



Department of Economics

Athens University of Economics and Business

WORKING PAPER no. 13-2026

The Persistence of Corporate Overinvestment and Underinvestment

Dimitris K. Christopoulos and Elias Tzavalis

October 2026

The Working Papers in this series circulate mainly for early presentation and discussion, as well as for the information of the Academic Community and all interested in our current research activity.

The authors assume full responsibility for the accuracy of their paper as well as for the opinions expressed therein.

The Persistence of Corporate Overinvestment and Underinvestment [†]

Dimitris K. Christopoulos

Elias Tzavalis

Athens University of Economics and Business^{*}

October 2026

Abstract. This paper examines the persistence and asymmetric adjustment of abnormal corporate investment. Conventional residual-based measures cannot distinguish persistent deviations from transitory shocks. We develop a nonlinear latent-state framework that decomposes benchmark-investment residuals into persistent and transitory components and allows persistence to differ between positive and negative states. Using U.S. public firms, we find stronger local propagation of positive than negative abnormal-investment deviations. Regime-continuation probabilities are high and relatively similar across states, showing that local propagation and regime continuation capture distinct dimensions of investment adjustment. Persistence asymmetry is pronounced among smaller and less profitable firms, while its relation to cash flow varies across samples. Persistent positive deviations also predict greater subsequent sales of property, plant, and equipment than equally large negative deviations, providing evidence on the economic relevance of the latent state. Overall, the findings are consistent with agency, financing, and adjustment mechanisms rather than a simple liquidity-based explanation.

JEL Classification: C32, C58, D22, E22, G31.

Keywords: Corporate investment; Underinvestment; Overinvestment; Latent-state models; Asymmetric investment dynamics; Filtering; Regime-dependent dynamics.

[†]We thank Mong Shan Ee and George Leledakis for helpful comments and suggestions on an earlier version of the paper. The research project is implemented in the framework of the H.F.R.I. call “3rd Call for H.F.R.I.’s Research Projects to Support Faculty Members & Researchers” (H.F.R.I. Project Number: 25948).

^{*}Patission 76, 104 34 Athens, Greece; dchristop@aueb; etzavalis@aueb.gr.

1. Introduction

A central question in corporate investment is not only why firms invest above or below levels implied by observable fundamentals, but also how persistent these deviations are and whether their adjustment differs across positive and negative investment states. This issue is closely related to theories of financing constraints, agency conflicts, adjustment costs, and corporate investment dynamics. Limited access to external finance may delay the correction of negative investment deviations, whereas managerial incentives, free-cash-flow problems, weak external discipline, or costly commitments to ongoing projects may sustain positive deviations. Jensen (1986), Fazzari, Hubbard, and Petersen (1988), Kaplan and Zingales (1997), Richardson (2006), and Hennessy and Whited (2007) emphasize different financial and agency mechanisms that can affect corporate investment. These mechanisms suggest that abnormal investment may exhibit asymmetric persistence, but they do not imply a unique ordering between overinvestment and underinvestment.

Empirically, overinvestment and underinvestment are commonly measured using residuals from benchmark investment regressions. Following Richardson (2006), a positive residual is typically interpreted as overinvestment and a negative residual as underinvestment. Owing to its transparency and relatively modest data requirements, this approach has become widely used in the corporate-finance literature. Residual-based measures have been employed to study the effects of financial reporting quality, corporate governance, liquidity, ownership structure, and other firm characteristics on investment efficiency; see, among others, Harford, Mansi, and Maxwell (2008), Biddle, Hilary, and Verdi (2009), Jiang et al. (2018), and Ee et al. (2022).

A broader literature suggests, however, that corporate investment has an important dynamic component. Past investment contains substantial information about current investment, consistent with dynamic investment models linking investment decisions and firms' financial policies (see Bond and Meghir (1994)). Investment adjustment can also be nonlinear because fixed costs, irreversibility, and other non-convex adjustment costs can generate periods of limited activity followed by larger investment episodes (see, e.g., Caballero (1999), and Cooper and Haltiwanger (2006)). Financing frictions can further affect the timing of adjustment. Whited (2006), for example, shows that constraints on external finance influence

the intertemporal pattern of investment and can lengthen the time between major investment episodes. Bloom (2009) shows that uncertainty shocks may induce firms to postpone investment temporarily and subsequently resume it as uncertainty declines. The literature therefore suggests that corporate investment combines persistent adjustment with short-lived responses, providing an economic motivation for separating persistent from transitory abnormal investment.

Ignoring this dynamic component can have several consequences. A benchmark-regression residual may combine a persistent deviation from benchmark investment with temporary demand shocks, short-lived financing conditions, measurement noise, and other transitory disturbances. Treating the residual itself as the underlying investment state can therefore lead to incorrect classification when its contemporaneous sign differs from that of the persistent component. It also treats short-lived and persistent deviations in the same way and provides little information about how slowly or rapidly an investment gap is corrected. In addition, imposing a common adjustment process can conceal differences in the propagation of positive and negative deviations. Importantly, persistence in the magnitude of an abnormal-investment state is conceptually distinct from the probability that a firm remains in the same state. Finally, firm characteristics may affect not only the contemporaneous magnitude of abnormal investment but also the strength of its propagation over time.

The paper makes three main contributions. First, it documents a pronounced asymmetry in the dynamics of abnormal investment. Positive persistent deviations exhibit substantially stronger local propagation than negative deviations, with the estimated persistence coefficients implying markedly different conditional rates of adjustment. This ordering is consistent with mechanisms that can sustain investment once firms move above their benchmark, including managerial agency, weak external discipline, continuing access to internal resources, and costly continuation or reversal of ongoing investment programs. The reduced-form evidence does not, however, separately identify these mechanisms.

Second, the paper treats persistence as a distinct dimension of corporate investment behavior and separates local propagation from regime continuation. Our empirical results show that a pronounced asymmetry in local state propagation can coexist with relatively similar regime-continuation probabilities. Thus, the magnitude of an abnormal-investment deviation and

the probability of remaining on the same side of the benchmark capture distinct dimensions of corporate investment adjustment.

Third, the paper develops a nonlinear latent-state framework for identifying and estimating these dynamic differences. The benchmark-investment residual is decomposed into a persistent latent component and a transitory disturbance. Persistence is allowed to differ between positive and negative latent states, while regime membership is inferred from the full nonlinear filtering distribution rather than from the contemporaneous sign of the residual. Positive persistent deviations are referred to as overinvestment and negative persistent deviations as underinvestment, following Richardson (2006) and the subsequent investment-efficiency literature. These labels are used in a reduced-form sense: the model identifies persistent deviations from benchmark investment rather than structural welfare losses. Firm-level bootstrap inference accounts for the generated nature of the first-stage residuals. Complementary Monte Carlo evidence shows that dynamic filtering substantially improves latent-state recovery and regime classification relative to raw investment residuals, accurately recovers regime-specific persistence, and delivers well-sized and high-power inference on persistence asymmetry, even with relatively short within-firm histories.

The empirical results reinforce these contributions. Using a common nonlinear investment benchmark that accounts for long-run firm heterogeneity and industry-specific time variation, we strongly reject symmetric persistence and find a pronounced asymmetry in the local propagation of abnormal investment. Positive abnormal-investment deviations exhibit substantially stronger propagation than negative deviations, with conditional half-lives of approximately 8–9 years compared with about 2 years, respectively.

Regime continuation, however, exhibits a much smaller asymmetry. Continuation probabilities are high and relatively similar for overinvestment and underinvestment, with the positive regime displaying only slightly higher continuation probabilities in both samples. Thus, the substantial difference in the propagation of investment deviations is not accompanied by a comparable difference in the probability of remaining in the same regime. These findings highlight local propagation and regime continuation as distinct dimensions of corporate investment adjustment.

Persistence asymmetry is particularly pronounced among smaller and less profitable firms,

primarily because negative deviations exhibit weaker propagation in these groups, while positive-state persistence remains high. Its association with cash flow is less uniform across the two samples. These patterns suggest that firm characteristics associated with financing capacity are related to the dynamics of abnormal investment, but do not support a simple liquidity-based explanation of the estimated persistence asymmetry.

We also assess the economic relevance of the estimated latent state using subsequent sales of property, plant, and equipment, an outcome that is not used in constructing the benchmark or estimating the latent-state model. Persistent positive abnormal-investment deviations predict greater subsequent asset sales, and this association is significantly stronger than that associated with equally large negative deviations. The finding is robust across the two long-history samples and over both the subsequent year and a cumulative three-year horizon, linking the estimated asymmetric investment dynamics to subsequent corporate capital reallocation.

The remainder of the paper is organized as follows. Section 2 develops the nonlinear latent-state framework, presents the estimation and inference procedures, and summarizes its finite-sample performance, with the full Monte Carlo analysis reported in Appendix F. Section 3 discusses the economic mechanisms that can generate asymmetric persistence and states the empirical hypotheses. Section 4 applies the framework to corporate investment and examines asymmetric persistence, regime dynamics, financing-capacity heterogeneity, predictive persistence, and the economic relevance of the latent state for subsequent capital reallocation. Section 5 concludes.

2. Latent-state model of abnormal investment

2.1 Benchmark investment equation and interpretation

The empirical literature commonly measures abnormal corporate investment as the residual from a benchmark investment regression. Let firm investment satisfy

$$I_{it} = X'_{it}\beta_0 + u_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T_i, \quad (1)$$

where I_{it} denotes investment by firm i in year t , and X_{it} contains an intercept and observable determinants of investment opportunities. The benchmark can include the standard firm-level controls, industry and year effects, and any additional variables used in the empirical specification. Let $\hat{\beta}$ denote the first-stage estimator and

$$\hat{u}_{it} = I_{it} - X'_{it}\hat{\beta}, \quad (2)$$

the estimated abnormal-investment residual.

Following Richardson (2006) and the subsequent investment-efficiency literature (e.g., Biddle et al. (2009), Chen et al. (2011) and Jiang et al. (2018)), positive and negative residuals are often interpreted as overinvestment and underinvestment. This interpretation is useful but incomplete. A regression residual can combine a persistent departure from benchmark investment with measurement error, temporary demand or financing shocks, and other short-lived disturbances. Its contemporaneous sign therefore need not reveal the sign of the persistent component.

The model below separates these components. Following the terminology established by Richardson (2006) and used throughout the subsequent literature, this paper refers to persistent positive abnormal investment as *overinvestment* and persistent negative abnormal investment as *underinvestment*. These are the paper’s primary regime labels. Formally, however, the model identifies persistent deviations from benchmark investment rather than deviations from a structurally identified optimal level of investment. The terms overinvestment and underinvestment should therefore be understood throughout the paper in this established reduced-form sense.

Persistent abnormal investment may reflect investment distortions associated with managerial incentives, financing frictions, or adjustment costs, but may also capture investment opportunities and other firm characteristics not fully represented in the benchmark equation. Importantly, both investment fundamentals and investment distortions may contain permanent firm-specific and time-varying components. Persistence in the benchmark residual therefore does not, by itself, establish the economic source of the underlying deviation. Our framework focuses on the dynamic propagation of persistent benchmark-relative abnormal investment. The empirical specification developed in Section 4 allows for additional sources

of firm and industry-time heterogeneity that may otherwise contribute to the estimated persistence of abnormal investment.

2.2 Asymmetric latent-state dynamics

The population residual is decomposed as

$$u_{it} = \eta_{it} + \varepsilon_{it}, \quad (3)$$

where η_{it} denotes the latent persistent state of abnormal investment and ε_{it} is a transitory disturbance. The latent state captures the component of abnormal investment that propagates over time, while ε_{it} captures short-lived variation. Positive values of the latent state, $\eta_{it} > 0$, define the overinvestment regime, whereas negative values, $\eta_{it} < 0$, define the underinvestment regime.

The distinction between persistent and transitory abnormal investment does not imply that every economically meaningful investment deviation must persist for several periods. A sufficiently large innovation to the latent state may generate a substantial deviation that subsequently declines rapidly or crosses the regime threshold. At the same time, some economically meaningful short-lived investment deviations may be captured by the transitory component. The proposed framework focuses on the propagation of persistent abnormal investment rather than on identifying every potentially inefficient investment decision.

The latent state evolves according to the threshold autoregressive process

$$\eta_{it} = \delta + \rho_+ \eta_{i,t-1} \mathbf{1}\{\eta_{i,t-1} \geq 0\} + \rho_- \eta_{i,t-1} \mathbf{1}\{\eta_{i,t-1} < 0\} + w_{it}, \quad (4)$$

Equation (4) can be written more compactly as

$$\eta_{it} = \delta + \rho(\eta_{i,t-1}) \eta_{i,t-1} + w_{it}, \quad (5)$$

where $\rho(x) = \rho_+ \mathbf{1}\{x \geq 0\} + \rho_- \mathbf{1}\{x < 0\}$. The intercept δ is necessary because the benchmark

residual is centered by construction, whereas an asymmetric zero-intercept state process need not have zero unconditional mean. Omitting δ would generally impose an unintended restriction linking ρ_+ and ρ_- whenever both regimes occur with positive probability.

For likelihood construction, the baseline specification assumes that w_{it} are $\text{iid}N(0, \sigma_w^2)$ and ε_{it} are $\text{iid}N(0, \sigma_\varepsilon^2)$. The two innovation sequences are mutually independent, independent across firms, and conditionally independent of the included regressors. The Gaussian assumption can be relaxed in robustness analysis, but the innovation distributions must be specified for likelihood-based estimation and for separating state from measurement variance.

The baseline parameter space imposes $0 \leq \rho_+ < 1$, $0 \leq \rho_- < 1$, $\sigma_w^2 > 0$, and $\sigma_\varepsilon^2 > 0$.

These restrictions give the persistence parameters their conventional economic interpretation and, together with standard innovation conditions, imply stability of the threshold process. Appendix A provides a formal stability result for the latent-state process. Allowing negative persistence is possible in principle, but it would generate sign-reversing dynamics and would make the usual half-life interpretation inappropriate. The estimated ordering can therefore be interpreted in light of the competing economic mechanisms developed in Section 3, although the reduced-form state equation does not separately identify the underlying structural channel.

The central empirical hypothesis is that persistence need not be symmetric across positive and negative abnormal-investment states, $\rho_+ \neq \rho_-$. We do not impose an *ex ante* ordering between the two persistence coefficients. Instead, the model allows the data to determine both the magnitude and the direction of persistence asymmetry. Section 3 develops the competing economic mechanisms that can rationalize either ordering and derives the corresponding directional hypotheses.

This specification differs from a conventional Markov-switching model. In a Markov-switching model, a discrete latent regime evolves according to an estimated transition matrix. Here, the latent state is continuous, and the slope of its conditional mean changes at the economically meaningful zero threshold. Persistence is therefore piecewise constant: it equals ρ_+ in the overinvestment region and ρ_- in the underinvestment region. The conditional mean remains continuous at zero because the two branches share the common intercept δ . Regime membership is determined by whether the latent state lies above or below the zero threshold

rather than by a separately evolving discrete state. This structure isolates the paper’s central object of interest, which is the asymmetry in the propagation of positive and negative persistent investment deviations.

The estimated persistence ordering can therefore be given an economic interpretation, but it should not be viewed as identifying a unique underlying mechanism. A finding that $\rho_- > \rho_+$ would be consistent with forces that slow the correction of negative investment deviations, whereas $\rho_+ > \rho_-$ would be consistent with forces that sustain positive deviations. Section 3 develops these competing mechanisms in greater detail. Because the state equation is deliberately reduced form, however, the empirical model identifies the direction and magnitude of persistence asymmetry rather than the particular structural channel generating it.

2.3 Nonlinear filtering, estimation, and inference

Because the transition slope depends on the unobserved sign of $\eta_{i,t-1}$, the filtering distribution is generally non-Gaussian. Likelihood evaluation must therefore integrate over both regions of the lagged-state distribution. In particular, it is not valid to select ρ_+ or ρ_- solely from the sign of a previously filtered point estimate.

Let

$$F_{it} = \sigma(\hat{u}_{i1}, \dots, \hat{u}_{it}, X_i)$$

denote the information available for firm i through period t , and let

$$\vartheta = (\delta, \rho_+, \rho_-, \sigma_w^2, \sigma_\varepsilon^2)'.$$

For notational convenience, in the density expressions below we suppress firm and time subscripts. Let x denote a generic realization of the lagged latent state $\eta_{i,t-1}$, and let η denote a generic realization of the current latent state η_{it} . The transition and measurement densities are

$$g_\vartheta(\eta \mid x) = \frac{1}{\sigma_w} \phi\left(\frac{\eta - \delta - \rho(x)x}{\sigma_w}\right),$$

and

$$h_{\vartheta}(u \mid \eta) = \frac{1}{\sigma_{\varepsilon}} \phi\left(\frac{u - \eta}{\sigma_{\varepsilon}}\right),$$

where $\phi(\cdot)$ denotes the standard normal density.

Given the filtered density $f_{i,t-1|t-1}(x)$, the one-step-ahead density of the latent state is

$$f_{it|t-1}(\eta) = \int g_{\vartheta}(\eta \mid x) f_{i,t-1|t-1}(x) dx.$$

The predictive density of the observed residual is

$$m_{it}(\hat{u}_{it}; \vartheta) = \int h_{\vartheta}(\hat{u}_{it} \mid \eta) f_{it|t-1}(\eta) d\eta,$$

and Bayes' rule gives the updated filtering density

$$f_{it|t}(\eta) = \frac{h_{\vartheta}(\hat{u}_{it} \mid \eta) f_{it|t-1}(\eta)}{m_{it}(\hat{u}_{it}; \vartheta)}.$$

These prediction and updating recursions follow the standard state-space filtering logic (see Durbin and Koopman (2012) for a general treatment). The firm-level and panel log-likelihoods are, respectively,

$$\begin{aligned} \ell_i(\vartheta; \hat{\beta}) &= \sum_{t=1}^{T_i} \log m_{it}(\hat{u}_{it}; \vartheta), \\ \ell_N(\vartheta; \hat{\beta}) &= \sum_{i=1}^N \ell_i(\vartheta; \hat{\beta}) \\ &= \sum_{i=1}^N \sum_{t=1}^{T_i} \log m_{it}(\hat{u}_{it}; \vartheta). \end{aligned} \tag{6}$$

The initial state density is taken to be the invariant density implied by the candidate parameter vector. A sufficiently long burn-in provides a practical alternative. In the baseline implementation, the integrals above are evaluated numerically using deterministic quadrature on a finite state grid. Numerical accuracy is assessed by refining and expanding the grid and verifying that the likelihood and parameter estimates remain stable. A particle filter provides an alternative numerical approximation but is not used in the baseline implementation.

