Capital Slack*

Isaac Baley†
UPF, CREI, BSE, CEPR

Miguel Bandeira‡
Insper, CFM

Andrés Blanco§
FRB Atlanta

Laura Castillo-Martínez¶
Duke, FRB Minneapolis, CEPR

Madalena Gaspar‖
UPF, BSE

Nicolás Oviedo**
UPF, BSE

February 17, 2026

Abstract

Investment is lumpy: firms remain inactive for long stretches and then adjust capital in bursts. This lumpiness drives a wedge between the capital firms would like to have and the capital they actually install. We recover this wedge—which we term *capital slack*—by estimating firms' latent desired capital from plant-level data. We microfound a nonlinear state-space representation using an Ss investment model with fixed adjustment costs and occasional free adjustment opportunities, and apply filtering and smoothing methods to infer reset capital in real time. Aggregating the recovered reset-capital series yields a capital slack index that leads the business cycle and predicts future movements in aggregate investment.

JEL: C11, D21, E22, E32

Keywords: frictionless investment, reset capital, investment slack, leading indicators, business cycles, lumpy investment, state-space models, Bayesian filtering and smoothing, microdata, capital adjustment frictions

---

*Baley acknowledges financial support from the Severo Ochoa Research Excellence Program (CEX2024- 001476-S) through Spain's State Research Agency (Agencia Estatal de Investigación), Spanish Ministry Grant PID2022-138443NB-I00, and the ERC Grant 101041334. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta, the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

†Universitat Pompeu Fabra, CREI, Barcelona School of Economics, and CEPR, isaac.baley@upf.edu

‡Insper and Centre for Macroeconomics (CFM), mbandeira@insper.edu.br

§Federal Reserve Bank of Atlanta and Emory University, julioablanco84@gmail.com

¶Duke University, Federal Reserve Bank of Minneapolis and CEPR, l.castillo-martinez@duke.edu

‖Universitat Pompeu Fabra and BSE, madalena.gaspar@upf.edu

**Universitat Pompeu Fabra and BSE, nicolas.oviedo@upf.edu

1 Introduction

Macroeconomists care about investment because it is central to business cycles. Yet investment is also hard to read in real time. The reason is microeconomic: capital adjustment is lumpy. Firms often remain inactive for long periods before adjusting in discrete episodes. In such environments, observed investment and observed capital stocks can substantially lag the evolution of fundamentals. The resulting wedge between what firms *install* and what they would *like to install* is a natural object for macro monitoring.

This paper formalizes that wedge as *capital slack*. The starting point is a latent firm-level object: reset capital $k^{\star}_{i,t}$, defined as the capital level a firm chooses when it adjusts. Reset capital tracks the evolution of fundamentals and desired scale, whereas observed capital $k_{i,t}$ is a sluggish, distorted measure of that target because fixed costs prevent continuous adjustment. We define capital slack as the gap between observed and reset capital (in logs),

$$s_{i,t} \equiv \log k_{i,t} - \log k^{\star}_{i,t}.$$

Slack accumulates during inaction because $k^{\star}_{i,t}$ continues to move with fundamentals while $k_{i,t}$ evolves mechanically. Slack is then fully released when firms adjust.

The central empirical challenge is that reset capital is unobserved: we observe only capital and sporadic adjustment episodes. Our approach treats this challenge as a state-space inference problem. Methodologically, we build a nonlinear state-space representation microfounded by an Ss investment model with fixed adjustment costs and occasional free adjustment opportunities (Baley and Blanco, 2021). The model implies a kinked measurement structure. When a firm is inactive, capital follows a deterministic depreciation law; when it adjusts, capital resets to the latent target $k^{\star}_{i,t}$. The latent state inherits the stochastic evolution of fundamentals, while the arrival of free adjustment opportunities introduces discrete regime variation. Crucially, inaction is informative: when the firm does not adjust, the Ss policy implies that the latent target must lie inside a band around undepreciated capital. These inequality restrictions are the main source of identification.

Following Bandeira, Castillo-Martınez and Wang (2024), we implement inference with a filtering–smoothing procedure tailored to kinked measurement equations. In the forward pass, we update beliefs about reset capital sequentially using Bayes' rule. In the backward pass, we refine these beliefs using the full sample, producing smoothed estimates of reset capital and associated uncertainty. These two objects play distinct roles in the results: filtered estimates capture what could be learned in real time from the data observed up to date t, while smoothed estimates provide a precision benchmark that clarifies the underlying signal.

Our empirical analysis uses plant-level manufacturing data from Chile. We estimate the model separately by subsector, allowing the inaction band and the arrival rate of adjustment opportunities to differ across industries while maintaining a common measurement structure. The posterior draws deliver sector-specific estimates of the inaction thresholds (L, U), the free adjustment opportunity rate λ, and the parameters governing the stochastic evolution of reset capital (μ, σ). We then recover plant-level paths of reset capital and slack, and we summarize them in a set of moments that make the economic content of the methodology transparent.

We first document the microeconomics of slack. The recovered slack distribution is highly non-degenerate and varies meaningfully over time: its median, dispersion, and tail behavior move with aggregate conditions. Importantly, slack is not just a cross-sectional curiosity. In the plant-level dynamics, slack builds up gradually during inaction spells and collapses when a plant adjusts, so the distribution of slack aggregates information about how much desired investment is currently "queued" across firms. We also show directly that slack predicts subsequent adjustment activity at the micro level: plants in the upper quartiles of the slack distribution are more likely to adjust in the following year, consistent with the Ss policy mapping from gaps to action.

To connect these objects to familiar investment facts, we summarize lumpy investment through its extensive and intensive margins. At the aggregate level, fluctuations in investment can reflect changes in the fraction of plants that adjust (the extensive margin) and changes in the average adjustment size conditional on adjustment (the intensive margin). We report both margins over time for Chilean manufacturing, showing that periods of high aggregate investment coincide with pronounced movements in adjustment frequency and adjustment size, consistent with large reallocations happening through episodic bursts rather than smooth marginal changes.

We then move from plant-level slack to sectoral and manufacturing slack. Using predetermined (lagged) capital weights, we construct sectoral aggregates of reset capital and slack, and we report analogous aggregates for manufacturing as a whole. The resulting aggregate slack series provides a compact, economically interpretable statistic: it measures the weighted gap between installed capital and the latent reset target that firms would choose upon adjustment. We also show that sectoral slack series co-move but are far from identical, highlighting that the buildup of latent investment pressure can be uneven across industries.

Finally, we evaluate slack as a leading indicator. We study predictive content using Granger-causality tests and forecasting regressions that map moments of the slack distribution to future investment over one- to four-year horizons in annual data. The key finding is directional: slack moments help forecast subsequent investment, whereas the reverse direction is weaker, consistent with slack behaving as a forward-looking state variable generated by micro-inaction rather than a mechanical transformation of realized investment. We also report out-of-sample performance metrics that quantify the improvement in forecast accuracy when slack moments are included, relative to investment-only benchmarks.

Contributions. The paper speaks to three strands of work. The first is the literature on lumpy investment and generalized Ss adjustment. Empirically, classic plant- and firm-level evidence documents infrequent, bursty investment and the relevance of fixed/nonconvex costs (e.g., Doms and Dunne, 1998; Cooper and Haltiwanger, 2006; Cooper *et al.*, 1999). On the theory and aggregation side, a large body of work studies irreversibility and nonconvex adjustment and their implications for aggregate investment dynamics (e.g., Pindyck, 1991; Bertola and Caballero, 1994; Caballero and Engel, 1993, 1999; Khan and Thomas, 2008), including quantitative general-equilibrium implementations that match micro lumpiness and its cyclical consequences (e.g., Thomas, 2002; Gourio and Kashyap, 2007; Bachmann *et al.*, 2013). Our approach is in this tradition, but with a different goal: rather than using the model to explain lumpiness, we leverage it as a *measurement device* to identify the latent reset policy and hence capital slack.

The second strand is the measurement of latent targets and "frictionless" counterfactuals in environments with adjustment frictions. The conceptual parallel with Bandeira, Castillo-Martınez and Wang (2024) is tight: menu costs create a wedge between observed prices and a latent reset target; fixed adjustment costs create a wedge between observed capital and a latent reset capital. The methodological novelty here is that the Ss investment policy produces a piecewise mapping from latent states to observables: whether the observed capital stock follows mechanical depreciation or jumps to the reset point is pinned down by the inaction thresholds (and, in our setup, occasional free adjustment opportunities).

The third strand is business-cycle monitoring using latent-index methods. A long macro tradition builds leading, and coincident indicators from multivariate information sets, including dynamic-factor and recession-probability approaches (e.g., Stock and Watson, 1989; Aruoba *et al.*, 2009; Chauvet and Piger, 2008). We contribute a complementary object that is explicitly microfounded: Slack is constructed from firm-level adjustment behavior and is mechanically linked to the Ss policy that governs when capital moves. This gives the aggregate slack index and its cross-sectional moments an interpretable structural content: they track the accumulation and release of investment pressure created by inaction, and therefore summarize the state of capital adjustment in the economy in a way that standard output-gap or utilization concepts do not.

2 A Parsimonious Investment Model

This section describes the economic environment that underlies our state-space representation of investment slack. We build a parsimonious Ss model of capital adjustment in which establishments face fixed costs of changing capital and, in addition, receive occasional free adjustment opportunities. The model is deliberately simple: its role is not to provide a fully articulated general-equilibrium theory of investment, but to deliver a transparent mapping from primitives to a kinked investment policy and, therefore, to an empirically useful measurement system.

2.1 Setup

Time is discrete and infinite. A continuum of ex ante identical firms produces output using capital subject to adjustment frictions. The economy is small and open, implying an exogenous constant real interest rate r.

Technology and shocks. Firm i produces output $y_{i,t}$ using capital $k_{i,t}$ according to a production function with decreasing returns to scale and no time-to-build:

$$y_{i,t} \;=\; u_{i,t}^{1-\alpha} k_{i,t}^{\alpha}, \qquad \alpha < 1. \tag{1}$$

Firm-level productivity $u_{i,t}$ follows a random walk with drift

$$\log u_{i,t} = \mu \;+\; \log u_{i,t-1} \;+\; \sigma\varepsilon_{i,t},\;\; \varepsilon_{i,t} \sim \mathcal{N}(0,1). \tag{2}$$

Shocks $\varepsilon_{i,t}$ are i.i.d. across plants i and time t. The capital stock, if uncontrolled, depreciates at a constant rate $\zeta > 0$.

Capital adjustment costs. The firm adjusts its capital stock by buying or selling capital. Any non-zero investment $i_{i,t} \equiv k_{i,t} - (1-\zeta)k_{i,t-1}$ entails an adjustment cost $\kappa_{i,t}$ proportional to idiosyncratic productivity $u_{i,t}$. Costs may differ for positive and negative investments. Free adjustment opportunities occur with probability λ.

Let $l_{i,t}$ indicate a free adjustment (and zero otherwise). We write the arrival of free adjustment opportunities using uniform shocks $v_{i,t}$, which are i.i.d. across plants i and time t:

$$l_{i,t} = \mathbb{1}\{v_{i,t} \leq \lambda\}, \quad v_{i,t} \;\sim\; \text{Uniform}[0,1]. \tag{3}$$

Capital adjustment costs are

$$\kappa_{i,t} \;\equiv\; \mathcal{K}(i_{i,t}, l_{i,t})u_{i,t} \tag{4}$$

where the function $\mathcal{K}$ captures asymmetric fixed costs and free adjustments:

$$\mathcal{K}(i_{i,t}, l_{i,t}) = \begin{cases} 0 & \text{if } i_{i,t} = 0 \;\text{ or }\; l_{i,t} = 1 \\ \kappa^{-} & \text{if } i_{i,t} > 0 \;\text{ and }\; l_{i,t} = 0 \\ \kappa^{+} & \text{if } i_{i,t} < 0 \;\text{ and }\; l_{i,t} = 0 \end{cases} \tag{5}$$

with $\kappa^{-} > 0$ and $\kappa^{+} > 0$ fixed.

