An Analytic Theory of Frictional Firm Dynamics*

Vasco M. Carvalho
University of Cambridge and CEPR

Hanbaek Lee
University of Cambridge

March 2026

Abstract

This paper develops an analytical general equilibrium theory of frictional firm dynamics. A parsimonious microfoundation yields a marginal value of capital that is affine in idiosyncratic productivity despite fixed adjustment costs, delivering closed-form adjustment thresholds. We decompose the distance between frictional and frictionless economies through three gap measures—investment, sorting, and output—and establish four results: macroeconomic neutrality is generically impossible; fixed costs amplify aggregate investment semi-elasticities rather than dampen them; firm-level uncertainty affects misallocation non-monotonically through a distributional channel; and the extensive margin generates state-dependent dynamics with endogenous negative skewness.

Keywords: Analytical firm dynamics, inaction band, misallocation, nonlinear aggregate fluctuations.

JEL Classification: D21, E22, E32

*The authors thank Isaac Baley, Anmol Bhandari, Andres Blanco, Simon Gilchrist, Axel Gottfries, Guido Menzio, Simon Mongey, Yongseok Shin, Stephen Terry, Maarten De Ridder, Jonathon Hazell, conference participants at SED, Mannheim Workshop on Firm Heterogeneity and Macroeconomics, Janeway Institute Cambridge Macro Conference, BSE Summer Forum, KER, ESWC 2025, and London junior macro workshop, and seminar participants at NYU, QMUL, Yonsei, and Jönköping for helpful suggestions and discussions. Carvalho gratefully acknowledges funding from ERC Consolidator Grant 101001221.

1 Introduction

How do firm-level investment frictions shape the business cycle? Microeconomic data consistently show that capital adjustment is lumpy and infrequent, yet the macroeconomic implications of this behavior remain debated. A significant barrier to resolving this debate is the computational complexity inherent in heterogeneous agent models. Because these models typically rely on numerical solutions, the transmission from micro-level inaction to aggregate fluctuations remains a "black box."

Our contributions are twofold. First, we develop a tractable framework for frictional firm dynamics in which the stationary recursive competitive equilibrium can be characterized analytically—an object conventionally obtained only through numerical methods. The source of tractability is a microfoundation that delivers Tobin's q—the shadow value of installed capital (Tobin, 1969; Hayashi, 1982)—as an affine function of idiosyncratic productivity. Under constant returns to scale production and a convex adjustment cost, the marginal benefit of investment is linear in idiosyncratic productivity and independent of the firm's own capital history. This yields an analytically tractable inaction region and closed-form adjustment thresholds. The recursive equilibrium can therefore be characterized in closed form, providing portable tools applicable beyond the specific investment problem studied here.

Second, we use this tractability to establish the macroeconomic implications of firm-level inaction. The analysis is organized around three gap measures—an investment gap Δ_I, a sorting gap Δ_{II}, and an output gap Δ_Y—that decompose the distance between frictional and frictionless economies. All three admit closed-form expressions as functions of structural parameters and the cross-sectional distribution within the inaction band. The gap structure delivers four results.

Impossibility of macroeconomic neutrality. We show that in the empirically relevant region ($\Delta_I > 0$), Jensen's inequality forces a strict ranking among the cross-sectional moments of inactive firms. All three gaps (Δ_I, Δ_{II}, Δ_Y) are therefore simultaneously strictly positive. Both channels of misallocation—the level of capital accumulation and the efficiency of its allocation—are simultaneously active, not substitutes. Even at the boundary case of perfect investment neutrality ($\Delta_I = 0$), the ranking ensures $\Delta_{II} > 0$ and hence $\Delta_Y > 0$: the investment and output gaps cannot vanish simultaneously. This is a structural consequence of the hierarchy of truncated moments within the inaction band, not a calibration-dependent finding. The impossibility sharpens the approximate symmetry conditions of Elsby and Michaels (2019) into an exact analytical result.

Elasticity amplification. We resolve an empirical puzzle in the investment elasticity literature. Zwick and Mahon (2017) document investment semi-elasticities to tax incentives of approximately 7.2, substantially larger than convex-adjustment-cost models predict. Winberry (2021) and Koby and Wolf (2020) show that convex costs *dampen* aggregate elasticity, yet the data suggest amplification. Our framework demonstrates that fixed adjustment costs amplify the aggregate interest-rate elasticity through an extensive-margin channel operating through Δ_I. Because inactive firms contribute zero, the same absolute investment shift produces a larger percentage response. Threshold crossings add a discrete layer of sensitivity absent from convex-cost models.

Non-monotonic uncertainty transmission. We provide the first analytical isolation of the frictional channel of uncertainty transmission. A large literature documents the macroeconomic effects of uncertainty shocks—through wait-and-see dynamics (Bloom, 2009), stochastic volatility in DSGE (Fernández-Villaverde et al., 2011), investment dispersion (Bachmann and Bayer, 2014), and state-dependent adjustment (Vavra, 2014)—yet isolating the transmission mechanism has proved difficult because effects are mediated by the entire cross-sectional distribution (Bloom et al., 2018). In our framework, frictionless investment is linear in productivity, so a mean-preserving spread in firm-level volatility has zero effect on the frictionless benchmark. All uncertainty effects therefore operate exclusively through $\Delta_I(\sigma_\epsilon)$—the mass and composition of firms within the inaction band. The resulting relationship between micro-level uncertainty and aggregate misallocation is an inverted U-shape. This non-monotonicity cannot emerge from precautionary-savings or risk-aversion channels, which predict monotonically increasing effects.

Endogenous business-cycle asymmetry. Business-cycle dynamics are inherently state dependent and asymmetric. Negative productivity shocks produce larger investment contractions than the expansions generated by symmetric positive shocks, generating endogenous negative skewness. The mechanism is visible in the geometry of the aggregate investment function: in recessions, the inaction region shifts into higher-density parts of the productivity distribution, widening Δ_I and trapping more firms in inaction. This increases the mass of near-threshold firms, making the economy more fragile precisely when conditions deteriorate. The resulting excess volatility, negative skewness, and lower ergodic mean constitute a "volatility tax" operating through a compositional channel, not risk aversion or incomplete markets.

These results rest on an analytic decomposition of aggregate misallocation into economically interpretable components. The investment gap Δ_I is the product of two forces: a quantity margin, determined by the mass of inactive firms M_0, and a quality margin, determined by the average productivity of those firms relative to the inaction midpoint. The output gap Δ_Y introduces an additional sorting channel that captures distortions in the capital-productivity covariance $cov(k,z)$ among inactive firms. The entire recursive equilibrium is characterized by a three-dimensional sufficient statistic—the capital stock K, the covariance $cov(k,z)$, and aggregate TFP Z—through which all cross-sectional information is channeled. Because all margins are available in closed form, the mechanisms can be traced to specific movements in the cross-sectional distribution rather than attributed to numerical artifacts.

Ultimately, these results offer a new perspective on the long-standing debate regarding the aggregate relevance of micro-level frictions. By deriving analytical solutions, we can revisit classic questions—from the conditions for neutrality to the dampening role of general equilibrium prices—without the ambiguity of numerical approximation. Our findings provide a unified framework for understanding the nonlinear transmission of shocks through the frictional firm dynamics documented empirically.

Related literature. Our paper is primarily related to the literature studying micro-level lumpiness and its aggregate implications. A central question is whether the firm-level (S,s) cycle is neutralized in aggregate. Caplin and Spulber (1987), Caballero and Engel (1993), and Elsby and Michaels (2019) progressively generalize sufficient conditions for neutrality; the last shows that distributional symmetry over the inaction band is key. We sharpen these conditions through exact analytical investment functions and prove that neutrality cannot hold for both investment and output simultaneously.

Regarding non-neutrality, Caballero and Engel (1991, 1999) analyze off-steady-state dynamics via analytical summary indices in generalized (S, s) frameworks; Cooper et al. (1999) emphasizes the role of investment-vintage distributions; and Baley and Blanco (2021) and Alvarez and Lippi (2022) develop sufficient-statistic and impulse-response approaches, respectively. Our framework complements this line by analytically characterizing the distance between frictional and frictionless dynamics in general equilibrium. A prominent strand debates whether GE forces neutralize lumpy investment. Thomas (2002), Khan and Thomas (2003, 2008), and House (2014) argue for approximate neutrality. Gourio and Kashyap (2007), Bachmann et al. (2013), and Winberry (2021)—calibrating to the semi-elasticity evidence of Zwick and Mahon (2017)—demonstrate non-neutral outcomes under alternative setups; Koby and Wolf (2020) argues convex costs discipline firm-level sensitivity. We provide the analytical sufficient condition for investment neutrality in GE and prove the impossibility of joint neutrality. Moreover, we show that fixed costs *amplify* rather than dampen the aggregate elasticity, resolving the tension between the large empirical semi-elasticities and the dampening prediction of convex-cost models.

Our analysis of uncertainty transmission contributes to the literature initiated by Bloom (2009) and developed along several dimensions: stochastic volatility in DSGE models (Fernández-Villaverde et al., 2011), investment dispersion dynamics (Bachmann and Bayer, 2014), state-dependent price adjustment under time-varying volatility (Vavra, 2014), and the quantitative decomposition of Bloom et al. (2018). A common challenge is that uncertainty effects are mediated by the entire cross-sectional distribution, making analytical decomposition intractable. We provide the first analytical isolation of the frictional channel: because frictionless investment is linear in productivity, all uncertainty effects operate exclusively through the investment gap. The resulting closed-form non-monotonic relationship complements the simulation-based decompositions in the existing literature.

On the empirical side, plant-level data reveal pervasive non-convexities (Caballero et al., 1995; Doms and Dunne, 1998) that drive the skewness and kurtosis of aggregate investment. These patterns are consistent with our model's prediction that fluctuations in the mass of active firms are the dominant channel for aggregate nonlinearity. Our methodological approach builds on Moll (2014), who exploits CRS technology for tractable saving rules under heterogeneity, and on the dynamic programming properties for homogeneous functions established by

Alvarez and Stokey (1998).[1]

2 Model

Time is discrete and infinite. A unit measure of ex-ante homogeneous firms becomes ex-post heterogeneous through a persistent firm-level productivity process. Each firm owns capital and decides investment one period before utilization.

In each period, an aggregate productivity shock Z_t is realized. The household and firms observe the aggregate productivity level Z_t and the distribution Φ_t of the individual firms. Therefore, the vector of aggregate state variables S_t is

$$S_t = \{Z_t, \Phi_t\}. \tag{1}$$

2.1 Technology and investment

Two features of the production technology drive the model's tractability: constant returns to scale and capital-augmenting productivity. Together, they ensure that operating profit is linear in effective capital zZk, so that investment incentives depend only on a firm's productivity, not on its scale.

Firm i operates using a production function F that is *homogeneous of degree one* with respect to capital stock k_{it}:

$$\lambda F(z_{it} Z_t k_{it}, n_{it}) = F(\lambda z_{it} Z_t k_{it}, \lambda n_{it}),\ \forall \lambda > 0. \tag{2}$$

where z_{it} and Z_t are *capital-augmenting* idiosyncratic and aggregate productivities; k_{it} is capital stock; n_{it} is labor demand. By homogeneity, the profit function satisfies

$$\Pi(z_{it} Z_t k_{it}; S_t) := \max_{n_{it}}\ F(z_{it} Z_t k_{it}, n_{it}) - w(S_t) n_{it} = \Pi(1; S_t) z_{it} Z_t k_{it}, \tag{3}$$

[1] More generally, our work relates to the literature on micro-level misallocation and macro dynamics (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009; Asker et al., 2014; David and Venkateswaran, 2019). We contribute by providing an analytical characterization of how misallocation fluctuates over the business cycle, identifying it as a key source of aggregate nonlinearity.

where S_t is the aggregate state; w is the wage. The firm-level output satisfies[2]

$$Y(z_{it}Z_tk_{it};S_t) = O(S_t)z_{it}Z_tk_{it}, \tag{4}$$

where $O(S_t) = \Pi(1;S_t) - w(S_t)\Pi_w(1;S_t)$.[3] Similarly, firm-level labor demand is

$$n^*_{it} = N(S_t)z_{it}Z_tk_{it} \tag{5}$$

where $N(S_t) = -\Pi_w(1;S_t)$.[4]

The firm-level productivity z_{it} and aggregate productivity Z_t follow AR(1) processes:

$$z_{it+1} = \rho z_{it} + \epsilon_{it+1}, \quad \epsilon_{it+1} \sim_{iid} G^z, \quad \mathbb{E}\epsilon_{it+1} = \bar{\epsilon} > 0, \quad Var(\epsilon_{it+1}) = \sigma^2_\epsilon \tag{6}$$

$$Z_{t+1} = \rho_{agg}Z_t + \epsilon_{t+1}, \quad \epsilon_{t+1} \sim_{iid} G^z_{agg}, \quad \mathbb{E}\epsilon_{t+1} = \bar{\epsilon}_{agg} > 0, \quad Var(\epsilon_{t+1}) = \sigma^2_{agg}. \tag{7}$$

We normalize $\bar{\epsilon} = 1 - \rho$ and $\bar{\epsilon}_{agg} = 1 - \rho_{agg}$ to anchor the unconditional average of both processes at unity. The three primitives of the idiosyncratic process—persistence ρ, innovation variance σ^2_ϵ, and mean $\bar{\epsilon}$—fully determine the stationary cross-sectional distribution (mean $\bar{\epsilon}/(1-\rho)$, variance $\sigma^2_\epsilon/(1-\rho^2)$). In the stationary analysis of Section 4.2, they enter the closed-form investment gap (47) exclusively through the distributional bracket, cleanly separated from the discount and depreciation parameters (β, δ) that govern $\mathcal{G}^{ss}$. In the nonstationary environment, the same primitives govern the time-varying cross-section. The gap then additionally depends on the aggregate state Z through the inaction midpoint $z_M(Z)$. Most analytical results require only the AR(1) structure and finite second moments, not the specific distribution G^z. We discuss the precise distributional requirements of each result in Section 3.