The nonlinear filter also provides the posterior mean of the latent state,

$$\hat{\eta}_{it} = E_{\hat{\vartheta}}(\eta_{it} \mid F_{it}),$$

and the posterior probability of a positive persistent deviation,

$$\hat{p}_{it}^+ = P_{\hat{\vartheta}}(\eta_{it} \geq 0 \mid F_{it}) = \int_0^\infty f_{it|t}(\eta) d\eta.$$

Thus, although the threshold defining the two regimes is fixed at zero, regime membership is latent and must be inferred from the filtering distribution rather than treated as directly observed. This distinction also explains why the raw benchmark-investment residual may provide an imperfect regime measure. Since the residual combines the persistent latent state with a transitory disturbance, a sufficiently large transitory shock can reverse its contemporaneous sign relative to that of the persistent component. Appendix B formalizes this residual-sign misclassification and shows how the probability of classification error depends on the magnitude of the latent state relative to transitory noise.

We therefore classify regime membership using the posterior probability rather than the sign of the raw residual. A firm-year observation is classified as overinvesting when $\hat{p}_{it}^+ > 1/2$ and as underinvesting when $\hat{p}_{it}^+ < 1/2$; an observation for which $\hat{p}_{it}^+ = 1/2$ is left unclassified. Because the posterior distribution need not be Gaussian, the sign of the posterior mean need not coincide with the regime having the larger posterior probability. Empirical classification is therefore based on $\hat{p}_{it}^+$, rather than on $\text{sign}(\hat{\eta}_{it})$. As a robustness check, observations with $0.30 < \hat{p}_{it}^+ < 0.70$ are treated as having uncertain regime membership. We additionally examine the distribution of posterior classification probabilities and recompute regime-transition probabilities using only adjacent observations for which both current and lagged states are confidently classified, defining $\hat{p}_{it}^+ > 0.70$ as a positive state and $\hat{p}_{it}^+ < 0.30$ as a negative state. The corresponding empirical results are reported in Appendix Table A.1. Appendix C establishes convergence of the feasible probability-based rule to the corresponding oracle classification and clarifies the irreducible uncertainty associated with recovering an unobserved latent state from noisy measurements.

Identification is based on the joint distribution of the observed residual history rather than

on direct observation of the latent state. The unit loading on η_{it} in Equation (3) fixes the scale of the latent state, while the threshold is normalized at zero because zero represents the benchmark-investment deviation. Conditional on the previous latent state,

$$E(\eta_{it} \mid \eta_{i,t-1} = x) = \begin{cases} \delta + \rho_+ x, & x \geq 0, \\ \delta + \rho_- x, & x < 0. \end{cases}$$

Thus, ρ_+ and ρ_- are the regime-specific slopes of the conditional state mean, while δ is their common intercept. Although the first-stage residual is centered by construction, the intercept δ is retained and estimated jointly because asymmetric state dynamics need not imply a zero-mean invariant distribution when the state equation is specified without an intercept. In the Monte Carlo design, the data-generating value of δ is calibrated to center the invariant latent-state distribution, while δ is estimated jointly with the remaining state-space parameters.

Identification requires the stationary distribution of the latent state to assign positive probability to both sides of zero and to contain sufficient variation within each regime. The state-innovation variance σ_w^2 and the transitory-variance component σ_ε^2 are distinguished by their different dynamic implications. Innovations to the latent state propagate to future observations, whereas transitory disturbances do not. Under the maintained distributional assumptions and regularity conditions, local identification of ϑ additionally requires the expected information matrix of the observed-data likelihood to be nonsingular. In the empirical analysis, identification can be assessed using regime occupancy, profile likelihoods, and the eigenvalues or condition number of the estimated information matrix. Appendix D provides formal identification results for the common intercept and the regime-specific persistence parameters and states the additional local-identification condition when the innovation variances are estimated jointly. Estimation proceeds in two steps. First, Equation (1) is estimated and the residuals defined in Equation (2) are constructed. Second, the state-space parameters are estimated by

$$\hat{\vartheta} = \arg \max_{\vartheta \in \Theta} \ell_N(\vartheta; \hat{\beta}),$$

subject to the parameter restrictions stated above. Because the likelihood is nonlinear, multiple starting values are used and convergence is assessed using the score, the Hessian,

and numerical stability of the likelihood.

Inference must account for the fact that the state-space model is estimated using generated first-stage residuals. Let $g_i(\beta)$ denote the firm-level estimating equations associated with the benchmark investment regression and let $s_i(\vartheta, \beta)$ denote the firm-level score of the state-space likelihood. Define

$$\Psi_i(\gamma) = \begin{pmatrix} g_i(\beta) \\ s_i(\vartheta, \beta) \end{pmatrix}, \quad \gamma = (\beta', \vartheta')'.$$

Under correct specification, cross-sectional independence or an appropriate cluster-dependence condition, suitably controlled T_i , and standard moment and smoothness conditions,

$$\sqrt{N}(\hat{\gamma} - \gamma_0) \xrightarrow{d} N(0, G^{-1}\Omega(G^{-1})'), \quad (7)$$

where

$$G = E \left[\frac{\partial \Psi_i(\gamma_0)}{\partial \gamma'} \right], \quad \Omega = E [\Psi_i(\gamma_0)\Psi_i(\gamma_0)'].$$

The covariance matrix of $\hat{\vartheta}$ is obtained from the corresponding block of the sandwich covariance matrix. This formulation accounts explicitly for estimation of the first-stage benchmark equation. Appendix E provides the corresponding two-step influence representation and shows explicitly how first-stage estimation uncertainty enters inference on the state-space parameters. In the empirical implementation, the preferred inference procedure is a firm-level bootstrap. Firms are resampled and both estimation stages are repeated, including construction of the benchmark residuals, nonlinear filtering, and likelihood maximization. This procedure propagates first-stage estimation uncertainty and preserves within-firm serial dependence. Hessian-based standard errors from the second stage alone may be reported as a diagnostic but are not used as the primary basis for inference.

Persistence symmetry is tested against the two-sided alternative

$$H_0 : \rho_+ = \rho_-, \quad H_1 : \rho_+ \neq \rho_-.$$

Let

$$d(\vartheta) = \rho_- - \rho_+.$$

The Wald statistic is

$$W = \frac{d(\hat{\vartheta})^2}{\widehat{\text{Var}}[d(\hat{\vartheta})]} \xrightarrow{d} \chi_1^2. \quad (8)$$

The variance used for the model-based Wald statistic in Equation (8) is obtained from the inverse Hessian of the second-stage likelihood. Because the benchmark-investment residuals are generated in the first stage, our preferred inference additionally uses a firm-level two-stage bootstrap that re-estimates both the benchmark equation and the latent-state model. The primary test of persistence symmetry is two-sided because the economic mechanisms discussed below can generate either direction of persistence asymmetry. In addition to the Wald test, we report a bootstrap confidence interval for $d(\vartheta) = \rho_- - \rho_+$. A positive value of d implies greater persistence of underinvestment, whereas a negative value implies greater persistence of overinvestment. Thus, conditional on rejecting persistence symmetry, the sign of the estimated persistence difference determines the direction of the asymmetry.

For $0 < \rho_s < 1$, the regime-specific conditional half-life is

$$h_s = \frac{\log(0.5)}{\log(\rho_s)}, \quad s \in \{+, -\}.$$

This quantity describes the rate at which a deviation decays while the latent state remains on the same side of the threshold. It should not be interpreted as the unconditional expected time spent in a regime, because innovations can move the latent state across zero.

2.4 Finite-sample performance

Because the empirical analysis relies on a nonlinear latent-state model estimated from generated benchmark-investment residuals, we evaluate the finite-sample performance of the procedure using Monte Carlo simulations. The experiments consider symmetric and asymmetric persistence designs and panel lengths chosen to span the within-firm histories used in the empirical application. They evaluate four features of the procedure: recovery

of the regime-specific persistence parameters, recovery of the latent state, classification of positive and negative investment regimes, and finite-sample inference for the null of symmetric persistence. Full details of the data- generating processes, implementation, and simulation results are reported in Appendix F.

The simulations support the empirical use of the estimator. Dynamic filtering substantially improves recovery of the persistent state relative to treating the raw investment residual as the state itself, reducing state-estimation mean squared error by approximately 62–75 percent across the designs considered. It also improves regime classification by several percentage points. The regime-specific persistence parameters are recovered accurately, with only modest finite-sample attenuation in the most demanding short-panel design. Importantly, the unrestricted model does not generate spurious persistence asymmetry under the symmetric data-generating process: the two-sided symmetry test has rejection frequencies close to its nominal 5 percent size, while power is approximately 96–100 percent under the asymmetric designs. These results are highly stable when the available within-firm history is shortened, providing support for the estimation and classification procedures used below.

3. Economic mechanisms and hypotheses

3.1 Sources of asymmetric persistence

The regime-specific coefficients ρ_+ and ρ_- measure the propagation of persistent abnormal investment in the overinvestment and underinvestment regions, respectively. Along a state path that remains in regime $s \in \{+, -\}$ for h periods, the local effect of a marginal change in the current latent state is

$$\frac{\partial \eta_{i,t+h}}{\partial \eta_{it}} = \rho_s^h. \quad (9)$$

A larger ρ_s therefore implies slower decay of an initial deviation while the latent state remains on the same side of the threshold. Because the model contains the intercept δ and innovations can move the latent state across zero, ρ_s should not be interpreted as determining the

unconditional time spent in a regime. Instead, it measures the local propagation of persistent abnormal investment within that regime.

Financing constraints provide a natural rationale for greater persistence of underinvestment. Firms with limited access to external finance may be unable to undertake profitable projects immediately when internal resources are insufficient or external capital is costly. Investment by financially constrained firms can therefore depend strongly on internally generated funds (see, e.g., Fazzari, Hubbard and Petersen (1988)), while collateral availability can further shape investment capacity (see Almeida and Campello (2007)). Information asymmetries may reinforce these effects by increasing the cost or limiting the availability of external finance (Stiglitz and Weiss (1981), and Myers and Majluf (1984), *inter alia*). When these financing frictions delay the expansion of investment, negative abnormal-investment states may decay more slowly, contributing to $\rho_- > \rho_+$.

A different set of mechanisms can generate greater persistence of overinvestment. Agency conflicts may induce managers to continue investing beyond levels justified by observable fundamentals when investment expands the resources under their control or generates private benefits (see Jensen (1986)). Abundant internal funds, weak external discipline, and persistent access to liquidity can allow such investment to continue (see, e.g., Richardson (2006), and Harford, Mansi and Maxwell (2008)). Commitments to ongoing projects or organizational objectives may reinforce these effects. When such forces dominate, positive abnormal-investment states may propagate more strongly, contributing to $\rho_+ > \rho_-$.

Adjustment costs and irreversibility can reinforce either direction of persistence asymmetry. Capital accumulation involves fixed costs, planning lags, installation costs, and partially irreversible commitments (see, e.g., Abel and Eberly (1994), and Dixit and Pindyck (1994)). When these costs primarily affect the initiation or expansion of investment, negative deviations may adjust slowly, contributing to $\rho_- > \rho_+$. Conversely, when investment programs involve multi-period commitments, sunk planning expenditures, complementary assets, or other costs of termination and reversal, positive deviations may persist, contributing to $\rho_+ > \rho_-$.

The distinction between the installed capital stock and the investment flow is important in this context. Although installed capital may be costly to dispose of, current investment can often be reduced by postponing or cancelling future projects. Such adjustment need

not be costless, however, when ongoing programs involve contractual, organizational, or technological commitments. Adjustment costs therefore do not imply a unique ordering of the two persistence parameters.

Economic uncertainty may also affect adjustment asymmetrically. Greater uncertainty can delay new investment by increasing the value of waiting or tightening financing conditions, while it may also induce firms to postpone, scale back, or cancel planned expenditures (see Bloom (2009)). Policy uncertainty has similarly been associated with weaker corporate investment (see Gulen and Ion (2016)). Its implications for the relative persistence of positive and negative deviations are therefore theoretically ambiguous.

Defining

$$d = \rho_- - \rho_+,$$

a positive value of d indicates greater persistence of underinvestment, whereas a negative value indicates greater persistence of overinvestment. Because the state-space model is deliberately reduced form, the mechanisms discussed above provide economic interpretations of the estimated ordering rather than separately identified structural channels.

3.2 Testable hypotheses

The preceding mechanisms imply that persistent abnormal investment need not adjust symmetrically across positive and negative states. They also imply that the direction of the asymmetry need not be universal. The empirical analysis therefore distinguishes between the existence of persistence asymmetry, its direction, and its heterogeneity across firms.

Hypothesis 1 (Asymmetric persistence). Persistent overinvestment and underinvestment exhibit different propagation dynamics. Defining the persistence difference as

$$d = \rho_- - \rho_+,$$

the primary hypothesis is

$$H_1 : \rho_- \neq \rho_+,$$

or equivalently,

$$H_1 : d \neq 0.$$

The corresponding null hypothesis of symmetric persistence, $\rho_- = \rho_+$, is evaluated using the two-sided Wald test in Equation (8). This is the paper's principal persistence hypothesis.

Hypothesis 2 (Direction of persistence asymmetry). Conditional on persistence being asymmetric, the mechanisms developed in Section 3.1 imply two competing directional outcomes. Financing frictions and costs of initiating or expanding investment can imply

$$H_{2a} : \rho_- > \rho_+,$$

or equivalently,

$$H_{2a} : d > 0.$$

By contrast, agency-related forces and costs of continuing or reversing ongoing investment programs can imply

$$H_{2b} : \rho_+ > \rho_-,$$

or equivalently,

$$H_{2b} : d < 0.$$

Neither directional inequality is imposed as the maintained alternative. The empirical sign of d , together with its confidence interval, determines the direction of the estimated persistence asymmetry.

Hypothesis 3 (Financing-capacity heterogeneity). The degree of persistence asymmetry may also vary systematically with firms' financing capacity. For relatively constrained and less-constrained firms, denoted by C and U , respectively, define

$$d^g = \rho_-^g - \rho_+^g, \quad g \in \{C, U\},$$

and let the magnitude of persistence asymmetry be

$$A^g = |d^g| = |\rho_-^g - \rho_+^g|.$$

The corresponding heterogeneity hypothesis is

$$H_3 : A^C \neq A^U.$$

Thus, Hypothesis 3 concerns whether financing capacity is systematically associated with the strength of asymmetric adjustment rather than presupposing a particular sign of the persistence difference. The sign of d^g is reported separately and indicates whether underinvestment or overinvestment is more persistent within each group.

Firm size, cash flow, and profitability are used as observable proxies for financing capacity. Because these variables are imperfect measures of financial constraints and may also be related to agency incentives, liquidity, investment opportunities, and organizational characteristics, the resulting cross-sectional evidence is interpreted as conditional heterogeneity in persistence rather than as causal identification of a specific financing-friction mechanism.

4. Asymmetric Persistence in Corporate Investment

We use the latent-state framework to address the paper’s central corporate-finance questions: whether persistent abnormal investment adjusts asymmetrically across overinvestment and underinvestment states, whether local state propagation differs from the probability of remaining in the same investment regime, and whether the strength of this asymmetry varies systematically with firms’ financing capacity. These questions correspond directly to the competing mechanisms and hypotheses developed in Sections 2 and 3 and allow the data to determine which investment state exhibits greater local persistence.

We examine these questions using a panel of U.S. publicly traded firms drawn from Compustat over the period 1991–2025. The empirical analysis is conducted using two long-history samples. The broader sample retains consecutive firm histories longer than 10 annual observations,

while the stricter sample retains histories longer than 20 annual observations. This design allows us to assess whether the estimated asymmetric investment dynamics are robust to the amount of within-firm time-series information available to the nonlinear filter.

The empirical analysis proceeds in three steps. First, we estimate benchmark investment using both a conventional linear specification and a more flexible generalized additive model (GAM), allowing us to separate abnormal investment from potentially nonlinear relationships between investment and observable firm fundamentals. Second, we apply the nonlinear latent-state model to the residuals from the preferred GAM benchmark and estimate the regime-specific persistence parameters. Third, we characterize the resulting overinvestment and underinvestment states in terms of regime incidence, investment intensity, transition dynamics, cross-sectional heterogeneity in persistence asymmetry, and their predictive content for future investment and regime membership. Throughout, positive and negative latent states are interpreted as persistent positive and negative deviations from benchmark investment, respectively, rather than as deviations from a structurally identified optimal level of investment.

4.1 Data, variables, and long-history samples

Firm-level accounting and market information are obtained from Compustat for 1991–2025. Investment, I_{it} , is measured as capital expenditures scaled by lagged total assets. The benchmark investment equation includes two lagged sales ratios, lagged Tobin’s Q , firm size, lagged profitability, lagged cash-flow intensity, lagged leverage, and lagged cash holdings. These variables capture observable investment opportunities, operating performance, internal financing capacity, capital structure, liquidity, and firm scale.

The inclusion of Tobin’s Q , following the investment- q tradition of Tobin (1969) and Hayashi (1982), is particularly useful because it provides a market-based measure of investment opportunities in addition to accounting variables. The sales controls capture realized product-market conditions, while profitability and cash flow summarize operating performance and internal fund generation. Leverage and cash holdings capture dimensions of capital structure and financial flexibility, and size controls for systematic differences in firm scale and access to financing.

An important consideration in constructing the benchmark investment equation is that firms

may differ persistently in their investment behavior because of unobserved characteristics, while investment opportunities and financing conditions may also vary systematically across industries over time. Such heterogeneity may remain in the benchmark residual and contribute to its estimated persistence. To account more fully for these sources of variation, we extend the conventional investment benchmark by including Mundlak-type firm-level averages of the explanatory variables and industry-by-year fixed effects. The former capture differences associated with firms' long-run observable characteristics, while the latter absorb common industry-specific investment conditions and shocks in each year. The resulting specification provides a more general benchmark for identifying the persistent component of abnormal investment.

The analysis uses two consecutive-history restrictions. The $T_i > 10$ sample contains 92,359 firm-year observations from 4,510 firms and 4,565 consecutive firm spells. The median consecutive history is 19 years. The $T_i > 20$ sample contains 54,073 firm-year observations from 1,960 firms and 1,960 consecutive spells, with a median history of 27 years. The broader and stricter samples therefore provide complementary evidence: the former retains a much larger cross section, whereas the latter provides a particularly long within-firm history for estimating persistent investment dynamics.

Table 1: Compustat Long-History Samples

Statistic	$T_i > 10$	$T_i > 20$
Sample period	1991–2025	1991–2025
Firm-year observations	92,359	54,073
Firms	4,510	1,960
Consecutive firm spells	4,565	1,960
Minimum consecutive history	11	21
Median consecutive history	19	27
Maximum consecutive history	35	35
Mean observations per spell	20.23	27.59

Note: The table summarizes the two long-history samples used in the empirical analysis. The $T_i > 10$ sample retains consecutive firm spells longer than 10 annual observations, while the $T_i > 20$ sample retains histories longer than 20 annual observations. The benchmark investment models are estimated once on the common first-stage sample. The resulting fitted values and abnormal-investment residuals are then carried unchanged into the $T_i > 10$ and $T_i > 20$ consecutive-history samples, within which the latent-state model is estimated separately. Both common-sample investment benchmarks include Mundlak-type firm-level means and SIC2-by-year fixed effects.