Investment problem. Let $V(k_{i,0}, u_{i,0})$ be the firm's value. Given the initial conditions $(k_{i,0}, u_{i,0})$, the firm chooses a sequence of capital adjustment dates $\{T_{i,h}\}_{h=1}^{\infty}$ and investments $\{i_{T_{i,h}}\}_{h=1}^{\infty}$, where h counts the number of adjustments, to maximize its expected discounted stream of profits. The sequential problem of the firm is described by

$$V(k_{i,0}, u_{i,0}) \;=\; \max_{\{T_{i,h}, i_{T_{i,h}}\}_{h=1}^{\infty}} \mathbb{E}\left[\sum_{t=0}^{\infty} \frac{1}{(1+r)^t} y_{i,t} \;-\; \sum_{h=1}^{\infty} \frac{1}{(1+r)^{T_{i,h}}} \left(\kappa_{i,T_{i,h}} + i_{T_{i,h}}\right)\right], \tag{6}$$

subject to the production function (1), idiosyncratic productivity (2), idiosyncratic adjustment costs in (3) and (4), and the law of motion for its capital stock

$$\log k_{i,t} \;=\; \log(1-\zeta) \;+\; \log k_{i,t-1} \;+\; \log\left(1 + \frac{i_{i,t}}{(1-\zeta)k_{i,t-1}}\right). \tag{7}$$

2.2 Optimal investment policy

Following Baley and Blanco (2021), we restate the firm's problem in terms of the (log) capital-to-productivity ratio,

$$\hat{k}_{i,t} \;\equiv\; \log k_{i,t} - \log u_{i,t}. \tag{8}$$

In this formulation, the optimal investment policy has an Ss structure and can be summarized by three numbers $\{\hat{k}^-, \hat{k}^\star, \hat{k}^+\}$: a lower inaction threshold $\hat{k}^-$, an upper inaction threshold $\hat{k}^+$, and a reset point $\hat{k}^\star$. The policy is time-invariant and common across firms. When the firm adjusts, it resets $\hat{k}_{i,t}$ to $\hat{k}^\star$; when it does not adjust, it lets capital evolve mechanically.

It is convenient to parameterize the inaction region relative to the reset point. Define constants $L \leq 0$ and $U \geq 0$ such that

$$\hat{k}^- \;=\; \hat{k}^\star + L, \qquad \hat{k}^+ \;=\; \hat{k}^\star + U, \tag{9}$$

where $L \equiv \hat{k}^- - \hat{k}^\star$ and $U \equiv \hat{k}^+ - \hat{k}^\star$. This policy representation naturally motivates the two observed objects in the data: an inaction dummy and the capital stock.

Inaction Dummy. Let $f(k_{i,t-1}) \equiv \log(1-\zeta) + \log k_{i,t-1}$ denote the log of undepreciated capital (i.e., capital after depreciation but before any new investment). The Ss policy implies that, absent a free adjustment opportunity ($l_{i,t} = 0$), a plant remains inactive when undepreciated capital lies inside the inaction band around reset capital. If a plant received a free adjustment opportunity ($l_{i,t} = 1$), it remains inactive only if $f(k_{i,t-1}) = \log k^\star_{i,t}$. This logic yields the inaction indicator:

$$d_{i,t} \;=\; \mathbb{1}_{\left\{f(k_{i,t-1}) \in [\log k^\star_{i,t} + L,\ \log k^\star_{i,t} + U]\right\}} \cdot (1 - l_{i,t}) \;+\; \mathbb{1}_{\left\{f(k_{i,t-1}) = \log k^\star_{i,t}\right\}} \cdot l_{i,t} \tag{10}$$

Given $d_{i,t}$, the capital stock evolves according to

$$\log k_{i,t} \;=\; f(k_{i,t-1}) \cdot d_{i,t} \;+\; \log k^{\star}_{i,t} \cdot (1 - d_{i,t}) \tag{11}$$

When $d_{i,t} = 1$ (inaction), observed capital equals undepreciated capital. When $d_{i,t} = 0$ (adjustment), observed capital equals reset capital.

Equations (10) and (11) make the "kink" explicit: the likelihood contribution of $(k_{i,t}, d_{i,t})$ changes discontinuously depending on whether $f(k_{i,t-1})$ falls inside or outside the inaction band.

Reset capital. We now recover the idiosyncratic reset capital in period t in levels, denoted by $k^{\star}_{i,t}$, as follows:

$$\log k^{\star}_{i,t} \;\equiv\; \hat{k}^{\star} + \log u_{i,t}. \tag{12}$$

Because $\hat{k}^{\star}$ is constant, reset capital inherits the law of motion of idiosyncratic productivity in (2):

$$\log k^{\star}_{i,t} \;=\; \mu \;+\; \log k^{\star}_{i,t-1} \;+\; \sigma\varepsilon_{i,t}, \qquad \varepsilon_{i,t} \sim \mathcal{N}(0,1). \tag{13}$$

The combination of (13) and the inaction band $[k^{\star}_{i,t} + L,\; k^{\star}_{i,t} + U]$ delivers a state-space representation with a kinked measurement rule: conditional on $d_{i,t} = 1$, observed capital follows mechanical depreciation dynamics; conditional on $d_{i,t} = 0$, observed capital jumps to the reset level.

Figure I illustrates these dynamics in a simulated example. Panel (a) plots observed and reset log capital, along with the inaction band around the reset target: observed capital jumps to the reset point when the plant adjusts and otherwise drifts mechanically with depreciation while the reset target continues to move with idiosyncratic shocks. Panel (b) plots the implied investment slack, $s_{i,t}$, which accumulates during inaction spells and resets upon adjustment. Panel (c) reports the inaction indicator, with shaded regions marking years classified as inaction ($d_{i,t} = 1$).

3 State-Space Representation

This section recasts the Ss investment model as a nonlinear state-space system. The key advantage of the state-space form is that it provides a sharp mapping from observed capital dynamics and inaction to latent reset capital: the Ss policy implies a kinked measurement equation, in which the inaction indicator selects between mechanical depreciation and a reset to the latent target.

3.1 Notation

We observe plants $i = 1, \ldots, n$ over periods $t = 1, \ldots, T$. For any plant-level series $x_{i,t}$, we write $x_{i,1:t} \equiv (x_{i,1}, \ldots, x_{i,t})$ for plant i's history up to date t, and $x_{i,1:T}$ for its full-sample history. For a

Figure I – Capital, Slack and Inaction for a Single Firm

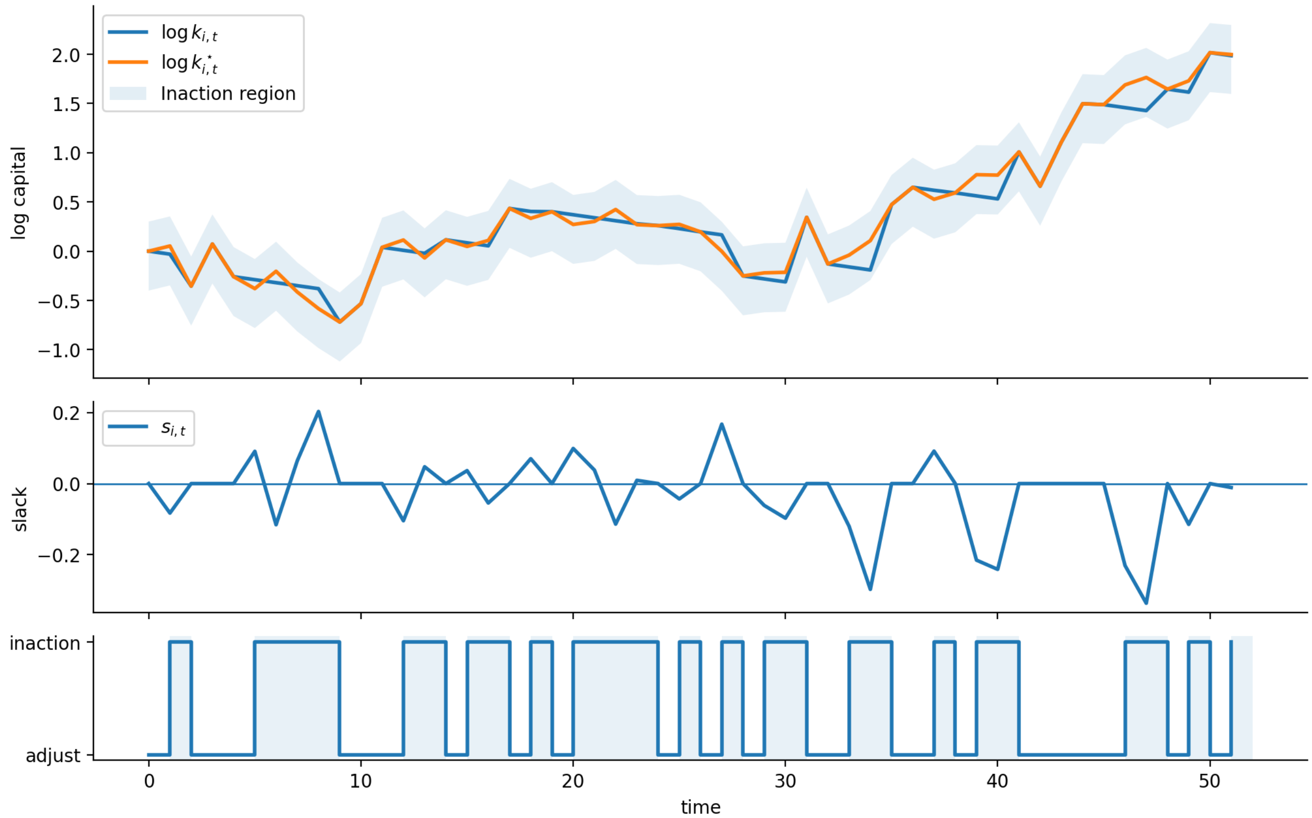

Note: Simulated paths from the Ss state-space model. The top panel plots observed log capital $\log k_{i,t}$ and latent reset capital $\log k^\star_{i,t}$; the shaded band is the inaction region $[\log k^\star_{i,t} + L,\ \log k^\star_{i,t} + U]$. The middle panel plots slack $s_{i,t} \equiv \log k_{i,t} - \log k^\star_{i,t}$. The bottom panel reports the inaction indicator $d_{i,t}$, with shaded areas marking periods of inaction ($d_{i,t} = 1$).

given date t, $x_{1:n,t} \equiv (x_{1,t}, \ldots, x_{n,t})$ denotes the cross-section across plants, while we denote with $x_{1:n,1:t} \equiv \{x_{i,\tau} : i = 1, \ldots, n,\ \tau = 1, \ldots, t\}$ the panel up to date t (and similarly for $x_{1:n,1:T}$).

The parameter vector we will estimate is $\theta \equiv [L,\ U,\ \lambda,\ \mu,\ \sigma]'$. Here (L, U) are the inaction thresholds recentered around the reset point, λ is the arrival rate of free adjustment opportunities, and (μ, σ) govern the drift and volatility of the latent reset-capital process. We treat the depreciation rate ζ as known and outside the estimation.

3.2 State-space

The state-space representation consists of two state-transition equations and two measurement equations. The two latent states are the reset capital $k^\star_{i,t}$, which follows (13), and the adjustment-opportunity indicator $l_{i,t}$, which follows (3). The observable variables are the inaction indicator $d_{i,t}$ in (10) and the log capital stock $\log k_{i,t}$ in (11). Together with an initial distribution for $(l_{i,0}, k^\star_{i,0})$, these elements define a nonlinear state-space model with a kinked measurement equation.

For convenience, we repeat the transition and measurement equations here:

$$l_{i,t} = \mathbb{1}\{v_{i,t} \le \lambda\}, \tag{14}$$

$$\log k^\star_{i,t} = \mu + \log k^\star_{i,t-1} + \sigma\varepsilon_{i,t}, \tag{15}$$

$$d_{i,t} = \mathbb{1}_{\left\{f(k_{i,t-1})\in[\log k^\star_{i,t}+L,\ \log k^\star_{i,t}+U]\right\}} \cdot (1-l_{i,t}) + \mathbb{1}_{\left\{f(k_{i,t-1})=\log k^\star_{i,t}\right\}} \cdot l_{i,t}, \tag{16}$$

$$\log k_{i,t} = f(k_{i,t-1}) \cdot d_{i,t} + \log k^\star_{i,t} \cdot (1-d_{i,t}), \tag{17}$$

where $v_{i,t} \sim_{iid} \text{Uniform}[0,1]$, $\varepsilon_{i,t} \sim_{iid} \mathcal{N}(0,1)$, and $(l_{i,0}, k^\star_{i,0})$ follow prespecified distributions.