Capital evolves as

$$k_{it+1} = (1-\delta)k_{it} + I_{it} \tag{8}$$

[2] The production technology is similar to Moll (2014), as it nests the CRS production function and assumes the capital-augmenting productivity process.

[3] The coefficient $O(\cdot)$ denotes the net-of-labor aggregate productivity and should not be confused with the Landau asymptotic notation used elsewhere.

[4] The result is immediate from Hotelling's lemma.

where I_{it} is investment.

Firm-level investment is subject to a unit period of time to build, convex adjustment cost $(\mu/2)I_{it}^2$, and fixed adjustment cost $\xi > 0$, which is incurred only when $I_{it} \neq 0$. The convex cost governs the intensive margin (how much to invest), while the fixed cost generates the extensive margin (whether to invest at all); Section 3 formalizes this separation. The form of the convex adjustment cost follows Abel (1983) and Cooper et al. (1999).

With probability $\varphi \in (0,1)$, a firm is hit by an exogenous disturbance that prevents capital adjustment. This *i.i.d.* shock captures residual sources of investment frictions, such as credit market frictions, beyond the explicit adjustment costs in our model.[5]

A distinctive feature of our model is that CRS operating profit implies no satiation point in capital accumulation. Without convex adjustment costs, firm size would be theoretically undefined. The cost $(\mu/2)I^2$ guarantees stationary investment dynamics by anchoring investment where marginal benefits and costs intersect.[6]

2.2 Recursive formulation of firms' problem

We denote future allocations with an apostrophe. At the beginning of each period, a firm with state $s = (k, z)$ observes aggregate states $S = (Z, \Phi)$. The recursive firm problem is

$$J(s; S) = \Pi(zZk; S) + (1-\varphi)\max\{R^*(s; S) - \xi, R^c(s; S)\} + \varphi R^c(s; S) \tag{9}$$

$$R^*(s; S) = \max_I -I - \frac{\mu}{2}I^2 + \mathbb{E}_{z,S}\Xi(S, S')J(k(1-\delta) + I, z'; S') \tag{10}$$

$$R^c(s; S) = \mathbb{E}_{z,S}\Xi(S, S')J(k(1-\delta), z'; S') \tag{11}$$

where J is the firm's value function; R^* and R^c are interim values for adjusters and non-adjusters. The stochastic discount factor $\Xi(S, S') \equiv \beta u_c(c(S'), l_H(S'))/u_c(c(S), l_H(S))$ is determined by the equity-owning household's marginal utility. The distribution evolves according to $\Phi' = M(S)$.

[5] The exogenous *i.i.d.* Poisson shock is similar to the Calvo price shock in the New Keynesian models. This shock does not play a key role in any of our results. It captures residual channels our model cannot account for in terms of extensive-margin investment.

[6] A firm can grow unboundedly with zero probability. The firm size distribution remains stable due to the stationary, mean-reverting nature of the idiosyncratic productivity process.

2.3 Household and equilibrium

The representative household owns equity in all firms, supplies labor, and chooses consumption to maximize lifetime utility subject to a budget constraint equating consumption to labor income plus aggregate dividends. The full household problem, including the objective, budget constraint, and optimality conditions, is presented in Appendix A.

A recursive competitive equilibrium is a set of policy and value functions such that: (i) household and firm problems are solved; (ii) labor and product markets clear,

$$
\begin{aligned}
g_{l_H}(S) &= \int g_n(s;S)d\Phi, \\
g_c(S) &= \int \Big(\Pi(1;S)zZk - g_I(s;S) - \tfrac{\mu}{2}g_I(s;S)^2 - \xi\mathbb{I}\{g_I(s;S) \neq 0\}\Big)\, d\Phi;
\end{aligned}
$$

where g_x denotes the equilibrium policy function for variable x; and (iii) the consistency condition $M_H(S) = M(S) = \Phi'$ holds.[7] The stationary recursive competitive equilibrium characterized in Section 4.2 is a straightforward extension of the recursive competitive equilibrium under constant aggregate productivity.[8]

3 Micro analytics

3.1 Affine marginal benefit

We first analytically characterize the marginal benefit of investment. Because profit is homogeneous of degree one, future marginal benefits are additively separable across periods. This property is unaffected by the presence of the fixed adjustment cost. To clarify this point, we define the frictionless benchmark value function $J^{NoFixed}$:

$$
J^{NoFixed}(s;S) = \max_I\ \Pi(zZk;S) - I - \tfrac{\mu}{2}I^2 + \mathbb{E}_{z,S}\Xi(S,S')J^{NoFixed}(k(1-\delta)+I,z';S'). \tag{12}
$$

[7] $\mathbb{K}$ and $\mathbb{Z}$ are the supports of the marginal distributions of capital and productivity induced from Φ.

[8] Existence and uniqueness of the stationary recursive competitive equilibrium follow from the monotonicity of the firm-level investment function (Hopenhayn and Prescott, 1992) and the explicitly derived aggregate demand and supply functions in factor markets, as established in Section 3.

In the following theorem, we show that (i) the marginal benefit of investment is independent of the fixed adjustment cost and (ii) it is affine in firm-level productivity z.

Theorem 1 (Marginal benefit)**.**
The marginal benefit of investment $q(k,z;S) := \frac{\partial \mathbb{E}_{z,S}\Xi(S,S')J(k,z';S')}{\partial k}$ *satisfies*

$$q(k,z;S) = \frac{\partial \mathbb{E}_{z,S}\Xi(S,S')J^{NoFixed}(k,z';S')}{\partial k} = \mathcal{G}(S)z + \mathcal{H}(S) \tag{13}$$

for some $\mathcal{G}, \mathcal{H} : \mathbb{S} \to \mathbb{R}$*, where* $\mathbb{S}$ *is the set of all possible aggregate states.*

Proof. See Appendix B. ∎

The economic content of Theorem 1 stems from the constant-returns-to-scale (CRS) assumption. Under CRS, the marginal product of capital depends only on the *contemporaneous* productivity realization, not on capital itself. The marginal benefit therefore decomposes into period-by-period future contributions, each depending only on future productivity and prices. This *additive separability* renders the marginal value *memoryless*. The return to investing today is the same regardless of whether the firm adjusts next period or remains inactive for several periods. By contrast, under decreasing returns the future marginal product depends on the *entire* future capital stock, requiring one to track all past investment decisions—making such models computationally intensive.

Our framework builds on a classical tradition. Lucas and Prescott (1971) showed that CRS ensures scale-independent investment rules, making industry equilibrium tractable with heterogeneous firms; Jovanovic (2025) highlights this as foundational for competitive industry dynamics. Our contribution is that CRS delivers tractability even with *fixed* adjustment costs. The marginal benefit is independent of the extensive-margin decision (adjust or not), separating the "how much" and "whether" dimensions into two tractable problems. As an immediate consequence, optimal investment size is orthogonal to ξ, enabling a closed-form inaction band (Theorem 2).

The marginal benefit coefficients $\mathcal{G}(S)$ and $\mathcal{H}(S)$ serve as *intermediary sufficient statistics* for the firm-level problem: all general-equilibrium feedback relevant for investment decisions is channeled through these two objects. In particular, they admit the following recursive characterization:

$$\mathcal{G}(S) = \mathbb{E}\left[\rho\Xi(S,S')\left(\Pi(1;S')Z' + (1-\delta)\mathcal{G}(S')\right)\right] \tag{14}$$

$$\mathcal{H}(S) = \left(\frac{1}{\rho} - 1\right)\mathcal{G}(S) + \mathbb{E}\left[\Xi(S,S')(1-\delta)\mathcal{H}(S')\right]. \tag{15}$$

$\mathcal{G}(S)$ captures how much a firm benefits from its *current* competitive position: the discounted future return from today's productivity advantage as it mean-reverts over time. $\mathcal{H}(S)$ captures the value of *future* opportunities: the discounted return from productivity innovations that the firm will receive regardless of its current state. Intuitively, $\mathcal{G}$ is pro-cyclical. When aggregate TFP is high, the future stream of marginal products is larger, making an additional unit of capital more valuable to a firm whose current productivity is z. $\mathcal{H}$ is pro-cyclical for a similar reason: higher aggregate TFP raises the expected value of future productivity innovations that the firm will receive.[9] Since $q = \mathcal{G}(S)z + \mathcal{H}(S)$ with both components pro-cyclical, the marginal benefit of investment q is itself pro-cyclical: firms' incentive to invest rises in booms and falls in recessions.

Both $\mathcal{G}$ and $\mathcal{H}$ are general-equilibrium objects embedding the full stream of stochastic discount factors and future prices. Without further assumptions on preferences, they generally lack closed-form expressions. Crucially, *explicit forms are not needed*. All results in Sections 3–5—including the inaction band, the gap decomposition, the impossibility theorem, the full equilibrium characterization, and the business-cycle implications—hold for *any* $\mathcal{G}$ and $\mathcal{H}$ satisfying (14)–(15). We exploit this separability in progressive steps: first deriving qualitative properties from the recursive forms alone; then obtaining explicit closed forms under the stationary equilibrium (Section 4.2); and finally characterizing the fully analytical nonstationary dynamics under GHH preferences.

Theorem 1 establishes a direct link to classical q-theory. The marginal benefit $q(z;S) = \mathcal{G}(S)z + \mathcal{H}(S)$ is the firm-level Tobin's q (Tobin, 1969): the shadow value of an additional unit of installed capital. As in Hayashi (1982), CRS production yields the equality of marginal and average q, $q(z;S) = q^a(z;S)$, so both $\mathcal{G}$ and $\mathcal{H}$ are empirically measurable from market-to-book ratios. In the classical q-theory

[9] The cyclical properties of $\mathcal{G}$ and $\mathcal{H}$ are based on a mild assumption that general equilibrium effects do not overturn the sign of the direct impact from TFP fluctuations, which is subject to the utility specification.

with convex costs alone (Abel, 1983), investment is monotonically increasing in q and the extensive margin is trivial. Our framework extends q-theory to accommodate fixed adjustment costs. The variable q continues to govern the intensive margin through the standard rule $I^* = (q-1)/\mu$ (Corollary 1), but now also governs the extensive margin: a firm adjusts if and only if the gain from investing at its q, net of convex costs, exceeds ξ (Theorem 2).

3.2 Optimal investment at the firm level

Given Theorem 1, the optimal investment size is determined by the point where the marginal benefit and the cost cross, as formalized in the following corollary.

Corollary 1 (Forward-looking affine investment in intensive margin)**.**
Firm-level optimal investment conditional on adjustment is:

$$I^*(z;S) = \frac{\mathcal{G}(S)z + \mathcal{H}(S) - 1}{\mu} \tag{16}$$

Proof. See Appendix B. ■

Equivalently, $I^*(z;S) = (q(z;S)-1)/\mu$: investment is linearly increasing in Tobin's q with slope $1/\mu$, the standard q-theory investment rule (Abel, 1983). The optimal investment is strictly increasing in the idiosyncratic productivity z and strictly decreasing in the convex adjustment cost parameter μ. Moreover, it is independent of the existing capital stock, as shown in Corollary 1. Similarly, the extensive-margin decision—whether to invest—is also independent of a firm's capital stock, so idiosyncratic productivity is the sole relevant individual state variable. The following theorem characterizes the threshold rule of a firm's investment with respect to the productivity level.

Theorem 2 (The inaction region of productivity)**.**
A firm does not invest when the productivity z lies in the inaction interval $\Omega = [z_L, z_H]$ specified as below:

$$\Omega(S) := [z_L(S), z_H(S)] = \left[\frac{1-\sqrt{2\mu\xi}-\mathcal{H}(S)}{\mathcal{G}(S)}, \frac{1+\sqrt{2\mu\xi}-\mathcal{H}(S)}{\mathcal{G}(S)}\right] \tag{17}$$

The intuition is transparent. A firm invests when the net present value of adjustment—the product of marginal lifetime profit $q = \mathcal{G}z + \mathcal{H}$ and optimal investment $I^* = (q-1)/\mu$—exceeds the fixed cost ξ. Substituting the optimal investment into the indifference condition yields a quadratic equation in the threshold productivity $\bar{z}$:

$$(\mathcal{G}(S)\bar{z} + \mathcal{H}(S) - 1)^2 = 2\mu\xi. \tag{18}$$

This equation crosses zero twice: the two roots $\bar{z} = \left(1 \pm \sqrt{2\mu\xi} - \mathcal{H}\right)/\mathcal{G}$ define z_L and z_H, and all productivity levels between them characterize the inaction region. The left-hand side is the squared deviation of q from the replacement cost; the right-hand side is the fixed cost scaled by the convex cost parameter. A firm adjusts when this deviation is large enough—in either direction—to justify paying ξ.

The inaction band is not firm-specific and varies over the business cycle. Substituting $q = \mathcal{G}z + \mathcal{H}$ into Theorem 2 yields a time-invariant representation in q-space:

Corollary 2 (The inaction q-band)**.**
A firm does not invest when its Tobin's q lies in the inaction q-band:

$$q(z;S) \in \Omega_q := [1 - \sqrt{2\mu\xi},\ 1 + \sqrt{2\mu\xi}]. \tag{19}$$

The q-band is symmetric around $q = 1$ with constant half-width $\sqrt{2\mu\xi}$.