4.2 Linear benchmark investment model

We begin with a linear benchmark because it provides a transparent link to the corporate-investment literature and makes the economic content of the first-stage variables explicit. Following the residual-based investment literature associated with Richardson (2006) and Biddle, Hilary, and Verdi (2009), investment is modeled as a function of observable investment opportunities and firm fundamentals. The specification also includes Tobin’s Q , in the spirit of the investment- q literature, to reduce the risk that persistent market-valued growth opportunities remain in the abnormal investment residual.

We estimate the linear benchmark once on the common first-stage sample, allowing for long-run firm heterogeneity and industry-specific time variation:

$$I_{it} = \beta_0 + \beta_1 SR_{i,t-1} + \beta_2 SR_{i,t-2} + \beta_3 Q_{i,t-1} + \beta_4 SIZE_{i,t-1} + \beta_5 PROF_{i,t-1} + \beta_6 PROF_{i,t-2} + \beta_7 LEV_{i,t-1} + \beta_8 CF_{i,t-1} + \beta_9 CASH_{i,t-1} + \bar{\mathbf{X}}_i' \boldsymbol{\gamma} + \lambda_{j(i),t} + u_{it}^L, \quad (10)$$

where SR denotes the sales ratio, Q Tobin’s Q , $PROF$ profitability, CF cash-flow intensity, LEV leverage, $CASH$ cash holdings, and $SIZE$ firm size.

The vector $\bar{\mathbf{X}}_i$ contains the firm-specific means of the seven distinct investment determinants (the nine regressors include two lags each of the sales ratio and profitability), calculated over the available observations for each firm in the common first-stage estimation sample. Following a Mundlak-type correlated-effects approach, these averages account for the component of permanent firm heterogeneity associated with the included regressors.

The term $\lambda_{j(i),t}$ denotes industry-by-year fixed effects, where $j(i)$ identifies the industry of firm i at the two-digit level of the Standard Industrial Classification (SIC), hereafter referred to as SIC2. These effects absorb common investment conditions and shocks affecting firms within the same SIC2 industry in a given year, including industry-specific changes in investment opportunities and financing conditions. Together, the Mundlak controls and SIC2-by-year fixed effects reduce the possibility that persistent firm heterogeneity or industry-specific time variation is incorrectly attributed to the persistent component of abnormal investment.

The linear benchmark is estimated once on the common first-stage sample. The resulting fitted values and residuals are retained unchanged when constructing the $T_i > 10$ and $T_i > 20$

long-history samples, within which the latent-state model is estimated separately.

Table 2 reports the estimates from the common first-stage sample, accounting for long-run firm heterogeneity through Mundlak-type controls and industry-specific time variation through SIC2-by-year fixed effects. Tobin’s Q , profitability, and cash flow are positively associated with investment, whereas firm size and leverage enter negatively. The negative coefficient on leverage is consistent with the view that higher indebtedness may constrain investment through debt-overhang effects or increased external financing costs. Cash holdings are not statistically significant after controlling for the other investment determinants and the additional sources of heterogeneity. The two lagged sales coefficients have opposite signs, indicating different contributions of recent and more distant sales conditions to current investment.

The lower panel of Table 2 reports the latent-state persistence estimates obtained after carrying the same first-stage residuals unchanged into the $T_i > 10$ and $T_i > 20$ consecutive-history samples. The positive-state persistence coefficients are estimated at 0.916 and 0.911, respectively, whereas the corresponding negative-state persistence coefficients are 0.704 and 0.742. Thus, positive abnormal-investment deviations exhibit substantially stronger propagation than negative deviations in both samples. The persistence differences, $\hat{\rho}_- - \hat{\rho}_+$, are estimated at -0.212 and -0.168 , respectively, and likelihood-ratio tests strongly reject the symmetry restriction $H_0 : \rho_- = \rho_+$ at $p < 0.001$. We defer the substantive interpretation of the estimated persistence asymmetry to the GAM-based results below, since the nonlinear model provides the preferred first-stage investment specification.

The linear specification can be viewed as a restricted approximation to the GAM when the conditional effect of each covariate is approximately linear over its empirical support. Its coefficients are therefore useful for summarizing average directional associations between firm fundamentals and investment. When the underlying relationship is nonlinear, however, a single coefficient averages over potentially heterogeneous marginal effects and should not be interpreted as a globally constant response.

Table 2: Linear First-Stage Investment Model and Associated Latent-State Persistence

Panel A. Common-sample linear investment benchmark		
Variable	Linear Model	
Sales ratio $t - 1$	0.010***	
	(0.001)	
Sales ratio $t - 2$	-0.008***	
	(0.001)	
Tobin's Q	0.007***	
	(0.000)	
Size $t - 1$	-0.008***	
	(0.000)	
Profitability $t - 1$	0.006***	
	(0.001)	
Profitability $t - 2$	0.004***	
	(0.001)	
Leverage $t - 1$	-0.026***	
	(0.001)	
Cash flow $t - 1$	0.011***	
	(0.002)	
Cash holdings $t - 1$	0.002	
	(0.002)	
Mundlak firm means	Yes	
SIC2 $\times$ year fixed effects	Yes	
Observations	143,145	
Firms	15,207	
Panel B. Associated latent-state persistence		
Statistic	$T_i > 10$	$T_i > 20$
Second-stage observations	92,359	54,073
Firms	4,510	1,960
Consecutive spells	4,565	1,960
$\hat{\rho}_+$	0.9159	0.9105
$\hat{\rho}_-$	0.7042	0.7421
$\hat{\rho}_- - \hat{\rho}_+$	-0.2116	-0.1684
LR test of $\rho_- = \rho_+$	851.195	372.421
p -value	< 0.001	< 0.001

Note: The dependent variable is investment scaled by lagged total assets. Panel A reports estimates of the linear investment benchmark, estimated once on the common first-stage sample of 143,145 firm-year observations from 15,207 firms. All nine firm-level investment determinants enter linearly. The specification also includes firm-specific averages of these determinants (Mundlak-type controls) and a full set of 2-digit SIC-by-year fixed effects. Standard errors for the first-stage coefficients, reported in parentheses, are clustered at the firm (*gvkey*) level. Panel B reports persistence estimates obtained by applying the nonlinear latent-state model to residuals from the common linear investment benchmark. The residuals are carried unchanged into the $T_i > 10$ and $T_i > 20$ consecutive-history samples, within which the latent-state model is estimated separately. The parameters ρ_+ and ρ_- denote persistence in the positive and negative latent states, respectively. The persistence difference is defined as $d = \rho_- - \rho_+$.

4.3 A flexible investment benchmark: generalized additive model

The linear benchmark is useful for economic interpretation, but it may leave systematic nonlinearities in the first-stage residual. This issue is unusually important in our setting because the residual is not simply discarded after the first-stage regression. Instead, it becomes the measurement series used by the nonlinear latent-state model. If the benchmark investment equation misspecifies nonlinear relationships between investment and its observable determinants, the remaining nonlinear component could be incorrectly attributed to persistent abnormal investment.

We therefore estimate a generalized additive model (GAM), following Hastie and Tibshirani (1990) and Wood (2017), on the common first-stage sample. The GAM retains an additive and economically interpretable structure while allowing all continuous firm-level determinants of investment to enter through penalized smooth functions. We retain the Mundlak-type firm-level averages and industry-by-year fixed effects introduced in the linear benchmark:

$$\begin{aligned}
 I_{it} = & \alpha + s_1(SR_{i,t-1}) + s_2(SR_{i,t-2}) + s_3(Q_{i,t-1}) + s_4(SIZE_{i,t-1}) + s_5(PROF_{i,t-1}) \\
 & + s_6(PROF_{i,t-2}) + s_7(LEV_{i,t-1}) + s_8(CF_{i,t-1}) + s_9(CASH_{i,t-1}) + \bar{\mathbf{X}}_i' \boldsymbol{\gamma}_G + \lambda_{j(i),t} + u_{it}^G.
 \end{aligned} \tag{11}$$

The functions $s_j(\cdot)$ are estimated from the data subject to smoothness penalties. The GAM thus allows the conditional investment response to each continuous firm-level fundamental to be nonlinear, while accounting for the additional sources of firm and industry-time heterogeneity described in Section 4.2.

The degree of flexibility used by each estimated smooth is summarized by its effective degrees of freedom (EDF). An EDF close to one corresponds approximately to a linear fitted relationship, whereas an EDF materially above one indicates that the estimated function uses additional nonlinear flexibility. The reported reference degrees of freedom (*Ref.df*) are used to construct the approximate F -test for the significance of a smooth term. The F -test evaluates whether the corresponding smooth contributes to explaining investment; it is not, by itself, a formal test of linearity against nonlinearity. Further details on the spline

representation, penalization, EDF, reference degrees of freedom, and approximate significance tests are provided in Appendix G.

Table 3 reports the smooth-term diagnostics from the common-sample GAM, which incorporates Mundlak-type firm means and SIC2-by-year fixed effects, together with the associated latent-state persistence estimates. All reported smooth terms are strongly significant. The EDF values are substantially larger than one for every continuous covariate, indicating that the fitted investment relationships use considerable nonlinear flexibility rather than collapsing to simple straight lines. The GAM achieves an adjusted R^2 of 0.406 and explains 41.6% of the deviance in the common first-stage sample.

The associated latent-state estimates also reveal substantial persistence asymmetry. The positive-state persistence coefficients are estimated at 0.928 and 0.919 for the $T_i > 10$ and $T_i > 20$ samples, respectively, whereas the corresponding negative-state coefficients are 0.689 and 0.731. The estimated persistence differences, $\hat{\rho}_- - \hat{\rho}_+$, are therefore -0.239 and -0.187 , respectively. Likelihood-ratio tests strongly reject the symmetry restriction $H_0 : \rho_- = \rho_+$ at $p < 0.001$ in both samples. Thus, the asymmetric propagation of positive and negative abnormal-investment deviations is also present under the more flexible nonlinear investment benchmark.

Figure 1 displays the estimated smooth functions from the common-sample GAM investment benchmark, incorporating Mundlak-type firm-level means and SIC2-by-year fixed effects. Since the GAM is estimated once on the common first-stage sample, the same estimated functions are used to construct abnormal-investment residuals for both long-history samples. The estimated relationships reveal considerable heterogeneity in the nonlinear responses of investment to its conventional determinants. Tobin's Q exhibits a broadly increasing relationship with investment, consistent with its role as a measure of investment opportunities, although its slope varies over the observed range. Cash flow displays a pronounced nonlinear pattern, with several local fluctuations and a substantial decline before rising sharply at higher cash-flow values. Leverage is associated with a generally declining investment contribution, following a relatively flat profile at low leverage levels. Firm size exhibits an initial increase followed by a sustained decline over most of its empirical support.

The remaining covariates also exhibit nonlinearities of varying magnitude. The two sales lags

Table 3: GAM First-Stage Smooths and Associated Persistence Estimates

Panel A. Common-sample GAM investment benchmark			
Smooth term	EDF (Ref.df)	F	p
$s(SR_{t-1})$	5.142 (7)	10.239	< 0.001
$s(SR_{t-2})$	4.992 (7)	6.994	< 0.001
$s(Q_{t-1})$	6.740 (7)	296.118	< 0.001
$s(SIZE_{t-1})$	5.772 (7)	180.616	< 0.001
$s(PROF_{t-1})$	6.939 (7)	503.227	< 0.001
$s(PROF_{t-2})$	6.020 (7)	28.003	< 0.001
$s(LEV_{t-1})$	5.264 (7)	60.831	< 0.001
$s(CF_{t-1})$	6.638 (7)	980.021	< 0.001
$s(CASH_{t-1})$	5.216 (11)	37.352	< 0.001
Mundlak firm means		Yes	
SIC2 $\times$ year fixed effects		Yes	
Observations		143,145	
Firms		15,207	
Adjusted R^2		0.406	
Deviance explained		41.6%	

Panel B. Associated latent-state persistence		
Statistic	$T_i > 10$	$T_i > 20$
Second-stage observations	92,359	54,073
Firms	4,510	1,960
Consecutive spells	4,565	1,960
$\hat{\rho}_+$	0.9281	0.9188
$\hat{\rho}_-$	0.6894	0.7314
$\hat{\rho}_- - \hat{\rho}_+$	-0.2387	-0.1874
LR test of $\rho_- = \rho_+$	1,025.072	443.179
p -value	< 0.001	< 0.001

Note: Panel A reports smooth-term diagnostics for the GAM investment benchmark, estimated once on the common first-stage sample of 143,145 firm-year observations from 15,207 firms. The specification includes smooth functions of the nine time-varying investment determinants, firm-specific averages of these determinants (Mundlak-type controls), and a full set of 2-digit SIC-by-year fixed effects. EDF denotes effective degrees of freedom and Ref.df the reference degrees of freedom used in the approximate significance test. EDF values materially above one indicate that the fitted relationship uses nonlinear flexibility. The reported F -statistics test whether the corresponding smooth term contributes to explaining investment; they are not direct tests of linearity against nonlinearity. Panel B reports persistence estimates obtained by applying the nonlinear latent-state model to residuals from the common-sample GAM investment benchmark. The residuals are carried unchanged into the $T_i > 10$ and $T_i > 20$ consecutive-history samples, within which the latent-state model is estimated separately. Positive latent states are labeled overinvestment and negative latent states underinvestment, in the reduced-form sense defined in Section 2. The persistence difference is defined as $d = \rho_- - \rho_+$, so the negative estimates imply $\rho_+ > \rho_-$. The LR statistic tests the null hypothesis $H_0 : \rho_- = \rho_+$.

display relatively modest variation, with the first exhibiting a shallow non-monotonic pattern and the second increasing gradually after a small initial decline. Profitability at $t - 1$ shows a generally increasing contribution over much of its observed range, followed by a marked decline near the upper end, whereas the second profitability lag exhibits a comparatively mild nonlinear pattern. Cash holdings display a relatively flat, mildly inverted-U relationship with investment. Overall, the estimated smooths indicate that the relationships between investment and several conventional firm fundamentals vary substantially across their empirical support.

From an economic perspective, these nonlinear patterns are broadly consistent with established corporate-investment mechanisms. The positive relation between Tobin’s Q and investment accords with its traditional interpretation as a measure of investment opportunities in the q -theory of investment (see, e.g., Hayashi (1982)). At the same time, financial frictions can weaken the otherwise tight mapping between Q and investment and generate heterogeneous investment responses (see Cao, Lorenzoni and Walentin (2019)). The sharp increase in the cash-flow contribution at higher levels is consistent with the role of internal funds emphasized by Fazzari, Hubbard and Petersen (1988). Particularly relevant for our nonlinear specification, Borri, Martins-Filho and Kalatzis (2022) document a markedly nonlinear relationship between investment and internal funds using a semiparametric framework.¹

The generally declining leverage profile is also economically interpretable. While leverage is associated with relatively limited variation in investment at low levels, its contribution becomes increasingly negative as indebtedness rises. This pattern is consistent with the possibility that higher debt burdens reduce financial flexibility and constrain investment through debt-overhang or refinancing-risk mechanisms. Such an interpretation accords with the classic debt-overhang mechanism of Myers (1977) and the empirical evidence of Lang, Ofek and Stulz (1996), as well as more recent evidence on the relationship between corporate indebtedness and investment (see Kalemli-Ozcan, Laeven and Moreno (2022), and Blickle and Santos (2024), among others).

The generally declining size profile, particularly beyond the lower end of its empirical support,

¹Investment–cash-flow sensitivities should not be interpreted mechanically as causal evidence of financing constraints. This qualification follows the long-standing debate initiated by Kaplan and Zingales (1997) and Erickson and Whited (2000); the nonlinear evidence of Borri, Martins-Filho, and Kalatzis (2022) provides a more recent illustration of the heterogeneous nature of the investment–internal-funds relationship.

is consistent with systematic differences in investment intensity across firm scale and maturity. More generally, the nonlinear profitability and sales relationships indicate that standard investment determinants retain familiar economic interpretations while their conditional contributions vary across firms. Taken together, the estimated smooths support the use of a flexible investment benchmark that accommodates nonlinear responses to firm fundamentals while accounting for long-run firm heterogeneity and industry-specific time variation.

Table 4 compares the GAM and linear investment benchmarks using identical evaluation samples. Both specifications are estimated once on the common first-stage sample, incorporating Mundlak-type firm-level means and SIC2-by-year fixed effects, and their fitted values and residuals are subsequently evaluated within the two long-history samples. The GAM provides a better fit in both samples. For $T_i > 10$, the evaluation-sample R^2 increases from 0.414 under the linear model to 0.452 under the GAM, while RMSE falls from 0.0492 to 0.0476. For $T_i > 20$, the corresponding R^2 increases from 0.440 to 0.481 and RMSE falls from 0.0466 to 0.0448. The fixed-residual deviance comparisons imply reductions of approximately 6.4% and 7.2%, respectively. The common-sample AIC and BIC comparisons also favor the GAM.

Importantly, the improvement is not confined to the evaluation of the original first-stage estimates. A separate firm-level out-of-sample validation exercise excludes 20% of eligible firms from benchmark estimation and evaluates predictions on their held-out observations. The GAM reduces out-of-sample RMSE by approximately 3.3% in the $T_i > 10$ sample and 3.8% in the $T_i > 20$ sample. The corresponding reductions in MAE are approximately 2.8% and 3.4%. Although these gains are more modest than the differences in evaluation-sample fit, they indicate that the additional nonlinear flexibility also improves prediction for firms excluded from benchmark estimation.

4.4 Asymmetric latent-state estimates

We next examine the dynamic properties of the persistent component extracted from the GAM residuals. Define

$$d = \rho_- - \rho_+.$$

The null of symmetric persistence is $d = 0$. As discussed in Section 2, the economic literature permits competing directional mechanisms, so the main empirical restriction is asymmetry




Figure 1: Nonlinear Components of the GAM Investment Benchmark

Note: The figure displays the estimated smooth contributions of the nine continuous investment determinants from the GAM estimated on the common first-stage sample of 143,145 firm-year observations from 15,207 firms. The specification includes Mundlak-type firm-level means and SIC2-by-year fixed effects. Dashed curves denote pointwise uncertainty bands, while rug marks indicate the empirical support of each covariate. The same estimated GAM is used to construct abnormal-investment residuals for both long-history samples.

