56.
- Berk JB, Green RC, Naik V (1999) Optimal investment, growth options, and security returns. *The Journal of Finance* 54(5):1553–1607.

- Cong LW, Xie D, Zhang L (2021) Knowledge accumulation, privacy, and growth in a data economy. *Management Science* 67(10):6480–6492.
- Dixit AK, Pindyck RS (1994) *Investment Under Uncertainty* (Princeton University Press, Princeton, NJ).
- Ewald CO, Taub B (2022) Real options, risk aversion and markets: A corporate finance perspective. *Journal of Corporate Finance* 72:102164.
- Farboodi M, Matray A, Veldkamp L, Venkateswaran V (2022a) Where has all the data gone? *The Review of Financial Studies* 35(7):3101–3138.
- Farboodi M, Matray A, Veldkamp L, Venkateswaran V (2022b) Valuing financial data. Working Paper 29894, National Bureau of Economic Research, Cambridge, MA.
- Gawer A, Henderson R (2008) *Platform Leadership: How Intel, Microsoft, and Apple Built Dynamic Capabilities* (Harvard Business School Press, Boston, MA).
- Jones CI (1995) R&D-based models of economic growth. *Journal of Political Economy* 103(4):759–784.
- Jones CI, Tonetti C (2020) Nonrivalry and the economics of data. *American Economic Review* 110(9):2819–2858.
- Kogan L, Papanikolaou D (2012) A theory of firm characteristics and stock returns: The role of investment-specific shocks. Working Paper 17975, National Bureau of Economic Research, Cambridge, MA.
- Kortum S (1997) Research, patenting, and technological change. *Econometrica* 65(6):1389–1419.
- Kuosmanen T, Zhou X (2021) Shadow prices and marginal abatement costs: Convex quantile regression approach. *European Journal of Operational Research* 289:666–675.
- Lambrecht BM (2017) Real options in finance. *Journal of Banking & Finance* 81:166–171.
- Lu J, Lv L, Wang Y, Zhu Y (2025) The real effect of monetary policy under uncertainty: Evidence from the change in corporate financing purposes. *Journal of Banking & Finance* 172:107381.
- Pakes A (1986) Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica* 54(4):755–784.
- Revankar N (1971) A class of variable elasticity of substitution production functions. *Econometrica* 39(1):61–71.
- Romer PM (1990) Endogenous technological change. *Journal of Political Economy* 98(5):S71–S102.
- The Royal Society, The British Academy (2017) Data management and use: Governance in the 21st century. The Royal Society, London.

- Segerstrom PS (1998) Endogenous growth without scale effects. *American Economic Review* 88(5):1290–1310.
- Veldkamp L (2023) Valuing data as an asset. *Review of Finance* 27(5):1545–1562.
- Veldkamp L, Chung C (2024) Data and the aggregate economy. *Journal of Economic Literature* 62(2):458–484.

In the appendix, we provide supplementary proofs, derivations, and robustness checks.

A Related Properties of the Innovation Kernel G

This appendix establishes properties of the innovation kernel used in the main text. Section A.1 proves Lemma 1. Section A.2 provides sufficient conditions for Assumption **A5**. Section A.3 provides a sufficient (small-gain) condition for Assumption **A6**. Throughout, we suppress time subscripts and write (I, D, A) for (I_t, D_t, A_t) .

A.1 Proof of Lemma 1

Proof of Lemma 1. For notational convenience, define

$$\Sigma := A^\alpha + \theta D^\alpha.$$

Then for $I > 0$, the innovation kernel is given by

$$G(I, D, A) = \beta \exp\left(aI\Sigma^{-1/\alpha}\right) I^{1-b} A^b D^b.$$

Using the chain rule, we compute the derivative of G with respect to D :

$$\begin{aligned} \frac{\partial G}{\partial D} &= G \left(\frac{b}{D} + aI \frac{\partial}{\partial D} (\Sigma^{-1/\alpha}) \right) \\ &= G \left(\frac{b}{D} - aI \frac{1}{\alpha} \Sigma^{-1/\alpha-1} \cdot \theta_\alpha D^{\alpha-1} \right) \\ &= \frac{G}{D} \left[b - \theta aI \frac{D^\alpha}{(A^\alpha + \theta D^\alpha)^{1+1/\alpha}} \right]. \end{aligned}$$

Since $I, D, A \in \mathbb{R}_{++}$, we have $G(I, D, A) > 0$ and $G(I, D, A)/D > 0$. Thus, $\partial_D G > 0$ if and only if

$$b - \theta aI \frac{D^\alpha}{(A^\alpha + \theta D^\alpha)^{1+1/\alpha}} > 0,$$

equivalently,

$$b(A^\alpha + \theta D^\alpha)^{(\alpha+1)/\alpha} > \theta aID^\alpha,$$

which is the stated condition. For $I = 0$ we have $G(0, D, A) = 0$ and hence $\partial_D G(0, D, A) = 0$, so the equivalence below is stated for $I > 0$. □

A.2 Sufficient conditions for Assumption A5

Proposition 7 (Sufficient conditions for Assumption **A5**). *Let $\mathcal{S}_0 \subset \mathbb{R}_{++}^3$ be the compact set in Assumption **A4**, and define $\mathcal{S}_0^{DA} := \{(D, A) : \exists K \text{ s.t. } (K, D, A) \in \mathcal{S}_0\}$. Fix any finite bound $\bar{y} > 0$ and consider the domain $[0, \bar{y}] \times \mathcal{S}_0^{DA}$.*

For

$$G(I, D, A) = \beta \exp(aI\phi(D, A)) I^{1-b} A^b D^b, \quad \phi(D, A) := (A^\alpha + \theta D^\alpha)^{-1/\alpha},$$

with $a, \beta > 0$, $b \in (0, 1)$, $\alpha \in (0, 1)$, and $\theta \in (0, 1]$, the following hold:

- (i) *Regularity.* $G \in C^2((0, \bar{y}) \times \mathcal{S}_0^{DA})$ and its partial derivatives admit continuous extensions to $(0, \bar{y}] \times \mathcal{S}_0^{DA}$. Moreover, G admits a continuous extension at $I = 0$ with $G(0, D, A) = 0$.
- (ii) *Concavity in I on $[0, \bar{y}]$.* Let $\phi_{\max} := \sup_{(D, A) \in \mathcal{S}_0^{DA}} \phi(D, A) < \infty$. For every $(D, A) \in \mathcal{S}_0^{DA}$, the map $I \mapsto G(I, D, A)$ is concave on $[0, \bar{y}]$ provided that

$$\bar{y} \leq \underline{I}^{\text{conc}} := \frac{\sqrt{1-b} - (1-b)}{a\phi_{\max}}. \quad (32)$$

If $\bar{y} < \underline{I}^{\text{conc}}$, then the concavity is strict on $(0, \bar{y}]$.

- (iii) *Continuity moduli.* G is Hölder continuous in I on $[0, \bar{y}]$ with exponent $(1-b)$ and Lipschitz continuous in (D, A) uniformly in I .

Proof of Proposition 7. Since $\mathcal{S}_0$ is compact, its projection $\mathcal{S}_0^{DA}$ is compact. Because ϕ is continuous on $\mathbb{R}_{++}^2$, there exist $0 < \phi_{\min} \leq \phi(D, A) \leq \phi_{\max} < \infty$ on $\mathcal{S}_0^{DA}$.

(i) *Regularity.* For $I > 0$, G is a product/composition of smooth functions in (I, D, A) on $(0, \bar{y}) \times \mathcal{S}_0^{DA}$, hence $G \in C^2$ there, and the derivatives extend continuously to $(0, \bar{y}] \times \mathcal{S}_0^{DA}$. Moreover, $G(0, D, A) = 0$ and $I^{1-b} \rightarrow 0$ as $I \downarrow 0$, so G extends continuously to $I = 0$.

(ii) *Concavity in I .* Fix $(D, A) \in \mathcal{S}_0^{DA}$ and define $c := a\phi(D, A) > 0$, $d := 1 - b \in (0, 1)$, and $C := \beta A^b D^b > 0$. Then for $I > 0$,

$$G(I, D, A) = Ce^{cI}I^d.$$

Differentiating twice yields

$$\partial_I^2 G(I, D, A) = Ce^{cI}I^{d-2}q(I), \quad q(I) := c^2I^2 + 2cdI + d(d-1).$$

The quadratic q has a unique positive root

$$I^{(D, A)} = \frac{-d + \sqrt{d}}{c} = \frac{\sqrt{1-b} - (1-b)}{a\phi(D, A)},$$

and satisfies $q(I) \leq 0$ for all $I \in (0, I^{(D, A)}]$. If (32) holds, then $\bar{y} \leq I^{(D, A)}$ for all $(D, A) \in \mathcal{S}_0^{DA}$ because $\phi(D, A) \leq \phi_{\max}$ implies $I^{(D, A)} \geq \frac{\sqrt{1-b} - (1-b)}{a\phi_{\max}}$. Hence $\partial_I^2 G(I, D, A) \leq 0$ for all $I \in (0, \bar{y}]$ and all $(D, A) \in \mathcal{S}_0^{DA}$, so $I \mapsto G(I, D, A)$ is concave on $(0, \bar{y}]$. Since $\lim_{I \downarrow 0} G(I, D, A) = 0 = G(0, D, A)$, the concavity inequality extends to $[0, \bar{y}]$ by taking $\varepsilon \downarrow 0$ in the concavity inequality between ε and I .

(iii) *Continuity moduli.* We establish (a) uniform Hölder continuity in I on $[0, \bar{y}]$ and (b) uniform Lipschitz continuity in (D, A) on $\mathcal{S}_0^{DA}$.

(a) *Hölder continuity in I .* Fix $(D, A) \in \mathcal{S}_0^{DA}$ and write $c := a\phi(D, A)$ and $d := 1 - b \in (0, 1)$. For any $I_1, I_2 \in [0, \bar{y}]$,

$$\begin{aligned} |e^{cI_1}I_1^d - e^{cI_2}I_2^d| &\leq e^{c\bar{y}}|I_1^d - I_2^d| + \bar{y}^d|e^{cI_1} - e^{cI_2}| \\ &\leq e^{c\bar{y}}|I_1 - I_2|^d + ce^{c\bar{y}}\bar{y}^d|I_1 - I_2|. \end{aligned}$$

Using $c = a\phi(D, A) \leq a\phi_{\max}$, we have $e^{c\bar{y}} \leq e^{a\phi_{\max}\bar{y}}$ and $ce^{c\bar{y}} \leq a\phi_{\max}e^{a\phi_{\max}\bar{y}}$. Since $|I_1 - I_2| \leq \bar{y}^{1-d}|I_1 - I_2|^d$ on $[0, \bar{y}]$, the right-hand side is bounded by $C_1|I_1 - I_2|^d$ for a constant C_1 that does not depend on (D, A) . Multiplying by the bounded factor $\beta A^b D^b$ on $\mathcal{S}_0^{DA}$ yields

$$|G(I_1, D, A) - G(I_2, D, A)| \leq C_I |I_1 - I_2|^{1-b}, \quad \forall I_1, I_2 \in [0, \bar{y}],$$

uniformly in $(D, A) \in \mathcal{S}_0^{DA}$.

(b) *Lipschitz continuity in (D, A) .* The partial derivatives $\partial_D G$ and $\partial_A G$ are continuous in (I, D, A) on the compact set $[0, \bar{y}] \times \mathcal{S}_0^{DA}$ (note that $\partial_D G, \partial_A G \rightarrow 0$ as $I \downarrow 0$ because $I^{1-b} \rightarrow 0$). Hence there exists

$$L_{DA} := \sup_{(I, D, A) \in [0, \bar{y}] \times \mathcal{S}_0^{DA}} \|\nabla_{(D, A)} G(I, D, A)\| < \infty.$$

By the mean-value theorem, for all $I \in [0, \bar{y}]$,

$$|G(I, D_1, A_1) - G(I, D_2, A_2)| \leq L_{DA}(|D_1 - D_2| + |A_1 - A_2|),$$

which proves uniform Lipschitz continuity in (D, A) . $\square$

A.3 Cross-complementarity under a small-gain condition

This subsection provides an easily verifiable *sufficient* condition for cross-complementarity. In the main text we maintain Assumption **A6** to allow for broader parameterizations beyond the small-gain region.

Lemma 2 (Cross-complementarity under a small-gain condition). *Let $\mathcal{S}_0^{DA} \subset \mathbb{R}_{++}^2$ be compact and fix $\bar{y} > 0$. There exists a constant $\bar{a} > 0$ (depending only on $\mathcal{S}_0^{DA}$ and the parameters b, α, θ) such that*

$$0 \leq a\bar{y} \leq \bar{a} \implies G_{AD}(I, D, A) \geq 0, \quad \forall (I, D, A) \in (0, \bar{y}] \times \mathcal{S}_0^{DA}.$$

By continuity, the inequality extends to $I = 0$ with $G_{AD}(0, D, A) = 0$.

Proof. For $I > 0$, write

$$\log G(I, D, A) = \log \beta + aI\phi(D, A) + (1-b)\log I + b\log A + b\log D, \quad \phi(D, A) = (A^\alpha + \theta D^\alpha)^{-1/\alpha}.$$

Then

$$\frac{G_{AD}}{G} = (\partial_A \log G)(\partial_D \log G) + \partial_{AD} \log G.$$

A direct calculation gives

$$\partial_A \log G = \frac{b}{A} + aI\phi_A, \quad \partial_D \log G = \frac{b}{D} + aI\phi_D, \quad \partial_{AD} \log G = aI\phi_{AD},$$

so

$$\frac{G_{AD}}{G} = \frac{b^2}{AD} + aI \left(\frac{b}{A}\phi_D + \frac{b}{D}\phi_A + \phi_{AD} \right) + (aI)^2 \phi_A \phi_D. \quad (33)$$

Because ϕ is smooth on $\mathbb{R}_{++}^2$ and $\mathcal{S}_0^{DA}$ is compact, the quantities $\phi_A, \phi_D, \phi_{AD}$, and $\phi_A\phi_D$ are bounded on $\mathcal{S}_0^{DA}$. Let

$$A_{\max} := \sup_{\mathcal{S}_0^{DA}} A, \quad D_{\max} := \sup_{\mathcal{S}_0^{DA}} D,$$

and define finite constants

$$M_1 := \sup_{(D,A) \in \mathcal{S}_0^{DA}} \left| \frac{b}{A} \phi_D + \frac{b}{D} \phi_A + \phi_{AD} \right|, \quad M_2 := \sup_{(D,A) \in \mathcal{S}_0^{DA}} |\phi_A \phi_D|.$$

Using $I \leq \bar{y}$ in (33), we obtain the lower bound

$$\frac{G_{AD}}{G} \geq \frac{b^2}{A_{\max} D_{\max}} - (a\bar{y})M_1 - (a\bar{y})^2 M_2, \quad \forall (I, D, A) \in (0, \bar{y}] \times \mathcal{S}_0^{DA}.$$

Let $c_0 := \frac{1}{2} \frac{b^2}{A_{\max} D_{\max}} > 0$. Choose $\bar{a} > 0$ such that for all $0 \leq a\bar{y} \leq \bar{a}$,

$$(a\bar{y})M_1 + (a\bar{y})^2 M_2 \leq c_0.$$

One admissible explicit choice is, for $M_1, M_2 > 0$,

$$\bar{a} := \min \left\{ \frac{c_0}{2M_1}, \sqrt{\frac{c_0}{2M_2}} \right\},$$

with the obvious simplifications when $M_1 = 0$ and/or $M_2 = 0$. Then $\frac{G_{AD}}{G} \geq c_0 > 0$ on $(0, \bar{y}] \times \mathcal{S}_0^{DA}$. Since $G > 0$ for $I > 0$, this implies $G_{AD} \geq 0$ there. Moreover, (G_{AD}/G) is bounded on $(0, \bar{y}] \times \mathcal{S}_0^{DA}$ and $G(I, D, A) \rightarrow 0$ as $I \downarrow 0$, hence $G_{AD}(I, D, A) \rightarrow 0$, which yields the continuous extension at $I = 0$ with $G_{AD}(0, D, A) = 0$. Therefore, the inequality extends to $I = 0$. $\square$

B Appendix: Proofs for Section 3

Throughout this appendix we work on interior (cap-slack) regions where $\tilde{G}_t = G(I_t, D_t, A_t)$, so the derivative-based formulas for G_I, G_D apply. Boundary/kink cases where the cap binds can be handled with standard directional-derivative/subgradient arguments.

B.1 Proof of Proposition 1

Proof of Proposition 1. Fix $A_t \equiv A$ and $S_t = (K_t, D_t)$. Under the static benchmark, J_t affects D_{t+1} only through (8) and does not enter the innovation kernel G_t at date t . Hence, conditional on S_t and w_t , the objective Y_{t+1} depends on (I_t, J_t) only through I_t via $G(I_t, D_t)$, and the problem reduces to choosing I_t subject to feasibility.

Note. The argument holds pointwise for any $w_t > 0$; the deterministic case $w_t \equiv w$ is a special case.

Write the innovation kernel as

$$G(I_t, D_t) = \beta \exp\left(a I_t \phi(D_t, A)\right) I_t^{1-b} A^b D_t^b, \quad \phi(D, A) := (A^\alpha + \theta D^\alpha)^{-1/\alpha}.$$

Let $c := a \phi(D_t, A) > 0$ and $d := 1 - b \in (0, 1)$. For every $I_t > 0$,

$$\frac{\partial G}{\partial I}(I_t, D_t) = \beta e^{c I_t} A^b D_t^b (c I_t^d + d I_t^{d-1}) = G(I_t, D_t) \left(c + \frac{d}{I_t}\right) > 0.$$

Since $w_t > 0$ by construction, the chain rule yields

$$\frac{dY_{t+1}}{dI_t} = \xi A^{1-\xi} \left((1-\delta)K_t + w_t G(I_t, D_t)\right)^{\xi-1} \cdot w_t \cdot \frac{\partial G}{\partial I}(I_t, D_t) > 0,$$

for all $I_t > 0$. At the boundary $I_t = 0$, $G(0, D_t) = 0$ and the right derivative is $+\infty$ (because $d - 1 < 0$), so $I \mapsto Y_{t+1}(I; S_t, w_t)$ is strictly increasing on any compact feasible interval.

It remains to incorporate the admissibility constraint $G(I_t, D_t) \leq D_t$. Because $I \mapsto G(I, D_t)$ is continuous and strictly increasing on $(0, \infty)$, the feasible set for I_t is the (possibly truncated) interval

$$\mathcal{I}(S_t) := \{I \in [0, \tau Y_t] : G(I, D_t) \leq D_t\} = [0, \min\{\tau Y_t, \bar{I}(D_t)\}],$$

where $\bar{I}(D_t) := \sup\{I \geq 0 : G(I, D_t) \leq D_t\} \in (0, \infty]$. Since Y_{t+1} is strictly increasing in I over $\mathcal{I}(S_t)$, the unique maximizer is

$$I_t^* = \min\{\tau Y_t, \bar{I}(D_t)\}, \quad J_t^* = \tau Y_t - I_t^*.$$

Under the maintained benchmark state-space/admissibility condition used in Section 3.1 (which ensures $G(\tau Y_t, D_t) \leq D_t$ on the relevant region), one has $\bar{I}(D_t) \geq \tau Y_t$ and thus $I_t^* = \tau Y_t$, implying $J_t^* = 0$ as stated.