3.3 Filtering and parameter estimation

Our empirical target is the latent reset-capital sequence $\{k^\star_{i,t}\}$, which corresponds to the capital firms would choose in the absence of adjustment frictions. However, recovering $\{k^\star_{i,t}\}$ requires knowledge of θ, which is unobserved. We therefore estimate the parameters and latent states jointly.

Formally, for each sector, we aim to characterize the joint posterior

$$p\big(\theta,\ \{k^\star_{i,t}, l_{i,t}\}_{i,t} \,\big|\, \{\log k_{i,t}, d_{i,t}\}_{i,t}\big) \tag{18}$$

The posterior reflects two distinct sources of information in the data. The time series of $\log k_{i,t}$ pins down when and by how much capital moves, while the inaction indicator $d_{i,t}$ imposes inequality restrictions implied by the Ss policy: when $d_{i,t} = 1$, undepreciated capital must lie inside the band around $\log k^\star_{i,t}$; when $d_{i,t} = 0$, observed capital must coincide with $\log k^\star_{i,t}$. These restrictions are precisely what make the measurement equation informative, even though $k^\star_{i,t}$ is latent.

To estimate states and parameters, we implement a Gibbs sampler algorithm that alternates between state and parameter blocks, which we describe next.[1]

State block: Given draws of the parameter vector θ, we recover plant-level latent states using a filtering–smoothing algorithm for nonlinear state-space models with kinked measurement equations. For each draw of θ, the *forward filter* constructs real-time posterior distributions,

$$p\big(k^\star_{i,t}, l_{i,t} \mid \log k_{i,1:t}, d_{i,1:t}, \theta\big)\,,$$

conditioning only on information available up to date t. The *backward smoother* then refines inference using the full sample,

$$p\big(k^\star_{i,t}, l_{i,t} \mid \log k_{i,1:T}, d_{i,1:T}, \theta\big)\,,$$

[1] Appendix A provides the full conditional distributions, the filtering and backward-sampling recursions, and implementation details.

delivering posterior means and credible intervals for the entire latent path. This system is both non-linear (due to the ”action/inaction” duality) and non-Gaussian (due to the random free adjustment opportunities).

Parameter block: Given the current states $(k^{\star}_{i,t},\, l_{i,t})$ and data $(\log k_{i,t}, d_{i,t})$, update θ and construct the posterior distribution

$$p\big(\theta \mid k^{\star}_{i,t}, l_{i,t}, \log k_{i,1:T}, d_{i,1:T}\big)$$

The parameter subset (μ, σ, λ) governs the state transitions and admits standard conjugate updates under normal, inverse-gamma, and beta priors, respectively. The thresholds (L, U) enter through the band restrictions in (10) and are updated with a Metropolis–Hastings algorithm that proposes new bands and accepts them based on the implied likelihood of the observed inaction/adjustment pattern.[2]

4 Empirical Application

This section implements our methodology on plant-level manufacturing data from Chile. The empirical application has three goals. First, we describe the data environment and document the lumpy adjustment patterns that motivate an Ss investment structure. Second, we construct the key inputs for estimation—capital stocks and inaction spells. Third, we estimate the model parameters for each manufacturing sector using our Bayesian state-space framework and recover the latent reset-capital series underlying our capital slack index.

4.1 Data

We use yearly data on manufacturing plants in Chile from the Annual National Manufacturing Survey (*Encuesta Nacional Industrial Anual*) for the period 1979 to 2011. Chilean National Accounts and Penn World Tables provide information on the depreciation rates and price deflators used to construct the capital series. We focus on the total capital stock. We consider plants that appear in the sample for at least 10 years (more than 60% of the sample) and have more than 10 workers. Appendix B describes the sample selection, variable construction, and analysis of each capital category separately: structures, vehicles, machinery, and equipment.

Capital stocks. We construct real capital stocks using the perpetual inventory method (PIM):

(19) $$k_{i,t} \;=\; (1-\zeta)\,k_{i,t-1} \;+\; \frac{I_{i,t}}{D_t},$$

[2]Appendix A provides additional information on the priors and the posterior distribution approximation.

Figure II – Capital Changes and Inaction Dummies

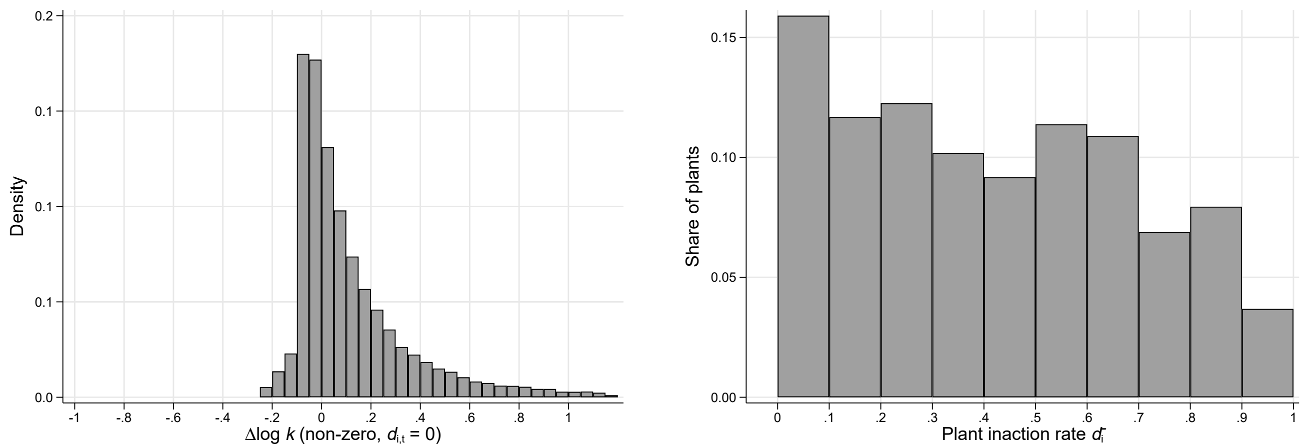

(a) Non-zero capital changes at adjustment **(b)** Inaction heterogeneity across plants

Note: Panel (a) depicts a positively skewed distribution of non-zero capital changes among adjustment observations ($d_{i,t} = 0$). Panel (b) depicts the distribution of plant inaction rates $\bar{d}_i$.

where ζ is the depreciation rate, D_t is the investment deflator, and $I_{i,t}$ is gross nominal investment (purchases plus upgrades/improvements minus sales). Initial capital $k_{i,0}$ is the plant's first non-negative reported book value, deflated by D_{t_0}.

Inaction dummies. For each plant i and year t, we construct an inaction dummy $d_{i,t} \in \{0, 1\}$ that records whether the plant is inactive in that year. Following Cooper and Haltiwanger (2006), we classify a plant-year observation as inaction when the investment rate is below 1% in absolute value, and set

$$d_{i,t} = \mathbb{1}_{\{|\iota_{i,t}| \leq 0.01\}}, \tag{20}$$

where $\iota_{i,t}$ denotes investment divided by lagged capital. We then define an adjustment date as any year with $d_{i,t} = 0$ and use the resulting sequence of dummies as the primary input to our state-space measurement equation, which conditions on whether capital evolves mechanically through depreciation (when $d_{i,t} = 1$) or jumps due to an adjustment (when $d_{i,t} = 0$). We also define a plant's average inaction rate as $\bar{d}_i \equiv T_i^{-1} \sum_t d_{i,t}$, where T_i denotes the number of observations of plant i. Finally, we truncate the investment-rate distribution at the 2nd and 98th percentiles to eliminate outliers.

Figure II summarizes the two key empirical inputs. Panel (a) plots the cross-sectional distribution of non-zero changes in (log) capital among adjustment observations, i.e., $\Delta \log k_{i,t}$ conditional on $d_{i,t} = 0$. The distribution is asymmetric and positively skewed. Panel (b) summarizes the inaction dummies by plotting the distribution of plants' inaction rates $\bar{d}_i$ (equivalently, the fraction of years with $d_{i,t} = 1$), which reveals substantial heterogeneity in inaction behavior across

Figure III – Time-series Behavior of Extensive and Intensive Margins

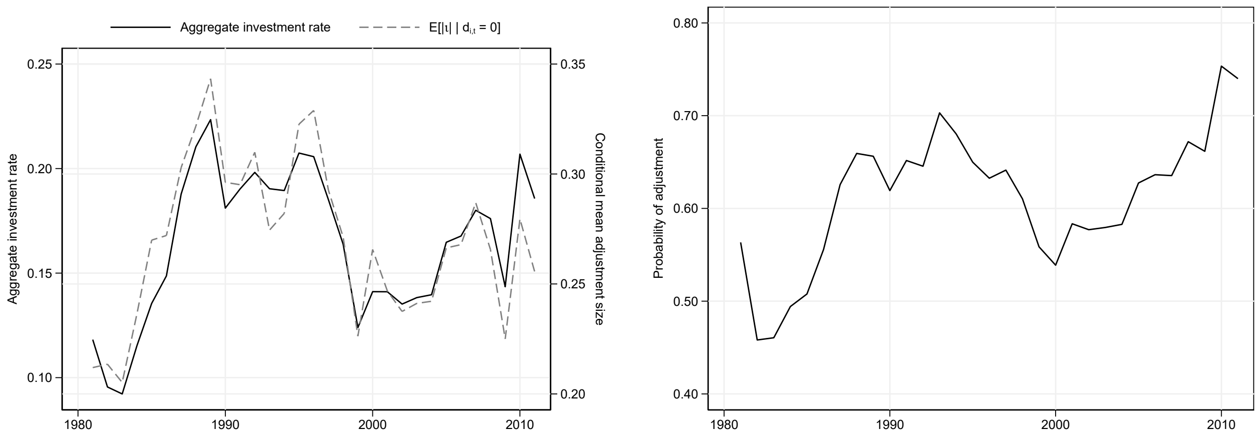

(a) Aggregate investment and conditional mean adjustment size

(b) Extensive margin

Note: Panel (a) plots aggregate investment and the mean adjustment size conditional on adjustment. Panel (b) plots the adjustment probability $\Pr(d_{i,t} = 0)$ (one minus the inaction rate).

plants. Below, we interpret these patterns through the lens of the model and exploit the kinked measurement structure implied by $d_{i,t}$ to recover latent reset capital.

4.2 Time-Series Statistics

To complement the cross-sectional evidence, we summarize the time-series behavior of the two margins of lumpy investment. The extensive margin is the fraction of plants that adjust in a given year, $\Pr(d_{i,t} = 0)$ (one minus the inaction rate). The intensive margin is the average adjustment size conditional on adjustment, $\mathbb{E}[|\iota_{i,t}| \mid d_{i,t} = 0]$. Figure III reports these objects for the Chilean manufacturing sector as a whole. Panel (a) plots aggregate investment together with the conditional mean adjustment size, while Panel (b) plots the extensive margin $\Pr(d_{i,t} = 0)$.

4.3 Sectoral Statistics

We partition manufacturing into eight major sectors: (1) Food and beverages; (2) Textiles, clothing and leather; (3) Wood products and furniture; (4) Paper and printing; (5) Chemistry, petroleum, rubber and plastic; (6) Non-metallic mineral products; (7) Basic metals; and (8) Metal products, machinery and equipment. Throughout the empirical analysis, we estimate the model separately by sector and assume that the structural parameters are common to all plants within a sector. This assumption is both practical and economically motivated: it allows the inaction thresholds, adjustment-opportunity rate, and productivity dynamics to reflect sector-specific technologies, resale markets, and adjustment frictions, while still exploiting rich within-sector plant-level variation to identify latent reset capital.

Table I – Sectoral Contributions within Manufacturing

Sector	Value added share (%)	Labor share (%)	Investment share (%)
1 Food	26.1	27.1	28.3
2 Textiles	9.7	19.9	4.2
3 Wood	3.5	7.6	4.9
4 Paper	12.5	9.0	23.3
5 Chemistry	25.1	13.9	10.7
6 Mineral	6.9	4.5	13.2
7 Metal	6.5	4.4	8.6
8 Machinery	9.6	13.5	6.8
All	100.0	100.0	100.0

Notes: Value added, labor, and investment shares are computed as each sector's average share of the manufacturing total over 1979–2011, based on sector-year aggregates of value added, sales, employment, and gross investment.