Table 4: GAM versus Linear Investment Benchmark

Statistic	$T_i > 10$		$T_i > 20$	
	GAM	Linear	GAM	Linear
Panel A. Fit in the long-history samples				
Observations	92,359	92,359	54,073	54,073
Firms	4,510	4,510	1,960	1,960
Consecutive spells	4,565	4,565	1,960	1,960
Evaluation-sample R^2	0.452	0.414	0.481	0.440
RMSE	0.0476	0.0492	0.0448	0.0466
MAE	0.0285	0.0294	0.0271	0.0281
Panel B. Agreement between GAM and linear benchmarks				
Kendall residual correlation	0.772		0.775	
Residual-sign agreement	0.877		0.877	
Fitted-value correlation	0.958		0.961	
RMSE of residual difference	0.0124		0.0119	
MAE of residual difference	0.0085		0.0082	
Panel C. Fixed-residual deviance comparison				
Linear residual deviance	223.710		117.211	
GAM residual deviance	209.387		108.718	
Deviance reduction	14.323		8.494	
Deviance reduction (%)	6.402		7.246	
Panel D. Common-sample information criteria				
AIC (GAM / Linear)	−434,582.431 / − 425,494.462			
BIC (GAM / Linear)	−410,614.153 / − 402,098.739			
Panel E. Firm-level out-of-sample prediction				
Test firms	900	900	392	392
Test observations	18,312	18,312	10,721	10,721
Out-of-sample RMSE	0.0473	0.0489	0.0460	0.0478
RMSE improvement (%)	3.329		3.828	
Out-of-sample MAE	0.0288	0.0296	0.0283	0.0294
MAE improvement (%)	2.806		3.437	

Note: The GAM and linear investment benchmarks are estimated once on the common first-stage sample of 143,145 firm-year observations from 15,207 firms. Both specifications include firm-specific means of all nine time-varying investment determinants (Mundlak controls) and a full set of SIC2-by-year fixed effects. The resulting fitted values and abnormal-investment residuals are carried unchanged into the $T_i > 10$ and $T_i > 20$ consecutive-history samples; neither benchmark is re-estimated after imposing the history restriction for Panels A–C or for the subsequent latent-state analysis. Panel A reports fit diagnostics of these fixed common-sample benchmarks within each long-history sample. Panel B reports concordance between their fitted values and residuals. Kendall’s τ measures rank association between the two sets of residuals. Panel C compares Gaussian residual loss using the fixed first-stage residuals; residual deviance equals the residual sum of squares. The reported deviance reductions are descriptive rather than sample-specific analysis-of-deviance tests. Panel D reports AIC and BIC from the common-sample first-stage model objects and therefore does not vary across the long-history samples. Panel E reports a separate firm-level out-of-sample validation exercise. For each long-history sample, 20% of eligible firms are randomly held out, and the GAM and linear benchmarks are re-estimated using all common first-stage observations from the remaining firms. For computational efficiency in this auxiliary exercise, the GAM is estimated using `gam` with `fREML` and discrete fitting, while the linear specification absorbs SIC2-by-year fixed effects

itself rather than an imposed sign for d .

Table 5 reports the full latent-state estimates obtained from the common-sample GAM investment benchmark incorporating Mundlak-type firm-level means and SIC2-by-year fixed effects. The results are highly consistent across the two long-history samples. In the $T_i > 10$ sample,

$$\hat{\rho}_+ = 0.928, \quad \hat{\rho}_- = 0.689,$$

implying $d = -0.239$. In the $T_i > 20$ sample,

$$\hat{\rho}_+ = 0.919, \quad \hat{\rho}_- = 0.731,$$

implying $d = -0.187$. Likelihood-ratio tests strongly reject the symmetry restriction $H_0 : \rho_+ = \rho_-$ at $p < 0.001$ in both samples.

Table 5: Nonlinear Latent-State Estimates

Statistic	$T_i > 10$	$T_i > 20$
ρ_+ : positive-state persistence	0.928	0.919
ρ_- : negative-state persistence	0.689	0.731
$d = \rho_- - \rho_+$	-0.239	-0.187
State innovation SD, σ_w	0.021	0.019
Transitory-error SD, σ_ε	0.029	0.026
State intercept, δ	-0.003	-0.002
Unrestricted log likelihood	167,645.729	102,569.631
Restricted common ρ	0.818	0.832
Restricted log likelihood	167,133.193	102,348.041
LR test of $\rho_+ = \rho_-$	1,025.072	443.179
p -value	< 0.001	< 0.001

Note: The table reports estimates of the nonlinear latent-state model using residuals from the common-sample GAM investment benchmark with Mundlak-type firm-level means and SIC2-by-year fixed effects. The GAM benchmark is estimated once on the common first-stage sample, and the resulting residuals are carried unchanged into the $T_i > 10$ and $T_i > 20$ consecutive-history samples, within which the latent-state model is estimated separately. Positive latent states are labeled overinvestment and negative latent states underinvestment in the reduced-form sense defined in Section 2. The persistence difference is defined as $d = \rho_- - \rho_+$. The LR test compares the unrestricted model with the restricted specification imposing $H_0 : \rho_+ = \rho_-$.

The finding $\hat{\rho}_+ > \hat{\rho}_-$ implies that positive persistent abnormal-investment deviations exhibit substantially stronger local propagation than negative deviations. Unlike specifications in

which the positive-state persistence coefficient approaches the unit boundary, the estimated coefficients in Table 5 allow an economically meaningful interpretation of the conditional half-lives.

Using the regime-specific half-life formula introduced in Section 2, $h_s = \frac{\log(0.5)}{\log(\rho_s)}$, $s \in \{+, -\}$, the estimated positive-state coefficients imply conditional half-lives of approximately 9.3 years for $T_i > 10$ and 8.2 years for $T_i > 20$. The corresponding negative-state half-lives are considerably shorter, at approximately 1.9 and 2.2 years, respectively. Thus, the effect of an initial positive abnormal-investment deviation decays substantially more slowly than that of a negative deviation, conditional on the latent state remaining within the corresponding regime.

These conditional half-lives should not be interpreted as the expected time firms spend in the overinvestment or underinvestment regime. The persistence coefficient ρ_s measures the local propagation of the magnitude of the continuous latent state, whereas a regime-continuation probability measures the probability of remaining on the same side of zero. The latter also depends on the state intercept, innovation distribution, and distribution of the latent state. The distinction between local state propagation and regime continuation is therefore economically important and is examined further below.

Economically, the estimated ordering supports the H_{2b} direction of the mechanisms discussed in Section 3. Positive abnormal-investment states exhibit substantially stronger local propagation than negative states. This pattern is consistent with mechanisms that may sustain investment once a firm is above its benchmark, including managerial agency, weak external discipline, continuing access to internal resources, and commitments or reversal costs associated with ongoing investment programs. Conversely, the aggregate estimates do not support the H_{2a} prediction that financing frictions make underinvestment the more persistent state. Because the model is reduced form, these mechanisms should be viewed as economic interpretations of the estimated ordering rather than as separately identified causal channels.

4.5 Regime incidence, investment intensity, and transition dynamics

Table 6 characterizes the estimated latent states. The hard classification assigns 52.0% of firm-year observations to the positive state in the $T_i > 10$ sample and 52.6% in the $T_i > 20$

sample. Positive abnormal-investment states are therefore slightly more frequent in both samples, although the regime shares remain relatively balanced. Posterior-weighted shares are similar to the hard classification shares, indicating that the estimated distribution of firm-year observations across the two regimes is not substantially altered when classification uncertainty is taken into account.

The two regimes also differ considerably in observed investment intensity. Mean investment is approximately 7.1% of lagged assets in the positive state and 2.9% in the negative state for $T_i > 10$, compared with 7.0% and 3.1% for $T_i > 20$. The raw regime gap is therefore approximately 3.9–4.1 percentage points, while the within-firm investment difference is approximately 3.5–3.8 percentage points. These statistics provide an economically transparent characterization of the latent classification, although they should not be interpreted as independent validation of investment inefficiency because the regimes are themselves estimated from abnormal-investment residuals.

A notable feature of Table 6 is that the positive state also exhibits a slightly higher continuation probability. In the $T_i > 10$ sample, $P(O_t | O_{t-1}) = 0.866$, compared with $P(U_t | U_{t-1}) = 0.847$; in the stricter sample, the corresponding probabilities are 0.872 and 0.849. Thus, positive abnormal-investment deviations exhibit both stronger local propagation in magnitude and a slightly higher probability of remaining positive from one year to the next.

Importantly, however, the differences in regime-continuation probabilities are relatively small compared with the differences in the estimated persistence coefficients. This reinforces the distinction between local propagation of the continuous latent state and the probability of remaining in the same regime. The former characterizes the rate at which the effect of an investment deviation decays, conditional on remaining within a regime, whereas the latter also depends on the state intercept, innovation distribution, and distribution of the latent state. The model therefore distinguishes two related but conceptually different dimensions of abnormal-investment dynamics.

Table 7 examines whether the estimated transition dynamics are sensitive to uncertainty in latent-state classification. In the $T_i > 10$ and $T_i > 20$ samples, respectively, 47.1% and 43.8% of observations have posterior positive-state probabilities between 0.30 and 0.70. This interval is deliberately broad, however: only 12.8% and 11.0% of observations lie close to

Table 6: Regime Incidence, Investment Outcomes, and Transition Dynamics

Statistic	$T_i > 10$	$T_i > 20$
Panel A. Regime incidence		
Hard share: overinvestment	0.520	0.526
Hard share: underinvestment	0.480	0.474
Posterior-weighted share: overinvestment	0.518	0.519
Posterior-weighted share: underinvestment	0.482	0.481
Panel B. Transition dynamics		
$P(O_t \mid O_{t-1})$	0.866	0.872
$P(U_t \mid U_{t-1})$	0.847	0.849
Under-minus-over continuation gap	-0.020	-0.023
Panel C. Investment outcomes		
Mean $I/A \mid O$	0.071	0.070
Mean $I/A \mid U$	0.029	0.031
Raw investment gap: $O - U$	0.041	0.039
Within-firm investment difference	0.038	0.035

Note: Hard regime shares are based on posterior classification using $\hat{p}_{it}^+ > 1/2$ for the positive latent state and $\hat{p}_{it}^+ < 1/2$ for the negative latent state. Posterior-weighted shares average the estimated regime probabilities. Transition probabilities are computed from the classified latent states. O and U denote the overinvestment and underinvestment regimes, respectively, in the reduced-form sense defined in Section 2. The variable I/A denotes firm investment scaled by lagged total assets. Accordingly, the reported conditional means are the average investment-to-assets ratios for observations classified in the positive and negative regimes, respectively. The raw investment gap is the difference between these conditional means, while the within-firm investment difference is computed for firms observed in both regimes. All statistics are descriptive and are conditional on the common-sample GAM investment benchmark incorporating Mundlak-type firm-level means and SIC2-by-year fixed effects, and on the associated latent-state classification. Minor discrepancies between reported differences and differences calculated from displayed values reflect rounding.

the 0.5 classification threshold, within $0.45 \leq \hat{p}_{it}^+ \leq 0.55$. When transition probabilities are recomputed using only adjacent observations for which both current and lagged states are confidently classified, the positive-state continuation probabilities are 0.980 and 0.977, while the corresponding negative-state probabilities are 0.957 and 0.958. Thus, the high degree of regime continuation documented in Table 6 remains when the analysis is restricted to observations with relatively sharp posterior classifications. The results in Table 7 therefore indicate that the transition patterns reported in Table 6 are not driven by observations lying close to the classification boundary.

Table 7: Posterior Classification Uncertainty and Confident-State Transitions

Panel A. Posterior classification uncertainty					
Sample	N	$\hat{p}_{it}^+ < 0.30$	$0.30 \leq \hat{p}_{it}^+ \leq 0.70$	$\hat{p}_{it}^+ > 0.70$	$0.45 \leq \hat{p}_{it}^+ \leq 0.55$
$T_i > 10$	92,359	25.78%	47.08%	27.13%	12.78%
$T_i > 20$	54,073	27.29%	43.77%	28.94%	10.96%

Panel B. Transition probabilities among confidently classified adjacent observations			
Sample	Valid transitions	$P(O_t = O \mid O_{t-1} = O)$	$P(U_t = U \mid U_{t-1} = U)$
$T_i > 10$	37,441	0.980	0.957
$T_i > 20$	23,994	0.977	0.958

Notes: $\hat{p}_{it}^+$ denotes the posterior probability that the latent investment state is positive. An observation is classified as confidently positive when $\hat{p}_{it}^+ > 0.70$ and confidently negative when $\hat{p}_{it}^+ < 0.30$. Panel B retains only adjacent observations within the same original spell for which both the current and lagged states are confidently classified.

Figure 2 provides a time-series characterization of investment intensity across the estimated regimes. Average investment scaled by assets is consistently higher in the positive state than in the negative state throughout 1991–2025, and this ordering does not reverse in either long-history sample. Both series exhibit a long-run decline relative to their levels in the 1990s, although the reduction is more pronounced in the positive regime. Consequently, the investment-intensity gap between the two regimes narrows over time but remains economically substantial. The time-series patterns are remarkably similar across the $T_i > 10$ and $T_i > 20$ samples, illustrating the evolution of investment intensity associated with the estimated latent regimes and the stability of this descriptive pattern across the two history restrictions.




Figure 2: Investment-to-Assets by Latent Regime, 1991–2025

Note: The figure plots annual mean investment scaled by lagged total assets for firm-year observations classified in the overinvestment and underinvestment regimes. The two panels incorporated in the figure correspond to the $T_i > 10$ and $T_i > 20$ consecutive-history samples, respectively. Regime classification is based on posterior probabilities from the nonlinear latent-state model applied to abnormal-investment residuals obtained from the common-sample GAM benchmark incorporating Mundlak-type firm-level means and SIC2-by-year fixed effects. Positive and negative latent states are labeled overinvestment and underinvestment in the reduced-form sense defined in Section 2.

4.6 Heterogeneity in asymmetric persistence

The theoretical discussion also allows the magnitude of persistence asymmetry to vary systematically with firms' financing capacity. Because the new estimates imply $\rho_+ > \rho_-$, it is useful to characterize heterogeneity using the absolute persistence gap

$$A_g = |\rho_-^g - \rho_+^g|,$$

where g indexes firm groups. This definition separates the magnitude of asymmetry from its direction.

Table 8 reports subgroup estimates based on firm-level median characteristics. Two patterns are consistent across the long-history samples: persistence asymmetry is more pronounced among less profitable firms and smaller firms. In the $T_i > 10$ sample, the absolute persistence gap is 0.289 for low-profitability firms, compared with 0.206 for high-profitability firms, and

0.363 for small firms, compared with 0.155 for large firms. The corresponding contrasts in the $T_i > 20$ sample are 0.236 versus 0.163 for profitability and 0.284 versus 0.108 for firm size.

The relationship between cash flow and persistence asymmetry is less uniform. The asymmetry is moderately larger among low-cash-flow firms in the $T_i > 10$ sample (0.257 versus 0.219), but is approximately identical across cash-flow groups in the $T_i > 20$ sample (0.182 in both groups). For Tobin’s Q , the asymmetry is somewhat larger among high- Q firms in both samples, particularly under the stricter history requirement.

Table 8: Heterogeneity in Persistence Asymmetry

Split	Group	$T_i > 10$			$T_i > 20$		
		ρ_+	ρ_-	A_g	ρ_+	ρ_-	A_g
Cash flow	Low cash flow	0.902	0.645	0.257	0.884	0.702	0.182
	High cash flow	0.938	0.720	0.219	0.931	0.750	0.182
Profitability	Low profitability	0.913	0.624	0.289	0.911	0.675	0.236
	High profitability	0.941	0.735	0.206	0.927	0.764	0.163
Firm size	Small firms	0.937	0.574	0.363	0.935	0.651	0.284
	Large firms	0.927	0.771	0.155	0.905	0.797	0.108
Tobin’s Q	Low Q	0.915	0.681	0.235	0.874	0.722	0.153
	High Q	0.940	0.691	0.249	0.947	0.732	0.216

Note: The table reports subgroup estimates of the asymmetric latent-state model based on abnormal-investment residuals from the common-sample GAM benchmark incorporating Mundlak-type firm-level means and SIC2-by-year fixed effects. For each long-history sample, firms are divided using the cross-sectional median of their firm-specific median characteristic, so that each firm belongs to one subgroup. The latent-state model is re-estimated separately within each subgroup. The parameters ρ_+ and ρ_- denote persistence in the positive (overinvestment) and negative (underinvestment) latent states, respectively. The measure $A_g = |\rho_-^g - \rho_+^g|$ denotes the magnitude of persistence asymmetry within subgroup g . Reported estimates are from the final 181-point-grid unrestricted model. Minor differences between reported asymmetry measures and differences calculated from displayed persistence coefficients reflect rounding.

The subgroup estimates indicate that the stronger local propagation of positive abnormal-investment deviations is not confined to firms with abundant cash flow or high profitability. Positive-state persistence remains high across all the reported groups, although its magnitude varies across firm characteristics. The greater asymmetry among smaller and less profitable firms is driven primarily by weaker propagation of negative deviations. For example, in the $T_i > 10$ sample, negative-state persistence is 0.574 among small firms compared with 0.771 among large firms, whereas positive-state persistence is 0.937 and 0.927, respectively. A

similar pattern appears across profitability groups and in the stricter $T_i > 20$ sample.

These cross-sectional patterns qualify a simple free-cash-flow interpretation of the aggregate result. Financing-related firm characteristics are associated with differences in asymmetric adjustment, but the evidence does not suggest that liquidity alone explains the stronger propagation of positive abnormal-investment deviations. Agency incentives, organizational commitments, and adjustment or reversal costs remain plausible complementary mechanisms.

4.7 Dynamic and predictive validation

We finally examine the implications of the estimated latent-state dynamics for future investment behavior. The following exercises document how the current filtered state is associated with subsequent investment outcomes and regime membership, conditional on the specified lagged controls and common year effects. Since the latent state is constructed from a dynamic model that explicitly allows for persistent abnormal-investment deviations, these predictive relationships are partly implied by the estimated state process. Accordingly, we interpret Tables 9 and 10 as empirical illustrations of the model’s dynamic implications, rather than as independent validation of investment inefficiency or evidence of causal effects.

Table 9 reports regressions of future investment on the filtered latent abnormal-investment state. The coefficient on the current state is positive and highly significant at horizons $t + 1$, $t + 2$, and $t + 3$ in both long-history samples. In the $T_i > 10$ sample, the coefficient declines from 0.892 at $t + 1$ to 0.595 at $t + 3$; the corresponding coefficients in the stricter $T_i > 20$ sample are 0.889 and 0.602. The estimates are remarkably similar across the two samples, and the gradual decline in the coefficients illustrates the propagation of persistent abnormal-investment deviations documented in the preceding sections.