Proof. By Theorem 1, $q(z;S) = \mathcal{G}(S)z + \mathcal{H}(S)$. Evaluating at the thresholds of Theorem 2 gives $q(z_H;S) = 1 + \sqrt{2\mu\xi}$ and $q(z_L;S) = 1 - \sqrt{2\mu\xi}$, both independent of S. ■

The q-band has three notable properties. First, it provides a transparent interpretation: a firm invests only when $|q-1| > \sqrt{2\mu\xi}$, i.e., when the net present value of adjustment justifies the fixed cost. The half-width $\sqrt{2\mu\xi}$ reveals an interaction—the convex cost μ amplifies the effective barrier from the fixed cost ξ. Second, the q-band is *invariant* to S: while the productivity thresholds $[z_L(S), z_H(S)]$ shift over the business cycle, the q-thresholds remain fixed, so dynamics operate by moving firms' q-values across a constant band. Third, the q-band reveals the source of business-cycle asymmetry (Section 5.3): since $\mathcal{G}$ and $\mathcal{H}$ are pro-cyclical, negative shocks push a mass of firms into the band from above, while positive shocks push firms into the low-density tail.

An increase in the adjustment friction parameters μ and ξ widens the inaction band, holding the general equilibrium effect fixed. Therefore, our analytical framework guides parameter identification by sharply distinguishing the roles of the two adjustment cost parameters: μ affects both intensive-margin investment and the inaction band width, while ξ affects only the extensive margin (inaction).

Figure 1 illustrates the optimal investment and the inaction band. The idiosyncratic productivity distribution remains unchanged over the business cycle. However, the unconstrained optimal investment function $I^*(z;S)$ and the inaction band shift with aggregate conditions, shaping nonlinear aggregate investment fluctuations.

Figure 1: The optimal firm-level investment as a function of productivity

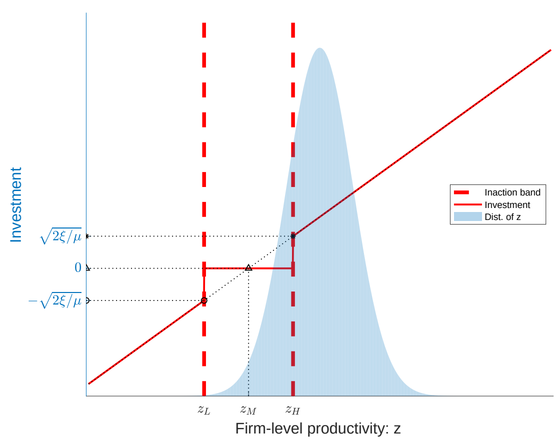

Notes: The thick line plots the constrained-optimal investment $\widehat{I}(z;S)$ as a function of idiosyncratic productivity z. The flat segment corresponds to the inaction band $\Omega(S)$, where the fixed adjustment cost makes zero investment optimal. The bell-shaped curve illustrates the stationary density of z. Parameters are set to the calibrated values in Section 6.

The inaction band leads to missing (dis-)investments that create a gap between frictionless and frictional aggregate investment. We define this wedge as the investment gap. It varies along two dimensions: quantity and quality. The length of the inaction band and the mass of the firms belonging to the band determine the quantity margin. The productivity midpoint z_M of the inaction band determines the quality margin.

Corollary 3 (The inaction midpoint)**.**

The midpoint of the inaction band z_M can be characterized as

$$z_M(S) = \frac{1 - \mathcal{H}(S)}{\mathcal{G}(S)}, \tag{20}$$

and the unconstrained optimal investment at the midpoint is zero: $I^*(z_M; S) = 0$.

Proof. See Appendix B. ■

Combining the theoretical results, the firm-level frictional investment $\widehat{I}(z; S)$ is

$$\widehat{I}(z; S) = \begin{cases} \frac{\mathcal{G}(S)z + \mathcal{H}(S) - 1}{\mu} & \text{if } z \notin \Omega(S), \text{ with probability } 1 - \varphi \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

We close the firm-level analysis by linking micro-level investment to macroeconomic dynamics through the inaction band. Macroeconomic conditions affect the firm-level incentive to invest through $\mathcal{G}$ and $\mathcal{H}$. In particular, under the fixed adjustment cost, $\mathcal{G}$ and $\mathcal{H}$ shape the dynamics of the inaction band, producing two distinct effects. The quantity effect operates through the (de-)synchronization of investment timings across firms. The quality effect operates through the selection driven by the endogenous choice of inaction.

Proposition 1 (Synchronization: quantity margin)**.**
If a shift in the aggregate productivity shock leads to $\Delta\mathcal{G}(S) > 0$, then the inaction band shrinks.

Proof. It is immediate from $\frac{\partial |z_H - z_L|}{\partial \mathcal{G}(S)} < 0$. ■

Intuitively, a higher $\mathcal{G}$ steepens the investment function, compressing the productivity range over which inaction is optimal. A larger fraction of firms simultaneously adjusts capital, amplifying the aggregate response to any given shock. A complementary channel operates through the location of the band:

Proposition 2 (Aggregate-level selection effect: quality margin)**.**
If a shift in the aggregate productivity shock leads to $\Delta\mathcal{G}(S) > 0$ and $\Delta\mathcal{H}(S) > 0$, then the inaction midpoint decreases.

Proof. See Appendix B. ■

Since $z_M = (1 - \mathcal{H})/\mathcal{G}$, the midpoint falls when both coefficients rise, shifting the band toward lower-productivity firms. The marginal firms newly entering investment are therefore of lower average quality—a negative selection effect that partially offsets the quantity gain from synchronization. Figure 2 illustrates

Figure 2: Shift of inaction bands and the firm-level investment policy

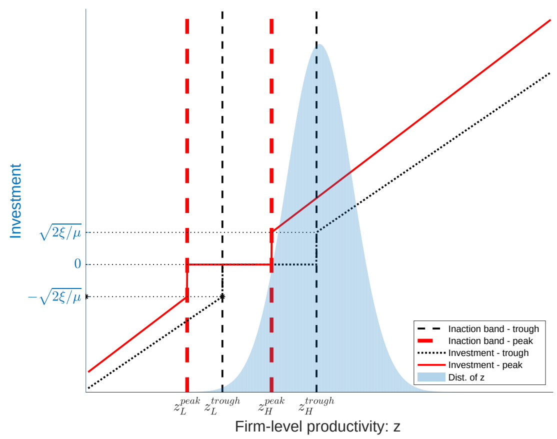

Notes: The solid and dashed lines plot the constrained-optimal investment in peak and trough aggregate states, respectively. The inaction band widens and shifts rightward in the trough, consistent with Propositions 1 and 2. The bell-shaped curve is the stationary density of z, which is invariant to the aggregate state. Parameters are calibrated as in Section 6.

the constrained optimal investment at the firm level for peak and trough periods.[10] These two channels create an amplification mechanism. In a boom, the band shrinks and shifts leftward: more firms invest, but marginal entrants are of lower productivity (negative selection). In a recession, the band widens and shifts rightward, concentrating inaction among better firms while fewer invest. This counter-cyclical selection means recessions are disproportionately damaging: the economy loses investment from the firms that would benefit most from reallocation. Later, under a fully analytical setup, we show that $d\mathcal{G}(S)/dZ > 0$ and $d\mathcal{H}(S)/dZ > 0$.

4 Equilibrium aggregation

4.1 Aggregate characterization

We decompose the recursive competitive equilibrium into analytically tractable components. The gap structure captures full global dynamics and permits an analytic treatment of the generalized impulse responses in Section 5.3.

Definition 1 (Misallocation)**.**

[10]The functions are based on the calibrated parameters in Section 6.

Investment gap Δ_I and output gap Δ_Y are defined as

$$\Delta_I(S) := I^{NoFixed}(S) - I(S), \quad \Delta_Y(S_{-1}, S) := Y^{NoFixed}(S_{-1}, S) - Y(S), \tag{22}$$

where

$$Y^{NoFixed}(S_{-1}, S) = \int F(zZ(k_{-1}(1-\delta) + I^*(s_{-1}; S_{-1})); S) d\Phi_{-1} \tag{23}$$

$$Y(S) = \int F(zZk; S) d\Phi. \tag{24}$$

The subscript -1 indicates a one-period lag. I^* is the firm-level frictionless optimal investment. These gaps arise solely from the presence of the fixed adjustment cost, which we interpret as friction-induced misallocation.

Definition 2 (Moments within the inaction band)**.**
$M_n(S)$ is defined as the n^{th} (non-central) moment of idiosyncratic productivity z over the inaction band $\Omega(S)$.

$$M_n(S) := \int_{z \in \Omega(S)} z^n d\Phi_z \tag{25}$$

Based on these definitions, we can analytically characterize the aggregate investment I and the frictionless benchmark $I^{NoFixed}$ as in the following proposition:

Proposition 3 (Aggregate investment characterization)**.**
The frictionless and frictional aggregate investments $(I^{NoFixed}, I)$ are

$$I^{NoFixed}(S) = \frac{1-\varphi}{\mu}\mathcal{G}(S)\left(\frac{\bar{\epsilon}}{1-\rho} - z_M(S)\right), \quad I(S) = I^{NoFixed}(S) - \Delta_I(S), \tag{26}$$

where the misallocation term $\Delta_I(S)$ is

$$\Delta_I(S) = \frac{1-\varphi}{\mu}\left(\mathcal{G}(S)M_1(S) + \left(\mathcal{H}(S) - 1\right)M_0(S)\right). \tag{27}$$

Proof. See Appendix B. ■

Proposition 3 shows that the gap Δ_I between aggregate investment and the frictionless benchmark can be expressed in closed form. It depends on the marginal-benefit terms $\mathcal{G}$ and $\mathcal{H}$ and the moments of the productivity distribution within the inaction region. Similarly, we analytically characterize the output gap in the following proposition:

Proposition 4 (Output gap characterization)**.**
The output gap Δ_Y is characterized as

$$\Delta_Y(S_{-1}, S) := O(S)Z(\rho\Delta_{II}(S_{-1}) + \bar{\epsilon}\Delta_I(S_{-1})), \tag{28}$$

where $\Delta_I(S)$ is as in Proposition 3 and the sorting gap is

$$\Delta_{II}(S) := \frac{1-\varphi}{\mu}\left(\mathcal{G}(S)M_2(S) + (\mathcal{H}(S) - 1)\, M_1(S)\right). \tag{29}$$

Proof. See Appendix B. ■

The output gap decomposes into a sorting gap and a level gap, whose macroeconomic implications are analyzed in Section 4.3:

$$\Delta_Y = O(S)Z\left(\underbrace{\rho\Delta_{II}(S_{-1})}_{\text{sorting gap}} + \underbrace{\bar{\epsilon}\Delta_I(S_{-1})}_{\text{level gap}}\right), \quad \bar{\epsilon} = 1-\rho, \tag{30}$$

The level gap contributes positively to Δ_Y when the inaction band sits left of the peak. Inactive firms near the upper threshold are high-productivity firms whose foregone positive investment dominates. The sorting gap captures how inaction weakens the capital-productivity covariance, increasing Δ_Y. Over the business cycle, both directions of sorting distortion are realized.[11]

4.2 Analytical equilibrium characterization

Despite firm-level heterogeneity and nonconvex adjustment costs, the entire RCE is characterized by a three-dimensional aggregate state: capital K, the capital-productivity covariance $cov(k, z)$, and aggregate TFP Z. Under GHH utility with linear consumption, all equilibrium objects admit closed-form steady-state expressions.

Proposition 5 (Analytical equilibrium characterization)**.**
Given $\mathcal{G}$,$\mathcal{H}$, $\sigma_\epsilon \geq 0$, and $\xi \geq 0$, the RCE allocations are completely characterized by a

[11] The comparison is against the frictionless benchmark; there are no externalities that disturb the first welfare theorem.

vector of aggregate variables S:

$$S = [K, cov(k, z), Z]. \tag{31}$$

Proof. See Appendix B. ∎

Aggregate output depends on the capital stock, the capital-productivity covariance $cov(k, z)$, and the net-of-labor aggregate productivity $O(S) := \Pi(1; S) - w(S)\Pi_w(1; S)$:

$$Y(S) = \underbrace{O(S)}_{\text{Net-of-labor TFP}} Z(K + cov(k, z)) \tag{32}$$

$$O(S) = \Pi(1; S) - w(S)\Pi_w(1; S). \tag{33}$$

Aggregate investment deviates from the frictionless benchmark by Δ_I, and aggregate consumption is the residual after deducting investment and total adjustment costs:

$$I(S) = I^{NoFixed}(\mathcal{G}, \mathcal{H}) - \Delta_I(\mathcal{G}, \mathcal{H}; \sigma_\epsilon, \xi) \tag{34}$$

$$I^{NoFixed}(\mathcal{G}, \mathcal{H}) = \frac{1-\varphi}{\mu} (\mathcal{G} + \mathcal{H} - 1) \tag{35}$$

Aggregate consumption $C(S)$ is the residual of output after deducting investment and adjustment costs $AC(S)$:

$$C(S) = Y(S) - I(S) - AC(S) \tag{36}$$

$$AC(S) = AC^{NoFixed}(S) - \tfrac{1}{2} (\mathcal{G}\Delta_{II} + (\mathcal{H} - 1)\Delta_I) \tag{37}$$

$$AC^{NoFixed}(S) = \tfrac{\mu}{2(1-\varphi)} (I^{NoFixed})^2 + \tfrac{1-\varphi}{2} \tfrac{\mathcal{G}^2 \sigma_\epsilon^2}{\mu(1-\rho^2)} \tag{38}$$

The capital-productivity covariance evolves under the tension between natural decay and endogenous sorting:

$$cov(k, z)' = \rho(1 - \delta) cov(k, z) \tag{39}$$

$$+ \rho \left[\tfrac{1-\varphi}{\mu} \tfrac{\sigma_\epsilon^2}{1-\rho^2} \mathcal{G} + \overbrace{\Delta_I(\mathcal{G}, \mathcal{H}; \sigma_\epsilon, \xi)}^{\text{sorting enhancing}} - \overbrace{\Delta_{II}(\mathcal{G}, \mathcal{H}; \sigma_\epsilon, \xi)}^{\text{sorting gap}} \right] \tag{40}$$

$$K' = K(1 - \delta) + I(S) \tag{41}$$

Since $cov(k, z)$ enters (32) as a first-order determinant of effective capital, the covariance dynamics in (40) directly shape aggregate output. The sorting gap Δ_{II} captures misallocation as inactive firms' productivity drifts while capital remains fixed; the level gap Δ_I acts as a sorting-enhancing force by concentrating reallocation at the extremes. Under our calibration $\Delta_{II}^{ss} > \Delta_I^{ss}$, so misallocation dominates; their differential response to shocks drives the nonlinear dynamics in Section 4.3.