The same conclusion holds if the objective is the conditional expectation $\mathbb{E}[Y_{t+1} \mid S_t]$: for any $0 \leq I_1 < I_2 \leq \tau Y_t$ and any realization with $w_t > 0$, the mapping $I \mapsto Y_{t+1}(I; w_t)$ is strictly increasing on the feasible interval $\mathcal{I}(S_t)$, hence $Y_{t+1}(I_2; w_t) > Y_{t+1}(I_1; w_t)$ pointwise. Taking conditional expectations yields $\mathbb{E}[Y_{t+1}(I_2) \mid S_t] > \mathbb{E}[Y_{t+1}(I_1) \mid S_t]$, so the conditional expectation is strictly increasing on $\mathcal{I}(S_t)$ and the unique maximizer remains the upper feasible bound, which equals $I_t = \tau Y_t$ under the maintained benchmark admissibility condition. $\square$

B.2 Proof of Theorem 1

Proof of Theorem 1. Let $\mathcal{S}_0 \subset \mathbb{R}_{++}^2$ be the compact, forward-invariant state space for (K, D) (Assumption A4), and let $\mathcal{B} := C_b(\mathcal{S}_0)$ denote the Banach space of bounded continuous functions on $\mathcal{S}_0$ endowed with the sup norm $\|\cdot\|_\infty$.

Under invariant technology $A_t \equiv A$, define the Bellman operator on pre-decision states by

$$(TV)(K, D) := \max_{\substack{I, J \geq 0, \\ I+J=\tau Y(K), \\ G(I, D) \leq D}} \left\{ (1-\tau)Y(K) + \rho \mathbb{E}[V(K', D') \mid K, D, I] \right\},$$

where

$$K' = (1 - \delta)K + w_t G(I, D), \quad D' = (1 - \delta_D)(D - G(I, D)) + J,$$

and w_t is the i.i.d. implementation shock with bounded support $0 < \underline{w} \leq w_t \leq \bar{w} < \infty$ and $\mathbb{E}[w_t] = 1$. Because $J = \tau Y(K) - I$, the maximization can equivalently be written over the one-dimensional feasible set

$$I \in \{I \in [0, \tau Y(K)] : G(I, D) \leq D\}.$$

Step 1 (Well-definedness: $\mathcal{T} : \mathcal{B} \rightarrow \mathcal{B}$). By Assumptions A1–A4, the state space $\mathcal{S}_0$ is compact and forward invariant, the feasible set is nonempty and compact at each (K, D) , and $Y(K)$ and $G(I, D)$ are continuous in their arguments by their explicit forms. Moreover, w_t has bounded support, and the admissibility constraint $G(I, D) \leq D$ guarantees $D' \geq 0$. By the forward-invariance property of $\mathcal{S}_0$ (Assumption A4), $(K', D') \in \mathcal{S}_0$ for all feasible choices.

Fix $V \in \mathcal{B}$. For each (K, D, I) , the integrand $V(K', D')$ is bounded by $\|V\|_\infty$, hence $\mathbb{E}[V(K', D')]$ is well-defined and continuous in (K, D, I) by dominated convergence. Berge's Maximum Theorem then implies that $\mathcal{T}V$ is continuous and bounded on $\mathcal{S}_0$, i.e., $\mathcal{T} : \mathcal{B} \rightarrow \mathcal{B}$.

Step 2 (Contraction). Take any $V_1, V_2 \in \mathcal{B}$ and any $(K, D) \in \mathcal{S}_0$. For any feasible (I, J) ,

$$\left| (1 - \tau)Y + \rho \mathbb{E}[V_1(K', D')] - ((1 - \tau)Y + \rho \mathbb{E}[V_2(K', D')]) \right| = \rho \left| \mathbb{E}[(V_1 - V_2)(K', D')] \right|.$$

By Jensen's inequality and boundedness of $V_1 - V_2$,

$$\left| \mathbb{E}[(V_1 - V_2)(K', D')] \right| \leq \mathbb{E}[|V_1 - V_2|(K', D')] \leq \|V_1 - V_2\|_\infty.$$

Taking the maximum over feasible controls and then the supremum over (K, D) gives

$$\|\mathcal{T}V_1 - \mathcal{T}V_2\|_\infty \leq \rho \|V_1 - V_2\|_\infty,$$

so $\mathcal{T}$ is a ρ -contraction on $(\mathcal{B}, \|\cdot\|_\infty)$.

Step 3 (Existence, uniqueness, and convergence). By Banach's fixed-point theorem, there exists a unique $V^* \in \mathcal{B}$ such that $V^* = \mathcal{T}V^*$; moreover, for any $V^0 \in \mathcal{B}$, the value iteration $V^{n+1} = \mathcal{T}V^n$ converges to V^* at a geometric rate no slower than ρ . $\square$

B.3 Interior FOC and the switching condition

Proof. Standing regularity and slackness. Throughout, we focus on interior states where the data-usage constraint $G(I_t, D_t) \leq D_t$ is slack; otherwise its Kuhn–Tucker multiplier enters the first-order condition in the standard way. We also work under the maintained local regularity on the compact on-policy region: V is continuously differentiable (or locally Lipschitz with measurable supergradient selections), and differentiation can be interchanged with conditional expectation by dominated convergence. Finally, Assumption A8 postulates (strong) concavity

of the one-dimensional objective in I on $[0, M_t]$ on the relevant region, where $M_t := \tau Y_t$. In the sequel, for notation convenience, we never distinguish V_D and p^D .

Interior FOC. On $(0, M_t)$, $I_t \mapsto G(I_t, D_t)$ is C^1 by construction. For any interior maximizer $0 < I_t < M_t$,

$$0 = \frac{\partial}{\partial I_t} \rho \mathbb{E}[V(S_{t+1}) \mid S_t] = \rho \mathbb{E}\left[V_K(S_{t+1}) \frac{\partial K_{t+1}}{\partial I_t} + V_D(S_{t+1}) \frac{\partial D_{t+1}}{\partial I_t} \mid S_t\right].$$

Using

$$\frac{\partial K_{t+1}}{\partial I_t} = w_t G_I(I_t, D_t), \quad \frac{\partial D_{t+1}}{\partial I_t} = -(1 - \delta_D) G_I(I_t, D_t) - 1,$$

and rearranging gives

$$\mathbb{E}\left[V_K(S_{t+1}) w_t G_I(I_t, D_t) \mid S_t\right] = \mathbb{E}\left[V_D(S_{t+1}) (1 + (1 - \delta_D) G_I(I_t, D_t)) \mid S_t\right]. \quad (34)$$

Boundary analysis and switching. Define the left directional derivative at the upper budget corner:

$$\Delta Q(M_t; S_t) := \left. \frac{\partial}{\partial I_t} \rho \mathbb{E}[V(S_{t+1}) \mid S_t] \right|_{I_t = M_t^-}.$$

For each state S_t , the feasible set for I_t is the compact interval $[0, M_t]$, and the map $I_t \mapsto \rho \mathbb{E}[V(S_{t+1}) \mid S_t]$ is continuous on $[0, M_t]$ and continuously differentiable on $(0, M_t)$. By Assumption A8, the one-dimensional objective is concave in I on $[0, M_t]$ on the relevant region. Hence the usual one-dimensional optimality conditions imply:

- if $\Delta Q(M_t; S_t) > 0$, the objective is locally increasing as $I_t \uparrow M_t$, so any maximizer is at the corner $I_t = M_t$ (and thus $J_t = 0$);
- if $\Delta Q(M_t; S_t) < 0$, the corner $I_t = M_t$ cannot be optimal, so any maximizer lies in the interior $0 < I_t < M_t$ and therefore satisfies (34);
- if $\Delta Q(M_t; S_t) = 0$, the corner point is included in the argmax; such states lie on the switching frontier between corner and interior regimes.

By continuity of the objects involved, the zero-level set $\{S_t : \Delta Q(M_t; S_t) = 0\}$ forms a frontier that separates the interior region $\{\Delta Q < 0\}$ from the corner region $\{\Delta Q > 0\}$. $\square$

B.4 Monotone comparative statics of the optimal policy

Fix a pre-decision state $S = (K, D)$ and consider the one-step objective

$$\Psi(I; S) := (1 - \tau)Y(S) + \rho \mathbb{E}[V(S') \mid S, I], \quad I \in [0, M(S)], \quad M(S) := \tau Y(S),$$

where the successor $S' = (K', D')$ follows the law of motion in Section 3. On the interior $(0, M(S))$, write Ψ_I for the partial derivative with respect to I . Throughout, we work under the maintained compactness and boundedness conditions in the invariant- A regime (bounded support of w_t and compact state/control sets), so that differentiation under the conditional expectation is justified by dominated convergence (equivalently, standard envelope/interchange results such as Benveniste–Scheinkman or Milgrom–Segal apply on the relevant compact on-policy region).

Standing primitives and notation. In this subsection, the innovation kernel $G = G(I, D)$ is C^2 with $G_I \geq 0$ and $G_{ID} \geq 0$ (data–innovation complementarity). Let $w_t > 0$ a.s. denote the i.i.d. implementation shock with $\mathbb{E}[w_t] = \bar{w} \in (0, \infty)$ and bounded support $0 < \underline{w} \leq w_t \leq \bar{w} < \infty$.

Lemma 3 (Monotonicity of V). *Suppose (i) the one-period reward $(1-\tau)Y(S)$ is nondecreasing in each state component (in the invariant- A regime, $Y_K \geq 0$ and $Y_D \equiv 0$); and (ii) for any fixed control, the transition is order-preserving:*

$$\frac{\partial K'}{\partial K} = 1 - \delta > 0, \quad \frac{\partial K'}{\partial D} = w_t G_D(I, D) \geq 0, \quad \frac{\partial D'}{\partial D} = (1 - \delta_D)(1 - G_D(I, D)) \geq 0,$$

(e.g., $0 \leq G_D \leq 1$ on the admissible set, or any small-gain condition ensuring $\partial D'/\partial D \geq 0$). Then the Bellman operator is isotone on $C_b(\mathcal{S}_0)$, and the fixed point V is nondecreasing in K and D . In particular, $V_K, V_D \geq 0$ wherever classical derivatives exist (otherwise interpret these as monotone superdifferential selections).

Assumption 1 (Marginal net-dominance (on-policy)). Along the optimal policy, for all feasible (S, I) in the on-policy region,

$$\mathbb{E}[w_t V_K(S') - (1 - \delta_D) V_D(S') \mid S, I] \geq 0.$$

Comment. Under an interior FOC $\Psi_I = 0$,

$$\mathbb{E}[w_t V_K(S') - (1 - \delta_D) V_D(S') \mid S, I] = \frac{1}{G_I(I, D)} \mathbb{E}[V_D(S') \mid S, I] \geq 0,$$

provided $V_D \geq 0$ (Lemma 3) and $G_I > 0$. Here $G_I(I, D)$ depends only on the current state and control, hence can be taken outside the expectation.

Assumption 2 (Curvature and state cross-effects). Along optimal paths, V is twice differentiable with

$$V_{KK} \leq 0, \quad V_{DD} \leq 0, \quad V_{KD} \geq 0 \quad (\text{increasing differences}).$$

Moreover, $Y_K \geq 0$.

Lemma 4 (Cross-partials of the marginal objective). *On the interior $(0, M(S))$,*

$$\Psi_I(I; S) = \rho \mathbb{E}[w_t V_K(S') G_I(I, D) - \{1 + (1 - \delta_D) G_I(I, D)\} V_D(S') \mid S, I],$$

where $S' = (K', D')$ follows the invariant- A transition and $S = (K, D)$. Then:

(a) **Decomposition in D and sufficient conditions.** Assume G is C^2 and V is twice differentiable on the relevant (on-policy) region. For any interior (S, I) ,

$$\partial_D \Psi_I(I; S) = \underbrace{\rho \mathbb{E}[(w_t V_K(S') - (1 - \delta_D) V_D(S')) G_{ID}(I, D) \mid S, I]}_{\text{direct effect } \text{DE}_D(I; S)} + \rho \text{RE}_D(I; S), \quad (35)$$

$$\text{RE}_D(I; S) := \mathbb{E}[w_t G_I(I, D) \partial_D V_K(S') - \{1 + (1 - \delta_D) G_I(I, D)\} \partial_D V_D(S') \mid S, I]. \quad (36)$$

Consequently, if $G_{ID} \geq 0$ and Assumption 1 holds, then $\text{DE}_D(I; S) \geq 0$.

Moreover, in the invariant-A regime one has $M_D(S) = 0$ and

$$\frac{\partial K'}{\partial D} = w_t G_D(I, D), \quad \frac{\partial D'}{\partial D} = (1 - \delta_D)(1 - G_D(I, D)).$$

Under Assumption 2 and the order-preserving bound $0 \leq G_D(I, D) \leq \bar{g} < 1$ on the relevant set, the remainder term can be written as

$$\text{RE}_D(I; S) = \mathbb{E}[(1 - \delta_D)(1 - G_D) B_D(S', I, w_t) - w_t G_D B_K(S', I, w_t) \mid S, I], \quad (37)$$

where (suppressing arguments of G_I and G_D for brevity)

$$B_D(S', I, w_t) := w_t G_I V_{KD}(S') + \{1 + (1 - \delta_D)G_I\}(-V_{DD}(S')) \geq 0,$$

$$B_K(S', I, w_t) := w_t G_I(-V_{KK}(S')) + \{1 + (1 - \delta_D)G_I\}V_{KD}(S') \geq 0.$$

A sufficient (small-gain) dominance condition ensuring $\text{RE}_D(I; S) \geq 0$ uniformly on the relevant set is

$$(1 - \delta_D)(1 - \bar{g}) \underline{B}_D \geq \bar{w} \bar{g} \bar{B}_K, \quad (38)$$

where $\bar{w} := \sup w_t < \infty$ and

$$\underline{B}_D := \inf B_D(S', I, w_t), \quad \bar{B}_K := \sup B_K(S', I, w_t),$$

with the infimum/supremum taken over (S, I, w_t) in the admissible (on-policy) region. Under (38), one has $\partial_D \Psi_I(I; S) \geq 0$ because both terms in (35) are then nonnegative.

- (b) **A sufficient condition for $\partial_K \Psi_I \leq 0$.** Assume Assumption 2 holds and that Y enters the transition only through the residual budget $J = M(S) - I$ with $M_K = \tau Y_K$. Then, for any interior (S, I) ,

$$\begin{aligned} \partial_K \Psi_I(I; K, D) = \rho \mathbb{E} \Big[& w_t G_I(I, D) ((1 - \delta) V_{KK}(S') + \tau Y_K(K, D) V_{KD}(S')) \\ & - \{1 + (1 - \delta_D)G_I(I, D)\} ((1 - \delta) V_{DK}(S') + \tau Y_K(K, D) V_{DD}(S')) \mid S, I \Big]. \end{aligned} \quad (39)$$

In particular, $\partial_K \Psi_I(I; K, D) \leq 0$ is ensured if the budget-slope term τY_K is not too large relative to the negative curvature terms. One sufficient small-gain condition is

$$\tau Y_K \cdot \sup \left\{ |w_t G_I V_{KD}| + |\{1 + (1 - \delta_D)G_I\}(-V_{DD})| \right\} \leq (1 - \delta) \cdot \inf \left\{ \{1 + (1 - \delta_D)G_I\} V_{DK} - w_t G_I(-V_{KK}) \right\}, \quad (40)$$

where the supremum/infimum are taken over the same admissible region as above.

Note. $G_I = G_I(I, D)$ is independent of K , so $\partial_K G_I \equiv 0$.

Assumption 3 (Strict concavity in I and interiority). There exists $\mu > 0$ such that $\Psi_{II}(I; S) \leq -\mu < 0$ on the interior, and the optimal policy satisfies $I^*(S) \in (0, M(S))$ for all S in a compact $\mathcal{S}_0$ (cf. Assumption A8).

Proposition 8 (Monotone comparative statics via the IFT). *Under Assumptions 1–3, $G_{ID} \geq 0$, and the sufficient conditions (38) and (40), the unique interior maximizer $I^*(K, D)$ solves $\Psi_I(I^*(K, D); K, D) = 0$ and is differentiable with*

$$\frac{\partial I^*}{\partial D} = -\frac{\partial_D \Psi_I}{\Psi_{II}} \geq 0, \quad \frac{\partial I^*}{\partial K} = -\frac{\partial_K \Psi_I}{\Psi_{II}} \leq 0,$$

whenever derivatives exist. Consequently, $J^*(S) = M(S) - I^*(S)$ satisfies

$$\partial_D J^* \leq 0, \quad \partial_K J^* \geq 0,$$

because $M_K = \tau Y_K \geq 0$ and $M_D = 0$ in the invariant- A regime.

Corollary (weak monotonicity with corners). If $\Psi_{II} \leq 0$ and the small-gain bound (40) holds, then there exists a measurable selection of maximizers $S \mapsto I^\dagger(S)$ that is nondecreasing in D and nonincreasing in K (Milgrom and Shannon, 1994).

B.5 Technical assumptions for steady-state analysis

A8 (Interior optimum and strong concavity relative to the budget). Let $M(S) := \tau Y(S)$ and $\mathcal{S}_0 \subset \mathbb{R}_+^2$ be a compact state set with strictly positive lower bounds. There exist $\mu > 0$ and $\varepsilon \in (0, \inf_{S \in \mathcal{S}_0} M(S))$ such that for every $S \in \mathcal{S}_0$ the continuation payoff

$$Q(I; S) := \rho \mathbb{E}[V(S') \mid S, I], \quad I \in [0, M(S)],$$

satisfies: (i) $Q \in C^2((0, M(S)))$ with $Q_{II}(I; S) \leq -\mu < 0$ on the interior; (ii) the unique maximizer $I(S)$ is interior and uniformly separated from the state-dependent boundary:

$$I(S) \in [\varepsilon, M(S) - \varepsilon] \quad \forall S \in \mathcal{S}_0.$$

Consequently $J(S) := M(S) - I(S) \geq \varepsilon$ for all $S \in \mathcal{S}_0$. *Note.* With $Q(I; S) := \rho \mathbb{E}[V(S') \mid S, I]$ and $\Phi(I; S) := \mathbb{E}[\cdot \mid S, I]$ (the FOC kernel), one has $Q_I = \rho \Phi$ and $Q_{II} = \rho \partial_I \Phi$.