Table 10 provides a complementary regime-based exercise in which future underinvestment is regressed on an indicator for current underinvestment. In the $T_i > 20$ sample, the estimated odds ratio declines from approximately 37.6 at $t + 1$ to 14.7 at $t + 2$ and 9.3 at $t + 3$; the corresponding ratios in the broader $T_i > 10$ sample are 35.1, 13.6, and 8.5. Although the predictive association weakens with the horizon, current regime membership remains strongly related to future membership three years ahead. This pattern is consistent with the high one-step underinvestment continuation probabilities in Table 6 and complements the

Table 9: Predictive Power of the Filtered Latent State

Horizon	$T_i > 10$		$T_i > 20$	
	Coefficient	R^2	Coefficient	R^2
$t + 1$	0.892 (0.021)	0.277	0.889 (0.030)	0.285
$t + 2$	0.680 (0.023)	0.187	0.688 (0.032)	0.199
$t + 3$	0.595 (0.024)	0.156	0.602 (0.033)	0.167

Note: The dependent variable is future investment at horizons $t + 1$, $t + 2$, and $t + 3$. The principal explanatory variable is the filtered posterior mean of the latent abnormal-investment state obtained from the common-sample GAM investment benchmark incorporating Mundlak-type firm-level means and SIC2-by-year fixed effects. The predictive regressions additionally control for the first and second lags of the sales ratio and include year fixed effects. Standard errors in parentheses are clustered at the firm level. The reported R^2 is the conventional coefficient of determination. All reported coefficients have $p < 0.001$. The exercise illustrates the predictive implications of the estimated latent-state dynamics rather than providing independent validation of investment inefficiency.

latent-state propagation results in Table 9.

Table 10: Persistence of Future Underinvestment

Horizon	$T_i > 10$			$T_i > 20$		
	Coefficient	Odds ratio	McFadden R^2	Coefficient	Odds ratio	McFadden R^2
$t + 1$	3.557 (0.027)	35.067	0.410	3.626 (0.038)	37.576	0.422
$t + 2$	2.608 (0.028)	13.578	0.260	2.691 (0.039)	14.746	0.274
$t + 3$	2.138 (0.029)	8.482	0.189	2.225 (0.041)	9.250	0.203

Note: The dependent variable equals one if the firm is classified in the underinvestment regime at horizons $t + 1$, $t + 2$, and $t + 3$. The principal explanatory variable is an indicator for current underinvestment based on the posterior regime classification obtained from the GAM-based nonlinear latent-state model. The logit regressions additionally control for the first and second lags of the sales ratio, lagged cash flow, lagged firm size, and year fixed effects. Standard errors in parentheses are clustered at the firm level. All reported coefficients have $p < 0.001$. Odds ratios equal the exponential of the corresponding logit coefficients, and McFadden's R^2 is the pseudo- R^2 from the corresponding logit model. Inference is conditional on the estimated latent-state classification.

4.8 Economic relevance of the latent state: subsequent asset sales

In this section, we assess the economic relevance of the estimated latent state by examining whether persistent abnormal-investment deviations predict future sales of property, plant, and equipment (SPPE), an outcome that is not used in constructing the investment benchmark or estimating the latent-state model. Asset sales provide an observable measure of ex post

capital reallocation. If persistent positive abnormal-investment states capture economically meaningful investment dynamics, firms experiencing larger positive deviations may subsequently reverse, restructure, or dispose of part of the capital accumulated during these investment episodes. Such a pattern is consistent with the mechanisms discussed in Section 3, whereby managerial incentives, weak external discipline, continued access to internal funds, and commitments or reversal costs may sustain investment above its benchmark before adjustment eventually occurs. SPPE has relatively broad coverage in the Compustat panel, leaving a substantial number of firm-year observations available in both long-history samples and at the horizons considered.

The exercise is interpreted as predictive evidence on the economic content of the latent state rather than as a structural test of investment inefficiency. Asset sales may arise for several reasons, including efficient restructuring, changes in investment opportunities, financing needs, or portfolio reallocation. Accordingly, a positive association between persistent positive abnormal investment and future asset sales does not imply that the original investment was welfare reducing. Rather, the exercise asks whether the persistent component identified by the latent-state model is systematically related to a subsequent corporate action that lies outside the system used to construct that state.

We consider both asset sales in the subsequent year and cumulative asset sales over the following three years. Specifically, we estimate

$$AS_{it}^{(H)} = \alpha_i + \lambda_{j(i),t} + \beta_+ \hat{\eta}_{it}^+ + \beta_- \hat{\eta}_{it}^- + \boldsymbol{\gamma}' \mathbf{Z}_{it} + e_{it}^{(H)}, \quad H \in \{1, 3\}, \quad (12)$$

where

$$AS_{it}^{(1)} = \frac{SPPE_{i,t+1}}{AT_{it}}, \quad AS_{it}^{(3)} = \frac{\sum_{s=1}^3 SPPE_{i,t+s}}{AT_{it}},$$

and the filtered latent state is decomposed into

$$\hat{\eta}_{it}^+ = \max(\hat{\eta}_{it}, 0), \quad \hat{\eta}_{it}^- = \max(-\hat{\eta}_{it}, 0).$$

Thus, $\hat{\eta}_{it}^+$ measures the magnitude of a positive persistent abnormal-investment deviation, whereas $\hat{\eta}_{it}^-$ measures the absolute magnitude of a negative persistent abnormal-investment

deviation. Both variables are non-negative. The control vector is

$$\mathbf{Z}_{it} = (SR_{it}, SR_{i,t-1}, Q_{it}, SIZE_{it}, PROF_{it}, PROF_{i,t-1}, LEV_{it}, CF_{it}, CASH_{it})'.$$

These controls are drawn from the firm-level investment determinants used in the benchmark specifications in Equations (10) and (11) and capture investment opportunities, firm scale, operating performance, internal financing capacity, and leverage. All controls enter linearly. The specification also includes firm-specific intercepts, α_i , and a full set of SIC2-by-year effects, $\lambda_{j(i),t}$, which absorb both aggregate time variation and industry-specific shocks in each year. Standard errors are clustered at the firm level. Because the firm-specific intercepts absorb time-invariant firm heterogeneity, additional Mundlak-type firm means are unnecessary. To limit the influence of extreme asset-sale observations, the SPPE outcomes are winsorized at the 1st and 99th percentiles separately by sample and horizon.

The first hypothesis examines whether persistent positive abnormal investment predicts subsequent asset sales:

$$H_0 : \beta_+ = 0 \quad \text{against} \quad H_1 : \beta_+ > 0.$$

A positive β_+ implies that larger persistent positive deviations from benchmark investment are followed by greater subsequent sales of PP&E, conditional on observable firm fundamentals, permanent firm heterogeneity, and industry-specific time conditions. Economically, such a finding would be consistent with subsequent reversal, restructuring, or disposal of capital following a persistent positive investment episode.

The second and central asymmetry test is

$$H_0 : \beta_+ = \beta_- \quad \text{against} \quad H_1 : \beta_+ > \beta_-.$$

Because both $\hat{\eta}_{it}^+$ and $\hat{\eta}_{it}^-$ are non-negative magnitudes, this test compares the subsequent asset-sales association with positive and negative persistent deviations of the same absolute size. Evidence that $\beta_+ > \beta_-$ would imply that future asset disposal is more strongly associated with persistent positive abnormal-investment states than with equally large persistent negative states.

Table 11: Subsequent Asset Sales and the Filtered Latent State

	$T_i > 10$		$T_i > 20$	
	$t + 1$	$t + 1 : t + 3$	$t + 1$	$t + 1 : t + 3$
Positive filtered state, $\hat{\eta}_{it}^+$	0.0277*** (0.0048)	0.0591*** (0.0155)	0.0272*** (0.0065)	0.0618*** (0.0212)
Negative filtered state, $\hat{\eta}_{it}^-$	-0.0116 (0.0095)	-0.0506 (0.0366)	-0.0225* (0.0126)	-0.0868* (0.0465)
SR_{it}	0.0006*** (0.0002)	0.0016** (0.0007)	0.0005 (0.0003)	0.0023** (0.0009)
$SR_{i,t-1}$	0.0002 (0.0002)	0.0006 (0.0006)	0.0005* (0.0003)	0.0013* (0.0008)
Q_{it}	-0.0002*** (0.0000)	-0.0004** (0.0001)	-0.0003*** (0.0001)	-0.0006*** (0.0002)
$SIZE_{it}$	-0.0004*** (0.0001)	-0.0027*** (0.0006)	-0.0005*** (0.0002)	-0.0023*** (0.0007)
$PROF_{it}$	-0.0019*** (0.0005)	-0.0054*** (0.0019)	-0.0016* (0.0008)	-0.0018 (0.0027)
$PROF_{i,t-1}$	-0.0008** (0.0004)	-0.0015 (0.0011)	-0.0005 (0.0006)	-0.0014 (0.0017)
LEV_{it}	0.0019*** (0.0005)	0.0037* (0.0020)	0.0018*** (0.0007)	0.0028 (0.0027)
CF_{it}	-0.0007 (0.0006)	-0.0011 (0.0020)	-0.0014 (0.0009)	-0.0063* (0.0032)
$CASH_{it}$	-0.0020*** (0.0005)	-0.0065*** (0.0020)	-0.0015* (0.0008)	-0.0055** (0.0027)
$\hat{\beta}_+ - \hat{\beta}_-$	0.0393	0.1096	0.0497	0.1486
One-sided p : $H_1 : \beta_+ > 0$	< 0.001	< 0.001	< 0.001	0.002
One-sided p : $H_1 : \beta_+ > \beta_-$	< 0.001	0.002	< 0.001	0.001
Firm fixed effects	Yes	Yes	Yes	Yes
SIC2 $\times$ year effects	Yes	Yes	Yes	Yes
R^2	0.440	0.610	0.468	0.607
Observations	66,738	53,560	38,518	31,933
Firms	4,238	3,994	1,849	1,772

Notes: The dependent variable is future sales of property, plant, and equipment (SPPE), scaled by total assets at time t . The $t + 1$ specification uses $SPPE_{i,t+1}/AT_{it}$, whereas the three-year specification uses cumulative SPPE over $t + 1$ through $t + 3$, scaled by AT_{it} . The dependent variables are winsorized at the 1st and 99th percentiles separately by sample and horizon. The positive and negative filtered states are defined as $\hat{\eta}_{it}^+ = \max(\hat{\eta}_{it}, 0)$ and $\hat{\eta}_{it}^- = \max(-\hat{\eta}_{it}, 0)$, respectively. The control vector corresponds to the full set of time-varying investment determinants in Equations (10) and (11), shifted to period t : SR_{it} , $SR_{i,t-1}$, Q_{it} , $SIZE_{it}$, $PROF_{it}$, $PROF_{i,t-1}$, LEV_{it} , CF_{it} , and $CASH_{it}$. All controls enter linearly. All specifications include firm fixed effects and a full set of SIC2-by-year effects. Standard errors clustered at the firm level are reported in parentheses. Stars refer to conventional two-sided coefficient tests: * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$. The reported hypothesis-test p -values are directional one-sided tests of $H_1 : \beta_+ > 0$ and $H_1 : \beta_+ > \beta_-$.

Table 11 provides strong support for the first hypothesis. The coefficient on the positive filtered state is positive in every specification and remarkably similar across the two long-history samples. At the one-year horizon, the estimates are 0.0277 for $T_i > 10$ and 0.0272 for $T_i > 20$, with one-sided p -values below 0.001 in both cases. For cumulative asset sales over the subsequent three years, the corresponding coefficients rise to 0.0591 and 0.0618 and remain statistically significant, with one-sided p -values below 0.001 and 0.002, respectively. Thus, persistent positive abnormal-investment deviations are systematically followed by greater subsequent asset sales, and the association is not confined to the immediately following year.

The magnitudes are also economically interpretable. Because the latent investment deviation and the asset-sales outcome are expressed as ratios, a one-percentage-point increase in the positive latent-state magnitude is associated with approximately 0.027 percentage points higher asset sales in the following year and approximately 0.059–0.062 percentage points higher cumulative asset sales over the subsequent three years. The close correspondence of these estimates across the two long-history samples further indicates that the relation is not sensitive to the amount of within-firm history used to estimate the latent state.

The results also strongly support the second, asymmetric hypothesis. The estimated differences $\hat{\beta}_+ - \hat{\beta}_-$ equal 0.0393 and 0.1096 in the $T_i > 10$ sample and 0.0497 and 0.1486 in the $T_i > 20$ sample for the one-year and cumulative three-year outcomes, respectively. The null of equal associations is rejected in favor of $\beta_+ > \beta_-$ in all four specifications, with one-sided p -values ranging from below 0.001 to 0.002. The negative-state coefficients are non-positive throughout and are only weakly significant in the stricter $T_i > 20$ sample. We therefore place greater emphasis on the direct asymmetry test than on the separate interpretation of β_- .

Overall, the asset-sales results provide further evidence on the economic relevance of the estimated latent state. Persistent positive abnormal-investment deviations predict greater subsequent sales of property, plant, and equipment than equally large negative deviations, and this asymmetry remains virtually unchanged after conditioning on the full set of time-varying investment determinants in Equations (10) and (11). The pattern is consistent with subsequent capital reallocation following persistent positive investment episodes and with the agency, continuation, and costly-reversal mechanisms discussed in Section 3. Because asset sales may also reflect efficient restructuring or changes in investment opportunities, however,

the evidence remains predictive and reduced form rather than a structural test of investment inefficiency.

5. Conclusion

This paper develops a dynamic latent-state framework for measuring persistent abnormal corporate investment. Unlike conventional residual-based measures, the approach separates persistent deviations from transitory disturbances and allows persistence to differ between positive and negative investment states. It therefore adds a dynamic dimension to the measurement of investment efficiency by distinguishing the persistence of abnormal investment from its contemporaneous magnitude.

Monte Carlo evidence further supports the finite-sample reliability of the latent-state procedure: dynamic filtering improves latent-state recovery and regime classification relative to raw investment residuals, regime-specific persistence is recovered accurately, and the two-sided test of persistence asymmetry is well sized and highly powerful even with relatively short within-firm histories.

Our empirical analysis provides four main findings. First, abnormal investment exhibits strong asymmetric persistence. Using the preferred nonlinear investment benchmark, positive abnormal-investment deviations display substantially stronger local propagation than negative deviations, with estimated persistence coefficients of approximately 0.92 and 0.69–0.73, respectively, across the two long-history samples. These estimates imply conditional half-lives of approximately 8–9 years for positive deviations, compared with about 2 years for negative deviations. Symmetric persistence is strongly rejected. The greater propagation of positive investment deviations is consistent with managerial agency, weak external discipline, continued access to internal resources, and organizational or contractual commitments that can make ongoing investment programs costly to reverse, although the reduced-form evidence does not separately identify these mechanisms.

Second, the results reveal a clear distinction between two dimensions of corporate investment persistence. Despite the pronounced asymmetry in local state propagation, regime-continuation probabilities are high and relatively similar across positive and negative states,

with only a modest difference in favor of the positive regime. Thus, the much slower conditional decay of positive abnormal-investment deviations is not accompanied by a comparably large difference in the probability of remaining in the same regime. This finding illustrates why local state propagation and regime continuation should be treated as distinct dimensions of corporate investment adjustment.

Third, persistence asymmetry displays substantial cross-sectional heterogeneity. It is particularly pronounced among smaller and less profitable firms, primarily because negative abnormal-investment deviations exhibit weaker propagation in these groups, while positive-state persistence remains high. The relationship between cash flow and persistence asymmetry is less uniform across the two long-history samples. These findings suggest that firm characteristics associated with financing capacity are related to the dynamics of abnormal investment, but do not support a simple liquidity-based explanation. Agency, organizational commitments, and adjustment or reversal costs remain plausible complementary mechanisms.

Fourth, the estimated latent state has economically meaningful predictive content for subsequent capital reallocation. Persistent positive abnormal-investment deviations are followed by greater subsequent sales of property, plant, and equipment, and this association is significantly stronger than that associated with equally large negative deviations. The result holds across both long-history samples and for asset sales measured over the subsequent year and cumulatively over the following three years. Importantly, the finding remains virtually unchanged after conditioning on the full set of time-varying investment determinants used in the benchmark specifications. Because subsequent asset sales may also reflect efficient restructuring or changes in investment opportunities, this evidence should be interpreted as support for the economic relevance of the latent state rather than as a structural test of investment inefficiency.

These findings have implications for empirical corporate-finance research. Two firms with similar contemporaneous abnormal-investment residuals may face very different subsequent adjustment paths if the persistent components of those residuals propagate at different rates. Studies of investment efficiency, financing constraints, governance, liquidity, or managerial incentives may therefore benefit from distinguishing the contemporaneous magnitude of abnormal investment from its persistence and from the probability of remaining in the same

investment regime. The latent-state framework provides a way to separate these economically distinct dimensions of corporate investment behavior.

Appendix

This appendix provides technical results and implementation details supporting the theoretical and empirical methodology of the paper. Appendices A–E present the formal results and proofs underlying Section 2, covering stability of the latent-state process, residual-based regime misclassification, probabilistic regime classification, identification of the state-space parameters, and two-step inference with generated residuals. Appendix F presents the Monte Carlo analysis in greater detail, assessing the finite-sample performance of the latent-state estimator and associated inference procedures. Appendix G provides implementation details for the GAM benchmark investment equation used in Section 4.

Appendix A. Stability of the latent-state process

For notational simplicity, the firm index is suppressed. The latent state evolves according to

$$\eta_t = \delta + \rho_+ \eta_{t-1} \mathbf{1}\{\eta_{t-1} \geq 0\} + \rho_- \eta_{t-1} \mathbf{1}\{\eta_{t-1} < 0\} + w_t.$$

Define

$$f(x) = \delta + \rho_+ x \mathbf{1}\{x \geq 0\} + \rho_- x \mathbf{1}\{x < 0\},$$

so that

$$\eta_t = f(\eta_{t-1}) + w_t.$$

Let

$$\rho^* = \max\{|\rho_+|, |\rho_-|\}.$$

Proposition A1 (Stability). Suppose that $\rho^* < 1$ and that w_t is iid with $E|w_t| < \infty$. Then the latent-state process admits a unique stationary and ergodic solution.

Proof. For any $x, y \in \mathbb{R}$, the common intercept δ cancels when $f(x) - f(y)$ is formed. Since the slopes on the two sides of the threshold are bounded in absolute value by ρ^* ,

$$|f(x) - f(y)| \leq \rho^* |x - y|.$$

Thus, f is globally Lipschitz with contraction coefficient $\rho^* < 1$. Consider two realizations of the process driven by the same innovation sequence but initialized at different values. Then

$$|\eta_t^{(1)} - \eta_t^{(2)}| \leq \rho^* |\eta_{t-1}^{(1)} - \eta_{t-1}^{(2)}|.$$

Iterating,

$$|\eta_t^{(1)} - \eta_t^{(2)}| \leq (\rho^*)^t |\eta_0^{(1)} - \eta_0^{(2)}|.$$

Since $\rho^* < 1$, the effect of the initial condition vanishes geometrically. Standard results for contractive stochastic recurrence equations therefore imply existence, uniqueness, and ergodicity of the stationary solution. $\square$

The baseline parameter space in Section 2.2 imposes $0 \leq \rho_+ < 1$ and $0 \leq \rho_- < 1$, and is therefore a special case of the stability condition above.