The sufficient statistic $S = [K, cov(k, z), Z]$ from Proposition 5 compresses the infinite-dimensional distribution into three aggregates, making the global nonlinear analysis of Section 4.3 tractable.

Finally, the system is closed by the recursive marginal benefit coefficients $\mathcal{G}(S)$ and $\mathcal{H}(S)$ characterized in (14)–(15).

Stationary Equilibrium Analysis To facilitate the analytical characterization, we assume GHH utility with infinite Frisch elasticity, which pins the wage to a constant $w(S) = \eta$. This fixes the aggregate coefficient $O(S) = O^{ss} = \frac{\Pi(1)}{\alpha}$ and mutes wage-driven general equilibrium effects. To ensure stationarity, we assume $Z = 1$.

$$S^{Repss} = \left[\tfrac{1-\varphi}{\delta\mu}\left(\mathcal{G}^{ss} + \mathcal{H}^{ss} - 1\right),\ 0,\ 1\right] \tag{42}$$

$$S^{NoFixedss} = \left[\tfrac{1-\varphi}{\delta\mu}\left(\mathcal{G}^{ss} + \mathcal{H}^{ss} - 1\right),\ \tfrac{(1-\varphi)\rho\sigma_\epsilon^2}{\mu(1-\rho(1-\delta))(1-\rho^2)}\mathcal{G}^{ss},\ 1\right] \tag{43}$$

$$S^{ss} = \left[\tfrac{1-\varphi}{\delta\mu}\left(\mathcal{G}^{ss} + \mathcal{H}^{ss} - 1\right) - \tfrac{\Delta_I^{ss}}{\delta},\ cov^{NoFixedss}(k, z) + \tfrac{\rho(\Delta_I^{ss} - \Delta_{II}^{ss})}{1-\rho(1-\delta)},\ 1\right] \tag{44}$$

where the covariance component uses $\bar{\epsilon}/(1-\rho) = 1$ by the normalization in Section 2, and

$$\mathcal{G}^{ss} = \Pi^{ss}(1)\left[\tfrac{\rho\beta\overline{Z}}{1-\rho\beta(1-\delta)}\right] \tag{45}$$

$$\mathcal{H}^{ss} = \Pi^{ss}(1)\left[\tfrac{\beta\bar{\epsilon}\overline{Z}}{(1-\rho\beta(1-\delta))(1-\beta(1-\delta))}\right]. \tag{46}$$

Equations (45)–(46) yield the first fully explicit characterization of $\mathcal{G}$ and $\mathcal{H}$. $\mathcal{G}^{ss}$ measures the lifetime return premium from persistent productivity; it is increasing in ρ. $\mathcal{H}^{ss}$ captures the value of future innovations orthogonal to current z; the extra factor $(1 - \beta(1-\delta))$ reflects that innovations accumulate one period later. In the non-stationary environment, $\mathcal{G}$ and $\mathcal{H}$ fluctuate with aggregate TFP through (14)–(15). Importantly, neither σ_ϵ nor σ_{agg} enters these expressions, cleanly separating the direct effect of TFP fluctuations from distributional effects of uncertainty, which operate through M_0, M_1, M_2.

The steady-state investment gap inherits a closed form from $\mathcal{G}^{ss}$, $\mathcal{H}^{ss}$, and the truncated moments (Proposition 6):

$$\Delta_I^{ss} = \frac{(1-\varphi)\mathcal{G}^{ss}}{\mu}\left[\left(\frac{\bar{\epsilon}}{1-\rho} - z_M^{ss}\right)\left(\Phi^{SN}(z_H^{ss*}) - \Phi^{SN}(z_L^{ss*})\right) - \frac{\sigma_\epsilon}{\sqrt{1-\rho^2}}\left(\phi^{SN}(z_H^{ss*}) - \phi^{SN}(z_L^{ss*})\right)\right] \tag{47}$$

where $z_j^{ss*} := (z_j^{ss} - \bar{\epsilon}/(1-\rho))/(\sigma_\epsilon/\sqrt{1-\rho^2})$ for $j \in \{H, L\}$ and $\bar{\epsilon} = 1-\rho$ (Section 2). The closed form reveals a clean factorization. The prefactor $(1-\varphi)\mathcal{G}^{ss}/\mu$ is the *intensive margin*: it measures how much a marginal unit of productivity is worth in lifetime profit, scaled by the convex cost parameter. This prefactor depends on persistence ρ and the discount rate β through $\mathcal{G}^{ss}$, but is entirely independent of the fixed cost ξ and idiosyncratic volatility σ_ϵ. All effects of ξ and σ_ϵ on misallocation are confined to the bracket—the *extensive margin*.

Within the bracket, two forces compete. The first term, $(\bar{\epsilon}/(1-\rho) - z_M^{ss})\big(\Phi^{SN}(z_H^{ss*}) - \Phi^{SN}(z_L^{ss*})\big)$, is the product of the inactive mass (the CDF difference) and the deviation of the inaction midpoint z_M^{ss} from the ergodic mean $\bar{\epsilon}/(1-\rho)$. When the band lies left of the peak ($z_M^{ss} < \bar{\epsilon}/(1-\rho)$), more inactive firms have above-average productivity and would invest positively under the frictionless rule; their missing investment raises Δ_I^{ss}. The second term, involving the PDF difference $\phi^{SN}(z_H^{ss*}) - \phi^{SN}(z_L^{ss*})$, corrects for threshold asymmetry—structurally isomorphic to a Heckman selection correction (Heckman, 1979)—reflecting that the densities at the upper and lower thresholds need not be equal.

This factorization organizes the macroeconomic analysis in Section 4.3. Because $\mathcal{G}^{ss}$ is isolated from ξ and σ_ϵ, each type of shock operates through a distinct channel in (47). An interest rate change shifts $\mathcal{G}^{ss}$ through discounting *and* moves the thresholds through the quadratic equation, so both the prefactor and the bracket respond. The gap ratio Δ_I^{ss}/I^{ss} then amplifies the aggregate semi-elasticity (Section 5.1). A dispersion shock σ_ϵ leaves $\mathcal{G}^{ss}$ entirely untouched and operates exclusively through the bracket, narrowing the standardized thresholds while amplifying the PDF prefactor $\sigma_\epsilon/\sqrt{1-\rho^2}$. This generates the non-monotone uncertainty response (Section 5.2). An aggregate TFP shock shifts the inaction midpoint z_M^{ss} through $\mathcal{H}$, translating the bracket through the bell-shaped density. The CDF and PDF terms respond asymmetrically because the density is curved rather than uniform, producing asymmetric impulse responses and negative skewness (Section 5.3). The sorting gap Δ_{II}^{ss} is defined analogously with M_1^{ss}, M_2^{ss} replacing M_0^{ss}, M_1^{ss} (Appendix A); Jensen's inequality ensures $\Delta_{II}^{ss} > \Delta_I^{ss}$ generically, driving the impossibility result of Section 4.3.

The transition from a representative-firm to a heterogeneous-firm framework reveals that output depends critically on the capital-productivity covariance $cov(k, z)$. Frictionless heterogeneity generates sorting gains ($C^{NoFixedss} > C^{Repss}$ when $2/(\alpha\beta) > 1$). Fixed costs introduce a tension: inaction reduces capital (Δ_I^{ss}) but enhances selection at adjustment thresholds. The output gap Δ_Y^{ss} in (48) captures the net effect:

$$\Delta_Y^{ss} = O^{ss}\overline{Z}\Big(\underbrace{(1-\rho)\,\Delta_I^{ss}}_{\text{level channel}} + \underbrace{\rho\,\Delta_{II}^{ss}}_{\text{sorting channel}} \Big). \tag{48}$$

The sorting gap enters with persistence weight ρ and the level gap with weight $1-\rho$. The level channel reflects foregone capital accumulation ($\bar{\epsilon} = 1-\rho$); the sorting channel reflects capital-productivity misallocation through Δ_{II}^{ss}, depressing $cov(k, z)$ and output via (32). Both terms are positive, so $\Delta_Y^{ss} > 0$. Higher ρ increases the sorting weight, though it also narrows the band through $\mathcal{G}^{ss}$; the net effect on Δ_Y^{ss} is positive under our calibration. Higher ξ widens the band unambiguously (holding $\mathcal{G}^{ss}$ fixed). Under our calibration $\Delta_{II}^{ss} > \Delta_I^{ss}$ (Jensen's inequality), so the sorting channel dominates.[12] The full welfare derivation is in Appendix B.

Analytical gaps The presence of fixed adjustment costs generates systematic deviations from the frictionless benchmark, which are characterized as investment, sorting, and output gaps. By utilizing the properties of the standard normal distribution, these gaps can be derived in closed form as functions of the normalized inaction cutoffs z_j^*, $j \in \{H, M, L\}$.

Proposition 6 (Analytical gaps)**.**
Under Gaussian productivity, the truncated moments M_0, M_1, M_2 in Propositions 3–4 admit closed-form expressions in terms of the standard normal CDF Φ^{SN} and PDF ϕ^{SN}, evaluated at the normalized cutoffs $z_j^* := \frac{z_j - \frac{\bar{\epsilon}}{1-\rho}}{\sqrt{\frac{\sigma_\epsilon^2}{1-\rho^2}}}$, $j \in \{H, M, L\}$. The explicit formulae

[12] The sorting-enhancing force of Δ_I^{ss} in (40) partially offsets Δ_{II}^{ss}, but the net effect $\Delta_I^{ss} - \Delta_{II}^{ss} < 0$ is negative.

are in Appendix A.

Proof. See Appendix B. ■

The closed-form expressions reveal two properties. First, each gap is proportional to $I^{NoFixed}$, so misallocation scales with the size of the economy. Second, the gaps depend on the cross-section only through the truncated moments (M_0, M_1/M_0, and variance within the band), providing low-dimensional sufficient statistics. We formally investigate the neutrality conditions in Section 4.3.

Using $\mathcal{H}(S) - 1 = -\mathcal{G}(S)z_M(S)$, the gaps from Propositions 3–4 can be equivalently written as

$$\Delta_I(S) = \tfrac{1-\varphi}{\mu}\mathcal{G}(S) \cdot M_0(S) \cdot \left(\tfrac{M_1(S)}{M_0(S)} - z_M(S)\right), \tag{49}$$

$$\Delta_{II}(S) = \tfrac{1-\varphi}{\mu}\mathcal{G}(S) \cdot M_1(S) \cdot \left(\tfrac{M_2(S)}{M_1(S)} - z_M(S)\right). \tag{50}$$

The selection-bias terms $M_1/M_0 - z_M$ and $M_2/M_1 - z_M$ are structurally isomorphic to the Heckman selection correction discussed after (47). Jensen's inequality ($M_2M_0 > M_1^2$) ensures $\Delta_{II} > \Delta_I$ generically.

Role of distributional assumptions The results form a hierarchy. The first tier—Theorems 1–2, the gap decomposition (Proposition 3), and the impossibility theorem (Proposition 8)—is distribution-free, requiring only CRS and AR(1) productivity. The second tier—elasticity amplification, the non-monotone uncertainty effect, and GIRF asymmetry—requires a continuous unimodal density but not Gaussianity. Only the third tier—the closed-form M_0, M_1, M_2 expressions and the Gaussian-derivative approximation—exploits the specific properties of the normal distribution. Under alternative specifications, the decomposition structure and qualitative mechanisms are preserved; only the explicit CDF/PDF expressions change.