Two sets of sectoral statistics organize the empirical environment. First, Table I reports each sector's average contribution to manufacturing value added, employment, and investment. These shares summarize the economic importance of each sector and provide natural weights for constructing sectoral and aggregate frictionless investment indexes. They also guide the interpretation of the aggregate results: sectors with large contributions to investment or value added mechanically exert a larger influence on economy-wide movements in desired capital and investment slack.

Second, Table II summarizes investment-rate moments by sector for the total capital stock. The patterns confirm substantial lumpiness across all sectors: high inaction frequencies coexist with occasional large adjustments (spikes), while serial correlation in investment rates is close to zero. At the same time, there is meaningful heterogeneity across sectors in average investment, inaction, and spike rates, which motivates estimating parameters separately by sector rather than imposing a single set of frictions for all manufacturing.

Table II – Investment Rate Statistics by Sector

Sector	Avg. inv. (%)	Pos. freq. (%)	Neg. freq. (%)	Inaction (%)	Spike (%)	Serial corr.
1 Food	15.9	56.6	2.5	41.0	22.9	0.0
2 Textiles	13.7	50.8	3.7	45.5	18.9	0.0
3 Wood	18.1	57.1	4.5	38.4	24.5	0.1
4 Paper	17.6	57.7	3.9	38.5	24.5	0.0
5 Chemistry	19.6	68.1	2.9	29.0	28.2	0.0
6 Mineral	17.6	58.2	2.7	39.1	23.0	0.0
7 Metal	16.4	63.0	3.2	33.9	24.0	0.0
8 Machinery	16.3	58.2	3.3	38.5	23.4	0.0
All	16.9	58.7	3.3	38.0	23.7	0.0

Notes: Investment rates are computed as investment divided by lagged capital. Pos./Neg. freq. refer to $\iota > 1\%$ and $\iota < -1\%$, respectively; inaction refers to $|\iota| \leq 1\%$; spikes refer to $|\iota| > 20\%$. Serial correlation is the first-order autocorrelation of the investment rate.

Figure IV – Parameter Estimation for Sector 1 (Food and beverages)

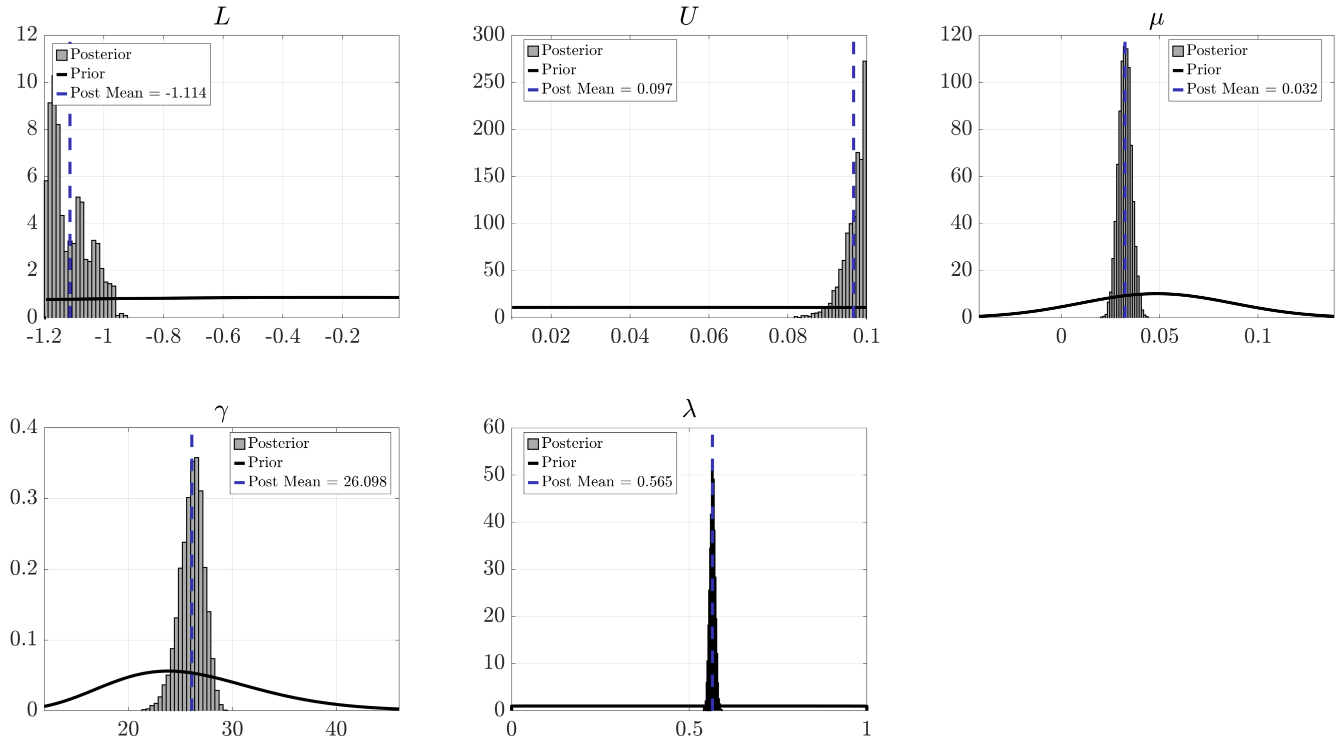

Note: Estimated parameters.

4.4 Parameter Estimates

We estimate the model parameters and states separately by manufacturing sector. Estimation targets five key parameters: the inaction region thresholds (L, U), the arrival rate of free adjustment opportunities λ, and the drift and volatility (μ, σ) governing the latent target process. Our Bayesian procedure combines these restrictions with the state-space transition of the latent target. For each sector, we specify priors over $(L, U, \lambda, \mu, \sigma)$ and compute the posterior using the likelihood implied by the kinked measurement equation and the latent-state transition.

We illustrate the Bayesian updating for one sector. Figure IV reports, for Sector 1 (Food and beverages), the prior (solid line) and posterior (histogram) distributions for the five key parameters, together with the posterior mean (vertical dashed line). The data are informative about the shape of the Ss policy: the posterior for the inaction thresholds (L, U) and the free-adjustment opportunity rate λ tightens markedly relative to the prior, reflecting that the prevalence and persistence of inaction pin down the width of the inaction band and how often firms adjust outside the band. The posterior for the target process parameters (μ, σ) is disciplined by the size and time-series dispersion of adjustment episodes, since observed capital jumps reveal information about the reset level and hence about the innovation variance and drift of $\log k^{\star}_{i,t}$. For this sector, the posterior means imply an estimated inaction region $[L, U] = [-1.114, 0.097]$, a free adjustment opportunity rate $\lambda = 0.565$, and productivityprocess parameters $(\mu, \sigma) = (0.032, 0.195)$.

Table III reports posterior means of the structural parameters by sector, with 95% credible

Table III – Parameter Estimates by Sector

Sector	L	U	μ	σ	λ
1 Food	-1.0301 [-1.1584,-0.9295]	0.1004 [0.0924,0.1032]	0.0364 [0.0302,0.0423]	0.1949 [0.1888,0.2032]	0.6424 [0.6288,0.6557]
2 Textiles	-1.0642 [-1.2237,-0.8863]	0.1766 [0.1583,0.1826]	-0.0014 [-0.0087,0.0059]	0.1850 [0.1774,0.1955]	0.5681 [0.5505,0.5856]
3 Wood	-1.1071 [-1.3187,-0.9216]	0.1772 [0.1510,0.1867]	0.0324 [0.0200,0.0442]	0.2259 [0.2156,0.2382]	0.6145 [0.5910,0.6373]
4 Paper	-1.1421 [-1.2977,-0.9438]	0.1185 [0.0950,0.1271]	0.0441 [0.0302,0.0577]	0.2295 [0.2177,0.2431]	0.6297 [0.6022,0.6564]
5 Chemistry	-1.1320 [-1.2434,-0.9081]	0.1272 [0.1070,0.1344]	0.0534 [0.0443,0.0624]	0.1997 [0.1923,0.2073]	0.6826 [0.6629,0.7016]
6 Mineral	-1.2523 [-1.4558,-0.9948]	0.1125 [0.0828,0.1239]	0.0284 [0.0140,0.0420]	0.2023 [0.1910,0.2141]	0.6259 [0.5945,0.6560]
7 Metal	-0.9387 [-1.0324,-0.8042]	0.1247 [0.1033,0.1322]	0.0284 [0.0203,0.0365]	0.1782 [0.1715,0.1858]	0.6474 [0.6278,0.6664]
8 Machinery	-1.0864 [-1.1899,-0.9268]	0.1124 [0.1023,0.1159]	0.0297 [0.0219,0.0374]	0.2000 [0.1926,0.2092]	0.6096 [0.5927,0.6259]

Notes: Posterior means are reported. Brackets report 95% credible intervals.

intervals in brackets. The estimated inaction bands in gap space are markedly asymmetric across sectors: the lower threshold L is far below zero (roughly between -0.94 and -1.25), while the upper threshold U is comparatively tight (roughly between 0.10 and 0.20). Recall that the lower side ($L < 0$) corresponds to tolerating large negative deviations of capital from its reset level (under-capitalization), whereas the upper side ($U > 0$) governs how quickly plants correct positive deviations (over-capitalization). The estimated asymmetry implies that plants allow sizeable under-capitalization before undertaking an upward adjustment, but correct over-capitalization more quickly.

The opportunity rate λ is precisely estimated and varies from about 0.60 (Sector 8) to about 0.68 (Sector 5), indicating substantial cross-sector differences in the frequency of "free" adjustment episodes. Finally, the target-process parameters (μ, σ) are also tightly pinned down. Estimated drifts are modest (from slightly negative in Sector 2 to about 0.05 in Sector 5), while volatilities cluster around 0.18–0.23 with the highest dispersion in sectors 3–4 and the lowest in Sector 7.

Overall, Table III shows that both the Ss policy parameters and the stochastic evolution of reset capital differ systematically across sectors, motivating sector-specific estimation and aggregation in the subsequent empirical analysis.

5 New Capital Slack Measures

This section presents the empirical implications of the recovered reset-capital series and the frictionless investment index. We proceed in four steps. First, we document basic properties of firm-level reset capital and the wedge between observed and reset capital, which we interpret as "capital slack" created by inaction. We show how this capital slack evolves over the cycle, and how it comoves with adjustment activity. Second, we aggregate the recovered reset-capital series using capital weights to construct sectoral and manufacturing-wide capital slack indexes, and we provide descriptive evidence that these indexes lead standard measures of real activity. Third, we formally evaluate the leading-indicator content of capital slack using forecasting regressions and local-projection exercises, and compare its predictive performance with that of commonly used macroeconomic and financial indicators. Finally, we assess robustness to alternative aggregation schemes, inaction definitions, and parameterizations of the state-space system, and we discuss how estimation uncertainty in the recovered latent states maps into uncertainty bands for the aggregate index.

5.1 Firm-level Capital Slack

Reset capital delivers a direct measure of capital slack,

$$s_{i,t} \equiv \log k_{i,t} - \log k^{\star}_{i,t},$$

which summarizes how far installed capital lies from the reset level at each date.

Cross-sectional Capital Slack Distribution Figure V shows substantial time variation in all moments of the cross-sectional distribution of capital slack. Both measures of central tendency and dispersion fluctuate markedly over the sample, indicating that the shape of the distribution changes meaningfully over time. Across panels, the mean appears to comove positively with kurtosis, while variance tends to move together with skewness. Although we do not impose a structural interpretation on these patterns, they suggest that shifts in the average level of capital slack are often accompanied by changes in tail thickness, whereas movements in dispersion are associated with changes in asymmetry.