A9 (Shock regularity, independence, and small-gain bounds). The implementation shock (w_t) is i.i.d., independent of (S, I) , with $0 < \underline{w} \leq w_t \leq \bar{w} < \infty$ and $\bar{w} := \mathbb{E}[w_t] \in (0, \infty)$. There exist finite constants $L_G, L_I, L_J > 0$ such that:

1. $G(I, D)$ is Lipschitz on $\{(I, D) : I \in [\varepsilon, M(S)], S \in \mathcal{S}_0\}$ with constant L_G ;
2. $I(\cdot)$ and $J(\cdot) = M(\cdot) - I(\cdot)$ are Lipschitz on $\mathcal{S}_0$ with constants L_I, L_J ;
3. (*small-gain*) The constants satisfy $(1 - \delta) + \bar{w} L_G L_I < 1$ and $(1 - \delta_D)(1 + L_G) < 1$, and the cross-terms are small enough such that the spectral radius of the effective dependency matrix is less than 1.

B.6 Proof of Proposition 2

Proof of Proposition 2. We work with the *expected-drift (certainty-equivalent) state mapping* $\bar{F}(S) := \mathbb{E}[S_{t+1} \mid S_t = S]$ induced by the optimal feedback policy $(I(S), J(S))$ on the compact

on-policy region $\mathcal{S}_0$. Under the maintained independence and bounded-support assumptions for w_t , the conditional mean takes the explicit form below.

Step 1: Lipschitz policy via the implicit function theorem. Fix $S = (K, D) \in \mathcal{S}_0$ and write the interior FOC

$$\Phi(I; S) := \mathbb{E}\left[V_K(S') w_t G_I(I, D) - V_D(S')(1 + (1 - \delta_D)G_I(I, D)) \mid S\right] = 0,$$

where S' is the successor state. By A8, $Q_{II}(I; S) \leq -\mu < 0$, hence $\partial\Phi/\partial I = Q_{II}/\rho \leq -\mu/\rho < 0$, and the implicit function theorem delivers a unique C^1 selection $I(\cdot)$ on $\mathcal{S}_0$. Compactness then yields a global Lipschitz bound. Since $J(S) = M(S) - I(S)$ and $M(\cdot)$ is smooth on $\mathcal{S}_0$, $J(\cdot)$ is Lipschitz as well.

Step 2: Lipschitz property of G on the relevant domain. Because $I(S) \in [\varepsilon, M(S) - \varepsilon]$ uniformly and $\mathcal{S}_0$ is compact with positive lower bounds, G has bounded first derivatives on the compact set $\{(I, D) : I \in [\varepsilon, M(S)], S \in \mathcal{S}_0\}$; hence there exists $L_G < \infty$ with

$$|G(I_1, D_1) - G(I_2, D_2)| \leq L_G(|I_1 - I_2| + |D_1 - D_2|).$$

Step 3: Expected-state mapping and block Lipschitz bounds. Because w_t is i.i.d. and independent of (S, I) with $0 < \underline{w} \leq w_t \leq \bar{w} < \infty$,

$$\mathbb{E}[w_t G(I(S), D) \mid S] = \bar{w} G(I(S), D), \quad \bar{w} := \mathbb{E}[w_t].$$

Throughout we take $\|S\| := \max\{|K|, |D|\}$. The expected-drift mapping induced by the optimal feedback (I, J) is

$$\bar{F}(S) = \begin{pmatrix} (1 - \delta) K + \bar{w} G(I(S), D) \\ (1 - \delta_D)(D - G(I(S), D)) + J(S) \end{pmatrix}.$$

Using the Lipschitz bounds for G , I , and J , there exist finite constants such that for all $S_1, S_2 \in \mathcal{S}_0$,

$$\begin{aligned} |\bar{K}(S_1) - \bar{K}(S_2)| &\leq ((1 - \delta) + \bar{w} L_G L_I) \|S_1 - S_2\| + \bar{w} L_G |D_1 - D_2|, \\ |\bar{D}(S_1) - \bar{D}(S_2)| &\leq (1 - \delta_D)(1 + L_G) |D_1 - D_2| + ((1 - \delta_D)L_G L_I + L_J) \|S_1 - S_2\|. \end{aligned}$$

Step 4: Contraction under a weighted sup norm. For $\eta > 0$, let $\|S\|_\eta := \max\{|K|, \eta|D|\}$. Then

$$\|\bar{F}(S_1) - \bar{F}(S_2)\|_\eta \leq \max\{c_1(\eta), c_2(\eta)\} \|S_1 - S_2\|_\eta,$$

with

$$c_1(\eta) = (1 - \delta) + \bar{w} L_G L_I + \frac{\bar{w} L_G}{\eta}, \quad c_2(\eta) = (1 - \delta_D)(1 + L_G) + \eta((1 - \delta_D)L_G L_I + L_J).$$

Note $c_1(\eta)$ is decreasing in η and $c_2(\eta)$ is increasing in η . By Assumption A9(3), the asymptotic values satisfy $\lim_{\eta \rightarrow \infty} c_1(\eta) < 1$ and $c_2(0) < 1$. Provided the coupling coefficients are sufficiently small (as ensured by the small-gain condition), there exists $\eta^* > 0$ such that

$\kappa := \max\{c_1(\eta^*), c_2(\eta^*)\} < 1$. Hence $\bar{F}$ is a contraction on $(\mathcal{S}_0, \|\cdot\|_{\eta^*})$.

Step 5: Unique fixed point and geometric convergence (mean dynamics). By Banach's fixed-point theorem, $\bar{F}$ has a unique fixed point $(\bar{K}, \bar{D}) \in \mathcal{S}_0$, and

$$\|S_t - (\bar{K}, \bar{D})\|_{\eta^*} \leq \kappa^t \|S_0 - (\bar{K}, \bar{D})\|_{\eta^*} \quad \forall t \geq 0,$$

for the iterates of the certainty-equivalent mean dynamics $S_{t+1} = \bar{F}(S_t)$. Interiority holds because A8 implies $I(\bar{S}) \in [\varepsilon, M(\bar{S}) - \varepsilon]$ and hence $J(\bar{S}) = M(\bar{S}) - I(\bar{S}) \geq \varepsilon > 0$, while $\bar{K}, \bar{D} > 0$ by construction of $\mathcal{S}_0$. $\square$

B.7 Extension to an exogenous technology index

Let the aggregate index (Z_t) be exogenous Markov on a compact $\mathcal{Z}$ (or satisfy a drift condition admitting a weighted sup norm). Working on $\mathcal{S} := \mathcal{S}_0 \times \mathcal{Z}$ and defining a weight $w(S, z) = w_0(S)\omega(z)$ with $\omega \geq 1$ such that

$$\sup_{z \in \mathcal{Z}} \frac{\mathbb{E}[\omega(Z_{t+1}) \mid z]}{\omega(z)} \leq \alpha_Z < \frac{1}{\rho},$$

the assumptions of Theorem 4 in Appendix C are satisfied for the enlarged state. Consequently the Bellman operator $\mathcal{T}$ remains a contraction on the weighted space $(C_b^w(\mathcal{S}), \|\cdot\|_w)$ with modulus at most $\rho\alpha_Z < 1$, and existence/uniqueness together with value-iteration convergence continue to hold with stochastic Z .

C Appendix: A Contraction Theorem for the Bellman Operator with Stochastic Z

This appendix gives an operations-research-style proof of existence, uniqueness, and geometric convergence for the Bellman equation with a possibly unbounded state (in particular, to accommodate a stochastic technology factor Z). We cover both the bounded-state (sup-norm) and the unbounded-state (weighted-norm) settings, and we treat policy-induced and maximization operators.

C.1 Setting, notation, and function spaces

Let $\mathcal{S}$ be a Polish space endowed with its Borel σ -algebra $\mathcal{B}(\mathcal{S})$. Let $\mathcal{A}$ be a Polish action space with its Borel σ -algebra, and let $\mathcal{A}(\cdot)$ be a nonempty feasible-action correspondence.

$$\text{Gr}(\mathcal{A}) := \{(s, a) : a \in \mathcal{A}(s)\} \subset \mathcal{S} \times \mathcal{A}.$$

Per-period payoffs are $\Pi : \text{Gr}(\mathcal{A}) \rightarrow \mathbb{R}$. A controlled Markov kernel $P(\cdot \mid s, a)$ governs transitions $S' \sim P(\cdot \mid s, a)$. Fix a discount factor $\rho \in (0, 1)$.

Operators. For $V : \mathcal{S} \rightarrow \mathbb{R}$ define

$$(\mathcal{P}_a V)(s) := \mathbb{E}[V(S') \mid S = s, a] = \int_{\mathcal{S}} V(s') P(ds' \mid s, a).$$

We consider

$$\begin{aligned} (\mathcal{T}_{\text{pol}} V)(s) &:= \Pi(s, \pi(s)) + \rho(\mathcal{P}_{\pi(s)} V)(s) \quad (\text{policy-induced}), \\ (\mathcal{T} V)(s) &:= \sup_{a \in \mathcal{A}(s)} \left\{ \Pi(s, a) + \rho(\mathcal{P}_a V)(s) \right\} \quad (\text{maximization}). \end{aligned}$$

Function spaces. Let $C_b(\mathcal{S})$ be bounded continuous functions with norm $\|V\|_{\infty} := \sup_{s \in \mathcal{S}} |V(s)|$.

For the unbounded case, let $w : \mathcal{S} \rightarrow [1, \infty)$ be continuous and define

$$C_b^w(\mathcal{S}) := \left\{ V : \|V\|_w := \sup_{s \in \mathcal{S}} \frac{|V(s)|}{w(s)} < \infty, \quad V/w \in C_b(\mathcal{S}) \right\}.$$

The map $J : C_b^w(\mathcal{S}) \rightarrow C_b(\mathcal{S})$, $J(V) := V/w$, is an isometric isomorphism ($\|V\|_w = \|J(V)\|_{\infty}$). Moreover, for any $f \in C_b(\mathcal{S})$, $J^{-1}(f) := fw$ lies in $C_b^w(\mathcal{S})$ and $J(J^{-1}(f)) = f$, so J is onto. Hence $(C_b^w(\mathcal{S}), \|\cdot\|_w)$ is a Banach space.

In applications with an exogenous technology factor Z_t , we will often take a separable weight $w(s, z) = w_0(s)\omega(z)$; in that case the drift bound in Assumption (W1) below will be written with a constant α_Z (see the extension remark in Appendix B).

All expectations are taken under the probability measure relevant for the Bellman recursion; we absorb the measure into the controlled kernel $P(\cdot \mid s, a)$.

C.2 Assumptions

We separate the bounded and weighted cases.

(B1) (Controlled Feller, bounded state). $\mathcal{S}$ is compact; $\Pi(\cdot, \cdot)$ is continuous and bounded on $\text{Gr}(\mathcal{A})$; and the kernel is *controlled Feller*: for all $f \in C_b(\mathcal{S})$, the map $(s, a) \mapsto (\mathcal{P}_a f)(s)$ is continuous on $\text{Gr}(\mathcal{A})$.

(B2) (Feasible actions for maximization). For each s , $\mathcal{A}(s)$ is nonempty and compact; $s \mapsto \mathcal{A}(s)$ is upper hemicontinuous with measurable graph. By Berge's maximum theorem, if $(s, a) \mapsto \Pi(s, a) + \rho(\mathcal{P}_a V)(s)$ is continuous, the supremum is attained and continuous in s , and the argmax admits a measurable selector.

(Pol) (Policy regularity). For a fixed stationary policy π , both $s \mapsto \Pi(s, \pi(s))$ and $s \mapsto (\mathcal{P}_{\pi(s)} f)(s)$ are continuous for every $f \in C_b(\mathcal{S})$ (e.g., this holds if π is continuous and (B1) holds).

(W1) (Weighted drift). There exists a continuous weight $w \geq 1$ and a constant $\alpha \in (0, 1/\rho)$ such that

$$\sup_{a \in \mathcal{A}(s)} \frac{\mathbb{E}[w(S') \mid s, a]}{w(s)} \leq \alpha \quad \text{for all } s \in \mathcal{S}.$$

In the exogenous-technology application in Appendix B, where the unbounded dimension is the Markov factor Z_t and $w(s, z) = w_0(s)\omega(z)$, we denote this drift-ratio bound by α_Z ;

in other words, α_Z there plays exactly the role of α here, and the contraction modulus in Theorem 4(b) becomes $c = \rho\alpha_Z$.

For the GBM specification of Z_t used in Section 4.1 (under $\mathbb{Q}$), a convenient choice is $\omega(z) = 1 + z^\eta$ for $\eta > 0$. If $Z' = Z \exp((\mu_Z - \frac{1}{2}\sigma_Z^2)\Delta + \sigma_Z\sqrt{\Delta}\varepsilon)$, then $\mathbb{E}[Z'^\eta \mid Z = z] = z^\eta \exp(\eta\mu_Z\Delta + \frac{1}{2}\eta(\eta-1)\sigma_Z^2\Delta)$, so

$$\sup_{z>0} \frac{\mathbb{E}[\omega(Z') \mid z]}{\omega(z)} \leq \max \left\{ 1, \exp(\eta\mu_Z\Delta + \frac{1}{2}\eta(\eta-1)\sigma_Z^2\Delta) \right\} =: \alpha_Z.$$

Hence (W1) holds provided $\rho\alpha_Z < 1$.

(W2) (Weighted controlled Feller and payoff regularity). $\Pi(\cdot, \cdot)/w(\cdot) \in C_b(\text{Gr}(\mathcal{A}))$ and, for all $V \in C_b^w(\mathcal{S})$, the map $(s, a) \mapsto (\mathcal{P}_a V)(s)/w(s)$ belongs to $C_b(\text{Gr}(\mathcal{A}))$.

(W3) (Weighted policy regularity). For a fixed stationary policy π , the functions $s \mapsto \Pi(s, \pi(s))/w(s)$ and $s \mapsto (\mathcal{P}_{\pi(s)} V)(s)/w(s)$ belong to $C_b(\mathcal{S})$ for all $V \in C_b^w(\mathcal{S})$ (e.g., this holds if π is continuous and (W2) holds).

C.3 Main theorem and corollaries

Theorem 4 (Contraction and well-posedness). *Fix $\rho \in (0, 1)$.*

- (a) Bounded state (sup-norm). *Under (B1) and (Pol), the policy-induced operator $\mathcal{T}_{\text{pol}}$ maps $C_b(\mathcal{S})$ into itself and is a ρ -contraction on $(C_b(\mathcal{S}), \|\cdot\|_\infty)$. If (B2) also holds, the maximization operator $\mathcal{T}$ maps $C_b(\mathcal{S})$ into itself and is likewise a ρ -contraction.*
- (b) Unbounded state (weighted norm). *Under (W1) and (W3), the policy-induced operator $\mathcal{T}_{\text{pol}}$ maps $C_b^w(\mathcal{S})$ into itself and is a $\rho\alpha$ -contraction on $(C_b^w(\mathcal{S}), \|\cdot\|_w)$. If (W1)–(W2) and (B2) hold, the maximization operator $\mathcal{T}$ maps $C_b^w(\mathcal{S})$ into itself and is a $\rho\alpha$ -contraction.*
- (c) Existence, uniqueness, and value iteration. *In each case, a unique fixed point V^* exists in the respective space, and for any V^0 ,*

$$\|\mathcal{T}^n V^0 - V^*\| \leq c^n \|V^0 - V^*\|, \quad c = \rho \text{ in (a), } c = \rho\alpha \text{ in (b),}$$

where $\|\cdot\|$ denotes $\|\cdot\|_\infty$ in (a) and $\|\cdot\|_w$ in (b).

Corollary 1 (Error bounds and stopping rules). *Let $c \in (0, 1)$ be the contraction modulus in Theorem 4. Then for any $n \geq 1$,*

$$\|V^n - V^*\| \leq \frac{c}{1-c} \|V^n - V^{n-1}\|.$$

Thus a practical stopping rule is $\|V^n - V^{n-1}\| \leq \varepsilon(1-c)/c$ to ensure $\|V^n - V^\| \leq \varepsilon$.*

C.4 Proofs

Lemma 5 (Basic inequality). *For any real families $\{x_a\}$ and $\{y_a\}$ indexed by a , $\sup_a x_a - \sup_a y_a \leq \sup_a (x_a - y_a)$.*

Proof. For each a , $x_a \leq (x_a - y_a) + \sup_b y_b$. Taking sup over a yields $\sup_a x_a \leq \sup_a (x_a - y_a) + \sup_b y_b$, i.e. $\sup_a x_a - \sup_a y_a \leq \sup_a (x_a - y_a)$. $\square$

Lemma 6 (Completeness of the weighted space). $(C_b^w(\mathcal{S}), \|\cdot\|_w)$ is complete.

Proof. The map $J(V) := V/w$ is an isometric isomorphism from $C_b^w(\mathcal{S})$ onto $C_b(\mathcal{S})$. Since $C_b(\mathcal{S})$ is complete, so is $C_b^w(\mathcal{S})$. $\square$

Proof of Theorem 4. We treat policy-induced and maximization operators in parallel, separating mapping properties from contraction.