Fully analytical nonstationary setup We now characterize the fully analytical recursive competitive equilibrium under GHH preferences with risk aversion $\sigma \to 0$:

$$\lim_{\sigma \to 0} \frac{(c - \eta l_H)^{1-\sigma}}{1-\sigma} \tag{51}$$

From the CRS Cobb-Douglas production function, the marginal profit of an effective capital stock satisfies $\Pi(1) := \alpha\left(\frac{1-\alpha}{\eta}\right)^{\frac{1-\alpha}{\alpha}}$. The coefficients $\mathcal{G}$ and $\mathcal{H}$ of the marginal benefit $q = \mathcal{G}z + \mathcal{H}$ (Theorem 1) have the following explicit form:

$$\mathcal{G}(S) = \mathcal{G}(Z) = \frac{\alpha\left(\frac{1-\alpha}{\eta}\right)^{\frac{1-\alpha}{\alpha}} \beta\rho\left[\rho_{agg}Z + \frac{\bar{\epsilon}_{agg}}{1-\rho\beta(1-\delta)}\right]}{1-\rho_{agg}\rho\beta(1-\delta)} \tag{52}$$

$$\mathcal{H}(S) = \mathcal{H}(Z) = \frac{\alpha\left(\frac{1-\alpha}{\eta}\right)^{\frac{1-\alpha}{\alpha}} \beta\left[\rho_{agg}Z\bar{\epsilon} + \frac{\bar{\epsilon}_{agg}\bar{\epsilon}(1-\rho\rho_{agg}\beta^2(1-\delta)^2)}{(1-\beta(1-\delta))(1-\rho\beta(1-\delta))}\right]}{(1-\rho_{agg}\rho\beta(1-\delta))(1-\rho_{agg}\beta(1-\delta))} \tag{53}$$

Notably, the volatility parameter σ_{agg} does not affect the future marginal benefit terms $\mathcal{G}$ and $\mathcal{H}$, while the persistence of the aggregate productivity process ρ_{agg} affects both terms. In this setup, we can analytically sign the derivatives of $\mathcal{G}$ and $\mathcal{H}$ with respect to the TFP variations:

$$\frac{\partial\mathcal{G}(Z)}{\partial Z} > 0, \quad \frac{\partial\mathcal{H}(Z)}{\partial Z} > 0. \tag{54}$$

This confirms the pro-cyclicality of $q = \mathcal{G}z + \mathcal{H}$ asserted in Section 3 and delivers the promised analytical sign for the comparative statics in Propositions 1–2.

Importantly, $\mathcal{G}$ and $\mathcal{H}$ are affine functions of the aggregate productivity Z. Therefore, the frictionless aggregate investment $I^{NoFixed}$ is also an affine function of Z:

$$I^{NoFixed} = \frac{1-\varphi}{\mu}\mathcal{G}(Z) + \frac{1-\varphi}{\mu}\left(\mathcal{H}(Z) - 1\right) \tag{55}$$

$$= aZ + b, \tag{56}$$

for some $a, b \in \mathbb{R}$. The analytical forms of a and b are in Appendix B.

Given $\mathcal{G}$ and $\mathcal{H}$, the inaction band and midpoint follow from Theorem 2 and Corollary 3. Aggregate investment $I(Z)$ is then analytically characterized by

$$I(Z) = (aZ+b)\left(1 - \underbrace{\Phi^{SN}(z^*_H(Z)) - \Phi^{SN}(z^*_L(Z))}_{\Delta_I\text{: Quantity}} - \underbrace{\frac{\phi^{SN}(z^*_H(Z)) - \phi^{SN}(z^*_L(Z))}{z^*_M(Z)}}_{\Delta_I\text{: Asymmetry}}\right). \tag{57}$$

This is the *nonstationary* generalization of the steady-state gap in (47). At the steady state ($Z = 1$), the thresholds and midpoint freeze at their ergodic values and the expression collapses to (47). Away from the steady state, the aggregate state Z enters through the midpoint $z_M^*(Z)$, making both correction terms functions of the business cycle. The *quantity margin*, $\Phi^{SN}(z_H^*) - \Phi^{SN}(z_L^*)$, is the mass of inactive firms. A wider band or denser region of the productivity distribution increases this mass, amplifying the aggregate investment shortfall. The *asymmetry margin*, $(\phi^{SN}(z_H^*) - \phi^{SN}(z_L^*))/z_M^*$, captures the threshold density imbalance—precisely the inverse Mills ratio of the Heckman selection correction identified in (47). The key difference from the stationary case is that $z_M^*(Z)$ now varies with Z: negative TFP shocks push z_M^* toward zero (the distribution peak), inflating both margins simultaneously.

Figure 3: Aggregate investment function

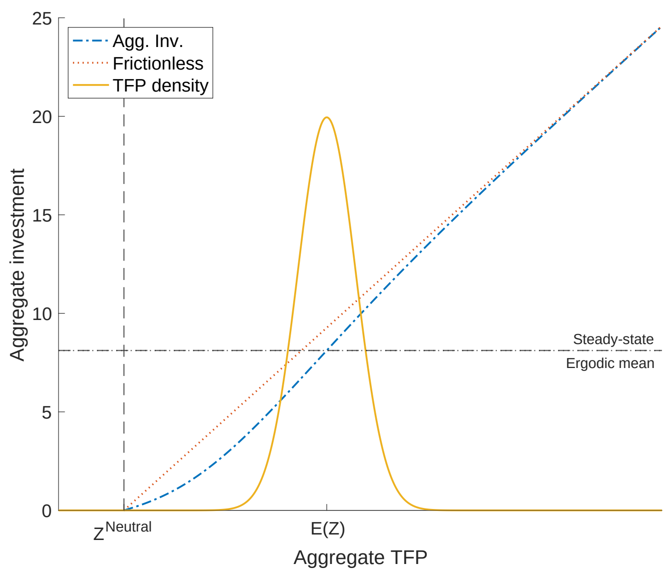

Notes: Dash-dotted line: frictional aggregate investment $I(Z)$. Dotted line: frictionless benchmark $I^{NoFixed}(Z)$. Solid curve: density of aggregate productivity Z. The vertical distance between the two investment lines is the investment gap $\Delta_I(Z)$. $Z^{Neutral}$ marks the productivity level at which $\Delta_I = 0$. Linear utility. Parameters are calibrated as in Section 6.

Figure 3 plots frictional aggregate investment (dash-dotted) and the frictionless benchmark (dotted) against aggregate productivity, with the TFP density overlaid (solid). The vertical distance between the two lines is the investment gap $\Delta_I(Z)$—the nonstationary counterpart of (47). At $Z = 1$ (the steady state), this distance equals the steady-state gap in (47). As Z varies, the distributional bracket responds through $z_M^*(Z)$ while the intensive-margin prefactor $\mathcal{G}(Z)/\mu$ comoves with TFP. Near $Z = 1$, the investment gap Δ_I is decreasing and *convex*—the source of the volatility tax and negative skewness in Section 4.3. The gap widens beyond the neutrality point $Z^{Neutral}$, where the symmetry condition (61) holds exactly. Below $Z^{Neutral}$, the frictionless benchmark is negative and misallocation operates in reverse. These features—convexity of the gap, widening gap, and the neutrality point—are the geometric foundations of the elasticity amplification (Section 5.1), non-monotone uncertainty (Section 5.2), and business-cycle asymmetry (Section 5.3). The neutrality point has the analytical form:

A closed-form neutrality point $Z^{Neutral}$ exists at which the investment gap is exactly zero. Moreover, the sign of Δ_I is equivalent to $Z > Z^{Neutral}$, $I^{NoFixed} > 0$, and $z^*_M < 0$—so the investment gap is positive whenever the economy is producing. The explicit expression and proof are in Appendix B.

While (47) provides a closed form at the steady state, the nonstationary gap involves CDF and PDF evaluations at both thresholds, which depend nonlinearly on the aggregate state through $z^*_M(Z)$. The following approximation reduces this to a single sufficient statistic while preserving the nonlinearity.

Proposition 7 (Gaussian-derivative approximation)**.**
In the analytical equilibrium, aggregate investment and the output gap admit the approximation

$$I(Z) \approx \underbrace{(aZ + b)}_{I^{NoFixed}(Z)} - \underbrace{c\,(\phi^{SN})'\left(\frac{d}{Z+e} + f\right)}_{\approx \Delta_I(Z)}, \tag{58}$$

$$\Delta_Y(Z_{-1}, Z) \approx gZ \cdot c\,(\phi^{SN})'\left(\frac{d}{Z_{-1}+e} + f\right), \tag{59}$$

for constants $a, \ldots, g \in \mathbb{R}$ determined by structural parameters, where $(\phi^{SN})'(x) = -x\phi^{SN}(x)$ is the derivative of the standard normal density. The investment gap Δ_I is evaluated at the lagged aggregate TFP Z_{-1}, reflecting that the output gap at date t arises from investment decisions at $t-1$. Two approximation constants ζ_1, ζ_2 are calibrated at the stationary equilibrium. The approximation error is $O(\xi)$ as $\xi \to 0$.

Proof. See Appendix B. ■

The Gaussian derivative $(\phi^{SN})'(x) = -x\phi^{SN}(x)$ is positive for $x < 0$ and negative for $x > 0$. Since $d/(Z+e) + f$ is decreasing and convex in Z, crossing zero near $Z^{Neutral}$, the investment gap inherits three properties: positivity when $Z > Z^{Neutral}$, a hump shape with a unique interior maximum, and concavity near its peak—generating the elasticity amplification, non-monotone uncertainty, and business-cycle asymmetry in Section 4.3. Under our calibration, the approximation error averages 0.47% for $I(Z)$ and 0.83% for Δ_Y. Coefficient expressions are in Appendix B.

4.3 Geometry of the gaps

In this section, we study under which conditions the misallocation terms become relevant for equilibrium dynamics. We establish an impossibility result: the investment and output gaps cannot vanish simultaneously—firm-level inaction *cannot* be neutral for aggregate investment and output. This is a structural consequence of Jensen's inequality, not a calibration-dependent finding. Models omitting fixed adjustment costs are qualitatively, not merely quantitatively, different from frictional economies. Recall from (49) that $\Delta_I(S) = \frac{1-\varphi}{\mu}\mathcal{G}(S)\, M_0(S)\big(M_1(S)/M_0(S) - z_M(S)\big)$, where the inaction midpoint z_M is characterized in Corollary 3. From this we deduce the following equivalence:

$$\Delta_I = 0 \iff \underbrace{\frac{M_1(S)}{M_0(S)}}_{\text{Average productivity among the inactive}} - \underbrace{z_M(S)}_{\text{Inaction midpoint}} = 0, \qquad (60)$$

when $M_0 > 0$. The investment gap is zero precisely when the average productivity among inactive firms equals the inaction midpoint. A sufficient condition is that Φ be symmetric around z_M on $\Omega(S)$, in which case

$$I(S) = I^{NoFixed}(S). \qquad (61)$$

The equivalence (60) makes the mechanism transparent. Symmetry forces the conditional mean M_1/M_0 to equal z_M: firms with productivity above the midpoint (who would invest more under the frictionless rule) are exactly offset by firms below it (who would invest less). Because $I^*(z;S)$ is linear in z (Corollary 1) and zero at z_M (Corollary 3), the positive and negative foregone investments integrate to zero (see Appendix B for the formal argument). This sharpens Elsby and Michaels (2019), who derive the same symmetry condition via analytical approximations; our result is exact and embeds in general equilibrium. The condition operates independently of GE smoothing, distinguishing it from Khan and Thomas (2008).

Crucially, symmetry neutralizes Δ_I but *cannot* simultaneously neutralize Δ_{II}: the productivity-weighted average M_2/M_1 exceeds M_1/M_0 by Jensen's inequality whenever the within-band distribution has positive variance. This asymmetry between the two gaps is the source of the impossibility result below.

The sorting channel in (30) satisfies an analogous equivalence:

$$\Delta_{II} = 0 \iff \underbrace{\frac{M_2(S)}{M_1(S)}}_{\text{Weighted average productivity among the inactive}} - \underbrace{z_M(S)}_{\text{Inaction midpoint}} = 0. \quad (62)$$

when $M_1 > 0$.[13] The economic content of (62) is as follows. The ratio M_2/M_1 is the average productivity among inactive firms when each firm is weighted by its own productivity level—high-z firms count more because they carry more capital and hence contribute more to the capital-productivity covariance. Symmetry of Φ around z_M equalizes the *unweighted* average M_1/M_0 with z_M (neutralizing Δ_I), but the productivity weighting tilts M_2/M_1 above z_M: high-productivity inactive firms are overrepresented in the sorting margin precisely because they are high-productivity. This is a direct consequence of Jensen's inequality ($M_2 M_0 > M_1^2$ whenever the within-band distribution has positive variance), and it ensures that the sorting gap Δ_{II} is strictly harder to neutralize than the investment gap Δ_I. The economic implication is that inaction distorts the allocation of *existing* capital across firms—the sorting channel in (40)—even when its effect on the *level* of aggregate investment vanishes. The Jensen ranking has a strong consequence. Note that the condition $\Delta_I \geq 0$ is not restrictive: it is equivalent to $I \geq 0$ (the inaction band lying left of the distribution peak), which is the only empirically relevant case—an inaction rate of approximately 30% (Zwick and Mahon, 2017) and rare disinvestment (Winberry, 2021) firmly place the economy in this region. The left-of-peak placement is also necessary for the steady state to be well defined:

[13] The weighted average productivity interpretation is directly from the following formulation:

$$\frac{M_2(S)}{M_1(S)} = \frac{\int_{\Omega(S)} z^2 d\Phi}{\int_{\Omega(S)} z d\Phi} = \int_{\Omega(S)} z \underbrace{\left(\frac{z}{\int_{\Omega(S)} z d\Phi}\right)}_{\text{Productivity-based weight}} d\Phi \quad (63)$$

if the band lies to the right of the peak, stationary aggregate investment becomes negative and the transversality condition is violated. In the closed form (47), $\Delta_I^{ss} > 0$ holds whenever z_M^{ss} lies below the ergodic mean $\bar{\epsilon}/(1-\rho)$, so the CDF term dominates.