5.2 Aggregate Capital Slack

We now move from plant-level latent objects to their sectoral and manufacturing counterparts. To summarize these objects at the sector level, we aggregate using predetermined capital weights. Let $\omega_{i,t}$ denote plant i's share of sector-g capital in year $t-1$ (so that weights are fixed before outcomes in year t are realized). We define sectoral reset capital and capital slack as

$$\log K^{\star}_{g,t} \equiv \sum_{i \in g} \omega_{i,t} \log k^{\star}_{i,t}, \qquad S_{g,t} \equiv \sum_{i \in g} \omega_{i,t} (\log k_{i,t} - \log k^{\star}_{i,t}), \tag{21}$$

and we report analogous aggregates for manufacturing as a whole. Using predetermined weights ensures that movements in $K^{\star}_{g,t}$ and $S_{g,t}$ reflect changes in plants' targets and gaps, rather than

Figure V – Capital Slack Moments over Time

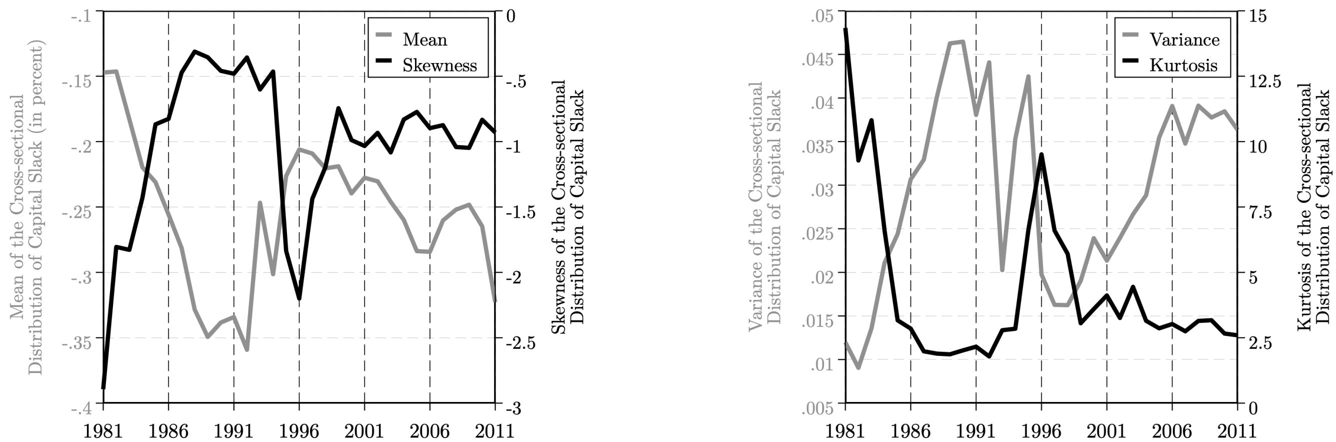

(a) Mean and Skewness over time

(b) Variance and Kurtosis over time

Note: Panel (a) shows the evolution of the mean (left axis) and skewness (right axis). Panel (b) shows the evolution of variance (left axis) and kurtosis (right axis) of the Capital Slack distribution.

mechanical reweighting induced by contemporaneous investment.

5.3 Leading-indicator for Business Cycles

We evaluate whether capital slack contains information about aggregate fluctuations beyond what is mechanically implied by current and lagged investment. The guiding idea is simple: when adjustment is lumpy, reset capital can move smoothly with fundamentals, while observed capital responds only when firms exit their inaction regions. The resulting wedge—capital slack—therefore aggregates "pent-up" adjustment pressure that may be informative about near-term investment dynamics and, more broadly, business-cycle conditions. We begin with a visual inspection, then proceed to reduced-form predictability diagnostics using Granger tests, and finally turn to out-of-sample forecasting and prediction exercises in the next subsection.

Capital Slack Distribution and Investment Rate. Figure VI plots the dynamics of the aggregate investment rate (black line, right axis) and the dynamics of the cross-sectional capital slack distribution (excluding zeros) for Chilean manufacturing. The shaded region summarizes the dispersion of plant-level capital slack in each year (interquartile range and the 5–95 percentile band), while the thick gray line reports the median of the capital slack distribution.

Three patterns stand out. First, aggregate investment comoves with the location of the capital slack distribution: years with higher investment tend to coincide with a median capital slack closer to zero (i.e., smaller average wedges between installed and reset capital), consistent with adjustment episodes closing accumulated gaps. Second, dispersion in capital slack is time-varying

Figure VI – Capital Slack Distribution and Investment Rate

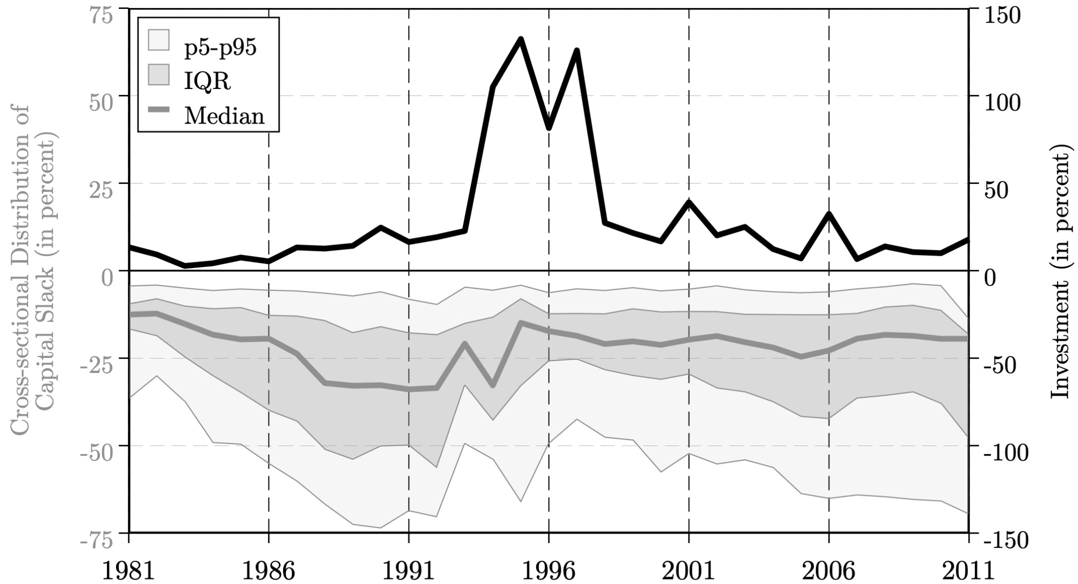

Note: Left axis: cross-sectional distribution of plant-level capital slack $s_{i,t}$ in each year (median, interquartile range, and 5–95 band). Right axis: aggregate investment rate (percent). Slack is computed from smoothed estimates of $k^{\star}_{i,t}$ and aggregated using predetermined $(t-1)$ capital weights.

and rises during periods when investment is high, indicating that aggregate investment episodes are associated not only with shifts in the distribution's center but also with changes in cross-sectional heterogeneity. Third, the percentile bands remain substantially asymmetric around the median throughout the sample, reflecting that large capital gaps are concentrated in one tail of the distribution and motivating our focus on distributional moments (not only means) as inputs for aggregate monitoring.

Event Study. Figure VII provides a simple event-style illustration of the mechanism linking capital slack to subsequent investment. We pick two years that are visually salient in the aggregate time-series evidence: one preceding a period of muted investment activity (1982), and one preceding the large investment episode highlighted in the aggregate series (1992, followed by the spike in the mid-1990s). Panel (a) compares the cross-sectional distribution of plant-level capital slack. In 1982, capital slack is close to zero and tightly concentrated. In 1992, the distribution shifted left and became more dispersed, indicating that a larger mass of plants is operating with installed capital substantially below its reset level. Panel (b) then summarizes what happens next in the aggregate: cumulative investment in 1993–94 is an order of magnitude larger than in 1983–84. The two panels together convey the basic interpretation of capital slack as "pent-up" reset capital: when capital slack is negative and large, subsequent investment is elevated as firms adjust and close the gap between installed and reset capital.

Figure VII – Capital Slack Distribution (Before and After an Aggregate Spike)

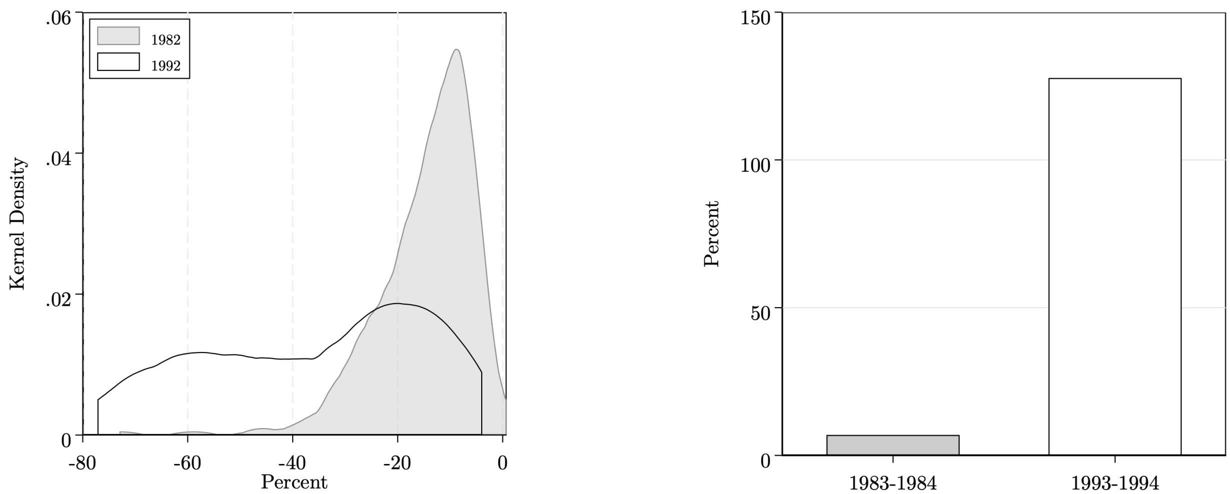

(a) Cross-sectional distribution of capital slack **(b)** Subsequent cumulative investment

Note: Panel (a) plots kernel densities of plant-level capital slack $s_{i,t}$ in 1982 and 1992 (in percent). Panel (b) reports cumulative aggregate investment over the subsequent two years (1983–84 and 1993–94).

5.3.1 Granger Causality

To assess whether capital slack systematically precedes investment, we exploit cross-sector and time variation to conduct panel Granger causality tests. For each sector g and year t, we summarize the cross-sectional distribution of plant-level capital slack $s_{i,t}$ using a vector of moments. Let $\mathcal{M}(s_{i,t})$ denote this collection of summary statistics,

$$\mathcal{M}_t^g \equiv \Big(\text{ZeroShare}_t^g, \text{ Mean}_t^g, \text{ Var}_t^g, \text{ Skew}_t^g, \text{ Kurt}_t^g\Big)',$$

where ZeroShare_t^g is the fraction of plants in sector g and year t with $s_{i,t} = 0$, and the remaining entries correspond to the cross-sectional mean, variance, skewness, and kurtosis of capital slack within sector g, computed excluding zero observations.