(a) Bounded state, policy-induced operator. *Mapping.* By (Pol), $\Pi(\cdot, \pi(\cdot)) \in C_b(\mathcal{S})$ and $s \mapsto (\mathcal{P}_{\pi(s)}V)(s) \in C_b(\mathcal{S})$ for all $V \in C_b(\mathcal{S})$, so $\mathcal{T}_{\text{pol}}V \in C_b(\mathcal{S})$. *Contraction.* For $U, V \in C_b(\mathcal{S})$,

$$\|\mathcal{T}_{\text{pol}}V - \mathcal{T}_{\text{pol}}U\|_\infty = \sup_{s \in \mathcal{S}} \rho \left| (\mathcal{P}_{\pi(s)}(V - U))(s) \right| \leq \rho \|V - U\|_\infty.$$

(a) Bounded state, maximization operator. *Mapping.* Fix $V \in C_b(\mathcal{S})$. By (B1), $(s, a) \mapsto (\mathcal{P}_aV)(s)$ is continuous, hence $(s, a) \mapsto \Pi(s, a) + \rho(\mathcal{P}_aV)(s)$ is continuous and bounded. By (B2) and Berge's maximum theorem, $s \mapsto (\mathcal{T}V)(s)$ is continuous and bounded, i.e., $\mathcal{T} : C_b(\mathcal{S}) \rightarrow C_b(\mathcal{S})$. *Contraction.* Using Lemma 5 and linearity of $\mathcal{P}_a$,

$$\begin{aligned} \|\mathcal{T}V - \mathcal{T}U\|_\infty &= \sup_{s \in \mathcal{S}} \left| \sup_{a \in \mathcal{A}(s)} \{\Pi(s, a) + \rho(\mathcal{P}_aV)(s)\} - \sup_{a \in \mathcal{A}(s)} \{\Pi(s, a) + \rho(\mathcal{P}_aU)(s)\} \right| \\ &\leq \sup_{s \in \mathcal{S}} \sup_{a \in \mathcal{A}(s)} \rho |(\mathcal{P}_a(V - U))(s)| \leq \rho \|V - U\|_\infty. \end{aligned}$$

(b) Weighted state, policy-induced operator. *Mapping.* Fix $V \in C_b^w(\mathcal{S})$. By $|V(S')| \leq \|V\|_w w(S')$ a.s. and (W1),

$$\frac{|(\mathcal{T}_{\text{pol}}V)(s)|}{w(s)} \leq \frac{|\Pi(s, \pi(s))|}{w(s)} + \rho \frac{\mathbb{E}[|V(S')| \mid s, \pi(s)]}{w(s)} \leq \frac{|\Pi(s, \pi(s))|}{w(s)} + \rho\alpha \|V\|_w.$$

Taking sup over s and using (W3) gives

$$\|\mathcal{T}_{\text{pol}}V\|_w \leq \|\Pi(\cdot, \pi(\cdot))/w(\cdot)\|_\infty + \rho\alpha \|V\|_w < \infty,$$

and $\mathcal{T}_{\text{pol}}V \in C_b^w(\mathcal{S})$. *Contraction.* For $U, V \in C_b^w(\mathcal{S})$,

$$\frac{|(\mathcal{T}_{\text{pol}}V - \mathcal{T}_{\text{pol}}U)(s)|}{w(s)} = \rho \frac{|(\mathcal{P}_{\pi(s)}(V - U))(s)|}{w(s)} \leq \rho \|V - U\|_w \frac{\mathbb{E}[w(S') \mid s, \pi(s)]}{w(s)} \leq \rho\alpha \|V - U\|_w.$$

(b) Weighted state, maximization operator. *Mapping.* Fix $V \in C_b^w(\mathcal{S})$. By (W2), $(s, a) \mapsto \Pi(s, a)/w(s) + \rho(\mathcal{P}_aV)(s)/w(s)$ is continuous and bounded on $\text{Gr}(\mathcal{A})$. By (B2) and Berge's maximum theorem, $s \mapsto (\mathcal{T}V)(s)/w(s)$ is continuous and bounded, hence $\mathcal{T} : C_b^w(\mathcal{S}) \rightarrow C_b^w(\mathcal{S})$. *Contraction.* By Lemma 5 and (W1),

$$\begin{aligned} \frac{|(\mathcal{T}V - \mathcal{T}U)(s)|}{w(s)} &\leq \sup_{a \in \mathcal{A}(s)} \rho \frac{|(\mathcal{P}_a(V - U))(s)|}{w(s)} \leq \rho \|V - U\|_w \sup_{a \in \mathcal{A}(s)} \frac{\mathbb{E}[w(S') \mid s, a]}{w(s)} \\ &\leq \rho\alpha \|V - U\|_w. \end{aligned}$$

In all cases the relevant function space is complete (Lemma 6 for weighted; standard for C_b), so Banach's fixed-point theorem yields existence, uniqueness, and geometric convergence. The error bound in Corollary 1 follows from the standard contraction argument. $\square$

C.5 Remarks on measurability, selectors, and policies

When $\mathcal{T}$ is a maximization operator, (B2)/(W2) imply the argmax correspondence is nonempty, compact-valued, and upper hemicontinuous; together with a measurable graph this ensures a measurable selector

$$\pi^*(s) \in_{a \in \mathcal{A}(s)} \{\Pi(s, a) + \rho(\mathcal{P}_a V^*)(s)\},$$

i.e., an optimal stationary policy. If maximizers are unique (e.g., strict concavity in a), the selector is unique and continuous under the same regularity.

Remark 6 (Endogenous A : no new difficulties). The well-posedness result with endogenous firm-specific technology A is an immediate corollary of Theorem 4 applied to the enlarged state (K, D, A) (or (K, D, A, Z)). The only modifications relative to the constant- A case are: (i) the transition for A enters the state kernel, which preserves the Feller/ w -Feller property; (ii) the constrained envelope for $p^A = V_A$ adds the budget-relaxation term $\lambda_t \tau Y_A$ because the constraint $I + J + R = \tau Y(K, A)$ depends on A ; and (iii) cross-partial of $G(I, D, A)$ (e.g. G_A) appear in the spillover term but do not affect the contraction modulus. Hence, under the same regularity conditions, no additional contraction proof is needed: the endogenous- A case is covered by Theorem 4 applied to the enlarged state (K, D, A) together with the envelope identities.

D Appendix: Proofs for Section 4

All derivatives and conditional expectations are evaluated at the optimal feedback $J_t^*(S_t)$; hence we suppress conditioning on J_t and write $\mathbb{E}_t^\mathbb{Q}[\cdot \mid S_t]$. With the single-control formulation $I_t = M_t - J_t$, the Bellman equation is

$$V(S_t) = \max_{0 \leq J_t \leq M_t} \left\{ (1 - \tau)Y_t - \frac{\kappa}{2}J_t^2 + \rho \mathbb{E}_t^\mathbb{Q}[V(S_{t+1}) \mid S_t, J_t] \right\}, \quad S_t = (K_t, D_t, Z_t),$$

where $M_t = \tau Y_t$. In the invariant-technology benchmark maintained in the main text, $Y_t = Y(K_t, Z_t)$ does not depend directly on D_t , hence M_t is locally constant in D_t . *Remark.* Throughout this appendix we allow a generic implementation factor w_t in the transition for K_{t+1} . The main text specializes to $w_t \equiv 1$ to focus on aggregate technology risk; setting $w_t \equiv 1$ below recovers (15)–(16) exactly.

Regularity note. Throughout we work on the interior region where the admissibility cap is slack, i.e., $\tilde{G}_t = G_t$ (equivalently $G_t < D_t$ and, if applicable, $\gamma G_t < A_t$). On this region the transition is C^1 and the derivatives G_I, G_D used below coincide with the derivatives of $\tilde{G}$. At kink points where the cap binds, the same identities hold with directional derivatives/subgradients, and the boundary characterization in the last paragraph applies.

On the relevant (on-policy) region, assume: (i) V is continuously differentiable in (K, D) , (ii) the transition mapping $S_{t+1} = S_{t+1}(S_t, J_t, w_t)$ is continuously differentiable in (D_t, J_t) , and

(iii) boundedness of w_t and local boundedness of the relevant derivatives ensure an integrable dominating function. Then the Benveniste–Scheinkman envelope theorem (or Milgrom–Segal) and dominated convergence justify differentiation under the conditional expectation and allow one to ignore the derivative of the optimal control with respect to the state in the envelope formulas.

Shadow-price recursion. Since the flow payoff does not depend directly on D_t , the envelope formula yields

$$p_t^D = \partial_{D_t} V(S_t) = \rho \mathbb{E}_t^{\mathbb{Q}} [V_K(S_{t+1}) \partial_{D_t} K_{t+1} + V_D(S_{t+1}) \partial_{D_t} D_{t+1} \mid S_t, J_t].$$

Using

$$K_{t+1} = (1 - \delta)K_t + w_t G(I_t, D_t), \quad D_{t+1} = (1 - \delta_D)(D_t - G(I_t, D_t)) + J_t,$$

and noting that, holding J_t fixed, $I_t = M_t - J_t$ satisfies $\partial I_t / \partial D_t = 0$ (because $M_t = \tau Y_t$ is independent of D_t), we obtain

$$\partial_{D_t} K_{t+1} = w_t G_D(I_t, D_t), \quad \partial_{D_t} D_{t+1} = (1 - \delta_D)(1 - G_D(I_t, D_t)).$$

Hence

$$\begin{aligned} p_t^D &= \rho \mathbb{E}_t^{\mathbb{Q}} [V_K(S_{t+1}) w_t G_D(I_t, D_t) + V_D(S_{t+1})(1 - \delta_D)(1 - G_D(I_t, D_t)) \mid S_t] \\ &= \rho(1 - \delta_D) \mathbb{E}_t^{\mathbb{Q}} [V_D(S_{t+1}) \mid S_t] + \rho \mathbb{E}_t^{\mathbb{Q}} [(w_t V_K(S_{t+1}) - (1 - \delta_D) V_D(S_{t+1})) G_D(I_t, D_t) \mid S_t] \\ &= \rho(1 - \delta_D) \mathbb{E}_t^{\mathbb{Q}} [p_{t+1}^D \mid S_t] + \rho \mathbb{E}_t^{\mathbb{Q}} [(w_t V_K(S_{t+1}) - (1 - \delta_D) p_{t+1}^D) G_D(I_t, D_t) \mid S_t], \end{aligned}$$

which reduces to (15) in the main text when $w_t \equiv 1$.

Interior FOC for J_t . At an interior optimum $0 < J_t^* < M_t$, the first-order necessary condition is

$$0 = \frac{\partial}{\partial J_t} \left(-\frac{\kappa}{2} J_t^2 + \rho \mathbb{E}_t^{\mathbb{Q}} [V(S_{t+1}) \mid S_t, J_t] \right) \Big|_{J_t=J_t^*} = -\kappa J_t^* + \rho \mathbb{E}_t^{\mathbb{Q}} \left[V_K(S_{t+1}) \frac{\partial K_{t+1}}{\partial J_t} + V_D(S_{t+1}) \frac{\partial D_{t+1}}{\partial J_t} \mid S_t \right].$$

Since $I_t = M_t - J_t$, we have $\partial I_t / \partial J_t = -1$ and therefore

$$\frac{\partial K_{t+1}}{\partial J_t} = w_t G_I(I_t, D_t) \frac{\partial I_t}{\partial J_t} = -w_t G_I(I_t, D_t), \quad \frac{\partial D_{t+1}}{\partial J_t} = 1 - (1 - \delta_D) G_I(I_t, D_t) \frac{\partial I_t}{\partial J_t} = 1 + (1 - \delta_D) G_I(I_t, D_t).$$

Therefore

$$\begin{aligned} \kappa J_t^* &= \rho \mathbb{E}_t^{\mathbb{Q}} [V_D(S_{t+1})(1 + (1 - \delta_D) G_I(I_t, D_t)) - V_K(S_{t+1}) w_t G_I(I_t, D_t) \mid S_t] \\ &= \rho \mathbb{E}_t^{\mathbb{Q}} [p_{t+1}^D - (w_t V_K(S_{t+1}) - (1 - \delta_D) p_{t+1}^D) G_I(I_t, D_t) \mid S_t], \end{aligned}$$

which reduces to (16) in the main text when $w_t \equiv 1$.

Eliminating $\mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D]$. From the interior FOC,

$$\mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] = \frac{\kappa}{\rho} J_t^* + \mathbb{E}_t^{\mathbb{Q}}[(w_t V_K(S_{t+1}) - (1 - \delta_D) p_{t+1}^D) G_I(I_t, D_t) \mid S_t].$$

Plugging this identity into the recursion above yields

$$p_t^D = (1 - \delta_D) \kappa J_t^* + \rho \mathbb{E}_t^{\mathbb{Q}}[(w_t V_K(S_{t+1}) - (1 - \delta_D) p_{t+1}^D) (G_D(I_t, D_t) + (1 - \delta_D) G_I(I_t, D_t)) \mid S_t],$$

which reduces to (17) when $w_t \equiv 1$.

Boundary characterization (directional derivatives). Let

$$Q(J_t; S_t) := -\frac{\kappa}{2} J_t^2 + \rho \mathbb{E}_t^{\mathbb{Q}}[V(S_{t+1}) \mid S_t, J_t], \quad J_t \in [0, M_t].$$

If $Q(\cdot; S_t)$ is concave (e.g., under standard curvature conditions on V and G), then the following are necessary and sufficient: (i) if $J_t^* \in (0, M_t)$, then $Q_J(J_t^*; S_t) = 0$; (ii) if $J_t^* = 0$, then $Q_J(0^+; S_t) \leq 0$; and (iii) if $J_t^* = M_t$, then $Q_J(M_t^-; S_t) \geq 0$. In our notation,

$$Q_J(J_t; S_t) = -\kappa J_t + \rho \mathbb{E}_t^{\mathbb{Q}}[V_D(S_{t+1})(1 + (1 - \delta_D) G_I(I_t, D_t)) - V_K(S_{t+1}) w_t G_I(I_t, D_t) \mid S_t],$$

with $I_t = M_t - J_t$. When $G_I(I, D)$ is unbounded as $I \downarrow 0$, the left derivative at $J_t = M_t$ should be interpreted as an extended-real directional derivative; in particular, any sufficient condition that makes $Q_J(M_t^-; S_t) = -\infty$ rules out the upper corner and implies $J_t^* < M_t$.

E Appendix: Option Valuation, Snell Envelope, and Trigger Equivalence

Throughout this appendix we work on the cap-slack interior region where $\tilde{G}_t = G(I_t, D_t, A_t)$ and derivatives such as G_I, G_D are well-defined; when the cap binds, the same conditions are interpreted via directional derivatives/subgradients.

We work in discrete time $t = 0, 1, 2, \dots$ on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \geq 0}, \mathbb{Q})$ under the same risk-neutral measure used in the main text. Let p_t^D be an adapted shadow-price process and fix $\rho \in (0, 1)$.

Integrability and Snell-envelope regularity. To ensure the infinite-horizon Snell-envelope construction is well-posed, assume $\{\rho^t X_t\}_{t \geq 0}$ is of class (D), i.e., $\{\rho^\tau X_\tau : \tau \in \mathcal{T}_0\}$ is uniformly integrable. A sufficient primitive condition is, for example, $\sup_t \mathbb{E}^{\mathbb{Q}}[|p_t^D|^p] < \infty$ for some $p > 1$. Throughout we also maintain $\sup_t \mathbb{E}^{\mathbb{Q}}[|p_t^D|] < \infty$, so all conditional expectations below are well-defined.

E.1 External (spot) acquisition as a Snell envelope

Fix a per-unit spot price $c_D > 0$. We interpret external acquisition as an on-the-spot increase in the current data stock, so the incremental gain from buying one unit at date t is $p_t^D - c_D$.

Define the gain process

$$X_t^{\text{ext}} := p_t^D - c_D.$$

Define $\mathcal{T}_t := \{\tau \geq t : \tau \text{ is an } \{\mathcal{F}_s\}\text{-stopping time and } \tau < \infty \text{ } \mathbb{Q}\text{-a.s.}\}$. The Snell envelope is

$$\mathcal{S}_t^{\text{ext}} := \max_{\tau \in \mathcal{T}_t} \mathbb{E}_t^{\mathbb{Q}}[\rho^{(\tau-t)} X_\tau^{\text{ext}}].$$

which is the smallest ρ -supermartingale dominating X^{ext} and satisfies the backward recursion

$$\mathcal{S}_t^{\text{ext}} = \max \left\{ X_t^{\text{ext}}, \rho \mathbb{E}_t^{\mathbb{Q}}[\mathcal{S}_{t+1}^{\text{ext}}] \right\}.$$

Moreover, the stopping time

$$\tau^* = \inf \left\{ t \geq 0 : X_t^{\text{ext}} \geq \rho \mathbb{E}_t^{\mathbb{Q}}[\mathcal{S}_{t+1}^{\text{ext}}] \right\}$$

is optimal and attains the value $\mathcal{S}_0^{\text{ext}}$.

Under a Markov and monotonicity structure, the exercise region is an upper set and, in the scalar case, admits a threshold representation

$$p_t^D \geq c_{\text{ext}}(S_t).$$

In the myopic/zero-continuation limit (e.g., as $\rho \downarrow 0$), the continuation term vanishes, $\mathcal{S}_t^{\text{ext}} \rightarrow (p_t^D - c_D)_+$, and hence $c_{\text{ext}}(S_t) \downarrow c_D$, recovering the heuristic rule “invest when $p_t^D \geq c_D$ ”.

E.2 Internal (make-to-order) accumulation as an instantaneous trigger

In our baseline transition, one unit of J_t increases D_{t+1} one-for-one. Conditional on S_t , the per-period objective over J_t can be written (up to J -independent terms) as

$$\max_{0 \leq J_t \leq M_t} \left\{ -\frac{\kappa}{2} J_t^2 - c_t^{\text{adj}} J_t + \rho \mathbb{E}_t^{\mathbb{Q}}[V(S_{t+1}) \mid S_t, J_t] \right\},$$

where $\kappa > 0$ is the quadratic adjustment-cost parameter and $c_t^{\text{adj}} \geq 0$ collects any per-unit wedge beyond the shadow value of internal funds.¹¹

Let $p_t^D := \partial_D V^*(S_t)$ and recall that $I_t = M_t - J_t$. Define the (discounted) shadow cost of relaxing the flow budget through the implementation margin,

$$\lambda_t := \rho \mathbb{E}_t^{\mathbb{Q}}[(w_t V_K(S_{t+1}) - (1 - \delta_D) p_{t+1}^D) G_I(I_t, D_t) \mid S_t],$$

which reduces to the expression in the main text when $w_t \equiv 1$.

Define the instantaneous (discounted) net payoff index

$$X_t^{\text{int}} := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t - c_t^{\text{adj}}.$$

To streamline notation, we suppress the conditioning on J_t and understand $\mathbb{E}_t^{\mathbb{Q}}[\cdot \mid S_t]$ as evaluated at the optimal choice $J_t^*(S_t)$. Then the KKT conditions for the constrained choice

¹¹In the special case of a constant per-unit internal fee c_D^{int} attached to J_t , set $c_t^{\text{adj}} \equiv c_D^{\text{int}}$.

$J_t \in [0, M_t]$ yield the projection (intensity) form

$$J_t^* = \min \left\{ M_t, \frac{(X_t^{\text{int}})_+}{\kappa} \right\},$$

together with the corner characterizations

$$J_t^* = 0 \iff X_t^{\text{int}} \leq 0, \quad 0 < J_t^* < M_t \iff 0 < X_t^{\text{int}} < \kappa M_t, \quad J_t^* = M_t \iff X_t^{\text{int}} \geq \kappa M_t.$$

Accordingly, the lower exercise frontier can be written as the KKT surface

$$\Gamma_J = \{S_t : X_t^{\text{int}} = 0\} = \{S_t : \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] = \lambda_t + c_t^{\text{adj}}\},$$

and, away from the upper budget corner, one has

$$J_t^* > 0 \iff X_t^{\text{int}} > 0.$$

Economically, the internal margin uses the next-period marginal value $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t]$ as the underlying and the state-dependent effective strike $\lambda_t + c_t^{\text{adj}}$ as the exercise cost, yielding an instantaneous (flow) option with linear intensity. When internal purchases are lumpy or carry fixed setup costs, the problem becomes a genuine optimal stopping problem with a Snell envelope analogous to §E.1.