Appendix B. Residual-Based Regime Misclassification

For notational simplicity, the firm index is suppressed. The observed abnormal-investment residual satisfies

$$u_t = \eta_t + \varepsilon_t, \tag{B.1}$$

where η_t is the persistent latent state and ε_t is the transitory disturbance. Define the latent regime as

$$s_t = \begin{cases} +, & \eta_t \geq 0, \\ -, & \eta_t < 0, \end{cases}$$

and the conventional residual-based classification as

$$\tilde{s}_t = \begin{cases} +, & u_t \geq 0, \\ -, & u_t < 0. \end{cases}$$

Residual-based misclassification occurs whenever $\tilde{s}_t \neq s_t$.

Proposition B1 (Residual-sign misclassification). Suppose that $\varepsilon_t \sim N(0, \sigma_\varepsilon^2)$, with $\sigma_\varepsilon^2 > 0$, and that ε_t is independent of η_t . Then

$$P(\tilde{s}_t \neq s_t \mid \eta_t) = \Phi\left(-\frac{|\eta_t|}{\sigma_\varepsilon}\right), \tag{B.2}$$

where $\Phi(\cdot)$ denotes the standard normal distribution function. Consequently,

$$P(\tilde{s}_t \neq s_t) = E\left[\Phi\left(-\frac{|\eta_t|}{\sigma_\varepsilon}\right)\right] > 0. \tag{B.3}$$

Proof. If $\eta_t > 0$, misclassification occurs when $\varepsilon_t < -\eta_t$, so that

$$P(\tilde{s}_t \neq s_t \mid \eta_t) = \Phi\left(-\frac{\eta_t}{\sigma_\varepsilon}\right).$$

If $\eta_t < 0$, misclassification occurs when $\varepsilon_t > -\eta_t$, yielding

$$P(\tilde{s}_t \neq s_t \mid \eta_t) = \Phi\left(\frac{\eta_t}{\sigma_\varepsilon}\right).$$

Both expressions equal $\Phi(-|\eta_t|/\sigma_\varepsilon)$. Integrating over the stationary distribution of η_t gives Equation (B.3). $\square$

Equation (B.2) shows that residual-sign classification is particularly uncertain when the persistent latent state lies close to the zero threshold. The probability of misclassification declines as $|\eta_t|$ increases, but remains positive for any finite $|\eta_t|$ under Gaussian measurement noise.

Appendix C. Probabilistic Regime Classification

For notational simplicity, the firm index is suppressed. Let

$$\gamma = (\beta', \vartheta)'$$

denote the full parameter vector entering the filtering procedure, where β determines the benchmark-investment residuals and ϑ contains the state-space parameters. Let γ_0 denote the true parameter vector and $\hat{\gamma}$ its estimator.

Let F_t denote the information set used by the filter through period t . Define the filtered probability of a positive persistent abnormal-investment state as

$$p_t^+(\gamma) = P_\gamma(\eta_t \geq 0 \mid F_t). \tag{C.1}$$

The feasible filtered probability is

$$\hat{p}_t^+ = p_t^+(\hat{\gamma}). \tag{C.2}$$

The corresponding probability-based classification rule is

$$\hat{s}_t = \begin{cases} +, & \hat{p}_t^+ > 1/2, \\ -, & \hat{p}_t^+ < 1/2, \\ \emptyset, & \hat{p}_t^+ = 1/2, \end{cases} \quad (\text{C.3})$$

where $\emptyset$ denotes an unclassified observation. The positive and negative regimes correspond, respectively, to persistent positive and persistent negative abnormal investment.

Assumption C1. The following conditions hold:

1. The estimator of the full parameter vector is consistent:

$$\hat{\gamma} \xrightarrow{P} \gamma_0.$$

2. For each t , the filtering probability $p_t^+(\gamma)$ is continuous in γ in a neighborhood of γ_0 .
3. The latent-state process is stationary and ergodic, and the nonlinear filter is stable with respect to its initialization.
4. The oracle posterior probability does not place positive probability exactly at the classification threshold:

$$P \left\{ p_t^+(\gamma_0) = \frac{1}{2} \right\} = 0.$$

Proposition C1 (Consistency of posterior probabilities and regime classification).

Under Assumption C1,

$$\hat{p}_t^+ - p_t^+(\gamma_0) = o_p(1). \quad (\text{C.4})$$

Consequently, the feasible probability-based classification converges to the oracle classification:

$$P(\hat{s}_t = s_t^*) \rightarrow 1, \quad (\text{C.5})$$

where

$$s_t^* = \begin{cases} +, & p_t^+(\gamma_0) > 1/2, \\ -, & p_t^+(\gamma_0) < 1/2. \end{cases} \quad (\text{C.6})$$

Proof. By consistency,

$$\hat{\gamma} \xrightarrow{p} \gamma_0.$$

Continuity of the filtering probability in γ and the continuous mapping theorem imply

$$p_t^+(\hat{\gamma}) \xrightarrow{p} p_t^+(\gamma_0).$$

Hence,

$$\hat{p}_t^+ - p_t^+(\gamma_0) = o_p(1),$$

which establishes Equation (C.4).

A discrepancy between $\hat{s}_t$ and s_t^* can occur only when estimation error moves the estimated posterior probability across the one-half classification threshold. For any $\epsilon > 0$, define

$$A_t(\epsilon) = P\left(\left|\hat{p}_t^+ - p_t^+(\gamma_0)\right| > \epsilon\right)$$

and

$$B_t(\epsilon) = P\left(\left|p_t^+(\gamma_0) - \frac{1}{2}\right| \leq \epsilon\right).$$

Then

$$P(\hat{s}_t \neq s_t^*) \leq A_t(\epsilon) + B_t(\epsilon). \quad (\text{C.7})$$

By Equation (C.4), $A_t(\epsilon) \rightarrow 0$ for every fixed $\epsilon > 0$. Assumption C1 implies that $B_t(\epsilon) \rightarrow 0$ as $\epsilon \downarrow 0$. Therefore,

$$P(\hat{s}_t \neq s_t^*) \rightarrow 0,$$

which establishes Equation (C.5). $\square$

Classification uncertainty. Proposition C1 establishes convergence of the feasible classification to the classification that would be obtained using the true model parameters. It does not imply that every realized latent regime can be recovered without error.

Define the realized latent regime as

$$s_t = \begin{cases} +, & \eta_t \geq 0, \\ -, & \eta_t < 0. \end{cases}$$

Conditional on F_t , the probability that the oracle classification differs from the realized latent regime is

$$P(s_t^* \neq s_t \mid F_t) = \min \{p_t^+(\gamma_0), 1 - p_t^+(\gamma_0)\}. \quad (\text{C.8})$$

Consequently, the unconditional oracle classification error is

$$R^* = E \left[\min \{p_t^+(\gamma_0), 1 - p_t^+(\gamma_0)\} \right]. \quad (\text{C.9})$$

This quantity is generally strictly positive when the transitory disturbance has nonzero variance. Parameter consistency therefore eliminates the additional discrepancy between the feasible and oracle classification rules, but it does not eliminate the irreducible uncertainty associated with recovering an unobserved latent state from noisy measurements.

Appendix D. Identification of the state-space parameters

For notational simplicity, the firm index is suppressed. The observed abnormal-investment measure satisfies

$$u_t = \eta_t + \varepsilon_t. \quad (\text{D.1})$$

The latent state evolves according to

$$\eta_t = \delta + \rho_+ \eta_{t-1} \mathbf{1}\{\eta_{t-1} \geq 0\} + \rho_- \eta_{t-1} \mathbf{1}\{\eta_{t-1} < 0\} + w_t. \quad (\text{D.2})$$

The loading of η_t in the measurement equation is normalized to one, fixing the scale of the latent state, and the regime threshold is normalized at zero.

Assumption D1. The following conditions hold:

1. The persistence parameters satisfy

$$|\rho_+| < 1, \quad |\rho_-| < 1.$$

2. The state innovation satisfies

$$E(w_t \mid \eta_{t-1}) = 0, \quad E(w_t^2) < \infty.$$

3. The measurement disturbances $\{\varepsilon_t\}$ are iid over time, independent of the latent-state process and the state innovations, with

$$E(\varepsilon_t) = 0, \quad E(\varepsilon_t^2) < \infty.$$

4. The distribution of ε_t is known or separately identified and its characteristic function

$$\phi_\varepsilon(s) = E[\exp(is\varepsilon_t)]$$

is nonzero for all $s \in \mathbb{R}$.

5. The stationary distribution of η_{t-1} assigns positive probability to both regimes and exhibits nondegenerate variation within each regime:

$$P(\eta_{t-1} \geq 0) > 0, \quad P(\eta_{t-1} < 0) > 0,$$

and

$$\text{Var}(\eta_{t-1} \mid \eta_{t-1} \geq 0) > 0, \quad \text{Var}(\eta_{t-1} \mid \eta_{t-1} < 0) > 0.$$

Proposition D1 (Identification of regime-dependent persistence). Under Assumption D1, the parameters δ , ρ_+ , and ρ_- are identified from the joint distribution of the observable abnormal-investment process $\{u_t\}$, conditional on identification of the measurement-disturbance distribution.

Proof. Consider the joint characteristic function of two consecutive observed outcomes:

$$\phi_{u_{t-1}, u_t}(s, r) = E[\exp\{i(su_{t-1} + ru_t)\}].$$

Using Equation (D.1) and the independence of the measurement disturbances,

$$\phi_{u_{t-1}, u_t}(s, r) = \phi_{\eta_{t-1}, \eta_t}(s, r) \phi_\varepsilon(s) \phi_\varepsilon(r). \quad (\text{D.3})$$

Since $\phi_\varepsilon(s) \neq 0$ for all s ,

$$\phi_{\eta_{t-1}, \eta_t}(s, r) = \frac{\phi_{u_{t-1}, u_t}(s, r)}{\phi_\varepsilon(s) \phi_\varepsilon(r)}. \quad (\text{D.4})$$

Hence, the observable joint distribution of (u_{t-1}, u_t) , together with the identified measurement-disturbance distribution, uniquely determines the joint distribution of (η_{t-1}, η_t) .

The identified joint distribution determines the conditional mean

$$m(x) = E(\eta_t \mid \eta_{t-1} = x),$$

which, from Equation (D.2), satisfies

$$m(x) = \begin{cases} \delta + \rho_+ x, & x \geq 0, \\ \delta + \rho_- x, & x < 0. \end{cases} \quad (\text{D.5})$$

Let $R_+ = \{\eta_{t-1} \geq 0\}$ and $R_- = \{\eta_{t-1} < 0\}$. Because $E(w_t \mid \eta_{t-1}) = 0$, the regime-specific slopes satisfy

$$\rho_s = \frac{\text{Cov}(\eta_t, \eta_{t-1} \mid R_s)}{\text{Var}(\eta_{t-1} \mid R_s)}, \quad s \in \{+, -\}. \quad (\text{D.6})$$

The denominators are strictly positive under Assumption D1. Hence, ρ_+ and ρ_- are identified from the identified joint distribution of the latent state. The common intercept is then identified from

$$\delta = E(\eta_t \mid R_s) - \rho_s E(\eta_{t-1} \mid R_s), \quad s \in \{+, -\}. \quad (\text{D.7})$$

Under the maintained model, Equation (D.7) gives the same value of δ from either regime. Therefore, δ , ρ_+ , and ρ_- are jointly identified. $\square$

Corollary D1 (Identification of persistence asymmetry). Because ρ_+ and ρ_- are separately identified, their difference

$$d = \rho_- - \rho_+$$

is also identified. Consequently, the null hypothesis

$$H_0 : \rho_+ = \rho_-$$

is testable from the observable distribution of abnormal investment.

Joint identification of the innovation variances. In the Gaussian state-space specification used in the empirical analysis, the complete parameter vector is

$$\vartheta = (\delta, \rho_+, \rho_-, \sigma_w^2, \sigma_\varepsilon^2)'.$$

When the innovation variances are estimated jointly with the state parameters, local identification of ϑ additionally requires the expected information matrix

$$\mathcal{I}(\vartheta_0) = -E \left[\frac{\partial^2 \ell_i(\vartheta_0)}{\partial \vartheta \partial \vartheta'} \right]$$

to be nonsingular at the true parameter value. This condition rules out local observational equivalence among changes in the state intercept, regime-specific persistence, and the signal-to-noise ratio.

Identification of the state-space parameters does not imply that every realized latent state is recovered without error. Individual states remain subject to filtering uncertainty and should therefore be summarized through their filtered or smoothed conditional distributions.

Appendix E. Generated Residuals and Two-Step Inference

The empirical procedure is implemented in two steps. In the first step, the benchmark investment equation is estimated and the generated residuals $\hat{u}_{it} = I_{it} - X'_{it}\hat{\beta}$ are constructed. In the second step, these residuals are used in the state-space likelihood to estimate ϑ .

Because $\hat{u}_{it}$ depends on the estimated first-stage parameter $\hat{\beta}$, inference on ϑ must in general account for first-stage estimation uncertainty. We therefore treat the two stages jointly through a stacked estimating-equation representation. Let $g_i(\beta)$ denote the firm-level estimating equations for the benchmark investment equation and let $s_i(\vartheta, \beta)$ denote the firm-level score of the state-space likelihood. Define

$$\Psi_i(\gamma) = \begin{pmatrix} g_i(\beta) \\ s_i(\vartheta, \beta) \end{pmatrix}, \quad \gamma = (\beta', \vartheta')'.$$

The two-step estimator satisfies the sample counterparts of these estimating equations.

Proposition E1 (Two-step asymptotic distribution). Suppose that firms are independent across i , or satisfy an appropriate cluster-dependence condition, that T_i is bounded or suitably controlled, and that the estimating equations satisfy standard identification, differentiability, and moment conditions. Then

$$\sqrt{N}(\hat{\gamma} - \gamma_0) \xrightarrow{d} N\left(0, G^{-1}\Omega(G^{-1})'\right),$$

where

$$G = E\left[\frac{\partial \Psi_i(\gamma_0)}{\partial \gamma'}\right], \quad \Omega = E[\Psi_i(\gamma_0)\Psi_i(\gamma_0)'].$$

Proof. A first-order expansion of the stacked sample estimating equations around γ_0 gives

$$0 = \frac{1}{\sqrt{N}} \sum_{i=1}^N \Psi_i(\gamma_0) + G\sqrt{N}(\hat{\gamma} - \gamma_0) + o_p(1).$$

Therefore,

$$\sqrt{N}(\hat{\gamma} - \gamma_0) = -G^{-1} \frac{1}{\sqrt{N}} \sum_{i=1}^N \Psi_i(\gamma_0) + o_p(1).$$

The result follows from a central limit theorem for the firm-level estimating equations. $\square$

Because the first-stage estimating equations do not depend on ϑ , the Jacobian has the block form

$$G = \begin{pmatrix} G_{\beta\beta} & 0 \\ G_{\vartheta\beta} & G_{\vartheta\vartheta} \end{pmatrix}.$$

Consequently, the influence representation for the second-stage estimator is

$$\sqrt{N}(\hat{\vartheta} - \vartheta_0) = -G_{\vartheta\vartheta}^{-1} \left[\frac{1}{\sqrt{N}} \sum_{i=1}^N s_i(\vartheta_0, \beta_0) - G_{\vartheta\beta} G_{\beta\beta}^{-1} \frac{1}{\sqrt{N}} \sum_{i=1}^N g_i(\beta_0) \right] + o_p(1).$$

The second term inside the brackets shows explicitly how uncertainty in the estimated benchmark investment equation enters inference on the state-space parameters. It disappears only under the special condition $G_{\vartheta\beta} = 0$; it is not assumed to vanish in the baseline analysis.

The corresponding covariance matrix for $\hat{\vartheta}$ is obtained from the relevant block of the full sandwich covariance matrix. In the empirical implementation, firm-level bootstrap inference provides a direct numerical counterpart to this result: firms are resampled and both the benchmark investment equation and the state-space model are re-estimated in every bootstrap replication.

This result concerns consistency and inference for the finite-dimensional parameter vector. It does not imply consistent pointwise recovery of each realized latent state as N increases. In particular, when the measurement disturbance has nonzero variance and the time dimension for each firm is bounded, filtering uncertainty generally remains even when the model parameters are known. As shown in Appendix D, posterior regime classification therefore remains subject to irreducible uncertainty.

Appendix F. Monte Carlo Analysis

This appendix evaluates the finite-sample performance of the proposed latent-state framework. The simulations examine four questions. First, can the model recover regime-specific persistence parameters when the persistent abnormal-investment state is observed only through noisy regression residuals? Second, does latent-state filtering improve state recovery and regime classification relative to the conventional residual-based approach? Third, does the

Wald test of symmetric persistence exhibit reliable finite-sample size and power? Fourth, how sensitive are these properties to the amount of within-firm time-series information available to the estimator?

The data-generating process follows the structure of the model developed in Section 2. For firm $i = 1, \dots, N$ and period $t = 1, \dots, T$, let

$$X_{it} = (1, S_{i,t-1}, S_{i,t-2})', \quad \beta_0 = (\alpha, \beta_1, \beta_2)'.$$

The benchmark investment equation in Equation (1), together with the decomposition in Equation (3), is specialized to

$$I_{it} = X_{it}'\beta_0 + u_{it} = \alpha + \beta_1 S_{i,t-1} + \beta_2 S_{i,t-2} + \eta_{it} + \varepsilon_{it}, \quad (\text{F1})$$

where S_{it} denotes sales, η_{it} is the persistent latent component of abnormal investment, and ε_{it} is the transitory disturbance. Sales evolve according to the stationary autoregressive process

$$S_{it} = \gamma S_{i,t-1} + \xi_{it}. \quad (\text{F2})$$

The latent state follows the asymmetric threshold process in Equation (4). Thus, persistence depends on the sign of the lagged latent state. Consistent with the interpretation in Section 2.2, positive values of η_{it} define the overinvestment regime and negative values define the underinvestment regime. Accordingly, ρ_+ measures persistence in the overinvestment region, whereas ρ_- measures persistence in the underinvestment region.

The state intercept δ is calibrated separately in each design so that the invariant distribution of the latent state satisfies $E(\eta_{it}) = 0$. Since $E(\varepsilon_{it}) = 0$, this also implies $E(u_{it}) = 0$, aligning the simulated abnormal-investment process with the centering induced by a benchmark regression containing an intercept. As discussed in Section 2.2, under asymmetric persistence $\delta = 0$ would not in general imply a zero-mean invariant distribution. The innovations are generated as

$$\xi_{it} \sim N(0, \sigma_\xi^2), \quad w_{it} \sim N(0, \sigma_w^2), \quad \varepsilon_{it} \sim N(0, \sigma_\varepsilon^2), \quad (\text{F3})$$

and are mutually independent across firms and time. Sales are initialized from their stationary Gaussian distribution. The latent state is initialized using a 750-period burn-in, providing a numerical draw from the invariant distribution implied by Equation (4).