First, we test whether moments of the capital slack distribution Granger-cause investment by estimating the following panel regression,

$$I_t^g = \alpha_g + \delta_t + \sum_{l=1}^{L} \gamma_l I_{t-l}^g + \sum_{l=1}^{L} \phi_l' \mathcal{M}_{t-l}^g + e_t, \tag{22}$$

where α_g denote sector fixed effects and δ_t denote time fixed effects. The lag length is set to $L = 2$. Panel (a) of Figure VIII reports the corresponding tests for Granger causality. The null that lagged capital slack moments have no predictive content for investment is rejected for the zero share, skewness, and kurtosis individually, as well as jointly across all moments. These results are consistent with the intuition from the case study in Figure VII, which shows that movements

Figure VIII – Granger Causality between Investment and Capital Slack Moments

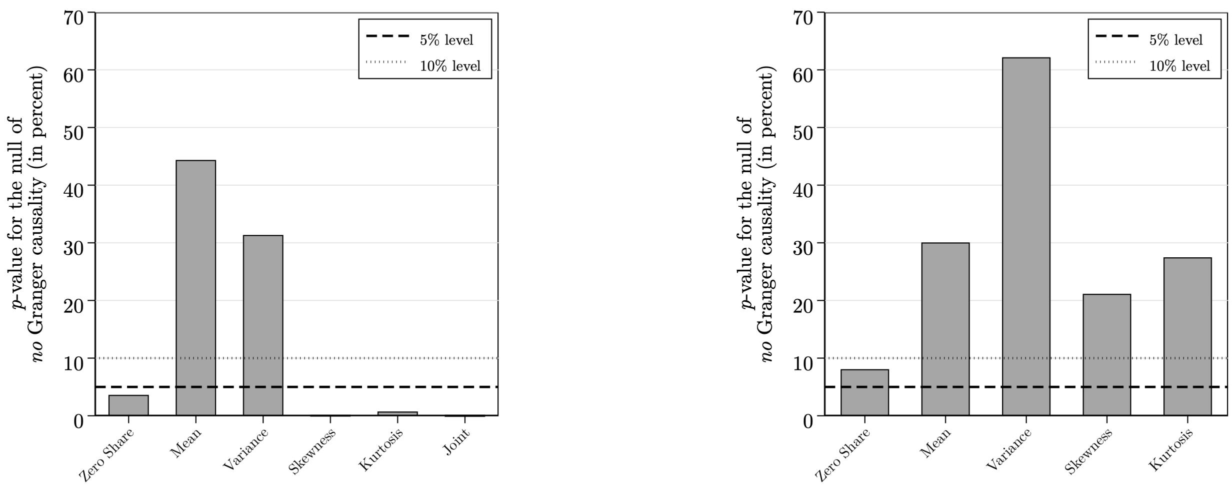

(a) Does capital slack predict investment? **(b)** Does investment predict capital slack?

Note: Bars report p-values (in percent) for the null of *no* Granger causality. Panel (a) tests whether capital slack moments predict investment; Panel (b) tests whether investment predicts capital slack moments. Dashed and dotted lines mark the 5% and 10% levels.

in the cross-sectional distribution of capital slack anticipate investment.

Second, to test whether Granger causality runs in the opposite direction, we estimate, for each moment in $\mathcal{M}(s_{i,t})$, the following panel regression,

$$\mathcal{M}^g_{j,t} = \alpha_g + \delta_t + \sum_{l=1}^{L} \Gamma_l I^g_{t-l} + \sum_{l=1}^{L} \mathbf{\Phi}'_l \mathcal{M}^g_{t-l} + \upsilon_t, \tag{23}$$

where $\mathcal{M}^g_{j,t}$ denotes the j-th element of $\mathcal{M}^g_t$. Panel (b) of Figure VIII reports that none of the capital slack moments is Granger-caused by investment at the 5% level. Taken together, the two panels rule out a pure "reaction" view—where capital slack adjusts only after realized investment—and instead support a "state" view in which capital slack reflects latent reset capital adjustments that precede observed investment.

5.3.2 Forecasting

While Granger tests provide an in-sample exclusion diagnostic, the following forecasting regressions quantify the magnitude and robustness of predictability, including out-of-sample performance.

We next quantify the predictive content of capital slack for aggregate investment using annual forecasting regressions. Let $\mathcal{I}^g_{t+1,t+h}$ denote cumulative investment in sector g between years $t+1$

Table IV – Forecasting Aggregate Investment with Capital Slack Moments

	$h=1$	$h=2$	$h=3$	$h=4$
Investment	0.01	0.05	0.01	0.05
	(0.02)	(0.02)	(0.04)	(0.07)
Zero Share	-0.18	-0.37	-0.75	-1.23
	(0.14)	(0.25)	(0.36)	(0.47)
Mean	-36.82	-71.39	-89.81	-134.38
	(28.07)	(38.58)	(25.88)	(43.43)
Variance	71.48	207.33	463.55	497.93
	(72.82)	(206.06)	(407.90)	(601.34)
Skewness	11.61	18.72	17.96	14.26
	(7.19)	(13.30)	(17.34)	(20.59)
Kurtosis	2.59	4.20	4.09	3.61
	(1.66)	(2.97)	(3.60)	(4.16)
Sector FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
p Inv	0.00	0.02	0.15	0.22
p Slack	0.00	0.20	0.00	0.01
R^2	0.53	0.55	0.60	0.64
Adj. R^2	0.43	0.45	0.52	0.55

Notes: Each column reports panel fixed-effects estimates where the dependent variable is cumulative investment $\sum_{k=1}^{h} I_{t+k}$ for horizon $h \in \{1,2,3,4\}$. The top panel reports the average coefficient for each variable from a regression that includes both current and one lagged value. Standard errors clustered by sector in parentheses. Zero Share is the fraction of plants with zero capital gap; the remaining slack moments (Mean, Variance, Skewness, Kurtosis) are computed excluding zero-gap plants. The rows p Inv and p Slack report p-values for the joint significance of all current and lagged investment coefficients and all current and lagged slack moment coefficients, respectively.

and $t+h$. For horizons $h \in \{1,2,3,4\}$ (one to four years), we estimate

$$\mathcal{I}^g_{t+1,t+h} = \alpha_g + \delta_t + \sum_{l=0}^{L} a_l I^g_{t-l} + \sum_{l=0}^{L} \mathbf{b}_l' \mathcal{M}^g_{t-l} + \epsilon_t, \tag{24}$$

where we include the contemporaneous value and one lag of each regressor. The cumulative left-hand side aligns the dependent variable with the interpretation of capital slack as pent-up adjustment pressure that may unwind gradually over several years.

Estimation results are reported in Table IV. The coefficients display economically sensible signs. Lagged investment is positive at all horizons, consistent with persistence. The zero share enters negatively, indicating that when more plants have no capital slack, subsequent investment is lower. The higher-order moments align with the heterogeneity mechanism: the mean of capital slack is negatively associated with future investment, while variance and kurtosis enter positively, consistent with greater dispersion generating more firms with large adjustment needs. Skewness is the only coefficient with a counterintuitive sign, but it is not statistically significant at the 5% level at any horizon. Joint tests show that lagged investment is significant at short horizons (one

Figure IX – Forecast Performance of Capital Slack Moments for Aggregate Investment

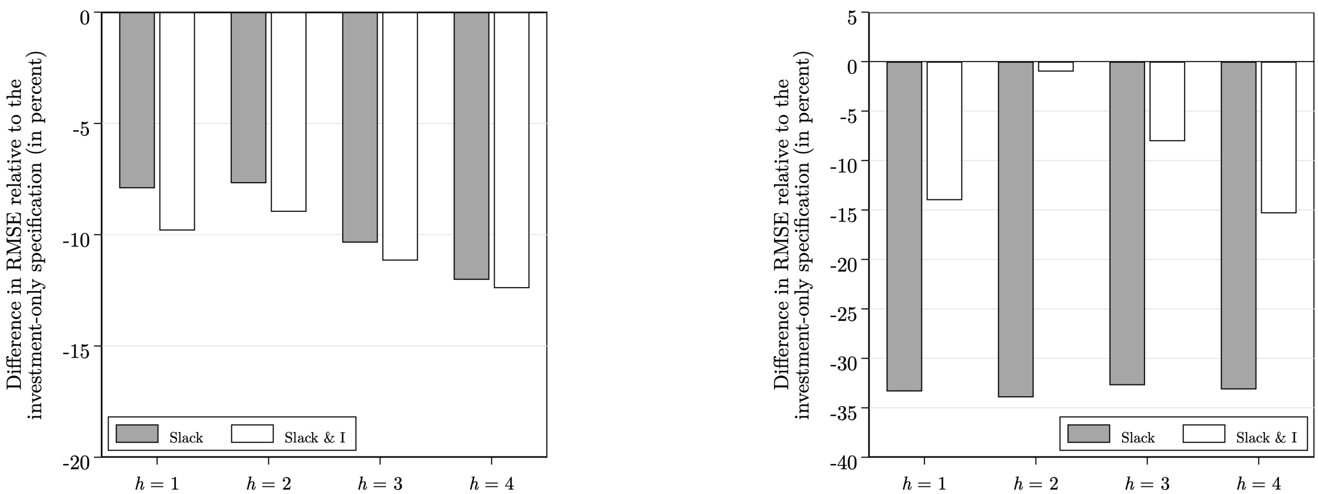

(a) In-sample fit **(b)** Out-of-sample fit

Note: Bars report the percent change in RMSE relative to an investment-only benchmark. "Capital slack" uses moments of the capital slack distribution as predictors; "Slack & I" augments these moments with lags of investment. Negative values indicate lower RMSE (better forecasts).

and two years), while capital slack moments are jointly significant at all horizons except two years.

Taken together, these results indicate that both lagged investment and cross-sectional capital slack moments help explain future investment, but they do not reveal which predictors drive the forecasting gains. To assess their relative importance, we estimate horse-race specifications of equation (24): one including only capital slack moments, one including both capital slack and lagged investment, and a benchmark including only lagged investment, and compare predictive performance across models.

Figure IX reports the forecasting comparison at horizons $h = 1, \ldots, 4$. Panel (a) shows that capital slack moments generate systematic in-sample RMSE reductions that increase with the horizon. Adding lagged investment yields only marginal additional improvements, particularly at short horizons. Panel (b) reports out-of-sample performance, where the model is estimated using the first ten years of data and forecasts are generated for the remaining sample with coefficients held fixed. Capital slack moments alone deliver substantial gains in predictive accuracy, whereas adding lagged investment attenuates these gains. While lagged investment slightly improves in-sample fit, its inclusion can worsen out-of-sample performance, suggesting that the richer specification overfits in small samples.

Overall, the results support a parsimonious forecasting takeaway: moments of the capital slack distribution contain incremental predictive information for aggregate investment, especially over multi-year horizons, but combining them with too many investment lags can hurt out-of-sample accuracy.

Real-time Feasibility: Filtering vs Smoothing. Because our leading-indicator exercises are inherently real-time, we replicate the main forecasting and monitoring results using *one-sided* filtered estimates, constructed from information available up to time t. We compare these to the two-sided smoothed estimates (which use the full sample) to quantify the loss of precision from real-time implementation and to verify that the predictive content does not rely on ex-post information.

Taken together, the evidence supports a simple interpretation: reset capital responds quickly to shocks, while observed capital adjusts intermittently, so the gap between the two contains forward-looking information about future investment dynamics.

5.4 Robustness (In Progress)

This section assesses whether our main empirical conclusions are sensitive to (i) how we aggregate firm-level objects, (ii) how we classify inaction in the micro data, (iii) the parametric structure imposed by the Ss model and state-space representation, and (iv) sample selection and panel structure. Throughout, we keep the baseline estimation and reporting conventions fixed and vary one ingredient at a time.

(i) Alternative aggregation and weighting. Our baseline sectoral and manufacturing indexes weight plants by predetermined capital shares (using $t-1$ capital). We recompute all aggregates under (a) equal weights, (b) value-added weights, (c) fixed weights anchored to an initial year, and (d) contemporaneous weights based on K_t instead of K_{t-1}. These exercises test whether the aggregate capital slack signal is driven mechanically by a small number of large plants or by compositional changes in the capital distribution.

(ii) Alternative inaction definitions and missing observations. The measurement equation uses the inaction indicator $d_{i,t}$, which we construct from a "small investment" threshold. We vary this threshold (e.g., 0.5%, 1%, 2% in absolute investment rates) and consider alternative treatments of missing or implausible investment reports, including (a) coding missing investment as missing $d_{i,t}$ (dropping the plant-year), (b) imputing investment as zero when accounting definitions suggest non-reporting corresponds to no purchases/sales, and (c) tightening/loosening trimming rules for extreme investment rates. These checks verify that the recovered reset capital and capital slack are not an artifact of a particular coding of inaction.

(iii) Alternative parameterizations of the latent process. We probe the role of parametric assumptions on the reset-capital process and adjustment opportunities. First, we allow for richer dynamics in $\log k^{\star}_{i,t}$, including AR(1) persistence and sector-specific drifts/volatilities. Second, we vary the adjustment-opportunity process, replacing the Bernoulli arrival model with (a) time-varying λ_t (common or sector-specific), or (b) a hazard that depends on observables (e.g., plant size bins). Third, we incorporate measurement error in observed capital (or in the investment rate used to form $d_{i,t}$) to test whether the capital slack distribution remains informative once we allow for noisy measurement in the observables.