E.3 Threshold representation under monotonicity

Assume a Markov state Z_t and that $p_t^D = p^D(Z_t)$ is nondecreasing in a relevant coordinate. Then the external exercise region

$$\{p^D(Z_t) \geq c_{\text{ext}}(S_t)\}$$

is an upper set; in the scalar-coordinate case it admits a one-dimensional threshold.

For the internal margin, if $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t]$, λ_t , and c_t^{adj} are monotone in the same coordinate(s), then X_t^{int} is monotone and the frontier

$$\Gamma_J = \left\{ \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D | S_t] = \lambda_t + c_t^{\text{adj}} \right\}$$

is an upper set. Locally (under a nondegeneracy condition on the gradient), it admits a C^1 graph representation.

Remark (continuous-time free boundaries). In continuous time, the external margin yields an American variational inequality of the form

$$\max \{(\mathcal{L} - r)U, p^D - c_D - U\} = 0$$

for the option value U , with a free boundary in the shadow price p^D . The internal margin yields

a free boundary associated with complementarity conditions of the form

$$J_t \geq 0, \quad \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t - c_t^{\text{adj}} \leq 0, \quad J_t(\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t - c_t^{\text{adj}}) = 0,$$

which are the discrete-time counterparts of continuous-time KKT/free-boundary characterizations.

F Discrete-to-Continuous Limit for invariant A: Stock–Flow, Bellman–HJB, and Shadow-Price Diffusions

We provide a self-contained derivation of the continuous-time benchmark reported in Section 4.4 and its connection to a vanishing-step discretization. Throughout, $S = (K, D, Z)$ denotes the state; the control is J with $I := M - J$ and $M = \tau Y$. When an object is written without the control argument it is evaluated at the optimal policy.

F.1 Standing Assumptions and Notation

Fix a step $\Delta t > 0$ and grid times $t_n = n\Delta t$. Discrete objects carry a superscript Δ (e.g., V^Δ , S^Δ , $p^{D,\Delta}$). Work on a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_{t_n}\}, \mathbb{Q})$.

A Δt -parametrized discretization. To connect the discrete model to the continuous-time benchmark as $\Delta t \downarrow 0$, we consider the standard “stock–flow” scaling in which (i) discounting and depreciation scale to first order in Δt , and (ii) flow quantities contribute at order Δt . Specifically, under a given admissible control J^Δ (interpreted as a *rate* control so that the per-step stock increment is $J^\Delta \Delta t$), the state transition over one step is

$$\begin{aligned} K_{t_{n+1}}^\Delta &= (1 - \delta \Delta t) K_{t_n}^\Delta + G(I_{t_n}^\Delta, D_{t_n}^\Delta) \Delta t, \\ D_{t_{n+1}}^\Delta &= (1 - \delta_D \Delta t) D_{t_n}^\Delta - G(I_{t_n}^\Delta, D_{t_n}^\Delta) \Delta t + J_{t_n}^\Delta \Delta t, \\ Z_{t_{n+1}}^\Delta &= Z_{t_n}^\Delta \exp\left((\mu_Z - \tfrac{1}{2}\sigma_Z^2)\Delta t + \sigma_Z \sqrt{\Delta t} \varepsilon_{n+1}\right), \quad \varepsilon_{n+1} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1), \end{aligned} \tag{41}$$

which matches the lognormal step (13) used in the main text. The one-step discount factor is $\rho^\Delta = e^{-r\Delta t}$.

H1 Primitive regularity on compacts. On any compact state–control set, G , Y , and the flow payoff Π are C^3 and their partial derivatives up to order 3 are bounded (equivalently, second derivatives are Lipschitz).

H2 Localization / tightness. For any finite horizon T , we work on a compact localization region (e.g., via truncation of Z as in Proposition 3, or via stopping times). On this region, the family of interpolated processes $\{S^\Delta\}_{\Delta t > 0}$ is tight and has uniformly bounded moments needed for the weak convergence statements below.

H3 Control regularity. Controls J^Δ are uniformly bounded on the localization region and are progressively measurable with respect to $\{\mathcal{F}_{t_n}\}$.

F.2 Stock–Flow Limit

From (41), the one-step increments satisfy

$$\begin{aligned} K_{t_{n+1}}^\Delta - K_{t_n}^\Delta &= (-\delta K_{t_n}^\Delta + G(I_{t_n}^\Delta, D_{t_n}^\Delta)) \Delta t, \\ D_{t_{n+1}}^\Delta - D_{t_n}^\Delta &= (-\delta_D D_{t_n}^\Delta - G(I_{t_n}^\Delta, D_{t_n}^\Delta) + J_{t_n}^\Delta) \Delta t, \\ Z_{t_{n+1}}^\Delta - Z_{t_n}^\Delta &= \mu_Z Z_{t_n}^\Delta \Delta t + \sigma_Z Z_{t_n}^\Delta \sqrt{\Delta t} \varepsilon_{n+1} + O(\Delta t), \end{aligned}$$

where the $O(\Delta t)$ term is uniform on the localization region (e.g., in L^p for any fixed $p < \infty$). Standard Markov-chain approximation arguments for diffusions, then imply that the continuous-time interpolations converge weakly (along a subsequence) to the Itô system

$$\begin{aligned} dK_t &= (-\delta K_t + G(I_t, D_t)) dt, \\ dD_t &= (-\delta_D D_t - G(I_t, D_t) + J_t) dt, \\ \frac{dZ_t}{Z_t} &= \mu_Z dt + \sigma_Z dW_t^\mathbb{Q}. \end{aligned} \tag{42}$$

(If an implementation shock is introduced at order $\sqrt{\Delta t}$ in the *increment* of K , then dK_t acquires an additional diffusion term; see Remark 7 below.)

F.3 Bellman–HJB Limit

Let the one-step payoff satisfy the flow scaling

$$\Pi^\Delta(S, J) = \Pi(S, J) \Delta t + o(\Delta t),$$

uniformly on the localization region, where Π is the continuous-time flow payoff in Section 4.4. The discrete Bellman equation is

$$V^\Delta(S) = \max_{0 \leq J \leq M(S)} \left\{ \Pi^\Delta(S, J) + e^{-r\Delta t} \mathbb{E}^\mathbb{Q}[V^\Delta(S') \mid S, J] \right\}.$$

Viscosity convergence (monotone scheme). The dynamic-programming operator is monotone and stable (by contraction and localization), and it is consistent with the HJB operator in the usual sense. In particular, under the same boundary clipping/localization used in the numerical implementation, the discrete operator remains order-preserving on the localized domain, so the Barles–Souganidis monotonicity requirement is satisfied. For any smooth test function $\varphi \in C^2$, the conditional expectation under (41) admits the local expansion

$$\mathbb{E}^\mathbb{Q}[\varphi(S') \mid S, J] = \varphi(S) + \Delta t \mathcal{L}^J \varphi(S) + o(\Delta t),$$

where $\mathcal{L}^J$ is the controlled generator associated with (42). Therefore, by the consistency–stability–compactness theorem (or the Markov-chain approximation framework), $V^\Delta \rightarrow V$ uniformly on compacts (along a subsequence), and the limit V is a viscosity solution of

$$rV = \max_{0 \leq J \leq M} \left\{ \Pi + V_K(-\delta K + G) + V_D(-\delta_D D - G + J) + \mathcal{L}_Z V \right\}, \tag{43}$$

with $\mathcal{L}_Z V = \mu_Z Z V_Z + \frac{1}{2} \sigma_Z^2 Z^2 V_{ZZ}$. If $V \in C^{1,2}$, (43) holds classically.

F.4 Envelope and the Shadow-Price PDE

Let $p^D := V_D$. Assume that Y is invariant in D (as in the main text) so that $\Pi_D = 0$ and $M_D = \tau Y_D = 0$; holding J fixed in the envelope argument, this implies $I_D = 0$. Differentiate (43) with respect to D and apply the envelope theorem (holding the optimizer J^* fixed so that $\partial J^*/\partial D$ drops out). Denoting

$$b_K := -\delta K + G(I, D), \quad b_D := -\delta_D D - G(I, D) + J,$$

we obtain

$$r p^D = \Pi_D + b_{K,D} V_K + b_K V_{KD} + b_{D,D} V_D + b_D V_{DD} + \mathcal{L}_Z p^D.$$

Here the terms $b_{K,D} V_K + b_K V_{KD} + b_{D,D} V_D + b_D V_{DD}$ arise from differentiating $b_K V_K + b_D V_D$ with respect to D . Since $\Pi_D = 0$, $b_{K,D} = G_D$, and $b_{D,D} = -\delta_D - G_D$, we can write the transport operator

$$\mathcal{L}f := b_K f_K + b_D f_D + \mathcal{L}_Z f,$$

to get

$$(r + \delta_D) p^D = (V_K - p^D) G_D + \mathcal{L}_Z p^D + b_K p_K^D + b_D p_D^D, \quad (44)$$

or, equivalently,

$$(r + \delta_D - \mathcal{L}) p^D = (V_K - p^D) G_D. \quad (45)$$

Equation (45) is the continuous-time analogue of the discrete shadow recursion and makes explicit the transport operator in (K, D) together with the Z -generator.

F.5 Itô Dynamics of the Shadow Price

Applying Itô's lemma to $p^D(K_t, D_t, Z_t)$ along (42) gives

$$dp_t^D = \mathcal{L}p^D(K_t, D_t, Z_t) dt + \sigma_Z Z_t p_Z^D(K_t, D_t, Z_t) dW_t^\mathbb{Q}.$$

Substituting $\mathcal{L}p^D = (r + \delta_D)p^D - (V_K - p^D)G_D$ from (45) simplifies the drift to

$$dp_t^D = \left[(r + \delta_D) p_t^D + (p_t^D - V_K(K_t, D_t, Z_t)) G_D(I_t, D_t) \right] dt + \sigma_Z Z_t p_Z^D dW_t^\mathbb{Q}, \quad (46)$$

which is the diffusion reported in the main text.

F.6 From Discrete Shadow Recursion to the PDE (First-Order Limit)

Start from the (interior) discrete shadow recursion under the risk-neutral measure $\mathbb{Q}$ and step size Δt . We work on the interior (cap-slack) region where the effective usage cap is slack so that $\tilde{G} = G$ and G_D is well-defined; boundary cases are handled by viscosity/subgradient arguments and do not affect the first-order limit on the localized domain. Under the stock-flow scaling

(41), the recursion takes the form

$$p_t^{D,\Delta} = \rho^\Delta (1 - \delta_D \Delta t) \mathbb{E}_t[p_{t+\Delta t}^{D,\Delta}] + \rho^\Delta \Delta t \mathbb{E}_t \left[(V_K^\Delta(S_{t+\Delta t}^\Delta) - (1 - \delta_D \Delta t) p_{t+\Delta t}^{D,\Delta}) G_D^\Delta \right] + o(\Delta t),$$

where $G_D^\Delta := G_D(I_t^\Delta, D_t^\Delta)$ is $\mathcal{F}_t$ -measurable and $\rho^\Delta = e^{-r\Delta t} = 1 - r\Delta t + o(\Delta t)$.

By smoothness on the localization region and the generator expansion for any C^2 test function φ ,

$$\mathbb{E}_t[\varphi(S_{t+\Delta t}^\Delta)] = \varphi(S_t^\Delta) + \Delta t \mathcal{L}\varphi(S_t^\Delta) + o(\Delta t),$$

with

$$\mathcal{L}\varphi = b_K \varphi_K + b_D \varphi_D + \mu_Z Z \varphi_Z + \frac{1}{2} \sigma_Z^2 Z^2 \varphi_{ZZ}, \quad b_K = -\delta K + G, \quad b_D = -\delta_D D - G + J.$$

Formally (or for smooth approximations of $p^{D,\Delta}$ on the localized domain), applying this to $\varphi = p^{D,\Delta}$ yields

$$\mathbb{E}_t[p_{t+\Delta t}^{D,\Delta}] = p_t^{D,\Delta} + \Delta t (\mathcal{L}p^{D,\Delta})(S_t^\Delta) + o(\Delta t).$$

Since G_D^Δ is $\mathcal{F}_t$ -measurable,

$$\mathbb{E}_t[p_{t+\Delta t}^{D,\Delta} G_D^\Delta] = G_D^\Delta \mathbb{E}_t[p_{t+\Delta t}^{D,\Delta}] = G_D^\Delta (p_t^{D,\Delta} + \Delta t \mathcal{L}p^{D,\Delta}) + o(\Delta t).$$

Moreover, $V_K^\Delta(S_{t+\Delta t}^\Delta) = V_K^\Delta(S_t^\Delta) + O(\Delta t)$ on the localization region, hence

$$\Delta t \mathbb{E}_t[V_K^\Delta(S_{t+\Delta t}^\Delta) G_D^\Delta] = \Delta t G_D^\Delta V_K^\Delta(S_t^\Delta) + o(\Delta t).$$

Substituting these expansions and $\rho^\Delta = 1 - r\Delta t + o(\Delta t)$ into the discrete recursion, collecting $O(\Delta t)$ terms, and cancelling $p_t^{D,\Delta}$, we obtain after dividing by Δt and letting $\Delta t \downarrow 0$,

$$0 = -(r + \delta_D) p^D + (V_K - p^D) G_D + \mathcal{L}p^D,$$

equivalently,

$$(r + \delta_D - \mathcal{L}) p^D = (V_K - p^D) G_D,$$

which is exactly the compact shadow-price PDE (45). Expanding $\mathcal{L}$ recovers (44):

$$(r + \delta_D) p^D = (V_K - p^D) G_D + b_K p_K^D + b_D p_D^D + \mu_Z Z p_Z^D + \frac{1}{2} \sigma_Z^2 Z^2 p_{ZZ}^D.$$

Remark 7 (On the implementation shock w_t in the continuous limit). In the discrete model one may include an implementation factor in the one-step capital increment. If one specifies the increment as

$$K_{t+\Delta t}^\Delta - K_t^\Delta = (-\delta K_t^\Delta + G(I_t^\Delta, D_t^\Delta)) \Delta t + \sigma_\omega G(I_t^\Delta, D_t^\Delta) \sqrt{\Delta t} \varepsilon_{t+1}^\omega,$$

then in the limit dK_t acquires the diffusion $+\sigma_\omega G(I_t, D_t) dW_t^\omega$ and, via Itô, dp_t^D gains the martingale term $\sigma_\omega G p_K^D dW_t^\omega$. The PDE structure (45) is unchanged; only the full generator $\mathcal{L}$ adds the K -diffusion component.

G Appendix: Static One-Period Problem with Endogenous Technology

G.1 Proof of Proposition 4

Proof. Fix the pre-decision state $S_t = (K_t, D_t, A_t)$ and define the one-period budget $B := \tau Y_t$. The myopic problem in (21) is

$$\max_{I_t, J_t, R_t \geq 0} Y_{t+1}(I_t, J_t, R_t; S_t) \quad \text{s.t.} \quad I_t + J_t + R_t = B,$$

where $Y_{t+1} = (A_{t+1}Z_{t+1})^{1-\xi}K_{t+1}^\xi$ and, in this benchmark, $Z_{t+1} \equiv 1$ and $w_{t+1} \equiv w$.

Step 1: Lagrangian and KKT. Define the Lagrangian

$$\mathcal{L}(I_t, J_t, R_t, \lambda) := Y_{t+1}(I_t, J_t, R_t; S_t) + \lambda(B - I_t - J_t - R_t), \quad \lambda \in \mathbb{R}.$$

The KKT conditions are

$$\begin{aligned} \partial_{I_t} Y_{t+1} - \lambda &\leq 0, \quad I_t \geq 0, \quad I_t(\partial_{I_t} Y_{t+1} - \lambda) = 0, \\ \partial_{J_t} Y_{t+1} - \lambda &\leq 0, \quad J_t \geq 0, \quad J_t(\partial_{J_t} Y_{t+1} - \lambda) = 0, \\ \partial_{R_t} Y_{t+1} - \lambda &\leq 0, \quad R_t \geq 0, \quad R_t(\partial_{R_t} Y_{t+1} - \lambda) = 0, \quad I_t + J_t + R_t = B. \end{aligned}$$

Step 2: Corner for data spending, $J_t^* = 0$. Under invariant Z and constant w , next-period output depends only on (K_{t+1}, A_{t+1}) . Moreover, under (4) (endogenous technology), J_t affects D_{t+1} and does not enter K_{t+1} nor A_{t+1} contemporaneously (since G depends on (D_t, A_t) under our timing and Y_{t+1} depends on (K_{t+1}, A_{t+1}) only). Hence $\partial_{J_t} Y_{t+1} \equiv 0$. Take any feasible (I_t, J_t, R_t) with $J_t > 0$ and consider the alternative feasible allocation

$$(I'_t, J'_t, R'_t) := (I_t, 0, R_t + J_t),$$

which satisfies $I'_t + J'_t + R'_t = B$. The induced one-step stocks satisfy

$$K_{t+1}(I'_t, 0, R'_t) = K_{t+1}(I_t, J_t, R_t), \quad A_{t+1}(I'_t, 0, R'_t) = A_{t+1}(I_t, J_t, R_t) + J_t,$$

because R_t enters A_{t+1} additively while K_{t+1} does not depend on R_t . Since $Y_A > 0$ for $K_{t+1} > 0$ and $A_{t+1} > 0$, we obtain

$$Y_{t+1}(I'_t, 0, R'_t; S_t) > Y_{t+1}(I_t, J_t, R_t; S_t),$$

contradicting optimality of any candidate with $J_t > 0$. Therefore the unique one-period optimizer satisfies

$$J_t^* = 0.$$

Step 3: Reduced (I_t, R_t) problem and interior FOC. With $J_t = 0$, the problem reduces to

$$\max_{I_t, R_t \geq 0} Y_{t+1}(I_t, 0, R_t; S_t) \quad \text{s.t.} \quad I_t + R_t = B.$$

Under the stated strict concavity on the feasible simplex, the solution (I_t^*, R_t^*) is unique. If the optimum is interior in (I_t, R_t) , the FOC equalizes marginal contributions:

$$\partial_{I_t} Y_{t+1}(I_t^*, 0, R_t^*; S_t) = \partial_{R_t} Y_{t+1}(I_t^*, 0, R_t^*; S_t).$$

Step 4: Chain rule for $\partial_{I_t} Y_{t+1}$ and $\partial_{R_t} Y_{t+1}$ (with $Z \equiv 1$). Recall

$$Y_{t+1} = A_{t+1}^{1-\xi} K_{t+1}^\xi, \quad K_{t+1} = (1 - \delta) K_t + w G(I_t, D_t, A_t),$$

$$A_{t+1} = (1 - \delta_A)(A_t - \gamma G(I_t, D_t, A_t)) + R_t.$$

Thus

$$\partial_{I_t} K_{t+1} = w G_I, \quad \partial_{I_t} A_{t+1} = -(1 - \delta_A) \gamma G_I, \quad \partial_{R_t} K_{t+1} = 0, \quad \partial_{R_t} A_{t+1} = 1,$$

where G_I is evaluated at (I_t, D_t, A_t) . The marginal products are

$$Y_K = \xi A_{t+1}^{1-\xi} K_{t+1}^{\xi-1} = \frac{\xi Y_{t+1}}{K_{t+1}}, \quad Y_A = (1 - \xi) A_{t+1}^{-\xi} K_{t+1}^\xi = \frac{(1 - \xi) Y_{t+1}}{A_{t+1}}.$$

Hence

$$\partial_{I_t} Y_{t+1} = Y_K \partial_{I_t} K_{t+1} + Y_A \partial_{I_t} A_{t+1} = Y_K w G_I - Y_A \gamma (1 - \delta_A) G_I, \quad \partial_{R_t} Y_{t+1} = Y_A.$$

Step 5: Rearrangement to (22). The interior condition $\partial_{I_t} Y_{t+1} = \partial_{R_t} Y_{t+1}$ is

$$Y_K w G_I = Y_A (1 + \gamma (1 - \delta_A) G_I).$$

Substituting $Y_K = \xi Y_{t+1} / K_{t+1}$ and $Y_A = (1 - \xi) Y_{t+1} / A_{t+1}$ yields

$$w G_I(I_t^*, D_t, A_t) = \frac{1 - \xi}{\xi} \frac{K_{t+1}}{A_{t+1}} (1 + \gamma (1 - \delta_A) G_I(I_t^*, D_t, A_t)),$$

where (K_{t+1}, A_{t+1}) are the one-step stocks induced by $(I_t^*, 0, R_t^*)$ via (4). This is exactly (22). Uniqueness then follows from the assumed strict concavity of $Y_{t+1}(I, 0, R; S_t)$ in (I, R) on the budget simplex. $\square$

G.2 Robustness of the One-Period Optimum to Parameter Perturbations

Let $\psi \in \Psi \subset \mathbb{R}^p$ collect structural parameters used in robustness checks (e.g., $w, \delta_A, \gamma, \alpha, \theta, b$). Fix the pre-decision state $S = (K, D, A)$ and consider

$$\max_{I, J, R \geq 0} Y_t(I, J, R; S, \psi) \quad \text{s.t.} \quad I + J + R = \tau Y_{t-1}.$$

By Proposition 4, the problem reduces to a one-dimensional program with $J^*(S, \psi) = 0$ and $R^* = \tau Y_{t-1} - I^*$.