The simulations use $N = 500$ firms and $R = 600$ Monte Carlo replications. To examine sensitivity to the within-firm time dimension, two panel lengths are considered: $T = 12$ and $T = 25$. Because the benchmark investment equation contains two lags of sales, the first two periods are used to construct the regressors. The $T = 12$ design therefore contains 10 usable observations per firm, or 5,000 usable firm-period observations per replication, whereas the $T = 25$ design contains 23 usable observations per firm, or 11,500 usable firm-period observations. The cross-sectional dimension is held fixed at $N = 500$ so that the experiment isolates the effect of shortening the within-firm history, which is the finite-sample dimension of primary interest here. All remaining DGP parameters are identical across the two panel lengths. The common DGP parameters are set to $\alpha = 0.20, \beta_1 = 0.60, \beta_2 = 0.20, \gamma = 0.50$, with $\sigma_\xi = 1.00, \sigma_\varepsilon = 0.80$. Three persistence designs are considered. The first is a symmetric benchmark,

$$\rho_+ = \rho_- = 0.90, \quad \sigma_w = 0.40.$$

This design is used to evaluate parameter recovery under symmetric persistence and the empirical size of the Wald test. The second design introduces moderate asymmetry,

$$\rho_+ = 0.75, \quad \rho_- = 0.90, \quad \sigma_w = 0.40.$$

The third design introduces stronger asymmetry,

$$\rho_+ = 0.55, \quad \rho_- = 0.90, \quad \sigma_w = 0.35.$$

In the two asymmetric designs, we set $\rho_- > \rho_+$ as a convenient normalization for organizing the Monte Carlo experiment rather than as a maintained directional prediction. Holding the persistence of one regime fixed while lowering that of the other generates moderate and

strong degrees of asymmetry and allows us to evaluate parameter recovery and test power as $|\rho_- - \rho_+|$ increases. The objective of the experiment is therefore to assess the ability of the estimator to recover regime-specific persistence and detect departures from symmetry, rather than to privilege one particular ordering of ρ_+ and ρ_- .

Under the symmetric Gaussian innovation structure used in the Monte Carlo design, reversing the ordering of the two persistence coefficients and reflecting the latent state around zero generates the corresponding mirror experiment in which overinvestment is more persistent than underinvestment. With the centering intercept adjusted accordingly, the positive and negative regime labels are interchanged, while the finite-sample implications for parameter recovery, latent-state recovery, regime classification, and the two-sided test of persistence symmetry are analogous. We therefore report one orientation of the asymmetric designs to avoid duplicating essentially mirror-image simulation results.

Each simulated panel is estimated in two stages, following the procedure described in Section 2.3. In the first stage, Equation (F1) is estimated by OLS using an intercept and the first two lags of sales. The abnormal-investment residual is then constructed as in Equation (2),

$$\hat{u}_{it} = I_{it} - X'_{it}\hat{\beta} = I_{it} - \hat{\alpha} - \hat{\beta}_1 S_{i,t-1} - \hat{\beta}_2 S_{i,t-2},$$

so that, up to first-stage estimation error,

$$\hat{u}_{it} \approx \eta_{it} + \varepsilon_{it}.$$

Three procedures are then used to estimate the latent state and classify the investment regime. The first is the conventional residual-based procedure. An observation is classified as overinvestment when $\hat{u}_{it} \geq 0$ and as underinvestment when $\hat{u}_{it} < 0$. The residual $\hat{u}_{it}$ is also treated directly as an estimate of the latent state.

The second procedure estimates a restricted latent-state model imposing $\rho_+ = \rho_-$. This symmetric filter retains the dynamic signal-extraction structure but rules out regime-dependent persistence. It therefore separates the gains arising from dynamic filtering from those attributable specifically to allowing asymmetric persistence.

The third procedure estimates the unrestricted asymmetric latent-state model in Equation (4). The parameter vector

$$\vartheta = (\delta, \rho_+, \rho_-, \sigma_w^2, \sigma_\varepsilon^2)'$$

is estimated jointly by maximum likelihood using the first-stage residuals as the observed measurement series. As in Section 2.3, the likelihood is evaluated using a nonlinear grid-based filter. The Monte Carlo implementation uses 61 state points over the interval $[-7, 7]$.

For both latent-state models, state estimation and regime classification follow the definitions in Section 2.3. The filtered posterior mean,

$$\hat{\eta}_{it} = E_{\hat{\vartheta}}(\eta_{it} \mid \mathcal{F}_{it}),$$

is used as the point estimate of the latent state, while classification is based on the posterior probability

$$\hat{p}_{it}^+ = P_{\hat{\vartheta}}(\eta_{it} \geq 0 \mid \mathcal{F}_{it}).$$

An observation is classified as overinvesting when $\hat{p}_{it}^+ > 1/2$ and as underinvesting when $\hat{p}_{it}^+ < 1/2$. Thus, consistent with Section 2.3, classification is based on posterior regime probabilities rather than on the sign of the filtered posterior mean.

Finite-sample performance is evaluated along four dimensions. First, parameter recovery is assessed using the Monte Carlo mean, bias, and root mean squared error of $\hat{\rho}_+$ and $\hat{\rho}_-$. Second, latent-state recovery is evaluated using the mean squared error

$$\text{MSE}(\hat{\eta}) = \frac{1}{N(T-2)} \sum_{i=1}^N \sum_{t=3}^T (\hat{\eta}_{it} - \eta_{it})^2. \quad (\text{F4})$$

The MSE of the raw regression residual is compared with those of the symmetric and asymmetric filtered-state estimates. Third, classification accuracy is defined as the proportion of usable firm-period observations for which the probability-based estimated regime agrees with the true sign of η_{it} . Accuracy is reported for the complete sample and for observations satisfying $|\eta_{it}| > 0.10$, thereby separating overall classification performance from performance near the zero threshold. Because the true latent state is observed by construction in the Monte

Carlo experiment, this comparison is a finite-sample diagnostic of classification performance; it does not imply error-free recovery of individual latent states in the empirical application.

Finally, inference on asymmetric persistence is evaluated using the Wald test in Equation (8),

$$H_0 : \rho_+ = \rho_-,$$

with

$$d(\vartheta) = \rho_- - \rho_+.$$

The Monte Carlo analysis focuses on the two-sided alternative

$$H_1 : \rho_+ \neq \rho_-,$$

or equivalently,

$$H_1 : d(\vartheta) \neq 0.$$

The empirical size of the Wald procedure is evaluated under the symmetric design, whereas its power is evaluated under the moderate- and strong-asymmetry designs. These two asymmetric designs correspond to different magnitudes of

$$|d(\vartheta)| = |\rho_- - \rho_+|,$$

and therefore allow us to assess how the power of the symmetry test changes as the degree of persistence asymmetry increases. As discussed above, the reported simulations use the normalization $\rho_- > \rho_+$, but the test itself is two-sided and does not impose a particular direction of asymmetry. Together, the experiments assess whether the proposed framework can recover regime-specific persistence, improve signal extraction and regime classification relative to raw regression residuals, provide reliable finite-sample inference on asymmetric adjustment, and remain effective when the available within-firm history is substantially shorter.

Table F1 reports regime-classification accuracy under the symmetric-persistence design, in which

$$\rho_+ = \rho_- = 0.90.$$

Because the data-generating process imposes common persistence across positive and negative latent states, this design provides a useful benchmark for evaluating the gains from dynamic signal extraction independently of persistence asymmetry.

Both latent-state filters improve classification relative to the conventional residual-sign procedure at both panel lengths. With $T = 12$, average classification accuracy increases from 77.1% for the first-stage residual to 81.4% under either latent-state filter. With $T = 25$, the corresponding increase is from 77.2% to 81.9%. The gain is also present away from the regime boundary. For observations satisfying $|\eta_{it}| > 0.10$, classification accuracy rises from 79.5% to 84.0% when $T = 12$, and from 79.5% to 84.6% when $T = 25$.

Table F1: Monte Carlo Classification Accuracy under Symmetric Persistence

Method	Classification accuracy	Accuracy away from threshold
Panel A. $T = 12$ (10 usable observations)		
Residual-sign classification	0.771	0.795
Symmetric filter	0.814	0.840
Asymmetric filter	0.814	0.840
Panel B. $T = 25$ (23 usable observations)		
Residual-sign classification	0.772	0.795
Symmetric filter	0.819	0.846
Asymmetric filter	0.819	0.846

Note: This table reports average regime-classification accuracy across 600 Monte Carlo replications under the symmetric-persistence design, $\rho_+ = \rho_- = 0.90$, with $N = 500$ firms. Panel A uses $T = 12$, providing 10 usable observations per firm after constructing the two sales lags. Panel B uses $T = 25$, providing 23 usable observations per firm. Classification accuracy is the fraction of observations for which the estimated regime agrees with the true sign of the simulated latent state η_{it} . “Away from threshold” restricts the calculation to observations satisfying $|\eta_{it}| > 0.10$. Residual-sign classification assigns an observation to the overinvestment regime when $\hat{u}_{it} \geq 0$ and to the underinvestment regime when $\hat{u}_{it} < 0$. The symmetric and asymmetric filters use the posterior regime probability $\hat{p}_{it}^+$, with an observation classified as overinvesting when $\hat{p}_{it}^+ > 1/2$ and as underinvesting when $\hat{p}_{it}^+ < 1/2$.

The comparison across panels shows that classification performance is highly stable when the available firm history is shortened. Moving from 23 to 10 usable observations reduces filtered classification accuracy by only about 0.5 percentage points in the full sample and by about 0.6 percentage points away from the threshold. The symmetric and asymmetric filters remain virtually identical under the symmetric DGP, as they should. Thus, the unrestricted

filter does not generate spurious classification gains when persistence is common, and the classification advantage of dynamic filtering does not depend on the relatively longer $T = 25$ design.

Having established the classification performance of the estimators under the symmetric benchmark, Table F2 evaluates their ability to approximate the magnitude of the simulated latent state. This is a more demanding criterion than regime classification because it evaluates the distance between the estimated and realized latent states rather than only whether their regime classifications agree.

Table F2 reports the mean squared error between the estimated and true latent states under the symmetric, moderate-asymmetry, and strong-asymmetry designs. For the residual-based procedure, the first-stage abnormal-investment residual $\hat{u}_{it}$ is treated directly as an estimate of η_{it} . For the symmetric and asymmetric latent-state models, the state estimate is the filtered posterior mean defined in Section 2.3,

$$\hat{\eta}_{it} = E_{\hat{\theta}}(\eta_{it} \mid \mathcal{F}_{it}).$$

The asymmetric designs in Table F2 are reported under the normalization $\rho_- > \rho_+$; analogous finite-sample conclusions obtain under the opposite ordering.

The results show that dynamic filtering produces large improvements in state-estimation accuracy at both panel lengths. In the $T = 12$ design, the asymmetric filter reduces MSE relative to the raw residual by approximately 62% under symmetric persistence, 66% under moderate asymmetry, and 74% under strong asymmetry. The corresponding reductions for $T = 25$ are approximately 64%, 67%, and 75%. Thus, even with only 10 usable observations per firm, most of the signal-extraction gain obtained in the longer panel is preserved.

As expected, the shorter panel produces somewhat larger filtered-state MSEs because less within-firm information is available to the filter. The increase, however, is modest: the asymmetric-filter MSE rises from 0.2320 to 0.2445 under symmetric persistence, from 0.2094 to 0.2171 under moderate asymmetry, and from 0.1600 to 0.1648 under strong asymmetry. The incremental gain from allowing regime-dependent persistence rather than imposing a symmetric filter remains comparatively small in both panel lengths. It is 0.16%, 1.07%, and

Table F2: Monte Carlo Latent-State Recovery

Design	Residual MSE	Symmetric MSE	Asymmetric MSE	Gain (%)
Panel A. $T = 12$ (10 usable observations)				
Symmetric: (0.90, 0.90)	0.6420	0.2449	0.2445	0.16
Moderate: (0.75, 0.90)	0.6409	0.2194	0.2171	1.07
Strong: (0.55, 0.90)	0.6396	0.1674	0.1648	1.56
Panel B. $T = 25$ (23 usable observations)				
Symmetric: (0.90, 0.90)	0.6405	0.2325	0.2320	0.18
Moderate: (0.75, 0.90)	0.6402	0.2112	0.2094	0.85
Strong: (0.55, 0.90)	0.6401	0.1633	0.1600	2.00

Note: The table reports mean squared errors (MSEs) for latent-state recovery, averaged across 600 Monte Carlo replications with $N = 500$. Panel A uses $T = 12$ and Panel B uses $T = 25$. The persistence parameters in parentheses are (ρ_+, ρ_-) . “Residual” treats the first-stage abnormal-investment residual $\hat{u}_{it}$ as the state estimate. “Symmetric” uses the filtered posterior mean $\hat{\eta}_{it}$ from the latent-state model imposing $\rho_+ = \rho_-$, whereas “Asymmetric” uses the filtered posterior mean from the unrestricted model in Equation (4). “Gain” is the percentage reduction in MSE of the asymmetric filter relative to the symmetric filter.

1.56% in Panel A and 0.18%, 0.85%, and 2.00% in Panel B. The main finite-sample gain therefore continues to come from dynamic signal extraction itself, while the asymmetric specification additionally permits the two persistence parameters to be estimated separately.

Table F3 compares classification accuracy across all three persistence designs. Both latent-state filters outperform residual-sign classification in every design and at both panel lengths. Under moderate asymmetry, the $T = 12$ classification rates are 73.2%, 77.0%, and 76.7% for the residual-sign procedure, symmetric filter, and asymmetric filter, respectively. The corresponding $T = 25$ rates are 73.2%, 77.5%, and 77.1%. Under strong asymmetry, the $T = 12$ rates are 68.4%, 71.2%, and 70.5%, compared with 68.4%, 71.6%, and 70.8% when $T = 25$.

The comparison across panels again indicates that classification is only mildly affected by the shorter within-firm history. The reduction in filtered classification accuracy is generally between 0.3 and 0.5 percentage points. Dynamic filtering therefore continues to improve regime classification by several percentage points relative to the raw residual even in the $T = 12$ design. The symmetric filter produces slightly higher classification rates than

Table F3: Monte Carlo Regime-Classification Accuracy

Design	T	Residual sign	Symmetric filter	Asymmetric filter
Symmetric	12	0.771	0.814	0.814
Moderate	12	0.732	0.770	0.767
Strong	12	0.684	0.712	0.705
Symmetric	25	0.772	0.819	0.819
Moderate	25	0.732	0.775	0.771
Strong	25	0.684	0.716	0.708

Note: The table reports average regime-classification accuracy across 600 Monte Carlo replications with $N = 500$. For $T = 12$ and $T = 25$, there are 10 and 23 usable observations per firm, respectively. The symmetric, moderate-, and strong-asymmetry designs use $(\rho_+, \rho_-) = (0.90, 0.90)$, $(0.75, 0.90)$, and $(0.55, 0.90)$. “Residual sign” classifies observations using $\text{sign}(\hat{u}_{it})$. The latent-state filters classify an observation as positive when $\hat{p}_{it}^+ > 1/2$ and negative when $\hat{p}_{it}^+ < 1/2$.

the asymmetric filter in the asymmetric designs. This does not imply that the symmetric model estimates the latent state or regime-specific persistence more accurately: classification accuracy records only whether the estimated regime agrees with the true sign of η_{it} , not the accuracy of the state magnitude or of the persistence parameters.

Table F4 examines recovery of the regime-specific persistence parameters and the finite-sample performance of the Wald test of symmetric persistence. Under the symmetric design, persistence recovery is essentially unchanged when the panel is shortened. With $T = 12$, the Monte Carlo mean estimates are $\hat{\rho}_+ = 0.900$ and $\hat{\rho}_- = 0.899$, and the Wald test rejects in 5.3% of replications. With $T = 25$, the corresponding estimates are again 0.900 and 0.899, with a rejection frequency of 5.5%. Thus, the unrestricted estimator does not generate spurious persistence asymmetry, and the Wald test remains close to its nominal 5% size even with only 10 usable observations per firm.

Under moderate asymmetry, the $T = 12$ mean estimates are 0.749 and 0.897, compared with true values of 0.75 and 0.90, while the $T = 25$ estimates are 0.749 and 0.899. Under strong asymmetry, shortening the panel has a more visible effect: the $T = 12$ mean estimates are 0.546 and 0.886, compared with true values of 0.55 and 0.90, whereas the $T = 25$ estimates are 0.552 and 0.897. The shorter panel therefore produces somewhat greater finite-sample attenuation in the more demanding strong-asymmetry design, but the average estimates remain close to their population values.

Table F4: Monte Carlo Recovery of Persistence Parameters and Wald-Test Performance

Design	True ρ_+	Mean $\hat{\rho}_+$	True ρ_-	Mean $\hat{\rho}_-$	Wald rejection
Panel A. $T = 12$ (10 usable observations)					
Symmetric persistence	0.900	0.900	0.900	0.899	0.053
Moderate asymmetry	0.750	0.749	0.900	0.897	0.957
Strong asymmetry	0.550	0.546	0.900	0.886	0.956
Panel B. $T = 25$ (23 usable observations)					
Symmetric persistence	0.900	0.900	0.900	0.899	0.055
Moderate asymmetry	0.750	0.749	0.900	0.899	1.000
Strong asymmetry	0.550	0.552	0.900	0.897	1.000

Note: This table reports true and Monte Carlo mean estimates of the regime-specific persistence parameters across 600 replications with $N = 500$. Panel A uses $T = 12$, corresponding to 10 usable observations per firm, and Panel B uses $T = 25$, corresponding to 23 usable observations per firm. The final column reports rejection frequencies for the two-sided Wald test of $H_0 : \rho_+ = \rho_-$ at the 5% significance level. The rejection frequency under the symmetric design measures empirical size, whereas those under the moderate- and strong-asymmetry designs measure statistical power.

The loss of power from shortening the panel is small. The rejection frequency falls from 100% in the $T = 25$ design to 95.7% under moderate asymmetry and 95.6% under strong asymmetry when $T = 12$. Thus, even with a short within-firm history, the Wald test retains high power against both degrees of persistence asymmetry while maintaining accurate size under the null.

Overall, the Monte Carlo results indicate that the finite-sample performance of the proposed latent-state framework is not dependent on a relatively long within-firm time dimension. Reducing the simulated panel from $T = 25$ to $T = 12$, corresponding to a decline from 23 to 10 usable observations per firm, produces only small changes in regime-classification accuracy and modest increases in latent-state MSE. The largest gains continue to arise from dynamic signal extraction: even in the shorter panel, filtering substantially improves latent-state recovery relative to the conventional investment residual. Regime-specific persistence is estimated accurately under the symmetric and moderate-asymmetry designs, while the strong-asymmetry design exhibits somewhat greater finite-sample attenuation when $T = 12$, as expected from the more limited within-firm information. Importantly, the Wald test remains well sized under the null and retains power of approximately 96% against both asymmetric designs considered. Overall, the simulations show that the main finite-sample

conclusions are robust to a substantially shorter panel dimension while preserving the central advantages of the latent-state approach.