(iv) Sensitivity to sample composition and panel structure. Finally, we assess whether results depend on the composition of plants and sectors in the sample. We repeat the analysis on (a) balanced panels (plants observed for at least T^* consecutive years) versus the baseline unbalanced panel, (b) alternative minimum-tenure thresholds (e.g., 5, 10, 15 years), and (c) alternative sector groupings and sector coverage. These exercises ensure that the capital slack signal is not driven by entry/exit patterns, short-lived plants, or a particular sectoral partition of manufacturing.

6 Conclusion

This paper introduces a simple, micro-founded way to measure *capital slack*: the wedge between observed capital and the latent reset capital that firms would choose upon adjustment. The key insight is that lumpy investment is not only a friction shaping aggregate dynamics, but also a *measurement device*. Empirically, the recovered distribution of capital slack provides a compact summary of the economy's pent-up investment pressure created by inaction. The results show that capital slack is systematically related to subsequent adjustment behavior at the micro level, and that aggregating the recovered objects yields informative sectoral and manufacturing-wide series.

Beyond this application, the approach is portable. Any environment in which policy rules generate thresholded mappings from latent targets to observables can be disciplined by the same state-space logic. In that sense, frictionless investment is a step toward a broader measurement agenda: using microeconomic kinks to recover latent state variables that are central for macroeconomic monitoring and policy analysis.

References

Aruoba, S. B., Diebold, F. X. and Scotti, C. (2009). *Real-Time Measurement of Business Conditions*. Working Paper 14349, National Bureau of Economic Research.

Bachmann, R., Caballero, R. J. and Engel, E. M. R. A. (2013). Aggregate implications of lumpy investment: New evidence and a DSGE model. *American Economic Journal: Macroeconomics*, **5** (4), 29–67.

Baley, I. and Blanco, A. (2021). Aggregate dynamics in lumpy economies. *Econometrica*, **89** (3), 1235–1264.

Bandeira, M., Castillo-Martinez, L. and Wang, S. (2024). *Frictionless Inflation*. Tech. rep., Tech. rep., Working Paper.

Bertola, G. and Caballero, R. J. (1994). Irreversibility and aggregate investment. *The Review of Economic Studies*, **61** (2), 223–246.

Caballero, R. J. and Engel, E. (1993). Microeconomic adjustment hazards and aggregate dynamics. *The Quarterly Journal of Economics*, **108** (2), 359–383.

— and Engel, E. M. R. A. (1999). Explaining investment dynamics in U.S. manufacturing: A generalized (S,s) approach. *Econometrica*, **67** (4), 783–826.

Chauvet, M. and Piger, J. (2008). A comparison of the real-time performance of business cycle dating methods. *Journal of Business & Economic Statistics*, **26** (1), 42–49.

Cooper, R., Haltiwanger, J. and Power, L. (1999). Machine replacement and the business cycle: Lumps and bumps. *American Economic Review*, **89** (4), 921–946.

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

Doms, M. and Dunne, T. (1998). Capital adjustment patterns in manufacturing plants. *Review of economic dynamics*, **1** (2), 409–429.

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

Khan, A. and Thomas, J. K. (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica*, **76** (2), 395–436.

Pindyck, R. S. (1991). Irreversibility, uncertainty, and investment. *Journal of Economic Literature*, **29** (3), 1110–1148.

Stock, J. H. and Watson, M. W. (1989). New indexes of coincident and leading economic indicators. In *NBER Macroeconomics Annual 1989*, MIT Press, pp. 351–394.

Thomas, J. K. (2002). Is lumpy investment relevant for the business cycle? *Journal of Political Economy*, **110** (3), 508–534.

A Gibbs Sampling Algorithm

Given the observed series of the capital stock $k_{i,t}$ and inaction $d_{i,t}$, together with the model parameters θ, our objective is to recover the distribution of the unobserved states $(k^\star_{i,t},\, l_{i,t})$ at each date t. The problem is thus to estimate the distribution of latent states conditional on the observed data and the parameters. We solve this using two standard steps in Bayesian state-space estimation: *forward filtering* and *backward sampling.*

Forward filtering asks: *What do the hidden states look like at each date, given the data so far?* Backward sampling answers: *Choose one coherent path for the hidden states over the entire sample, consistent with both the model dynamics and all observations.* In our MCMC procedure, each parameter draw θ is paired with latent-state draws from forward filtering and backward sampling, allowing us to compute cross-sectional moments and the cumulative impulse response needed to test the sufficient statistics formula.

A.1 State space in probabilistic form

State densities. Conditional on θ, these two components evolve independently, so that

$$p(X_{i,t} \mid X_{i,t-1}, \theta) \;=\; p(k^\star_{i,t} \mid k^\star_{i,t-1}, \theta) \times p(l_{i,t} \mid \theta). \tag{25}$$

This density captures the model-implied law of motion for the unobserved states, a key building block of the forward filtering recursion. Lemma 1 establishes that the state transition densities are independent across plants.

Lemma 1. *State transition densities are independent across plants:*

$$p(X_{1:n,t} \mid X_{1:n,t-1}, k_{1:n,1:t-1}, \theta) \;=\; \prod_{i=1}^{n} p(X_{i,t} \mid X_{i,t-1}, \theta). \tag{26}$$

Measurement densities. The measurement density $p(k_{i,t} \mid X_{i,t}, k_{i,t-1}, \theta)$ links the observed capital stock $k_{i,t}$ to the latent state $X_{i,t}$ and the previous period's observed capital $k_{i,t-1}$, as given by the measurement equations (10) and (11). When the inaction indicator $d_{i,t} = 1$, capital evolves deterministically via depreciation, so $k_{i,t}$ is a deterministic function of $k_{i,t-1}$. When $d_{i,t} = 0$, the plant adjusts to the reset level $k^\star_{i,t}$. Consequently, the measurement density is degenerate given the state: conditional on $(k^\star_{i,t}, l_{i,t})$ and $k_{i,t-1}$, the model pins down $k_{i,t}$ exactly. Moreover, the lumpy adjustment policies implied by the model bound the unobserved states, allowing for efficient grid-based Bayesian filtering and smoothing despite the kinks in the measurement equations.

This density is used in the update step of the forward filtering recursion to incorporate information about the observed capital stock into the distribution of latent states.

Lemma 2 establishes that the measurement densities are also independent across plants.

Lemma 2. *Measurement densities are independent across plants:*

$$p(k_{1:n,t} \mid X_{1:n,t}, k_{1:n,1:t-1}, \theta) \;=\; \prod_{i=1}^{n} p(k_{i,t} \mid X_{i,t}, k_{i,t-1}, \theta). \tag{27}$$

Initial densities Finally, we establish initial conditions. Conditional on parameter vector θ, the state-space representation is finalized with the initial state distributions:

$$p(X_{1:n,1}|k_{1:n,1}, \theta) \stackrel{\text{i.i.d.}}{=} \prod_{i=1}^{n} p(X_{i,1}|k_{i,1}, \theta). \tag{28}$$

A.2 Forward filtering

The goal of forward filtering is to compute, for each date t, the joint filtered distribution of the latent states $X_{i,t} \equiv (k^{\star}_{i,t}, l_{i,t})$ for all plants:

$$p(X_{1:n,t} \mid k_{1:n,1:t}, \theta).$$

The recursion starts from the initial distribution in (28) and proceeds forward from $t = 1$ to T, incorporating each new data point as it arrives.

We begin with the full joint distribution across all n plants. As shown in Lemma 1 (independence of state transitions) and Lemma 2 (independence of measurement densities), both the state transition densities and the measurement densities factorize across plants. These results imply that the joint filtered density can be written as the product of individual densities:

$$p(X_{1:n,t} \mid k_{1:n,1:t}, \theta) \;=\; \prod_{i=1}^{n} p(X_{i,t} \mid k_{i,1:t}, \theta). \tag{29}$$

Thus, the filtering problem reduces to n independent problems, one for each plant i.

Proposition 1 (**Recursive forward filtering**). *Given an initial density, the filtered density for each plant i can be computed recursively as follows:*

1. ***Prediction step**: Starting from the distribution at $t-1$, use the state transition in (25) to obtain the one-step-ahead predicted density (do not use yet the current observation $k_{i,t}$):*

$$p(X_{i,t} \mid k_{i,1:t-1}, \theta) \;=\; \int \underbrace{p(X_{i,t} \mid X_{i,t-1}, \theta)}_{\text{state transitions}} p(X_{i,t-1} \mid k_{i,1:t-1}, \theta)\, dX_{i,t-1}. \tag{30}$$

2. ***Update step**: Incorporate the new observation $k_{i,t}$ using the measurement equations (10)–*

(11) *and Bayes' theorem:*

$$p(X_{i,t} \mid k_{i,1:t}, \theta) \;=\; \frac{p(k_{i,t} \mid X_{i,t}, k_{i,t-1}, \theta)\, p(X_{i,t} \mid k_{i,1:t-1}, \theta)}{\int p(k_{i,t} \mid X_{i,t}, k_{i,t-1}, \theta)\, p(X_{i,t} \mid k_{i,1:t-1}, \theta)\, dX_{i,t}}. \tag{31}$$

A.3 Backward sampling

While forward filtering yields marginal distributions period-by-period, Bayesian inference requires draws from the *joint* distribution of the states over the whole sample:

$$p(X_{1:n,1:T} \mid k_{1:n,1:T}, \theta). \tag{32}$$

Backward sampling reconstructs coherent state histories that are consistent with both past and future observations.

We begin with the full joint distribution across all n plants. As in forward filtering, Lemma 1 and Lemma 2 imply that this joint distribution factorizes across plants:

$$p(X_{1:n,1:T} \mid k_{1:n,1:T}, \theta) \;=\; \prod_{i=1}^{n} p(X_{i,1:T} \mid k_{i,1:T}, \theta). \tag{33}$$

Thus, backward sampling can be carried out independently for each plant i.

Proposition 2 shows how to recover an entire history of the latent states by working backwards in time. Starting from the final period, we sequentially condition on the future state and the observed data to sample earlier states. This recursion relies on the Markov property of the state process and the fact that the filtered and predicted densities are already available from forward filtering. This recursion generates a full sequence $\{X_{i,t}\}_{t=1}^{T}$ for each plant i that is consistent with the model dynamics and all observed data.

Proposition 2 (Recursive backward sampling). *Given the filtered and predicted distributions from Proposition 1, the backward sampling recursion for each plant i is:*

1. ***Initialization at*** T*: Draw the final state from its filtered distribution:*

$$X_{i,T} \;\sim\; p(X_{i,T} \mid k_{i,1:T}, \theta). \tag{34}$$

2. ***Backward recursion****: For $t = T-1, T-2, \ldots, 1$, draw $X_{i,t}$ from*

$$p(X_{i,t} \mid k_{i,1:T}, X_{i,t+1}, \theta) \;=\; \frac{p(X_{i,t+1} \mid X_{i,t}, \theta)\, p(X_{i,t} \mid k_{i,1:t}, \theta)}{p(X_{i,t+1} \mid k_{i,1:t}, \theta)}, \tag{35}$$

   where:

   - $p(X_{i,t+1} \mid X_{i,t}, \theta)$ *is the state transition equation;*

- $p(X_{i,t} \mid k_{i,1:t}, \theta)$ *is the filtered distribution;*
- $p(X_{i,t+1} \mid k_{i,1:t}, \theta)$ *is the one-step-ahead predicted distribution.*

A.4 Parameter estimation

Conditional on the data and states panel, our parameter estimates are drawn from their posterior distributions. This step is also done in blocks, i.e. each parameter is sampled from its full conditional distribution (conditional on all states and other parameters). This section describes the assumed prior distributions and posterior draws.

Inaction bands L and U. The inaction bands are updated in a Metropolis-Hastings step. For L, we consider a truncated normal prior defined by the parameters μ_L, σ_L, $\underline{L}$, and $\bar{L}$. Similarly, for U we consider a truncated normal prior defined by the parameters μ_U, σ_U, $\underline{U}$, and $\bar{U}$.

We draw a new parameter proposals $L^\star$ and $U^\star$ from a proposal distribution with density $q(\cdot)$ and compute the acceptance ratio α

$$\alpha = \frac{Z(L^\star, U^\star)}{Z(L,U)} \frac{q(L,U|L^\star,U^\star)}{q(L^\star,U^\star|L,U)} \frac{p(L^\star,U^\star)}{p(L,U)} \tag{36}$$

where $Z(\cdot)$ is the likelihood of L and U given parameters, states and data, and $p(\cdot)$ is the prior density.