Proposition 8 (Impossibility of macroeconomic neutrality)**.**
For any aggregate state S with $M_0(S) > 0$ and $\Delta_I(S) \geq 0$ (equivalently, $I(S) \geq 0$):

(i) $\Delta_{II}(S) > 0$.

(ii) When $\Delta_I(S) > 0$, all three gaps are simultaneously strictly positive: $\Delta_I(S) > 0$, $\Delta_{II}(S) > 0$, and $\Delta_Y(S, S') > 0$ for any subsequent state S'.

In particular, the investment and output gaps cannot vanish simultaneously.

Proof. See Appendix B. ■

The impossibility rests on Jensen's hierarchy of truncated moments:

$$\frac{M_2(S)}{M_1(S)} > \frac{M_1(S)}{M_0(S)}, \tag{64}$$

which holds strictly whenever the within-band distribution is non-degenerate. Part (i): the condition $\Delta_I \geq 0$ ensures $M_1/M_0 \geq z_M$; combining with the ranking gives $M_2/M_1 > z_M$, so $\Delta_{II}(S) > 0$. Even at the boundary of perfect investment neutrality ($M_1/M_0 = z_M$, $\Delta_I = 0$), the ranking forces $\Delta_{II} > 0$ and hence $\Delta_Y > 0$ through the sorting channel. Part (ii): when $\Delta_I > 0$, the ranking implies $M_2/M_1 > M_1/M_0 > z_M$, so all three gaps are simultaneously strictly positive via $\Delta_Y(S, S') = O(S')Z'(\rho\,\Delta_{II}(S) + \bar{\epsilon}\,\Delta_I(S))$. This is structural, not a calibration artifact. Quantitatively, $\Delta_I/I = 8.22\%$ and $\Delta_Y/Y \approx 12\%$ at the steady state (Section 6).

Aggregate implications The left-of-peak location delivers three interrelated implications through the gap structure (47)–(48). First, $\Delta_I^{ss} > 0$: average investment is lower than median investment, implying a negatively skewed firm-level investment distribution.[14] The inaction band traps predominantly low-productivity firms that would make small positive investments; their missing contributions

[14] Because the idiosyncratic productivity distribution is symmetric and the frictionless investment rule is linear, $I^{\text{NoFixed}}(S)$ coincides with the median firm's investment.

pull the average below the median. This *cross-sectional* negative skewness translates into *time-series* negative skewness in aggregate investment over the business cycle.

Second, $\Delta_{II}^{ss} > 0$: inactive firms near the upper threshold carry below-target capital relative to their productivity, weakening $cov(k,z)$ through (40). Jensen's inequality ($M_2 M_0 > M_1^2$) ensures this productivity-weighted misallocation exceeds the simple average, so $\Delta_{II}^{ss} > \Delta_I^{ss}$.

Third, the *output gap* $\Delta_Y^{ss} > 0$ inherits contributions from both channels: lost capital accumulation (the $(1-\rho)\Delta_I^{ss}$ term) and capital misallocation (the $\rho\Delta_{II}^{ss}$ term). These operate through the sufficient statistic $S = [K, cov(k,z), Z]$: Δ_I^{ss} reduces the ergodic K directly, while the gap differential $\Delta_I^{ss} - \Delta_{II}^{ss} < 0$ depresses $cov(k,z)$ and hence output through (32).

5 Macroeconomic implications

5.1 Frictional investment semi-elasticity

A central empirical puzzle is the *magnitude* of firm-level investment responses to tax incentives and interest rate changes. Zwick and Mahon (2017) document that the investment semi-elasticity to tax incentives is approximately 7.2, substantially larger than standard convex-adjustment-cost models predict. Winberry (2021) shows that convex costs alone *dampen* aggregate elasticity, and Koby and Wolf (2020) find similar attenuation. A key unresolved question is: *can adjustment frictions amplify rather than dampen the aggregate elasticity?* Our framework delivers an affirmative answer. The mechanism is visible in the closed form (47): since $\mathcal{G}^{ss} \propto 1/R$, an interest rate increase simultaneously reduces the intensive-margin prefactor and shifts the inaction thresholds, compressing both margins of aggregate investment. The investment gap Δ_I^{ss} opens an extensive-margin amplification channel absent from convex-cost models. To analyze the elasticity to a one-shot change in the interest rate, we express $\mathcal{G}^{ss}$ and $\mathcal{H}^{ss}$ (see Section 4.2) as functions of the contemporaneous gross interest rate R and stationary interest rate $1/\beta$.[15]

[15] Specifically,

$$\mathcal{G}^{ss} = \frac{\Pi^{ss}(1)}{R}\left[\frac{\rho\overline{Z}}{1-\rho\beta(1-\delta)}\right], \quad \mathcal{H}^{ss} = \frac{\Pi^{ss}(1)}{R}\left[\frac{\bar{\epsilon}\overline{Z}}{(1-\rho\beta(1-\delta))(1-\beta(1-\delta))}\right]$$

The derivation is available in Appendix B.

We then characterize the sensitivity at both firm and aggregate levels.

Proposition 9 (Firm-level investment semi-elasticity)**.**
The firm-level investment elasticity at the steady state is

$$\frac{\partial log(\widehat{I}^{ss}(z))}{\partial R} = \begin{cases} -\frac{1}{R}\left(1+\frac{1}{\mu \widehat{I}^{ss}(z)}\right), & \textit{if the firm stays unconstrained} \\ -1, & \textit{if the firm becomes constrained} \\ 0, & \textit{if the firm stays constrained.} \\ +\infty, & \textit{if the firm becomes unconstrained.} \end{cases} \tag{65}$$

Proof. See Appendix B. ■

The proposition highlights stark heterogeneity in firm responses. Unconstrained firms exhibit finite negative elasticity through the standard intensive margin. Firms deep within the inaction band have zero elasticity. However, firms at the edges of the inaction band exhibit infinite local elasticity, as a small change in R triggers a discrete policy switch.

Proposition 10 (Aggregate investment semi-elasticity)**.**
The aggregate investment semi-elasticities for the baseline and frictionless benchmark at the steady state are given by the following, where z_H^*, z_L^* denote the standardized thresholds and $\sigma_z = \sigma_\epsilon / \sqrt{1-\rho^2}$:

$$\frac{\partial \log(I^{ss})}{\partial R} \tag{66}$$

$$= \left(1+\frac{\Delta_I}{I^{ss}}\right)\frac{\partial \log(I^{NoFixed,ss})}{\partial R} - \frac{1}{I^{ss}\sigma_z}\left(I^*(z_H^{ss})\phi^{SN}(z_H^{*,ss})\frac{\partial z_H^{ss}}{\partial R} - I^*(z_L^{ss})\phi^{SN}(z_L^{*,ss})\frac{\partial z_L^{ss}}{\partial R}\right)$$

$$\frac{\partial \log(I^{NoFixed,ss})}{\partial R} = -\frac{1}{R}\left(1+\frac{R(1-\rho)(1-\beta(1-\delta))}{\Pi^{ss}(1)\overline{Z\epsilon}-R(1-\rho)(1-\beta(1-\delta))}\right) \tag{67}$$

Proof. See Appendix B. ■

The semi-elasticity decomposes into two channels. A scaling effect: since $I^{ss} < I^{NoFixed,ss}$, the same absolute shift produces a larger percentage response. And a threshold effect from the movement of inaction boundaries with R. Whether these channels reinforce or offset determines a central question in the literature: does the frictional economy respond more or less to interest rates than the frictionless benchmark?

Corollary 4 (Elasticity amplification)**.**

If $\Delta_I^{ss} \geq 0$, the frictional aggregate investment semi-elasticity at the steady state is greater than the frictionless benchmark:

$$\left|\frac{\partial \log(I^{ss})}{\partial R}\right| > \left|\frac{\partial \log(I^{NoFixed,ss})}{\partial R}\right| \tag{68}$$

Proof. See Appendix B. ■

Fixed adjustment costs amplify the aggregate interest-rate elasticity of investment, resolving the puzzle raised by Winberry (2021) and Koby and Wolf (2020): while convex adjustment costs dampen elasticity by spreading adjustment smoothly, *fixed* costs amplify it by concentrating responses among active firms and creating discrete threshold crossings.

The mechanism operates through two reinforcing channels, decomposed in Figure 4:

$$\frac{\partial \log(I^{ss})}{\partial R} = \underbrace{\left(1 + \frac{\Delta_I^{ss}}{I^{ss}}\right)\frac{\partial \log(I^{NoFixed,ss})}{\partial R}}_{\text{amplified intensive margin}<0} - \underbrace{\frac{1}{I^{ss}}\frac{\partial \Delta_I^{ss}}{\partial R}}_{\text{extensive margin}>0} \tag{69}$$

The first term scales up the frictionless response by $1 + \Delta_I^{ss}/I^{ss}$: inaction reduces total investment, so the same absolute shift produces a larger percentage change. The second term is the extensive margin, whose structure is transparent from (47). Since $\mathcal{G}^{ss} \propto 1/R$, a higher R reduces the prefactor (the lifetime return to productivity). Simultaneously, the quadratic equation widens the inaction band, increasing the inactive mass and midpoint deviation in the distributional bracket. Both forces raise $\partial\Delta_I^{ss}/\partial R > 0$, reinforcing the intensive margin. Higher fixed cost ξ amplifies both channels: a wider band raises the gap ratio Δ_I^{ss}/I^{ss} (first term) and increases the density of firms near the thresholds who switch discretely when R changes (second term). Economies with larger ξ therefore exhibit stronger interest-rate amplification, a prediction testable across industries with different adjustment frictions.

The amplification extends beyond investment: through $S = [K, cov(k,z), Z]$, an interest rate increase both reduces K directly and alters the balance between Δ_I and Δ_{II} in (40), propagating to output via (32)—a GE feedback absent from convex-cost models.

Figure 4: Firm-level investment response to the interest rate change

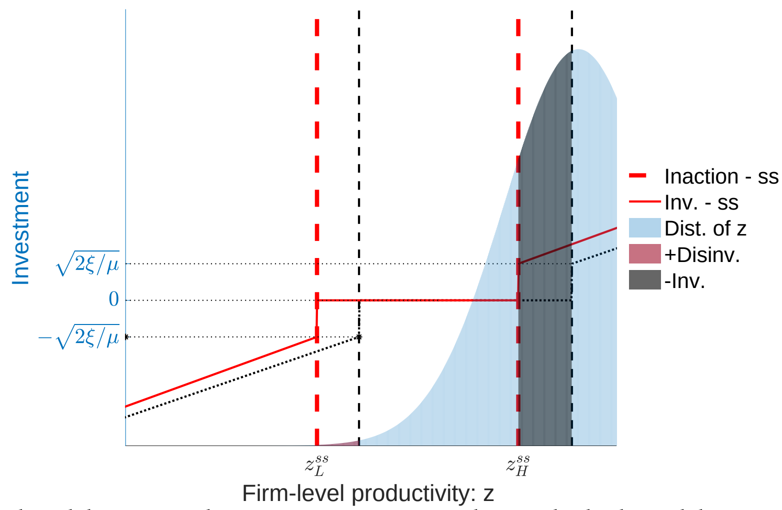

Notes: Solid red line: steady-state investment policy. Black dotted line: investment policy after an interest rate increase. The dark grey area ("−Inv.") captures firms switching from positive investment to inaction. The pink area ("+Disinv.") captures firms switching from inaction to disinvestment. The bell-shaped curve is the stationary density of z. Parameters are calibrated as in Section 6.

5.2 Non-monotonic impact of dispersion shocks

The macroeconomic effects of firm-level uncertainty shocks are well documented—through wait-and-see dynamics (Bloom, 2009), stochastic volatility (Fernández-Villaverde et al., 2011), investment dispersion (Bachmann and Bayer, 2014), state-dependent adjustment (Vavra, 2014), and quantitative heterogeneous-firm models (Bloom et al., 2018)—but analytically isolating the transmission mechanism has remained out of reach. The difficulty is fundamental: in standard models, uncertainty effects are mediated by the entire cross-sectional distribution of firms, so existing decompositions are necessarily simulation-based and model-specific.

Our framework resolves this by exploiting the clean separation visible in (47): σ_ϵ is entirely absent from the intensive-margin prefactor $\mathcal{G}^{ss}/\mu$, so all uncertainty effects are confined to the distributional bracket. More precisely, because frictionless investment is linear in productivity (Corollary 1), a mean-preserving spread in σ_ϵ has *zero* effect on the frictionless benchmark. All uncertainty effects operate exclusively through $\Delta_I(\sigma_\epsilon)$. The resulting relationship is *non-monotonic*: a critical volatility σ_ϵ^* maximizes aggregate misallocation. At low volatility, few firms populate the inaction region; at high volatility, the distribution flattens and within-band asymmetry vanishes. Our calibrated σ_ϵ lies near this peak.