Appendix G. Generalized Additive Model for the Benchmark Investment Equation

This appendix provides additional details on the generalized additive model (GAM) used to construct the preferred benchmark-investment residual. The GAM retains an additive specification in economically motivated covariates while allowing the conditional effect of each continuous firm-level fundamental to depart from global linearity. It also includes Mundlak-type firm-level means and industry-by-year fixed effects to account for persistent firm heterogeneity and industry-specific time variation.

Specification

Consistent with equation (11) in the main text, the GAM investment benchmark is

$$\begin{aligned}
I_{it} = & \alpha + s_1(SR_{i,t-1}) + s_2(SR_{i,t-2}) + s_3(Q_{i,t-1}) + s_4(SIZE_{i,t-1}) + s_5(PROF_{i,t-1}) \\
& + s_6(PROF_{i,t-2}) + s_7(LEV_{i,t-1}) + s_8(CF_{i,t-1}) + s_9(CASH_{i,t-1}) \\
& + \bar{\mathbf{X}}_i' \gamma + \lambda_{j(i),t} + u_{it}^G,
\end{aligned} \tag{G1}$$

where $\bar{\mathbf{X}}_i$ contains the firm-specific means of the seven distinct investment determinants, two of which enter with two lags, and $\lambda_{j(i),t}$ denotes SIC2-by-year fixed effects. The nine time-varying determinants enter through smooth functions, while the Mundlak controls and industry-by-year effects enter additively without smoothing penalties.

The GAM is estimated once on the common first-stage sample of 143,145 firm-year observations from 15,207 firms. Its fitted values and abnormal-investment residuals are carried unchanged into the $T_i > 10$ and $T_i > 20$ consecutive-history samples, within which the nonlinear latent-state model is estimated separately. Hence, differences in persistence estimates across the two samples do not reflect re-estimation of the benchmark equation.

Let

$$m(\mathbf{X}_{it}, \bar{\mathbf{X}}_i, j(i), t) = E(I_{it} \mid \mathbf{X}_{it}, \bar{\mathbf{X}}_i, j(i), t)$$

denote the conditional benchmark, where $\mathbf{X}_{it}$ collects the nine time-varying regressors. The abnormal-investment measurement series is

$$\hat{u}_{it}^G = I_{it} - \widehat{m}_G(\mathbf{X}_{it}, \bar{\mathbf{X}}_i, j(i), t).$$

If the conditional mean is nonlinear but a globally linear benchmark is imposed, systematic nonlinear components may remain in the estimated residual. Likewise, omitted persistent firm heterogeneity or industry-specific time variation may contaminate the residual. Because this residual is subsequently decomposed into persistent and transitory components, such misspecification could be attributed incorrectly to the latent abnormal-investment state. The GAM reduces this risk by flexibly estimating the conditional mean while accounting for these additional sources of heterogeneity before applying the latent-state decomposition.

Basis representation and penalization

Each smooth term has the finite-basis representation

$$s_j(x) = \sum_{k=1}^{K_j} b_{jk}(x) \theta_{jk}, \quad (\text{G2})$$

where $b_{jk}(x)$ is a basis function and θ_{jk} its coefficient. Following the standard penalized-regression-spline formulation of GAMs (see Wood (2017)), excessive curvature is penalized to limit overfitting.

For the Gaussian identity-link model, estimation can be represented schematically as

$$\min_{\alpha, \gamma, \{\lambda_{j,t}\}, \{\theta_j\}} \left\{ \sum_{(i,t) \in \mathcal{C}} \left[I_{it} - \alpha - \sum_{j=1}^9 s_j(X_{jit}) - \bar{\mathbf{X}}_i' \gamma - \lambda_{j(i),t} \right]^2 + \sum_{j=1}^9 \kappa_j \theta_j' S_j \theta_j \right\}, \quad (\text{G3})$$

where $\mathcal{C}$ denotes the common first-stage sample, S_j is the penalty matrix for smooth j , and κ_j its smoothing parameter. The Mundlak controls and SIC2-by-year fixed effects are estimated

jointly with the smooth coefficients but are not penalized. Larger κ_j values impose stronger curvature penalties and produce smoother fitted relationships, whereas smaller values allow greater local flexibility.

The smoothing parameters are selected by restricted maximum likelihood (REML), following Wood (2011). The resulting effective complexity of each smooth is summarized by its effective degrees of freedom, allowing the fitted relationships to adapt to the data while limiting excessive flexibility.

Effective degrees of freedom and approximate F -tests

The effective degrees of freedom (EDF) summarize the complexity of each fitted smooth after penalization. An EDF close to one corresponds approximately to a linear fitted relationship, whereas values materially above one indicate additional nonlinear flexibility. EDF need not be integers.

The GAM output also reports reference degrees of freedom, Ref.df, used in the approximate significance test for each smooth. The reported F -statistic tests whether the smooth contributes to explaining investment and should not be interpreted as a direct test of linearity against nonlinearity. Evidence supporting the GAM as a useful nonlinear benchmark therefore comes jointly from the EDF, the estimated smooth shapes and uncertainty bands, the improvement in evaluation-sample fit, and the firm-level out-of-sample predictive performance reported in Table 4.

References

- Abel, A. B., & Eberly, J. C. (1994). A unified model of investment under uncertainty. *American Economic Review*, 84(5), 1369–1384.
- Almeida, H., & Campello, M. (2007). Financial constraints, asset tangibility, and corporate investment. *Review of Financial Studies*, 20(5), 1429–1460.
- Biddle, G. C., Hilary, G., & Verdi, R. S. (2009). How does financial reporting quality relate to investment efficiency? *Journal of Accounting and Economics*, 48(2–3), 112–131.
- Blickle, K. S., & Santos, J. A. C. (2024). The costs of corporate debt overhang. *Journal of*

Financial Intermediation, 60, 101118.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3), 623–685.

Bond, S., & Meghir, C. (1994). Dynamic investment models and the firm's financial policy. *Review of Economic Studies*, 61(2), 197–222.

Borri, K. T., Martins-Filho, C., & Kalatzis, A. E. G. (2022). Exploring nonlinearities between investment and internal funds: Evidence of the U-shaped investment curve. *Economics Letters*, 218, 110713.

Caballero, R. J. (1999). Aggregate investment. In J. B. Taylor & M. Woodford (Eds.), *Handbook of Macroeconomics* (Vol. 1B, pp. 813–862). Elsevier.

Cao, D., Lorenzoni, G., & Walentin, K. (2019). Financial frictions, investment, and Tobin's q . *Journal of Monetary Economics*, 103, 105–122.

Chen, F., Hope, O.-K., Li, Q., & Wang, X. (2011). Financial reporting quality and investment efficiency of private firms in emerging markets. *The Accounting Review*, 86(4), 1255–1288.

Cooper, R., & Haltiwanger, J. (2006). On the nature of capital adjustment costs. *Review of Economic Studies*, 73(3), 611–633.

Dixit, A. K., & Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton University Press.

Durbin, J., & Koopman, S. J. (2012). *Time Series Analysis by State Space Methods* (2nd ed.). Oxford University Press.

Ee, M. S., Iftekhar, H., & Huang, H. (2022). Stock liquidity and corporate labor investment. *Journal of Corporate Finance*, 72, 102142.

Erickson, T., & Whited, T. M. (2000). Measurement error and the relationship between investment and q . *Journal of Political Economy*, 108(5), 1027–1057.

Fazzari, S. M., Hubbard, R. G., & Petersen, B. C. (1988). Financing constraints and corporate investment. *Brookings Papers on Economic Activity*, 19(1), 141–206.

- Gulen, H., & Ion, M. (2016). Policy uncertainty and corporate investment. *Review of Financial Studies*, 29(3), 523–564.
- Harford, J., Mansi, S. A., & Maxwell, W. F. (2008). Corporate governance and firm cash holdings in the U.S. *Journal of Financial Economics*, 87(3), 535–555.
- Hastie, T. J., & Tibshirani, R. J. (1990). *Generalized Additive Models*. Chapman & Hall.
- Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1), 213–224.
- Hennessey, C. A., & Whited, T. M. (2007). How costly is external financing? Evidence from a structural estimation. *Journal of Finance*, 62(4), 1705–1745.
- Jensen, M. C. (1986). Agency costs of free cash flow, corporate finance, and takeovers. *American Economic Review*, 76(2), 323–329.
- Jiang, F., Cai, W., Wang, X., & Zhu, B. (2018). Multiple large shareholders and corporate investment: Evidence from China. *Journal of Corporate Finance*, 50, 66–83.
- Kalemli-Özcan, Ş., Laeven, L., & Moreno, D. (2022). Debt overhang, rollover risk, and corporate investment: Evidence from the European crisis. *Journal of the European Economic Association*, 20(6), 2353–2395.
- Kaplan, S. N., & Zingales, L. (1997). Do investment-cash flow sensitivities provide useful measures of financing constraints? *Quarterly Journal of Economics*, 112(1), 169–215.
- Lang, L., Ofek, E., & Stulz, R. M. (1996). Leverage, investment, and firm growth. *Journal of Financial Economics*, 40(1), 3–29.
- Myers, S. C. (1977). Determinants of corporate borrowing. *Journal of Financial Economics*, 5(2), 147–175.
- Myers, S. C., & Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics*, 13(2), 187–221.
- Richardson, S. (2006). Over-investment of free cash flow. *Review of Accounting Studies*,

11(2–3), 159–189.

Stiglitz, J. E., & Weiss, A. (1981). Credit rationing in markets with imperfect information. *American Economic Review*, 71(3), 393–410.

Tobin, J. (1969). A general equilibrium approach to monetary theory. *Journal of Money, Credit and Banking*, 1(1), 15–29.

Whited, T. M. (2006). External finance constraints and the intertemporal pattern of intermittent investment. *Journal of Financial Economics*, 81(3), 467–502.

Wood, S. N. (2011). Fast stable restricted maximum likelihood and marginal likelihood estimation of semiparametric generalized linear models. *Journal of the Royal Statistical Society: Series B*, 73(1), 3–36.

Wood, S. N. (2017). *Generalized Additive Models: An Introduction with R* (2nd ed.), CRC Press.



**Department of Economics
Athens University of Economics and Business**

List of Recent Working Papers

2024

- 01-24 Market Timing & Predictive Complexity, Stelios Arvanitis, Foteini Kyriazi, Dimitrios Thomakos**
- 02-24 Multi-Objective Frequentistic Model Averaging with an Application to Economic Growth, Stelios Arvanitis, Mehmet Pinar, Thanasis Stengos, Nikolas Topaloglou**
- 03-24 State dependent fiscal multipliers in a Small Open Economy, Xiaoshan Chen, Jilei Huang, Petros Varthalitis**
- 04-24 Public debt consolidation: Aggregate and distributional implications in a small open economy of the Euro Area, Eleftherios-Theodoros Roumpanis**
- 05-24 Intangible investment during the Global Financial Crisis in the EU, Vassiliki Dimakopoulou, Stelios Sakkas and Petros Varthalitis**
- 06-24 Time will tell! Towards the construction of instantaneous indicators of different agent-types, Iordanis Kalaitzoglou, Stelios Arvanitis**
- 07-24 Norm Constrained Empirical Portfolio Optimization with Stochastic Dominance: Robust Optimization Non-Asymptotics, Stelios Arvanitis**
- 08-24 Asymptotics of a QLR-type test for optimal predictive ability, Stelios Arvanitis**
- 09-24 Strongly Equitable General Equilibrium Allocations, Angelos Angelopoulos**
- 10-24 The Greek macroeconomy: A note on the current situation and future outlook, Apostolis Philippopoulos**
- 11-24 Evolution of Greek Tax System, A Survey of Legislated Tax Changes from 1974 to 2018, Panagiotis Asimakopoulos**
- 12-24 Macroeconomic Impact of Tax Changes, The case of Greece from 1974 to 2018, Panagiotis Asimakopoulos**
- 13-24 'Pareto, Edgeworth, Walras, Shapley' Equivalence in a Small Economy, Angelos Angelopoulos**
- 14-24 Stimulating long-term growth and welfare in the U.S, James Malley and Apostolis Philippopoulos**
- 15-24 Distributionally Conservative Stochastic Dominance via Subsampling, Stelios Arvanitis**
- 16-24 Before and After the Political Transition of 1974. Institutions, Politics, and the Economy of Post-War Greece, George Alogoskoufis**
- 17-24 Gaussian Stochastic Volatility, Misspecified Volatility Filters and Indirect Inference Estimation, Stelios Arvanitis, Antonis Demos**
- 18-24 Endogenous Realtor Intermediation and Housing Market Liquidity, Miroslav Gabrovski, Ioannis Kospentaris, Victor Ortego-Marti**
- 19-24 Universal Choice Spaces and Expected Utility: A Banach-type Functorial Fixed Point, Stelios Arvanitis, Pantelis Argyropoulos, Spyros Vassilakis**

2025

- 01-25 Democracy, redistribution, and economic growth: Some evidence from post-1974 Greece, George C. Bitros
- 02-25 Assessing Downside Public Debt Risks in an Environment of Negative Interest Rates Growth Differentials, Yiannis Dendramis, Georgios Dimitrakopoulos and Elias Tzavalis
- 03-25 Wicksellian General Equilibrium, Angelos Angelopoulos
- 04-25 Institutions and Ethics, Angelos Angelopoulos
- 05-25 General Equilibrium in a Capitalist Economy, Angelos Angelopoulos
- 06-25 Public investment multipliers revisited: The role of production complementarities, Vasiliki Dimakopoulou, George Economides, and Apostolis Philippopoulos
- 07-25 The heterogeneous causal effects of the EU's Cohesion Fund, Angelos Alexopoulos, Ilias Kostarakos, Christos Mylonakis, and Petros Varthalitis
- 08-25 Gaussian Invariant Markov Chain Monte Carlo, Michalis K. Titsias, Angelos Alexopoulos, Siran Liu, and Petros Dellaportas
- 09-25 Financial Intermediation, Investment, and Monetary Policy, Ioannis Kospentaris and Florian Madison
- 10-25 Temperature as a premium factor. Insurance incorporating climate change. A comparative case study between Europe and North America, Nomikou-Lazarou Irene
- 11-25 How High will our Buildings Become? Angelos Angelopoulos
- 12-25 Emergency Room General Equilibrium, Angelos Angelopoulos
- 13-25 Updating Density Estimates using Conditional Information Projection: Stock Index Returns and Stochastic Dominance Relations, Stelios Arvanitis, Richard McGee, and Thierry Post
- 14-25 Measuring Textual Cohesion via Graph Curvature, Weighted Dominance, and Embedding-Based Coherence, Stelios Arvanitis
- 15-25 Universal Preference Spaces for Expected Utility via a Banach-type Functorial Fixed Point, Stelios Arvanitis, Pantelis Argyropoulos, and Spyros Vassilakis
- 16-25 Planet Earth's General Equilibrium, Angelos Angelopoulos
- 17-25 Electable and Stable Insiders' Coalition Governments, Tryphon Kollintzas and Lambros Pechlivanos
- 18-25 Long Memory from Cheeger Bottlenecks and Long Cycles in Network Dynamics, Stelios Arvanitis
- 19-25 Robust Multiobjective Model Averaging in Predictive Regressions, Stelios Arvanitis, Mehmet Pinar, Thanasis Stengos, and Nikolas Topaloglou
- 20-25 On robust Bayesian causal inference, Angelos Alexopoulos and Nikolaos Demiris
- 21-25 Leverage Effect, Volatility Feedback and the Influence of Jumps, Orestis Agapitos, Ioannis Papantonis, Leonidas Rompolis, and Elias Tzavalis
- 22-25 Leadership: The other deficit of post-1974 Greece, George C. Bitros

2026

- 01-26 The Tell-Tale Clock! Speed and Agent Composition in High Frequency Trading, Stelios Arvanitis
- 02-26 Sentiment Utopia Index: Rhetorical Structure and Corpus-Relative Financial Sentiment, Stelios Arvanitis
- 03-26 Draghi's report a year later: Beneficial but no remedy for EU challenges, George C. Bitros, Anastasios G. Malliaris, and Mark S. Rzepczynski

- 04-26 Tax Structure Reform and Macroeconomic Strategy in Greece: Econometric Evidence and Policy Implications, Panagiotis Asimakopoulos
- 05-26 A Perspective on International Monetary Leadership. From Silver and Gold to the Dollar Standard, George Alogoskoufis
- 06-26 Behavioral Balanced Scorecard Framework for Tax Administrations: Integrating Tax Morale, Institutional Trust, and Compliance Behavior, Panagiotis Asimakopoulos
- 07-26 Education and governance: The Russell–Hayek Conundrum through the lenses of ancient Sparta and the “Sparta of the East”, George C. Bitros
- 08-26 Overhauling Democracy: Layered State Authority, Corporate-Style Accountability, and the Free Market Economy, George C. Bitros and Steve Bakalis
- 09-26 ECONOMICS 200. The Economics of transition to a Steady State Economy, Theodore P. Lianos
- 10-26 Empirical Likelihood Tests for Pairwise Stochastic Dominance, Stelios Arvanitis, Richard McGee, and Thierry Post
- 11-26 Production Function Estimation under Input Distortions: Identification Using Heteroscedasticity-Based Instruments, Dimitris K. Christopoulos and Elias Tzavalis
- 12-26 Learning Dimension-Reduced State-Adaptive Model Averaging in Regime-Dependent Environments, Stelios Arvanitis and Thanasis Stengos
- 13-26 The Persistence of Corporate Overinvestment and Underinvestment, Dimitris K. Christopoulos and Elias Tzavalis



Department of Economics Athens University of Economics and Business

The Department is the oldest Department of Economics in Greece with a pioneering role in organising postgraduate studies in Economics since 1978. Its priority has always been to bring together highly qualified academics and top quality students. Faculty members specialize in a wide range of topics in economics, with teaching and research experience in world-class universities and publications in top academic journals.

The Department constantly strives to maintain its high level of research and teaching standards. It covers a wide range of economic studies in micro-and macroeconomic analysis, banking and finance, public and monetary economics, international and rural economics, labour economics, industrial organization and strategy, economics of the environment and natural resources, economic history and relevant quantitative tools of mathematics, statistics and econometrics.

Its undergraduate program attracts high quality students who, after successful completion of their studies, have excellent prospects for employment in the private and public sector, including areas such as business, banking, finance and advisory services. Also, graduates of the program have solid foundations in economics and related tools and are regularly admitted to top graduate programs internationally. Three specializations are offered: 1. Economic Theory and Policy, 2. Business Economics and Finance and 3. International and European Economics. The postgraduate programs of the Department (M.Sc and Ph.D) are highly regarded and attract a large number of quality candidates every year.

For more information:

<https://www.dept.aueb.gr/en/econ/>