Assumptions

We impose the following regularity conditions:

(R1) **Smoothness and strong concavity.** Define the reduced objective

$$\tilde{Y}_t(I; S, \psi) := Y_t(I, 0, \tau Y_{t-1} - I; S, \psi).$$

Assume $\tilde{Y}_t$ is C^2 in (I, ψ) on $I \in [0, \tau Y_{t-1}]$, and there exists $\mu > 0$ (independent of ψ in a neighborhood $\mathcal{N}(\psi_0)$) such that

$$\partial_{II} \tilde{Y}_t(I; S, \psi) \leq -\mu < 0, \quad \forall (I, \psi) \in [0, \tau Y_{t-1}] \times \mathcal{N}(\psi_0).$$

(R2) **Bounded cross-derivatives.** There exists $L_{\psi I} < \infty$ such that $\|\partial_{I\psi} \tilde{Y}_t(I; S, \psi)\| \leq L_{\psi I}$ on $[0, \tau Y_{t-1}] \times \mathcal{N}(\psi_0)$.

(R3) **Corner persistence for J (structural).** Under the maintained timing, J affects only D_t but does not enter (K_t, A_t) at the one-step horizon, and operating output depends on (K_t, A_t) only. Hence $\partial_J Y_t(I, J, R; S, \psi) \equiv 0$ for all arguments, and the dominance argument in Proposition 4 implies $J^*(S, \psi) = 0$ throughout a neighborhood $\mathcal{N}(\psi_0)$.

Proposition 9 (Policy continuity and local Lipschitz robustness). *Under (R1)–(R3), for each fixed state S there exists a neighborhood $\mathcal{N}(\psi_0)$ such that the unique optimizer $(I^*(S, \psi), R^*(S, \psi))$ with $R^* = \tau Y_{t-1} - I^*$ satisfies:*

1. **Continuity:** $\psi \mapsto I^*(S, \psi)$ is continuous on $\mathcal{N}(\psi_0)$.
2. **Local Lipschitz bound:** for all $\psi_1, \psi_2 \in \mathcal{N}(\psi_0)$,

$$|I^*(S, \psi_1) - I^*(S, \psi_2)| \leq \frac{L_{\psi I}}{\mu} \|\psi_1 - \psi_2\|, \quad |R^*(S, \psi_1) - R^*(S, \psi_2)| \leq \frac{L_{\psi I}}{\mu} \|\psi_1 - \psi_2\|.$$

Proof. Define the scalar FOC map

$$\Phi(I; S, \psi) := \partial_I \tilde{Y}_t(I; S, \psi).$$

By (R1), $\partial_I \Phi = \partial_{II} \tilde{Y}_t \leq -\mu < 0$ on $[0, \tau Y_{t-1}] \times \mathcal{N}(\psi_0)$. The Implicit Function Theorem yields a unique C^1 selection $I^*(S, \psi)$ satisfying $\Phi(I^*(S, \psi); S, \psi) = 0$ and

$$\frac{\partial I^*}{\partial \psi}(S, \psi) = -\frac{\partial_{I\psi} \tilde{Y}_t(I^*(S, \psi); S, \psi)}{\partial_{II} \tilde{Y}_t(I^*(S, \psi); S, \psi)}.$$

By (R1)–(R2), $\|\partial_{I\psi} \tilde{Y}_t\| \leq L_{\psi I}$ and $|\partial_{II} \tilde{Y}_t| \geq \mu$, hence $\|\partial I^* / \partial \psi\| \leq L_{\psi I} / \mu$ on $\mathcal{N}(\psi_0)$. The mean value theorem yields the Lipschitz bound; continuity follows. Since $R^* = \tau Y_{t-1} - I^*$, the same bound holds for R^* . Finally, (R3) gives $J^*(S, \psi) = 0$ on $\mathcal{N}(\psi_0)$. $\square$

Proposition 10 (Objective-value sensitivity bound). *Under (R1)–(R2) and compact feasibility, there exist constants $L_0 < \infty$ such that for all $\psi_1, \psi_2 \in \mathcal{N}(\psi_0)$,*

$$\left| \max_{I \in [0, \tau Y_{t-1}]} \tilde{Y}_t(I; S, \psi_1) - \max_{I \in [0, \tau Y_{t-1}]} \tilde{Y}_t(I; S, \psi_2) \right| \leq L_0 \|\psi_1 - \psi_2\| + \frac{L_{\psi I}^2}{2\mu} \|\psi_1 - \psi_2\|^2.$$

Sketch. Let $I_i^* := I^*(S, \psi_i)$ for $i = 1, 2$. Decompose

$$\tilde{Y}_t(I_1^*; \psi_1) - \tilde{Y}_t(I_2^*; \psi_2) = [\tilde{Y}_t(I_1^*; \psi_1) - \tilde{Y}_t(I_1^*; \psi_2)] + [\tilde{Y}_t(I_1^*; \psi_2) - \tilde{Y}_t(I_2^*; \psi_2)].$$

The first term is bounded by $L_0 \|\psi_1 - \psi_2\|$ by continuity of partial derivatives on a compact set. For the second term, strong concavity implies

$$\tilde{Y}_t(I_1^*; \psi_2) - \tilde{Y}_t(I_2^*; \psi_2) \leq -\frac{\mu}{2} (I_1^* - I_2^*)^2.$$

Using Proposition 9, $|I_1^* - I_2^*| \leq (L_{\psi I}/\mu) \|\psi_1 - \psi_2\|$, which yields the stated bound. $\square$

Lemma 7 (Corner stability for J). *Under (R3), there exists a neighborhood $\mathcal{N}(\psi_0)$ such that $J^*(S, \psi) = 0$ for all $\psi \in \mathcal{N}(\psi_0)$.*

Proof. In the reduced one-period problem, $\partial_J Y_t \equiv 0$. The dominance argument in the proof of Proposition 4 implies $J^*(S, \psi) = 0$ uniformly for ψ in a neighborhood of ψ_0 . $\square$

G.3 Comparative Statics in τ and θ

This appendix provides self-contained proofs for Propositions 5 and 6. All expectations are taken under the risk-neutral measure $\mathbb{Q}$ whenever Z_t is stochastic.

Assumption 4 (Comparative-statics primitives for τ and θ). In addition to the existence and contraction conditions in Appendix C and the continuous-time bounds in Appendix H, assume:

1. **Strict concavity of the inner problem in the budget level.** For each fixed pre-decision state S_t and continuation value function V , the one-step inner problem

$$\Phi(b) := \max_{I, J, R \geq 0, I+J+R=b} \Psi(I, J, R; S_t, V), \quad b = \tau Y_t(S_t),$$

is strictly concave in (I, J, R) on the simplex $\{I, J, R \geq 0, I+J+R=b\}$. Consequently, Φ is concave in b and the associated Lagrange multiplier $\lambda_t(b)$ exists and is weakly decreasing in b .

2. **θ increases data intensity in the innovation kernel.** The kernel $G(I, D, A; \theta)$ is C^1 in θ and satisfies $G_{D\theta} \geq 0$ on the admissible region. (When needed below, we also assume $G_{A\theta} \geq 0$.)
3. **Small-gain/positivity for the coupled shadow-price system.** There exists $\varepsilon > 0$ such that

$$\sup_{(I, D, A)} \{\gamma G_A(I, D, A; \theta) + G_D(I, D, A; \theta)\} \leq r + \min\{\delta_A, \delta_D\} - \varepsilon.$$

This ensures that the operator in (48) is invertible and its resolvent is positivity preserving.

4. **Monotone marginal continuation values (sufficient condition).** For the fixed point $V^*(\cdot; \tau, \theta)$, the discounted marginal continuation value of data $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D(\tau, \theta) \mid S_t]$ is weakly increasing in τ and in θ , statewise.

Proof of Proposition 5 (budget expansion)

Fix a pre-decision state S_t and a continuation value function V . Write the one-step inner problem as

$$\Phi(b) := \max_{I, J, R \geq 0, I+J+R=b} \Psi(I, J, R; S_t, V), \quad b := \tau Y_t(S_t).$$

Under Assumption 4(i), Φ is concave in b and the budget multiplier $\lambda_t(b)$ exists and is weakly decreasing in b . Since $b = \tau Y_t(S_t)$ and $Y_t(S_t)$ is fixed at the inner maximization stage, $\lambda_t(\tau)$ is weakly decreasing in τ .

Define the excess-return gaps

$$\Delta_J(S_t; \tau) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D(\tau) \mid S_t] - \lambda_t(\tau), \quad \Delta_R(S_t; \tau) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A(\tau) \mid S_t] - \lambda_t(\tau).$$

By Assumption 4(iv), the discounted marginal continuation values are weakly increasing in τ , while $\lambda_t(\tau)$ is weakly decreasing. Therefore both Δ_J and Δ_R are weakly increasing in τ , which shifts the switching surfaces $\Gamma_J = \{S : \Delta_J(S; \tau) = 0\}$ and $\Gamma_R = \{S : \Delta_R(S; \tau) = 0\}$ outward in the sense stated in Proposition 5. $\square$

Proof of Proposition 6 (data weight)

We provide a continuous-time argument based on the generator-form system in Theorem 3. Along the optimal policy, $(p^A, p^D) = (V_A, V_D)$ solves

$$\begin{aligned} (r + \delta_A - \mathcal{L}) p^A &= (V_K - p^D - \gamma p^A) G_A(I, D, A; \theta) + (1 - \tau + \lambda_t \tau) Y_A, \\ (r + \delta_D - \mathcal{L}) p^D &= (V_K - p^D - \gamma p^A) G_D(I, D, A; \theta), \end{aligned} \quad (47)$$

where $\mathcal{L}$ is the controlled generator in Theorem 3. Rearranging yields the linear ‘‘Metzler’’ system

$$\begin{pmatrix} r + \delta_A - \mathcal{L} + \gamma G_A & G_A \\ \gamma G_D & r + \delta_D - \mathcal{L} + G_D \end{pmatrix} \begin{pmatrix} p^A \\ p^D \end{pmatrix} = \begin{pmatrix} V_K G_A + (1 - \tau + \lambda_t \tau) Y_A \\ V_K G_D \end{pmatrix}. \quad (48)$$

Under Assumption 4(iii), the inverse operator is positivity preserving.

Consider two parameters $\theta_2 > \theta_1$. By Assumption 4(ii), $G_D(\theta_2) \geq G_D(\theta_1)$ pointwise on the admissible region (and, when imposed, $G_A(\theta_2) \geq G_A(\theta_1)$). Assumption 4(iv) implies that the discounted marginal continuation value of data is weakly increasing in θ , which implies (equivalently) that the statewise data shadow price is weakly increasing in θ under the resolvent representation induced by (48). Formally, since both the operator coefficients and the right-hand side in (48) are componentwise nonnegative and weakly increasing in θ , positivity of the

resolvent implies the comparison principle

$$p^D(\theta_2) \geq p^D(\theta_1) \quad \text{statewise.}$$

Finally, the interior KKT condition (25) gives

$$J_t^*(\theta) = \left(\frac{\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D(\theta) \mid S_t] - \lambda_t(\theta)}{\kappa} \right)_+.$$

Hence, whenever the gap $\rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D(\theta) \mid S_t] - \lambda_t(\theta)$ is weakly increasing in θ (in particular under Assumption 4(iv)), both J_t^* and the exercise region $\{J_t > 0\}$ are weakly increasing in θ , as claimed in Proposition 6. $\square$

G.4 Additional Theory: Threshold Structure and Monotonicity of Switching Surfaces

Generic threshold representation for data investment (sufficient conditions)

Fix a state $S_t = (K_t, D_t, A_t, Z_t)$ and consider the inner problem embedded in the Bellman operator. With quadratic adjustment costs, the KKT condition (25) implies the exact representation

$$J_t^*(S_t) = \left(\frac{\Delta_J(S_t)}{\kappa} \right)_+, \quad \Delta_J(S_t) := \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D \mid S_t] - \lambda_t.$$

Thus $J_t^* > 0$ if and only if $\Delta_J(S_t) > 0$, and the switching surface is $\Gamma_J = \{S : \Delta_J(S) = 0\}$.

If, holding (K_t, A_t, Z_t) fixed, the gap $\Delta_J(K_t, D_t, A_t, Z_t)$ is continuous and strictly decreasing in D_t , then there exists a threshold function $\kappa_D(K_t, A_t, Z_t)$ such that

$$J_t^* > 0 \iff D_t < \kappa_D(K_t, A_t, Z_t), \quad J_t^* = 0 \iff D_t \geq \kappa_D(K_t, A_t, Z_t).$$

Local monotonicity of switching surfaces (sufficient conditions)

On interior points of the switching surface, the condition $\Delta_J(K, D, A, Z) = 0$ implicitly defines $D = \kappa_D(K, A, Z)$ locally. If $\partial_D \Delta_J < 0$ (strict), then by the Implicit Function Theorem,

$$\partial_K \kappa_D = -\frac{\partial_K \Delta_J}{\partial_D \Delta_J}, \quad \partial_A \kappa_D = -\frac{\partial_A \Delta_J}{\partial_D \Delta_J}.$$

Hence, under increasing-differences conditions ensuring $\partial_K \Delta_J \geq 0$ and $\partial_A \Delta_J \geq 0$, we obtain $\partial_K \kappa_D \geq 0$ and $\partial_A \kappa_D \geq 0$.

G.5 Steady-State Characterization

We derive the steady-state relations stated in Section 5.2. Suppose an optimal stationary policy induces constant allocations $(\bar{I}, \bar{J}, \bar{R})$ and a constant innovation flow $\bar{G} := G(\bar{I}, \bar{D}, \bar{A})$, with $\bar{w} := \mathbb{E}[w_t] \in (0, \infty)$.

Taking expectations in (4), the expected updates are

$$\bar{K}_{t+1} = (1 - \delta)\bar{K}_t + \bar{w} \bar{G}, \quad \bar{D}_{t+1} = (1 - \delta_D)(\bar{D}_t - \bar{G}) + \bar{J}, \quad \bar{A}_{t+1} = (1 - \delta_A)(\bar{A}_t - \gamma \bar{G}) + \bar{R}.$$

At an expected steady state $(\bar{K}, \bar{D}, \bar{A})$, imposing $\bar{K}_{t+1} = \bar{K}_t = \bar{K}$, etc., yields

$$\delta \bar{K} = \bar{w} \bar{G}, \quad \delta_D \bar{D} = \bar{J} - (1 - \delta_D) \bar{G}, \quad \delta_A \bar{A} = \bar{R} - \gamma(1 - \delta_A) \bar{G}.$$

These are understood as expected steady-state identities under an optimal stationary policy.

Static benchmark. In the purely myopic benchmark of Section 5.1, the firm maximizes next-period output $Y_t = A_t^{1-\xi} K_t^\xi$ subject to $I + J + R = \tau Y_{t-1}$. Since J does not affect (K_t, A_t) contemporaneously, the dominance argument implies $J^* = 0$ (Appendix G.1). The associated budget multiplier is weakly positive and is not pinned down to zero by the corner $J^* = 0$.

G.6 Local Stability Condition

Let $\bar{F} : \mathbb{R}_+^3 \rightarrow \mathbb{R}_+^3$ denote the expected state mapping induced by an optimal stationary policy: $S_{t+1} = \bar{F}(S_t)$ for $S_t = (K_t, D_t, A_t)$, with $\bar{F}$ using $\bar{w}$. Linearizing around $(\bar{K}, \bar{D}, \bar{A})$ yields

$$\begin{pmatrix} K_t - \bar{K} \\ D_t - \bar{D} \\ A_t - \bar{A} \end{pmatrix} \approx \nabla \bar{F}(\bar{K}, \bar{D}, \bar{A}) \begin{pmatrix} K_{t-1} - \bar{K} \\ D_{t-1} - \bar{D} \\ A_{t-1} - \bar{A} \end{pmatrix}.$$

A sufficient condition for local asymptotic stability of the expected dynamics is

$$\varrho(\nabla \bar{F}(\bar{K}, \bar{D}, \bar{A})) < 1,$$

where $\varrho(\cdot)$ denotes spectral radius.