The investment gap is explicitly a function of this volatility:

$$\Delta_I(\sigma_\epsilon) = \frac{1-\varphi}{\mu}\mathcal{G}M_0(\sigma_\epsilon)\left(\frac{M_1(\sigma_\epsilon)}{M_0(\sigma_\epsilon)} - z_M\right) \tag{70}$$

The clean separation follows from the linearity of the frictionless investment rule:

Proposition 11 (Irrelevance of the uncertainty shock on frictionless investments)**.**
Given $\mathcal{G}$ and $\mathcal{H}$, for any aggregate state S,

$$\frac{\partial I^{NoFixed}(S)}{\partial \sigma_\epsilon} = 0. \tag{71}$$

Proof. Immediate from Corollary 1: $I^{NoFixed}(S)$ is affine in z with coefficients that depend on S but not on σ_ϵ. ■

All σ_ϵ effects are therefore compositional, operating through the mass and selection of firms within the inaction band without household precautionary motivation—a channel unavailable in models where uncertainty also affects investment through precautionary motives or nonlinear profit functions. As volatility rises from zero, firms enter the band and Δ_I grows. But the gap cannot grow indefinitely: the distribution eventually flattens, eroding within-band asymmetry. This implies a "worst-case" uncertainty:

Proposition 12 (Existence of the critical value)**.**
Suppose $\mu_z \notin [z_L, z_H]$, i.e., the unconditional mean of productivity lies outside the inaction band.[16] *If $\Delta_I > 0$ for some σ_ϵ, then there exists σ_ϵ^* such that*

(i) $\Delta_I(\sigma_\epsilon^*) \geq \Delta_I(\sigma_\epsilon)$ *for all* $\sigma_\epsilon \in (0, \infty)$*, and*

(ii) $\frac{\partial \Delta_I(\sigma_\epsilon)}{\partial \sigma_\epsilon}\Big|_{\sigma_\epsilon^*} = 0.$

Proof. See Appendix B. ■

The non-monotonicity arises from competing *quantity* and *quality* channels: the mass of inactive firms M_0 versus their average productivity relative to z_M. At σ_ϵ^*, these forces exactly offset. Both Δ_I and Δ_Y exhibit inverted U-shapes (see Appendix B), and our calibrated σ_ϵ lies near the peak, where micro-level uncertainty has its maximum distorting impact. Because the sufficient statistic S depends on

[16] This holds whenever the half-width of the inaction band $\sqrt{2\mu\xi}/\mathcal{G}$ is less than $|\mu_z - z_M|$. The condition is satisfied at the calibration.

σ_ϵ only through the gaps, the non-monotonicity in $\Delta_I(\sigma_\epsilon)$ propagates directly to $cov(k,z)$ through (40) and hence to all aggregate quantities.

5.3 Nonlinear aggregate investment dynamics

This section establishes three global results beyond the local elasticity analysis: (i) asymmetric impulse responses (negative shocks $\Rightarrow$ larger contractions); (ii) state dependence (recessions $\Rightarrow$ heightened fragility); and (iii) excess volatility and negative skewness. All are absent from the frictionless benchmark. They arise from a single cross-sectional asymmetry visible in $I(Z)$ (Figure 3): the inaction band sits left of the distribution's peak, making $\Delta_I(Z)$ convex in the empirically relevant region. The vertical distance between the linear frictionless benchmark and the nonlinear frictional function *is* $\Delta_I(Z)$; the three results follow directly from this gap's geometry. In the closed form (47), a TFP shock shifts the inaction midpoint z_M through $\mathcal{H}$, translating the distributional bracket through the Gaussian density. The CDF and PDF terms respond asymmetrically because the density is bell-shaped rather than uniform. Negative shocks push the band into the high-density region, widening the gap; positive shocks push it into the low-density tail.

To capture the inherently global nonlinearity of the extensive margin, we consider an instantaneous TFP shock $\Delta_{\mathcal{H}}$ that shifts the marginal benefit to $\mathcal{G}(S)z + (\mathcal{H}(S) + \Delta_{\mathcal{H}})$. By expanding the state to $(S, \Delta_{\mathcal{H}})$, we express aggregate investment and the investment gap as functions of the shock magnitude. The generalized impulse response function (GIRF) formalizes this comparison.

Definition 3 (Generalized impulse response function (GIRF))**.**
Generalized impulse response functions $g(S, \Delta_{\mathcal{H}})$ and $g^{NoFixed}(S, \Delta_{\mathcal{H}})$ of aggregate investment and frictionless benchmark are defined as

$$g(S, \Delta_{\mathcal{H}}) := I(S, \Delta_{\mathcal{H}}) - I(S, 0) \tag{72}$$

$$g^{NoFixed}(S, \Delta_{\mathcal{H}}) := I^{NoFixed}(S, \Delta_{\mathcal{H}}) - I^{NoFixed}(S, 0). \tag{73}$$

A key advantage of our analytical framework is that the GIRF admits an exact decomposition into intensive and extensive margins.

$$g(S,\Delta_{\mathcal{H}}) = \overbrace{\frac{1-\varphi}{\mu}\Delta_{\mathcal{H}}(1-M_0(S,\Delta_{\mathcal{H}}))}^{\text{Intensive margin}} - \overbrace{(\Delta_I(S,\Delta_{\mathcal{H}}) - \Delta_I(S,0))}^{\text{Extensive margin}} \tag{74}$$

$$g^{NoFixed}(S,\Delta_{\mathcal{H}}) = \frac{1-\varphi}{\mu}\Delta_{\mathcal{H}} \tag{75}$$

The intensive margin is scaled by the active mass $(1-M_0)$, creating an endogenous recession multiplier as M_0 rises. The extensive margin captures discrete threshold crossings—it equals the *change in the investment gap* $\Delta_I(S,\Delta_{\mathcal{H}}) - \Delta_I(S,0)$. In booms, Δ_I narrows as the band shifts into the low-density tail, so the extensive margin reinforces the intensive margin. In recessions, Δ_I widens sharply as the band enters the high-density region, and the extensive margin amplifies the contraction beyond what the intensive margin alone would produce. Together, this yields asymmetric and state-dependent propagation:

Proposition 13 (Asymmetric and state-dependent GIRF)**.**
The generalized impulse response of aggregate investment is asymmetric between positive and negative shocks and state-dependent:

(i) $|g(S,\Delta_{\mathcal{H}})| < |g(S,-\Delta_{\mathcal{H}})|$*, for inaction bands to the left of the peak of the firm-level productivity distribution.*[17]

(ii) *S affects g through the intensive margin:* $\frac{1-\varphi}{\mu}\Delta_{\mathcal{H}}(1-M_0(S,\Delta_{\mathcal{H}}))$*, decreasing in* M_0.

(iii) *S affects g through the extensive margin:* $\Delta_I(S,\Delta_{\mathcal{H}}) - \Delta_I(S,0)$.

Proof. See Appendix B. ∎

The asymmetry is rooted in density asymmetry: a negative shock shifts the band rightward toward the high-density region, trapping many firms; a positive shock shifts it into the low-density tail. The mechanism is self-reinforcing—recessions widen the band near the peak, raising M_0 and vulnerability to further shocks, while booms shrink the band into the tail. In contrast, the frictionless benchmark is symmetric:

[17] The opposite statement is true for inaction bands to the right of the peak of the firm-level productivity distribution.

Corollary 5 (Symmetric and state-independent frictionless GIRF)**.**
The generalized impulse response of frictionless aggregate investment is symmetric between positive and negative shocks and state-independent:

(i) $|g^{NoFixed}(S, \Delta_{\mathcal{H}})| = |g^{NoFixed}(S, -\Delta_{\mathcal{H}})|$

(ii) $|g^{NoFixed}(S_1, \Delta_{\mathcal{H}})| = |g^{NoFixed}(S_0, \Delta_{\mathcal{H}})|$*, for all* S_0, S_1.

Proof. See Appendix B. ■

The investment gap $\Delta_I(Z)$ is hump-shaped—large near the gap-maximizing TFP level Z^{Max} and vanishing in the tails. In the empirically relevant region, the economy operates predominantly above Z^{Max}, where $\Delta_I(Z)$ is decreasing and convex in Z (flattening toward zero). Since $I^{NoFixed}(Z)$ is affine, $\mathbb{E}I(Z) - I(\mathbb{E}Z) = -[\mathbb{E}\Delta_I(Z) - \Delta_I(\mathbb{E}Z)]$. Jensen's inequality applied to the locally convex Δ_I implies that TFP fluctuations raise the ergodic mean gap above its steady-state level, depressing ergodic mean investment—a "volatility tax" operating through the investment gap, not risk aversion or incomplete markets. The mechanism is visible in Figure 3: negative TFP shocks push toward the thick center of the hump (large gap increase) while positive shocks push into the thin tail (small gap decrease). Losses exceed gains on average. We formally sign the moments:

Proposition 14 (The nonlinearity in the business cycle)**.**
If $I^{NoFixed}(Z) > 0$ *and* $\mathbb{P}(Z < Z^{Max})$ *is sufficiently small, where* Z^{Max} *is the aggregate productivity level at which* $\Delta_I(Z)$ *is maximized,*

(i) $\mathbb{E}I(Z) < I(\mathbb{E}Z)$

(ii) $s.d.(I) > s.d.(I^{NoFixed})$

(iii) $skewness(I) < skewness(I^{NoFixed})$

Proof. See Appendix B. ■

The local concavity of $I(Z)$ in the empirically relevant region is the common source of all three results. Jensen's inequality directly yields part (i). The extensive margin creates excess volatility in part (ii) because $\Delta_I(Z)$ is negatively correlated with TFP: recessions widen the gap while booms narrow it, amplifying the variance of I beyond that of $I^{NoFixed}$. The same concavity generates part (iii): symmetric TFP fluctuations produce asymmetric investment responses, with larger contractions than expansions, tilting the ergodic distribution leftward. Through the sufficient statistic $S = [K, cov(k,z), Z]$, these moment distortions compound. The depressed mean lowers ergodic K; excess volatility destabilizes $cov(k,z)$ through (40); and the left skew ensures that recessionary misallocation exceeds expansionary efficiency gains.

6 Quantitative assessment

6.1 Calibration

This section quantitatively assesses the theoretical predictions. We calibrate four parameters of the fully analytical model: the convex adjustment cost μ, the fixed adjustment cost ξ, the labor disutility parameter η, and the aggregate TFP shock volatility σ_Z. For the remaining parameters, we follow standard values in the firm dynamics literature. The unit period is a quarter. The calibration proceeds in two steps: (1) calibration of the frictional parameters (μ, ξ, η) based on the stationary equilibrium, and (2) calibration of the TFP volatility parameter (σ_Z) based on the stochastic equilibrium. The volatility is obtained after applying the HP filter to both the model and the data. Table 1 reports the level of the calibrated and fixed parameters.

Table 2 reports targeted and untargeted equilibrium moments. The upper panel shows data targets and model-implied counterparts; the last column specifies data sources. The lower panel reports untargeted moments implied by the model. In the stationary equilibrium with the calibrated parameters, the investment gap is 8.22% of the equilibrium investment.

Under the calibrated parameters, 28.8% of firms are inactive (close to the data target of 30.2%), and the inaction region lies to the left of mean idiosyncratic productivity. That is, roughly 29% of firms do not invest because their productivity falls within the inaction band.[18]

The model also performs well on untargeted moments. The untargeted investment semi-elasticity of 9.81 is close to the empirical estimate of 7.20 from Zwick and Mahon (2017). The model's investment rate dispersion—$s.d.(\hat{I}/k)$ of 13.95—is close to the Compustat value of 16.00, capturing firm-level lumpiness. These untargeted moments provide out-of-sample validation of the framework.

[18] There is a negligible portion of firms making negative investment (0.02%), which is also consistent with the data.

Table 1: Calibrated and fixed parameters

Parameters	Description	Value
Stationary eq. parameters		
α	Capital share	0.3300
δ	Depreciation rate	0.0250
ρ	Firm-level productivity shock persistence	0.9750
σ_ϵ	Firm-level productivity shock volatility	0.0220
μ	Convex adjustment cost*	0.6950
ξ	Fixed adjustment cost*	0.0055
$\overline{\epsilon}$	Firm-level productivity shock average	$1-\rho$
η	Labor disutility*	1.8940
φ	Financial shock	0.1000
β	Household's discount factor	0.9900
Aggregate shock parameters		
ρ_{agg}	Aggregate productivity shock persistence	0.9500
$\overline{\epsilon}_{agg}$	Aggregate productivity shock persistence	$1-\rho_{agg}$
σ_Z	Aggregate productivity shock volatility*	0.0108

Notes: Parameters with an asterisk (*) are calibrated. The remaining parameters are set to standard values in the firm dynamics literature. The unit period is a quarter.

Table 2: Targeted and untargeted equilibrium moments

	Model	Data	Reference
Target moments (annualized)			
Inaction portion (%)	28.80	30.20	Zwick and Mahon (2017)
$mean(\hat{I}/k)$	10.12	10.40	Zwick and Mahon (2017)
$s.d.(log(Y))$	1.46	1.44	NIPA data (HP-filtered)
Labor hour	0.30	0.33	Bureau of Labor Statistics
Untargeted moments (annualized)			
$s.d.(\hat{I}/k)$	13.95	16.00	Zwick and Mahon (2017)
$s.d.(\hat{I})/mean(\hat{I})$	6.16	6.50	Compustat data
Investment semi-elasticity	9.81	7.20	Zwick and Mahon (2017)
Δ_I/I (%)	8.22		

Notes: The upper panel reports moments targeted in the calibration; the lower panel reports untargeted moments implied by the model. Data sources are listed in the rightmost column.

6.2 Quantitative assessment of the theoretical predictions

The semi-elasticity increases from 7.83 (frictionless) to 9.81 (frictional)—a 25% amplification driven by the extensive margin. The GIRFs exhibit strong asymmetry: a 2-s.d. shock produces a 34% investment deviation when negative versus

25% when positive. Figure 5 plots impulse responses for 2-s.d. TFP shocks, displaying stark asymmetry and nonlinearity in magnitude.

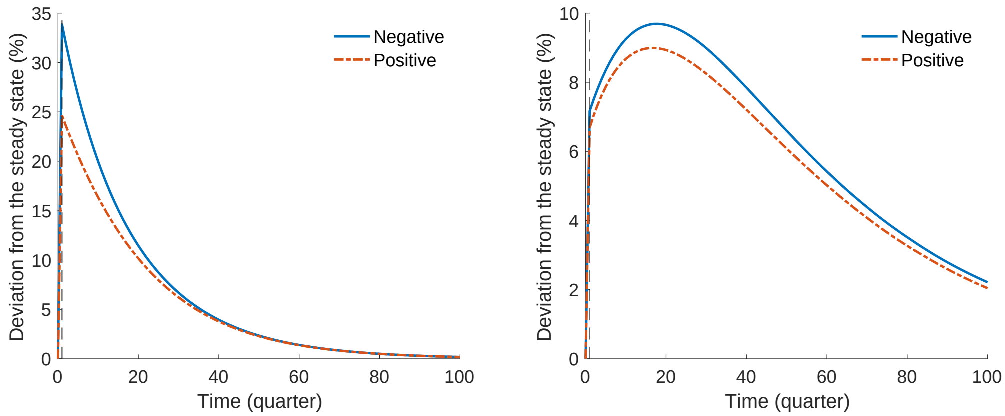

(a) I: 2-$s.d.$ shock (b) Y: 2-$s.d.$ shock

Figure 5: Asymmetric (absolute) impulse responses to an aggregate TFP shock

Notes: Solid line: response to a negative TFP shock. Dash-dotted line: response to a positive shock of equal magnitude. Left: aggregate investment. Right: aggregate output. All responses are in absolute value. Parameters are calibrated as in Section 6.

Furthermore, the response is highly state-dependent. The response to a 2-s.d. shock increases from 34% at steady state to 44% in a downturn (TFP=0.95), because the mass of firms near adjustment thresholds increases during recessions.

The general equilibrium effect of log utility The fully analytical model assumes linear utility, abstracting from inter-temporal smoothing. Khan and Thomas (2008) showed that general equilibrium effects can dampen the nonlinearity from firm-level inaction. Our Proposition 8 establishes that the nonlinearity cannot be perfectly neutralized, but the degree of attenuation is a quantitative question. We assess this by computing the general equilibrium under log utility using the repeated transition method of Lee (2025).

The log utility significantly flattens investment functions: the slope decreases by 20% (30% for the frictionless benchmark).[19] Nevertheless, the investment gap still widens during downturns, confirming that the nonlinearity from firm-level inaction survives the GE smoothing force.

Business cycle statistics Table 3 reports business cycle statistics from a 5,000-period simulation of the fully analytical model. Frictional dynamics significantly

[19]Aggregate TFP realizations beyond three standard deviations require prohibitively long simulations and are excluded.

amplify investment volatility by 29% relative to the frictionless benchmark (from 4.6% to 5.9%), with a modest increase in output volatility as well. The frictional model also generates more negative skewness—investment skewness is 32% more negative than the frictionless benchmark—consistent with the nonlinear investment gap mechanism illustrated in Figure 3.

Table 3: Business cycle statistics

	Unfiltered				HP-filtered			
	$log(I)$		$log(Y)$		$log(I)$		$log(Y)$	
Fixed cost	Yes	No	Yes	No	Yes	No	Yes	No
Volatility	0.1444	0.1123	0.1023	0.1018	0.0593	0.0460	0.0146	0.0144
Skewness	-0.3727	-0.2476	0.0780	0.0803	-0.3059	-0.2313	-0.0830	-0.0812
Kurtosis	3.0805	2.9795	2.7201	2.7208	3.3527	3.2259	3.1458	3.1464

Notes: All variables are in logs. Aggregate investment and output are HP filtered with a smoothing parameter of 1600. "Frictional" denotes the baseline model with the fixed adjustment cost; "Frictionless" denotes the benchmark without it. The model is simulated for 5,000 periods under linear utility.

Table 4 reports business cycle statistics under log utility. As under linear utility, fixed adjustment costs significantly amplify volatility and negative skewness in aggregate allocations, confirming that general equilibrium smoothing does not neutralize frictional nonlinearity.

Table 4: Business cycle statistics: log utility

	Unfiltered				HP-filtered			
	$log(I)$		$log(Y)$		$log(I)$		$log(Y)$	
Fixed cost	Yes	No	Yes	No	Yes	No	Yes	No
Volatility	0.1115	0.0881	0.1105	0.1100	0.0206	0.0163	0.0153	0.0152
Skewness	-0.1624	-0.0685	-0.1297	-0.1280	-0.1209	-0.0914	-0.0912	-0.0911
Kurtosis	2.7564	2.7604	2.5054	2.5053	3.1312	3.1096	3.0963	3.0961

Notes: Same specification as Table 3, except the household has log utility. The general equilibrium solution is computed using the repeated transition method of Lee (2025).

Figure 6 validates these predictions against U.S. data. Frictional investment (solid) is more volatile than the frictionless benchmark (dashed) in normalized terms (Panel a), and the gap Δ_I widens during downturns (Panel b). Panel (c) confirms this in the data: U.S. fixed investment diverges from its linear projection onto TFP during recessions, consistent with the nonlinearly increasing investment gap.

Figure 6: Investment dynamics: model and data

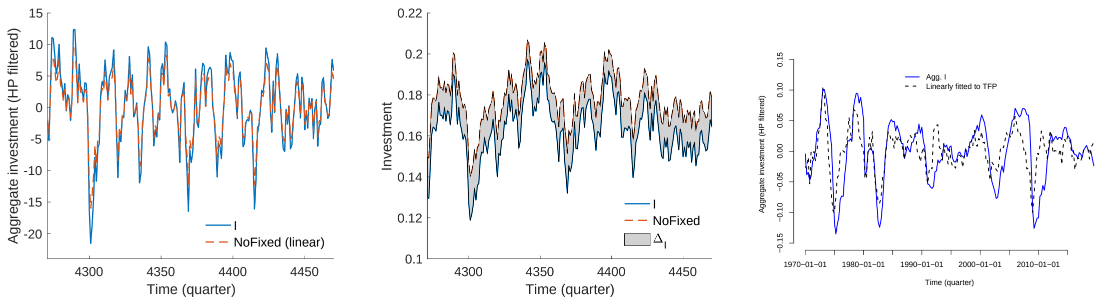

(a) Normalized (model) (b) Levels (model) (c) U.S. data

Notes: Panels (a)–(b): Solid line is the frictional model; dashed line is the frictionless benchmark. Panel (a) shows investment in % deviation from the mean; Panel (b) shows levels, where the vertical distance is Δ_I. Simulated for 5,000 periods under linear utility with calibrated parameters. Panel (c): Solid line is U.S. real private fixed investment (NIPA); dashed line is the linear projection onto utilization-adjusted TFP (Fernald, 2014). Both series HP filtered ($\lambda = 1600$). Shaded areas indicate NBER recession dates.

7 Concluding remarks

This paper develops a tractable general equilibrium framework that renders frictional firm dynamics analytically transparent. Closed-form solutions isolate three margins of aggregate misallocation—quantity, quality, and sorting—through which micro-level inaction shapes the business cycle. Firm-level inaction creates systematic wedges—investment and output gaps—that persist even under general equilibrium price adjustments.

Three insights emerge. First, in the empirically relevant region, all three gaps—investment, sorting, and output—are simultaneously strictly positive; joint neutrality is impossible by a Jensen's inequality ranking of truncated moments. Second, fixed costs amplify the economy's sensitivity to interest rates and generate a non-monotonic (inverted-U) relationship between firm-level uncertainty and aggregate misallocation, challenging the view that frictions simply dampen aggregate responses. Third, the extensive margin induces state dependence and asymmetry, making the economy more fragile during downturns.

Beyond frictional firm dynamics, the framework provides a methodological toolkit. By rendering heterogeneous-agent dynamics analytically transparent, it opens a path for future research to dissect the interplay between micro-level frictions and macro-level policy without sacrificing the clarity of closed-form theory.

Our paper opens several future research pathways. The elasticity amplification result connects directly to monetary policy transmission: embedding nominal rigidities would clarify how firm-level inaction shapes the investment channel of monetary policy. The closed-form misallocation characterization invites optimal-policy analysis—for instance, the design of investment tax credits that account for the extensive-margin response documented here. Finally, the sample-selection structure underlying the gap decomposition suggests testable predictions linking firm-level investment spikes and inaction rates to cross-sectional productivity moments.

References

Abel, A. B. (1983): "Optimal Investment Under Uncertainty," *The American Economic Review*, 73, 228–233.

Alvarez, F. and F. Lippi (2022): "The Analytic Theory of a Monetary Shock," *Econometrica*, 90, 1655–1680.

Alvarez, F. and N. L. Stokey (1998): "Dynamic Programming with Homogeneous Functions," *Journal of Economic Theory*, 82, 167–189.

Asker, J., A. Collard-Wexler, and J. De Loecker (2014): "Dynamic Inputs and Resource (Mis)Allocation," *Journal of Political Economy*, 122, 1013–1063.

Bachmann, R. and C. Bayer (2014): "Investment Dispersion and the Business Cycle," *American Economic Review*, 104, 1392–1416.

Bachmann, R., R. J. Caballero, and E. M. R. A. Engel (2013): "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model," *American Economic Journal: Macroeconomics*, 5, 29–67.

Baley, I. and A. Blanco (2021): "Aggregate Dynamics in Lumpy Economies," *Econometrica*, 89, 1235–1264.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

Caballero, R. J. and E. M. R. A. Engel (1991): "Dynamic (S, s) Economies," *Econometrica*, 59, 1659–1686.

——— (1993): "Heterogeneity and Output Fluctuations in a Dynamic Menu-Cost Economy," *The Review of Economic Studies*, 60, 95–119.

——— (1999): "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach," *Econometrica*, 67, 783–826.

Caballero, R. J., E. M. R. A. Engel, J. C. Haltiwanger, M. Woodford, and R. E. Hall (1995): "Plant-Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity*, 1995, 1.

Caplin, A. S. and D. F. Spulber (1987): "Menu Costs and the Neutrality of Money," *The Quarterly Journal of Economics*, 102, 703–726.

Cooper, R., J. Haltiwanger, and L. Power (1999): "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review*, 89, 32.

David, J. M. and V. Venkateswaran (2019): "The Sources of Capital Misallocation," *American Economic Review*, 109, 2531–67.

Doms, M. E. and T. Dunne (1998): "Capital Adjustment Patterns in Manufacturing Plants," *Review of Economic Dynamics*, 1, 409–429.

Elsby, M. W. and R. Michaels (2019): "Fixed adjustment costs and aggregate fluctuations," *Journal of Monetary Economics*, 101, 128–147.

Fernald, J. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," *Federal Reserve Bank of San Francisco, Working Paper Series*, 01–28.

Fernández-Villaverde, J., P. Guerrón-Quintana, J. F. Rubio-Ramírez, and M. Uribe (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101, 2530–2561.

Gourio, F. and A. K. Kashyap (2007): "Investment spikes: New facts and a general equilibrium exploration," *Journal of Monetary Economics*, 54, 1–22.

Hayashi, F. (1982): "Tobin's Marginal q and Average q: A Neoclassical Interpretation," *Econometrica*, 50, 213–224.

Heckman, J. J. (1979): "Sample Selection Bias as a Specification Error," *Econometrica*, 47, 153–161.

Hopenhayn, H. A. and E. C. Prescott (1992): "Stochastic Monotonicity and Stationary Distributions for Dynamic Economies," *Econometrica*, 60, 1387.

House, C. L. (2014): "Fixed costs and long-lived investments," *Journal of Monetary Economics*, 68, 86–100.

Hsieh, C.-T. and P. J. Klenow (2009): "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 124, 1403–1448.

Jovanovic, B. (2025): "Competitive Industry Dynamics under Constant Returns," *Journal of Political Economy*, forthcoming.

Khan, A. and J. K. Thomas (2003): "Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?" *Journal of Monetary Economics*, 50, 331–360.

——— (2008): "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics," *Econometrica*, 76, 395–436.

Koby, Y. and C. K. Wolf (2020): "Aggregation in Heterogeneous-Firm Models: Theory and Measurement," *Working Paper*, 66.

Lee, H. (2025): "Global Nonlinear Solutions in Sequence Space and the Generalized Transition Function," *Working paper*, 42.

Lucas, R. E. and E. C. Prescott (1971): "Investment Under Uncertainty," *Econometrica*, 39, 659–681.

Moll, B. (2014): "Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?" *American Economic Review*, 104, 3186–3221.

Restuccia, D. and R. Rogerson (2008): "Policy distortions and aggregate productivity with heterogeneous establishments," *Review of Economic Dynamics*, 11, 707–720.

Thomas, J. (2002): "Is Lumpy Investment Relevant for the Business Cycle?" *Journal of Political Economy*, 110, 508–534.

Tobin, J. (1969): "A General Equilibrium Approach To Monetary Theory," *Journal of Money, Credit and Banking*, 1, 15–29.

Vavra, J. (2014): "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *Quarterly Journal of Economics*, 129, 215–258.

Winberry, T. (2021): "Lumpy Investment, Business Cycles, and Stimulus Policy," *American Economic Review*, 111, 364–396.

Zwick, E. and J. Mahon (2017): "Tax Policy and Heterogeneous Investment Behavior," *American Economic Review*, 107, 217–248.