G.7 Rigorous KKT and Shadow-Price Derivations with Endogenous Technology

Dynamic program (risk-neutral evaluation)

Let the state be $S_t = (K_t, D_t, A_t, Z_t)$ and evaluate continuation values under the risk-neutral measure $\mathbb{Q}$. The Bellman equation is

$$V(S_t) = \max_{I_t, J_t, R_t \geq 0} \left\{ (1-\tau)Y(K_t, A_t, Z_t) - \frac{\kappa}{2}J_t^2 - \frac{\kappa'}{2}R_t^2 + \rho \mathbb{E}_t^{\mathbb{Q}}[V(S_{t+1})] \right\}, \quad I_t + J_t + R_t = \tau Y(K_t, A_t, Z_t), \quad (49)$$

with transitions

$$K_{t+1} = (1 - \delta)K_t + w_t G(I_t, D_t, A_t), \quad (50a)$$

$$D_{t+1} = (1 - \delta_D)(D_t - G(I_t, D_t, A_t)) + J_t, \quad (50b)$$

$$A_{t+1} = (1 - \delta_A)(A_t - \gamma G(I_t, D_t, A_t)) + R_t, \quad (50c)$$

and Z_{t+1} as in (13). All derivative-based KKT and envelope identities below are stated on the interior (cap-slack) region where $\tilde{G} = G$ and G_I, G_D, G_A are well-defined; boundary cases can be handled via directional derivatives/subgradients. Assume G is C^2 and the inner problem is strictly concave (e.g., $\kappa, \kappa' > 0$), ensuring a unique optimizer.

KKT conditions and full marginal effects

Let λ_t denote the multiplier on the budget. The Lagrangian is

$$\mathcal{L} = (1 - \tau)Y - \frac{\kappa}{2}J_t^2 - \frac{\kappa'}{2}R_t^2 + \rho \mathbb{E}_t^{\mathbb{Q}}[V(S_{t+1})] + \lambda_t(\tau Y - I_t - J_t - R_t).$$

Define shadow prices $p_t^K := V_K(S_t)$, $p_t^D := V_D(S_t)$, $p_t^A := V_A(S_t)$. Using (50),

$$\partial_{I_t} K_{t+1} = w_t G_I, \quad \partial_{I_t} D_{t+1} = -(1 - \delta_D)G_I, \quad \partial_{I_t} A_{t+1} = -\gamma(1 - \delta_A)G_I,$$

hence

$$-\lambda_t + \rho \mathbb{E}_t^{\mathbb{Q}}[w_t p_{t+1}^K - (1 - \delta_D)p_{t+1}^D - \gamma(1 - \delta_A)p_{t+1}^A] G_I(I_t, D_t, A_t) \leq 0, \quad I_t(\cdot) = 0. \quad (51)$$

Similarly,

$$-\kappa J_t + \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D] - \lambda_t \leq 0, \quad -\kappa' R_t + \rho \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A] - \lambda_t \leq 0,$$

with complementary slackness. These coincide with (24)–(27) in the main text.

G.8 Technology as a Second Real Option (Appendix)

Shadow-price recursions and constrained envelope

Applying the constrained envelope theorem to (49) yields the discrete-time recursions

$$p_t^D = \rho(1 - \delta_D) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D] + \rho \mathbb{E}_t^{\mathbb{Q}}[(w_t p_{t+1}^K - (1 - \delta_D)p_{t+1}^D - \gamma(1 - \delta_A)p_{t+1}^A) G_D(I_t, D_t, A_t)], \quad (52)$$

$$p_t^A = (1 - \tau) Y_A(K_t, A_t, Z_t) + \lambda_t \tau Y_A(K_t, A_t, Z_t) + \rho(1 - \delta_A) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A] + \rho \mathbb{E}_t^{\mathbb{Q}}[(w_t p_{t+1}^K - (1 - \delta_D)p_{t+1}^D - \gamma(1 - \delta_A)p_{t+1}^A) G_A(I_t, D_t, A_t)], \quad (53)$$

where $Y_A := \partial Y / \partial A$. The additional term $\lambda_t \tau Y_A$ arises because the feasible set depends on A_t through $Y(K_t, A_t, Z_t)$.

On interior sets where (51) holds with equality, $\lambda_t = \rho \mathbb{E}_t^{\mathbb{Q}}[\dots] G_I(I_t, D_t, A_t)$, and (53) admits the compact form

$$p_t^A = (1 - \tau) Y_A + \rho(1 - \delta_A) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^A] + \rho \mathbb{E}_t^{\mathbb{Q}}[(w_t p_{t+1}^K - (1 - \delta_D)p_{t+1}^D - \gamma(1 - \delta_A)p_{t+1}^A) (G_A + \tau Y_A G_I)]. \quad (54)$$

Quadratic adjustment costs affect (J_t, R_t) through the KKT conditions but do not enter the envelope recursions directly.

G.9 Derivation of the Shadow-Price Recursions

Data shadow price

Since the current payoff and the budget constraint do not depend on D_t , the envelope theorem gives

$$p_t^D = \rho \mathbb{E}_t^{\mathbb{Q}}[V_K(S_{t+1}) \partial_{D_t} K_{t+1} + V_D(S_{t+1}) \partial_{D_t} D_{t+1} + V_A(S_{t+1}) \partial_{D_t} A_{t+1}].$$

From (50),

$$\partial_{D_t} K_{t+1} = w_t G_D, \quad \partial_{D_t} D_{t+1} = (1 - \delta_D)(1 - G_D), \quad \partial_{D_t} A_{t+1} = -(1 - \delta_A)\gamma G_D,$$

so

$$p_t^D = \rho(1 - \delta_D) \mathbb{E}_t^{\mathbb{Q}}[p_{t+1}^D] + \rho \mathbb{E}_t^{\mathbb{Q}}[(w_t p_{t+1}^K - (1 - \delta_D)p_{t+1}^D - \gamma(1 - \delta_A)p_{t+1}^A) G_D],$$

which is (52).

Technology shadow price

Here the current payoff and the budget depend on A_t through $Y(K_t, A_t, Z_t)$. The constrained envelope theorem yields

$$p_t^A = (1 - \tau) Y_A + \lambda_t \tau Y_A + \rho \mathbb{E}_t^{\mathbb{Q}}[V_K \partial_{A_t} K_{t+1} + V_D \partial_{A_t} D_{t+1} + V_A \partial_{A_t} A_{t+1}].$$

From (50),

$$\partial_{A_t} K_{t+1} = w_t G_A, \quad \partial_{A_t} D_{t+1} = -(1 - \delta_D)G_A, \quad \partial_{A_t} A_{t+1} = (1 - \delta_A)(1 - \gamma G_A),$$

which gives (53). The compact form (54) follows by substituting the interior equality for λ_t from (51).

Corners and subgradients

At corners ($J_t = 0$ or $R_t = 0$), the KKT inequalities replace equalities and $p^\bullet$ should be interpreted as (right) supergradients at kinks. The envelope recursions remain valid under standard measurable selection and differentiation-under-the-expectation conditions.

H Appendix: Continuous-Time Limit of Shadow-Price Recursions

This appendix derives the continuous-time benchmark for the shadow prices of data and firm-specific technology, and establishes its first-order connection to the discrete-time model. The state is $S = (K, D, A, Z)$; controls are (I, J, R) with the *flow* budget constraint

$$I + J + R = \tau Y(K, A, Z),$$

and the innovation kernel $G = G(I, D, A)$ is C^2 . Throughout, $Y(K, A, Z) = (AZ)^{1-\xi} K^\xi$. Throughout the analytical derivations we work on the interior (cap-slack) region where the effective usage cap is slack so that $\tilde{G} = G$ and G_I, G_D, G_A are well-defined; boundary cases can be handled via directional derivatives/viscosity arguments.

H.1 Standing assumptions and notation

Fix a time step $\Delta t > 0$ and grid times $t_n = n\Delta t$. Discrete objects carry a superscript Δ (e.g., $V^\Delta, S^\Delta, p^{\bullet, \Delta}$). Unless a control argument is shown explicitly, all objects are evaluated at the (unique) optimal control.

H1 Regularity on compacts. Y, G , and the one-period payoff rate π are C^3 on compact state-control sets; all partial derivatives up to order 2 are uniformly bounded on such sets.

H2 Localization, tightness, and moment bounds. For any finite horizon T , we work with a localization (e.g., truncation of Z or stopping times) such that the stopped/interpolated processes $(K^\Delta, D^\Delta, A^\Delta, Z^\Delta)$ remain in a common compact set on $[0, T]$ a.s. On this localized region, the family has uniformly bounded second moments, and a Lindeberg condition holds for the i.i.d. shocks used below.

H3 Risk-neutral scaling and flow normalization. The discount factor satisfies $\rho^\Delta = e^{-r\Delta t}$ and depreciation is scaled as $\delta_j^\Delta := 1 - e^{-\delta_j\Delta t}$ for $j \in \{K, D, A\}$. The technology index follows a GBM-consistent Euler step

$$Z_{t_{n+1}}^\Delta = Z_{t_n}^\Delta \left(1 + \mu_Z \Delta t + \sigma_Z \sqrt{\Delta t} \zeta_{n+1}\right), \quad \mathbb{E}[\zeta_n] = 0, \quad \mathbb{E}[\zeta_n^2] = 1,$$

with $\{\zeta_n\}$ i.i.d. and $\mathbb{E}|\zeta_n|^3 < \infty$.

Flow normalization. Controls (I, J, R) and the payoff π are *rates*. Thus, the discrete problem over a period of length Δt uses the one-period payoff $\pi \Delta t$ and stock updates that are first-order in Δt (see below).

H.2 Stock-flow limit

Assume the Δ -model uses the (scaled) stock laws

$$K_{t+\Delta t}^\Delta = (1 - \delta_K^\Delta) K_t^\Delta + w G(I_t^\Delta, D_t^\Delta, A_t^\Delta) \Delta t + o(\Delta t), \quad (55)$$

$$D_{t+\Delta t}^\Delta = (1 - \delta_D^\Delta) D_t^\Delta - G(I_t^\Delta, D_t^\Delta, A_t^\Delta) \Delta t + J_t^\Delta \Delta t + o(\Delta t), \quad (56)$$

$$A_{t+\Delta t}^\Delta = (1 - \delta_A^\Delta) A_t^\Delta - \gamma G(I_t^\Delta, D_t^\Delta, A_t^\Delta) \Delta t + R_t^\Delta \Delta t + o(\Delta t). \quad (57)$$

Equivalently, in increment form,

$$\begin{aligned} K_{t+\Delta t}^\Delta - K_t^\Delta &= -\delta_K^\Delta K_t^\Delta + w G(I_t^\Delta, D_t^\Delta, A_t^\Delta) \Delta t + o(\Delta t), \\ D_{t+\Delta t}^\Delta - D_t^\Delta &= -\delta_D^\Delta D_t^\Delta - G(I_t^\Delta, D_t^\Delta, A_t^\Delta) \Delta t + J_t^\Delta \Delta t + o(\Delta t), \\ A_{t+\Delta t}^\Delta - A_t^\Delta &= -\delta_A^\Delta A_t^\Delta - \gamma G(I_t^\Delta, D_t^\Delta, A_t^\Delta) \Delta t + R_t^\Delta \Delta t + o(\Delta t), \\ Z_{t+\Delta t}^\Delta - Z_t^\Delta &= \mu_Z Z_t^\Delta \Delta t + \sigma_Z Z_t^\Delta \sqrt{\Delta t} \zeta_{n+1}. \end{aligned}$$

Using $\delta_j^\Delta = \delta_j \Delta t + o(\Delta t)$ and the martingale FCLT, the continuous-time interpolations converge weakly (under **H1–H3**) to the Itô system

$$\begin{aligned} dK_t &= (-\delta_K K_t + w G(I_t, D_t, A_t)) dt, \\ dD_t &= (-\delta_D D_t - G(I_t, D_t, A_t) + J_t) dt, \\ dA_t &= (-\delta_A A_t - \gamma G(I_t, D_t, A_t) + R_t) dt, \\ \frac{dZ_t}{Z_t} &= \mu_Z dt + \sigma_Z dW_t^\mathbb{Q}. \end{aligned} \tag{58}$$

(If implementation shocks diffuse into K —e.g. $w_t = 1 + \sigma_w dB_t$ —then dK_t acquires an extra diffusion term; the PDE structure below is unchanged except that the generator gains an additional second-order term in K .)

H.3 Bellman–HJB limit

Let the *flow* payoff rate be

$$\pi(S, J, R) := (1 - \tau)Y(S) - \frac{\kappa}{2}J^2 - \frac{\kappa'}{2}R^2, \quad S = (K, D, A, Z).$$

The discrete Bellman equation at step Δt is

$$V^\Delta(S) = \max_{I+J+R=\tau Y(S)} \left\{ \pi(S, J, R) \Delta t + e^{-r\Delta t} \mathbb{E}[V^\Delta(S') \mid S, (I, J, R)] \right\},$$

where $S' = (K', D', A', Z')$ is given by (55)–(57) and the Euler step for Z .

For any smooth test function $\varphi \in C^2$, a second-order Taylor expansion in (K, D, A, Z) yields

$$\mathbb{E}[\varphi(S') \mid S, (I, J, R)] = \varphi(S) + \Delta t \mathcal{L}^{(I, J, R)} \varphi(S) + o(\Delta t),$$

where $\mathcal{L}^{(I, J, R)}$ is the controlled generator implied by (58):

$$(\mathcal{L}^{(I, J, R)} f)(S) = b_K f_K + b_D f_D + b_A f_A + \mu_Z Z f_Z + \frac{1}{2} \sigma_Z^2 Z^2 f_{ZZ}, \tag{59}$$

with drifts

$$b_K = -\delta_K K + w G, \quad b_D = -\delta_D D - G + J, \quad b_A = -\delta_A A - \gamma G + R,$$

and $G = G(I, D, A)$.

Using $e^{-r\Delta t} = 1 - r\Delta t + o(\Delta t)$ and rearranging gives

$$0 = \max_{I+J+R=\tau Y} \left\{ \pi(S, J, R) - rV^\Delta(S) + \mathcal{L}^{(I, J, R)} V^\Delta(S) \right\} + o(1).$$

By the Barles–Souganidis consistency–stability–monotonicity framework, $V^\Delta \rightarrow V$ locally uniformly and V is the (unique) viscosity solution of the stationary HJB

$$rV = \max_{I+J+R=\tau Y} \left\{ \pi(S, J, R) + V_K b_K + V_D b_D + V_A b_A + \mathcal{L}_Z V \right\}, \tag{60}$$

where $\mathcal{L}_Z V := \mu_Z Z V_Z + \frac{1}{2} \sigma_Z^2 Z^2 V_{ZZ}$. If $V \in C^{1,2}$, (60) holds classically.

H.4 Envelope identities and shadow-price PDEs

Define $p^K := V_K$, $p^D := V_D$, $p^A := V_A$, and let λ be the multiplier on $I + J + R = \tau Y(K, A, Z)$. Note that

$$Y_A(K, A, Z) = (1 - \xi) Z^{1-\xi} K^\xi A^{-\xi}.$$

Let $\mathcal{L}$ denote the generator (59) evaluated at the optimal control.

Proposition 11 (Shadow-price PDEs under general G). *Under (60) and **H1–H3**, the shadow prices satisfy*

$$\begin{aligned} (r + \delta_D - \mathcal{L}) p^D &= (w p^K - p^D - \gamma p^A) G_D(I, D, A), \\ (r + \delta_A - \mathcal{L}) p^A &= (1 - \tau) Y_A(K, A, Z) + \lambda \tau Y_A(K, A, Z) + (w p^K - p^D - \gamma p^A) G_A(I, D, A), \end{aligned} \tag{61}$$

where (I, J, R) are evaluated at the optimal policy. At interior points,

$$\lambda = (w p^K - p^D - \gamma p^A) G_I(I, D, A). \tag{62}$$

Sketch. Differentiate (60) with respect to D and A . By the (constrained) envelope theorem, derivatives of the optimal controls do not appear; only explicit state dependence matters. Since D does not enter Y and does not affect the feasible set, there is no budget-relaxation term in the p^D equation.

Let $b = (b_K, b_D, b_A)$ be the drift vector under the optimal policy. Because ∂_D commutes with $\mathcal{L}_Z$ and

$$\mathcal{L} p^D = b_K V_{KD} + b_D V_{DD} + b_A V_{AD} + \mathcal{L}_Z V_D,$$

we obtain

$$r p^D = \mathcal{L} p^D + p^K \partial_D b_K + p^D \partial_D b_D + p^A \partial_D b_A.$$

From (58),

$$\partial_D b_K = w G_D, \quad \partial_D b_D = -\delta_D - G_D, \quad \partial_D b_A = -\gamma G_D.$$

Rearranging yields

$$(r + \delta_D - \mathcal{L}) p^D = (w p^K - p^D - \gamma p^A) G_D.$$

For p^A , the constrained envelope picks up the budget-relaxation term because the feasible set depends on A through Y :

$$\partial_A \left[\max_{I+J+R=\tau Y} \{\cdots\} \right] = (1 - \tau) Y_A + \lambda \tau Y_A + (\text{transport terms}).$$

Proceeding as above and using

$$\partial_A b_K = w G_A, \quad \partial_A b_D = -G_A, \quad \partial_A b_A = -\delta_A - \gamma G_A,$$

gives the second PDE in (61). Finally, the interior FOC for I gives (62). $\square$

H.5 Itô forms (shadow-price diffusions)

Applying Itô to $p^\bullet(K_t, D_t, A_t, Z_t)$ along (58),

$$dp_t^\bullet = (\mathcal{L}p^\bullet)(S_t) dt + \sigma_Z Z_t p_Z^\bullet(S_t) dW_t^\mathbb{Q}.$$

Substituting $\mathcal{L}p^\bullet$ from (61) yields

$$dp_t^D = \left[(r + \delta_D) p_t^D - (w p_t^K - p_t^D - \gamma p_t^A) G_D(I_t, D_t, A_t) \right] dt + \sigma_Z Z_t p_Z^D dW_t^\mathbb{Q}, \quad (63)$$

$$dp_t^A = \left[(r + \delta_A) p_t^A - ((1 - \tau) + \lambda_t \tau) Y_A(K_t, A_t, Z_t) - (w p_t^K - p_t^D - \gamma p_t^A) G_A(I_t, D_t, A_t) \right] dt + \sigma_Z Z_t p_Z^A dW_t^\mathbb{Q}. \quad (64)$$

H.6 First-order discrete-to-continuous limit

Start from the discrete shadow-price recursions under $\mathbb{Q}$ (consistently scaled with **H3**):

$$\begin{aligned} p_t^{D,\Delta} &= \rho^\Delta (1 - \delta_D^\Delta) \mathbb{E}_t[p_{t+\Delta t}^{D,\Delta}] + \rho^\Delta \Delta t \mathbb{E}_t \left[(w p_{t+\Delta t}^{K,\Delta} - (1 - \delta_D^\Delta) p_{t+\Delta t}^{D,\Delta} - \gamma (1 - \delta_A^\Delta) p_{t+\Delta t}^{A,\Delta}) G_D^\Delta \right] + o(\Delta t), \\ p_t^{A,\Delta} &= ((1 - \tau) Y_A^\Delta + \lambda_t^\Delta \tau Y_A^\Delta) \Delta t + \rho^\Delta (1 - \delta_A^\Delta) \mathbb{E}_t[p_{t+\Delta t}^{A,\Delta}] \\ &\quad + \rho^\Delta \Delta t \mathbb{E}_t \left[(w p_{t+\Delta t}^{K,\Delta} - (1 - \delta_D^\Delta) p_{t+\Delta t}^{D,\Delta} - \gamma (1 - \delta_A^\Delta) p_{t+\Delta t}^{A,\Delta}) G_A^\Delta \right] + o(\Delta t), \end{aligned}$$

where $G_\chi^\Delta := G_\chi(I_t^\Delta, D_t^\Delta, A_t^\Delta)$ is $\mathcal{F}_t$ -measurable.

Use $\rho^\Delta = 1 - r\Delta t + o(\Delta t)$ and $1 - \delta_j^\Delta = 1 - \delta_j\Delta t + o(\Delta t)$, and the generator expansion (for smooth limits)

$$\mathbb{E}_t[p_{t+\Delta t}^{\bullet,\Delta}] = p_t^{\bullet,\Delta} + (\mathcal{L}p^{\bullet,\Delta})(S_t) \Delta t + o(\Delta t),$$

together with $\mathbb{E}_t[p_{t+\Delta t}^{\bullet,\Delta} G_\chi^\Delta] = G_\chi^\Delta \mathbb{E}_t[p_{t+\Delta t}^{\bullet,\Delta}]$. Subtract $p_t^{\bullet,\Delta}$, divide by Δt , and let $\Delta t \downarrow 0$ to recover (61). In particular,

$$\rho^\Delta (1 - \delta_j^\Delta) = e^{-(r+\delta_j)\Delta t} = 1 - (r + \delta_j)\Delta t + o(\Delta t),$$

which is the source of the $(r + \delta_j)$ terms in (61).

H.7 Continuous-time optimal stopping (exercise regions)

With per-unit data acquisition cost $c_D > 0$, the perpetual option to acquire data has value $S(\cdot)$ characterized by the variational inequality

$$\max \{ (\mathcal{L} - r) S, p^D - c_D - S \} = 0,$$

where $\mathcal{L}$ is the generator of the Markov state (or of p^D when p^D is Markov in the state). Value matching and smooth pasting at the free boundary characterize the trigger. In the multidimensional state (K, D, A, Z) , the exercise set is a switching surface, consistent with the discrete-time Snell-envelope characterization.

H.8 Remarks on regularity and corners

If V is not $C^{1,2}$, interpret (61) in the viscosity sense. The Bellman scheme is monotone and consistent under **H1–H3** and hence converges to the unique viscosity solution. Quadratic adjustment costs affect optimal controls through the KKT conditions but do not enter the envelope identities directly. At corners ($J_t = 0$ or $R_t = 0$), the same PDEs hold with supergradients/subgradients.

H.9 Comparative statics via a Metzler small-gain bound

We now read monotone cross-effects and boundedness directly from (61).

Lemma 8 (Monotone cross-effects and non-explosive complementarity). *Assume:*

- (i) $G \in C^2$ with $G_D, G_A \geq 0$ and $G_{AD} \geq 0$ on a compact admissible set;
- (ii) $Y(K, A, Z) = (AZ)^{1-\xi} K^\xi$ so $Y_A \geq 0$, and along the optimal policy the Bellman operator preserves monotonicity in (D, A) (so $V_{KD}, V_{KA} \geq 0$ where defined);
- (iii) at interior points $\lambda = (w p^K - p^D - \gamma p^A) G_I$ with $G_I \geq 0$ (or, more generally, λ is nondecreasing in (D, A));
- (iv) the small-gain bound

$$\sup_{(I,D,A)} \{ \gamma G_A(I, D, A) + G_D(I, D, A) \} < r + \min\{\delta_A, \delta_D\}.$$

Then the unique bounded solution of (61) satisfies

$$\partial_D p^A \geq 0, \quad \partial_A p^D \geq 0,$$

and (p^A, p^D) is uniformly bounded (no explosive feedback).

Sketch. Differentiate (61) with respect to D and set $u := \partial_D p^A$ and $v := \partial_D p^D$. Collecting terms yields a linear system of the form

$$\mathcal{M}_D \begin{pmatrix} u \\ v \end{pmatrix} = g_D,$$

where the operator $\mathcal{M}_D$ is Metzler: its diagonal blocks are of the form $(r + \delta_j - \mathcal{L}) + (\text{nonnegative multiplication})$, while its off-diagonal blocks are nonnegative multiplications proportional to G_A and γG_D . Under (iv), the associated feedback operator has norm strictly below one, so $\mathcal{M}_D^{-1}$ admits a Neumann-series expansion and is positivity-preserving. Since g_D is componentwise nonnegative under (i)–(iii), we obtain $u, v \geq 0$, hence $\partial_D p^A \geq 0$. The same argument after differentiating with respect to A yields $\partial_A p^D \geq 0$. Applying the same decomposition directly to (61) gives uniform boundedness of (p^A, p^D) . $\square$

Table 2: Baseline Calibration

Parameter	Value	Description
Discount rate r	0.05	Annual risk-free rate
Time step Δ	0.25	Years per period (quarter)
Output elasticity ξ	0.35	Capital share in $Y = (AZ)^{1-\xi}K^\xi$
Reinvestment share τ	0.30	Fraction of output retained for investment
Depreciation of K δ	0.04	Annual physical-capital depreciation
Depreciation of D δ_D	0.06	Annual data obsolescence / decay
Depreciation of A δ_A	0.06	Annual technology depreciation
Revankar curvature α	0.60	Substitutability in $G(I, D, A; \theta)$
Scale β	1.00	Scale parameter in G
Effort weight a	0.50	Weight on implementation I in G
Technology weight b	0.30	Weight on A in G
Data weight θ	0.80	Relative weight of D in G
Knowledge-usage rate γ	0.50	Flow of G allocated to A
Adj. cost on J κ_J	6.0	Quadratic cost on data investment
Adj. cost on R κ_R	8.0	Quadratic cost on R&D
Z drift μ_Z	0.00	Annual drift of Z
Z volatility σ_Z	0.15	Annual volatility of Z

I Appendix: Numerical Implementation Details and Robustness Checks

This appendix documents the numerical solution, diagnostics, and robustness checks for Section 7. We describe (i) the state grids and quadrature, (ii) the discrete Bellman operator and solution method for the constant- A and dynamic- A regimes, (iii) numerical shadow prices, KKT gaps, and frontier extraction, and (iv) robustness to grid, quadrature, and time-step refinements. Tables 2 and 3 summarize the baseline calibration and numerical diagnostics.

I.1 Baseline calibration and numerical diagnostics

I.2 State space, grids, and quadrature

State and discretization. We solve the stochastic dynamic program on $S = (K, D, A, Z)$ with controls (I, J, R) subject to $I + J + R = \tau Y(K, A, Z)$ and transition laws in (4). The baseline discretization is quarterly, $\Delta = 0.25$, with discount factor $\rho = e^{-r\Delta}$.

Tensor-product grids. We use tensor-product grids with $(N_K, N_D, N_A, N_Z) = (60, 60, 60, 11)$ in the dynamic- A model and $(60, 60, 1, 11)$ in the constant- A benchmark (where $A \equiv 1$). The (K, D, A) grids are logarithmically spaced on compact bounds chosen to contain the baseline simulated paths and frontier overlays.

Aggregate shock and probability-form Gauss–Hermite. When Z is stochastic, it follows the lognormal step

$$Z' = Z \exp\left((\mu_Z - \tfrac{1}{2}\sigma_Z^2)\Delta + \sigma_Z\sqrt{\Delta}\varepsilon\right), \quad \varepsilon \sim \mathcal{N}(0, 1).$$

Table 3: Numerical Diagnostics (Baseline Calibration)

Statistic	Constant- A	Dynamic- A
State grid (N_K, N_D, N_A, N_Z)	(60, 60, 1, 11)	(60, 60, 60, 11)
GH nodes n_{GH}	5	5
Bellman residual (rel. sup)	$\approx 10^{-4}$	$\approx 10^{-4}$
Bellman residual (abs. sup)	$\approx 10^{-6}$	$\approx 10^{-6}$
Budget residual (max.)	$< 10^{-15}$	$< 10^{-15}$
KKT complementarity (max.)	$\lesssim 10^{-7}$	$\lesssim 10^{-7}$
Value diff. (grid refinement) [†]	$< 10^{-3}$	$< 10^{-3}$
Value diff. (time-step refinement) [‡]	10^{-3} – 10^{-2}	10^{-3} – 10^{-2}
Frontier shift (max.)	< 1 grid step	< 1 grid step

[†]Comparing baseline grids to finer grids (e.g., (60, 60, 60, 11) vs. (80, 80, 80, 15) in the dynamic- A case).

[‡]Quarterly $\Delta = 0.25$ vs. monthly $\Delta = 1/12$ after mapping annual primitives to period primitives.

Conditional expectations are computed with a probability-form Gauss–Hermite (GH) rule: weights $\{\omega_m\}$ are nonnegative and normalized so that $\sum_m \omega_m = 1$. Baseline uses $n_{\text{GH}} = 5$; increasing to 7 or 9 yields negligible changes in values and frontiers (Section I.9).

I.3 Interpolation and boundary handling

Multilinear interpolation. For off-grid continuation states S' , we evaluate $V(S')$ by multilinear interpolation over (K, D, A, Z) . For interior points, multilinear interpolation is a convex combination of neighboring grid values.

Baseline boundary rule (clipping). If a transition produces a coordinate slightly outside the grid, we apply *boundary clipping* (clamping) to the nearest boundary node before interpolation. This preserves the convex-combination structure that underlies the discrete nonexpansiveness bounds in Section I.7. As a robustness diagnostic, we also allow one-sided linear extrapolation for mildly out-of-grid points; reported objects remain unchanged within tolerances. Because extrapolation can break convexity, it is not used in the baseline computations.

I.4 Solution method: constant- A and dynamic- A models

We solve the discrete Bellman equations in both regimes by Gauss–Seidel (GS) value iteration. The constant- A problem has a one-dimensional choice set (effectively J), while the dynamic- A problem features the intra-period allocation (I, J, R) governed by the KKT system in Appendix G.7.

Intra-period allocation via the shadow value of funds. With quadratic adjustment costs, the KKT conditions imply, for a given shadow value of internal funds λ ,

$$J(\lambda) = \frac{(\rho \mathbb{E}^{\mathbb{Q}}[p^D(S') \mid S] - \lambda)_+}{\kappa_J}, \quad R(\lambda) = \frac{(\rho \mathbb{E}^{\mathbb{Q}}[p^A(S') \mid S] - \lambda)_+}{\kappa_R}, \quad I(\lambda) = \tau Y(S) - J(\lambda) - R(\lambda).$$

At each state S , we select $\lambda(S)$ by scalar bisection to satisfy the I -optimality condition (FOC_I) from Appendix G.7. A convenient bracket is

$$\lambda_{\min} = 0, \quad \lambda_{\max} = \max\{\rho \mathbb{E}^{\mathbb{Q}}[p^D(S') \mid S], \rho \mathbb{E}^{\mathbb{Q}}[p^A(S') \mid S]\},$$

so that $J(\lambda_{\max}) = R(\lambda_{\max}) = 0$ and $I(\lambda_{\max}) = \tau Y(S) \geq 0$. We stop bisection when the budget residual is below 10^{-4} .

Sequential update of shadow prices. Given a value iterate V^n , we compute shadow prices $(p^{D,n}, p^{A,n})$ via finite differences (Section I.5) and use them to evaluate the KKT allocation when applying the Bellman operator to obtain V^{n+1} . This staggered update stabilizes the KKT-based inner step while preserving the standard contraction logic of value iteration on the grid.

[t]

Grids for (K, D, A, Z) ; tolerances $(\varepsilon_V, \varepsilon_B, \varepsilon_{\text{KKT}})$; initial guess V^0 Fixed point V and policies (I^*, J^*, R^*) on the grid

$n \leftarrow 0$

stopping rules are not satisfied Finite-difference gradients on-grid Compute $p^{D,n}(S) \approx \partial_D V^n(S)$ and $p^{A,n}(S) \approx \partial_A V^n(S)$ on the grid

grid state S * [f] one GS sweep Discounted one-step expectations via GH + interpolation $\mathcal{E}^D(S) \leftarrow \rho \mathbb{E}^{\mathbb{Q}}[p^{D,n}(S') \mid S]$ $\mathcal{E}^A(S) \leftarrow \rho \mathbb{E}^{\mathbb{Q}}[p^{A,n}(S') \mid S]$

Find $\lambda(S)$ by bisection on $[\lambda_{\min}(S), \lambda_{\max}(S)]$ so that the I -FOC holds

$J^* \leftarrow (\mathcal{E}^D(S) - \lambda(S))_+ / \kappa_J$ $R^* \leftarrow (\mathcal{E}^A(S) - \lambda(S))_+ / \kappa_R$ $I^* \leftarrow \tau Y(S) - J^* - R^*$

Bellman update $V^{n+1}(S) \leftarrow \Pi(S; I^*, J^*, R^*) + \rho \mathbb{E}^{\mathbb{Q}}[V^n(S') \mid S, I^*, J^*, R^*]$

$n \leftarrow n + 1$ Check stopping rules (Section I.6)

I.5 Shadow prices, KKT gaps, and switching frontiers

Finite-difference gradients. Shadow prices are defined as derivatives of the fixed-point value function: $p^D = V_D$ and $p^A = V_A$. On the grid we compute centered differences in the interior and one-sided differences at boundaries. For example, along D ,

$$p^D(K, D, A, Z) \approx \frac{V(K, D + h, A, Z) - V(K, D - h, A, Z)}{2h},$$

with the analogous formula for p^A along the A axis. Representative directional derivatives are stable under grid refinement and time-step refinement, and agree with the envelope-recursion implementation up to numerical precision (order 10^{-4}).

Exercise gaps. Given $\lambda(S)$, the signed KKT gaps are

$$\Delta_J(S) = \rho \mathbb{E}^{\mathbb{Q}}[p^D(S') \mid S] - \lambda(S), \quad \Delta_R(S) = \rho \mathbb{E}^{\mathbb{Q}}[p^A(S') \mid S] - \lambda(S).$$

Switching surfaces are zero level sets $\{S : \Delta_J(S) = 0\}$ and $\{S : \Delta_R(S) = 0\}$. The (K, D) frontiers in Section 7 are extracted on a mid- Z slice.

I.6 Stopping rules and diagnostics

We terminate iterations when all of the following hold:

- Bellman residual: relative sup norm below 10^{-4} and absolute sup norm below 10^{-6} ;
- Budget feasibility: $\max_S |I(S) + J(S) + R(S) - \tau Y(S)| < 10^{-4}$;
- KKT complementarity residual: maximum violation below 10^{-3} (typically $\mathcal{O}(10^{-7})$).

Table 3 reports achieved residuals under the baseline calibration.

I.7 Discrete contraction of the numerical Bellman operator (sketch)

Let $\mathcal{T}_h$ denote the discrete Bellman operator on the grid, with interpolation $\mathcal{I}_h$ (multilinear with clipping) and quadrature expectation $\mathbb{E}_h$ (probability-form GH, weights sum to one). For bounded grid functions V, W :

Nonexpansiveness of $\mathcal{I}_h$. Under multilinear interpolation with clipping, $\mathcal{I}_h V(S')$ is a convex combination of neighboring grid values, hence

$$\|\mathcal{I}_h V - \mathcal{I}_h W\|_\infty \leq \|V - W\|_\infty.$$

Nonexpansiveness of $\mathbb{E}_h$. Because GH weights are nonnegative and sum to one,

$$\|\mathbb{E}_h[V(S')] - \mathbb{E}_h[W(S')]\|_\infty \leq \|V - W\|_\infty.$$

Blackwell discounting. Monotonicity and discounting imply

$$\|\mathcal{T}_h V - \mathcal{T}_h W\|_\infty \leq \rho \|V - W\|_\infty,$$

so value iteration converges geometrically at rate ρ on the grid. If one-sided extrapolation is used at boundaries, the convex-combination structure can fail; for this reason clipping is the baseline rule and extrapolation is used only as a robustness diagnostic.

I.8 Frontier extraction and overlay statistic

We extract frontiers on a mid- Z slice as zero-level sets of Δ_J and Δ_R on a regular (K, D) evaluation grid. We compute Δ_J, Δ_R by bilinear interpolation and apply a marching-contours routine to approximate the zero-level sets. When an entire K -column is same-signed, we fall back to one-dimensional interpolation in D to locate the sign change between nodes, preventing spurious “missing frontier” artifacts.

For overlay plots we report the average vertical displacement

$$\Delta D^* = \frac{1}{M} \sum_{m=1}^M (D^{\text{hi}}(K_m) - D^{\text{lo}}(K_m)),$$

where $D^{\text{hi}}(K_m)$ and $D^{\text{lo}}(K_m)$ are frontier ordinates under the highest and lowest parameter values in an overlay.

I.9 Global sensitivity analysis and robustness checks

Randomized global sensitivity design. We perturb $(\tau, \delta_A, \delta_D, \theta)$ over $N = 40$ Latin-hypercube draws with `seed=2025`:

$$\tau \in [0.25, 0.35], \quad \delta_A \in [0.04, 0.08], \quad \delta_D \in [0.04, 0.08], \quad \theta \in [0.7, 1.1].$$

For each draw we re-solve the dynamic program and compute summary objects including mean output $\bar{Y}$ and frontier locations. A standardized OLS surrogate of $\bar{Y}$ on standardized inputs yields the tornado plot in Figure 6.

Grid refinement. Increasing (N_K, N_D, N_A, N_Z) (e.g., $60 \rightarrow 80$ in (K, D, A) and $11 \rightarrow 15$ in Z) changes values, policies, and shadow prices only marginally. Switching frontiers move by less than one baseline grid step.

Quadrature refinement. Increasing n_{GH} from 5 to 7 (and to 9 in diagnostics) yields negligible changes in values and frontier locations at the plotting scale.

Time-step refinement. We re-solve on a monthly grid $\Delta = 1/12$ while preserving annual primitives through standard mappings:

$$\rho_\Delta = e^{-r\Delta}, \quad 1 - \delta_{j,\Delta} = (1 - \delta_{j,\text{annual}})^\Delta, \quad \sigma_\Delta = \sigma_{\text{annual}}\sqrt{\Delta}, \quad j \in \{K, D, A\}.$$

Sup-norm differences in values lie between 10^{-3} and 10^{-2} and frontier displacements are negligible; see Table 3.

Replication. The complete replication package (code, parameter files, and logs) is provided in the e-companion. Run logs record command-line arguments, diagnostics, and random seeds. Results replicate on a standard workstation; run times scale approximately linearly in grid size and the number of GH nodes.