Then, we draw a threshold r from a uniform $[0,1]$ distribution and use the following rule:

- If $r < \alpha$, accept the proposed parameters $(L^\star, U^\star)$.
- If $r \geq \alpha$, reject the proposed parameter and maintain (L, U).

Productivity drift μ. We consider a normal prior for the drift parameter μ with mean β_0 and variance Σ_0. Given the data generating process for the reset point (equation (15)) and the parameter σ, the posterior distribution of μ follows a normal distribution with the following parameters:

$$\beta^{\text{post}} = \left(\frac{nT-n}{\sigma^2} + \Sigma_0^{-1}\right)^{-1} \left(\sum_{i=1}^{n}\sum_{t=2}^{T} \frac{(\log k_{i,t}^\star - \log k_{i,t-1}^\star)}{\sigma^2} + \Sigma_0^{-1}\beta_0\right) \tag{37}$$

$$\Sigma^{\text{post}} = \left(\frac{nT-n}{\sigma^2} + \Sigma_0^{-1}\right)^{-1} \tag{38}$$

We draw a new guess for μ from this posterior distribution.

Variance σ^2. We consider a gamma prior for the inverse variance $\gamma \equiv 1/\sigma^2$ with shape parameter $\nu_0/2$ and rate parameter $\sigma_0^2(\nu_0/2)$. Given the data generating process for the reset point (equation (15)) and the parameter μ, the posterior distribution of γ is a gamma distribution with the following shape and rate parameters:

$$\text{shape}^{\text{post}} = \frac{\nu_0 + n \cdot T - n}{2} \tag{39}$$

$$\text{rate}^{\text{post}} = \frac{\nu_0 \sigma_0^2 + SSR}{2} \tag{40}$$

where $SRR = \sum_{i=1}^{n} \sum_{t=2}^{T} (\log k_{i,t}^{\star} - \log k_{i,t-1}^{\star} - \mu)^2$. We draw a new guess for γ from this posterior distribution and invert it to obtain the new guess for σ^2.

Arrival rate of free adjustment opportunities λ. We consider a beta prior for the parameter λ with parameters β_1^{prior} and β_2^{prior}. Given this prior and the data generating process of $l_{i,t}$ (equation (14)), the posterior distribution of λ is also a beta distribution with the following parameters:

$$\beta_1^{\text{post}} = \beta_1^{\text{prior}} + \sum_{i=1}^{n} \sum_{t=1}^{T} l_{i,t} \tag{41}$$

$$\beta_2^{\text{post}} = \beta_2^{\text{prior}} + N \cdot T - \sum_{i=1}^{n} \sum_{t=1}^{T} l_{i,t} \tag{42}$$

We draw a new guess for λ from this posterior distribution.

B Data description

This appendix describes the Chilean plant-level data, sample construction, and the capital series used in the empirical analysis.

B.1 Data source and sample

We use annual manufacturing-plant data from the *Encuesta Nacional Industrial Anual* (ENIA). We drop (i) plants permanently below 10 workers, (ii) observations with non-positive book capital, wage bill, or sales, (iii) plants with non-zero investment in fewer than 10% of observed years, and (iv) plants with fewer than 3 years of coverage. Plants that exit for more than three years and later re-enter are treated as new plants. In baseline results, we further restrict to plants observed for at least 10 years. Table V summarizes removals.

Table V – Data cleaning

Description	
Start year	1979
End year	2011
Avg. number of plants per year	543
Plant-year observations	154,591
Cleaning	**Number of removed obs.**
Less than 10 employees	3,984
Non-positive wage bill, capital, or sales	5,604
Frequency of non-zero investment less than 10%	12,819
Less than 3 years of coverage	6,221
Remaining observations	125,963
% of total	81
With more than 10 years of data	98,360
% of remaining observations	78

Notes: Authors' calculations using establishment-level survey data from Chile (ENIA). "Less than 10 employees" refers to plants with fewer than 10 employees in all years they appear in the sample.

B.2 Capital stocks and investment rates

We construct real capital stocks using the perpetual inventory method (PIM). For plant i, capital type j, and year t,

$$K_{i,j,t} = (1-\zeta_j)K_{i,j,t-1} + \frac{I_{i,j,t}}{D_{j,t}}. \tag{43}$$

We consider $j \in \{\text{structures}, \text{ machinery \& equipment}, \text{ vehicles}\}$ and also total capital.

Investment. Gross nominal investment is purchases plus upgrades/reforms/improvements minus sales:

$$(44) \qquad I_{i,j,t} = \text{purchases}_{i,j,t} + \text{reforms}_{i,j,t} + \text{improvements}_{i,j,t} - \text{sales}_{i,j,t}.$$

Initial capital is the plant's first non-negative reported book value for that asset type, deflated by D_{j,t_0}.

Depreciation and deflators. We use constant, type-specific geometric depreciation rates (Table VI) and type-specific investment deflators from the Penn World Tables (PWT). We trim investment rates to exclude those below the 2nd percentile and those above the 98th percentile.

Table VI – Depreciation rates

Structures	Machinery & Equipment	Vehicles
3%	11%	15%

B.3 Validation and auxiliary checks

Reported vs. PIM capital. Figure X compares aggregate reported book capital to aggregate PIM capital (by type and total); the series track each other closely.

Figure X – Chile: Reported Book Value vs Perpetual Inventory

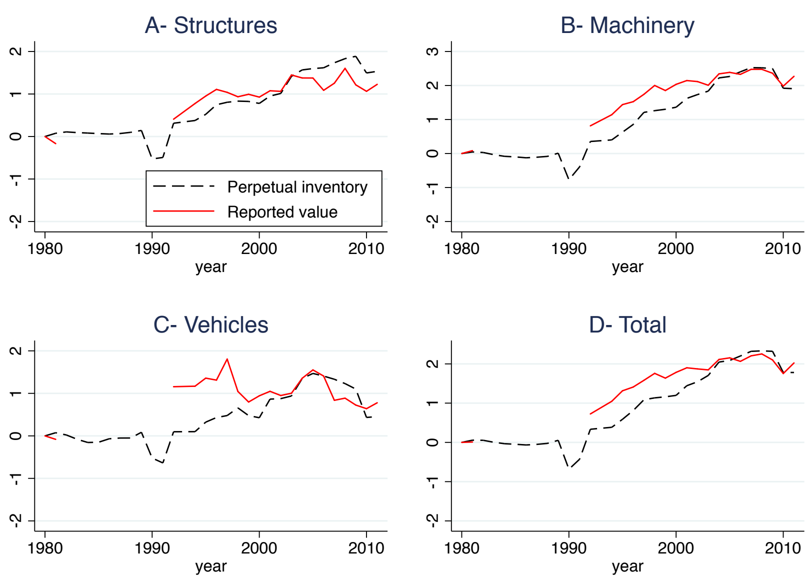

Notes: Aggregate manufacturing capital reported by plants and computed via PIM. Variables are in logs and real terms, normalized to zero in 1980.

Micro vs. National Accounts. Figure XI compares manufacturing investment aggregated from ENIA to the National Accounts series; Table VII reports investment composition by capital type when available.

Figure XI – **Chile**: Micro-data vs. National Accounts

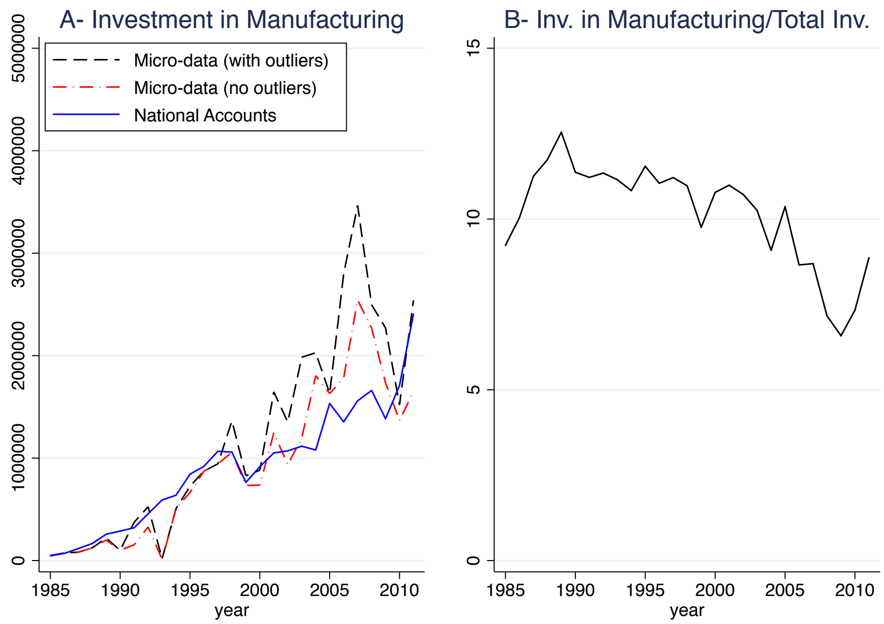

Notes: Panel A compares nominal manufacturing investment from ENIA (black dashed) and National Accounts (blue solid). The red dashed-dotted line drops extreme micro outliers (investments larger than 5% of aggregate investment). Panel B shows manufacturing investment as a share of total investment.

Table VII – Chile: Distribution of Investment Across types of Capital

	Structures	Machinery	Vehicle
National Accounts	35.4	51.4	13.1
ENIA	29.1	68.6	2.1

Notes: Describes average percentage of investments across different types of capital in the ENIA and national accounts.

B.4 Investment Rates Statistics

Investment. Figure XII plots aggregate investment by type (real, logs, normalized in 1995).

Figure XII – Chile: Total Investment from ENIA

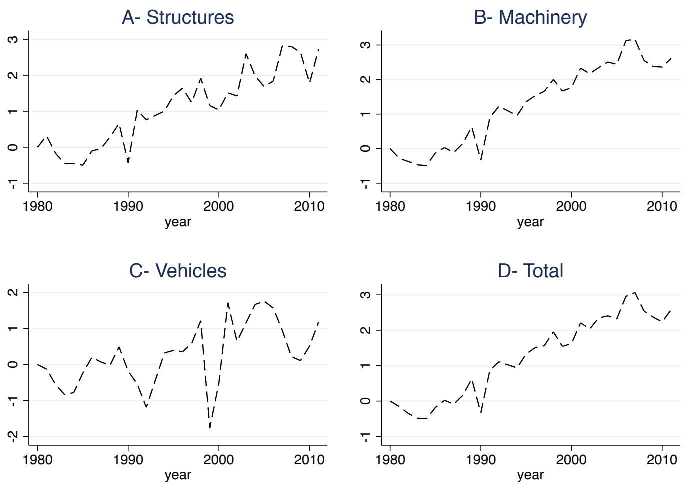

Notes: Aggregate manufacturing investment in structures, machinery, vehicles, and total. Variables are in logs and real terms.

Investment/Capital ratios. Figure XIII plots aggregate investment-to-capital ratios by type (capital via PIM).

Figure XIII – Chile: Investment to Capital Ratio

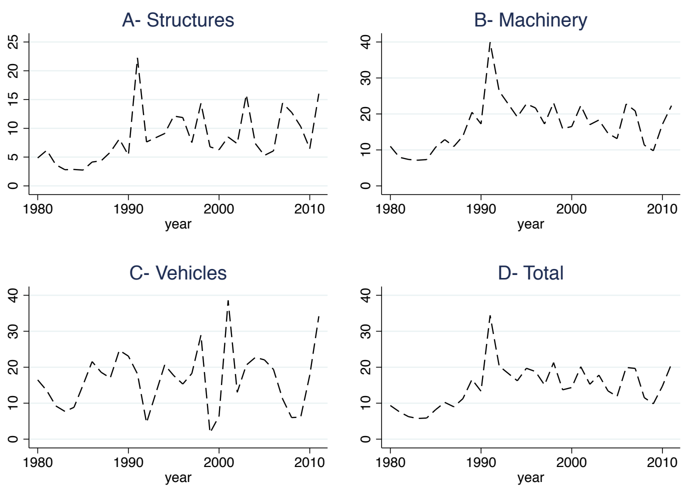

Notes: Aggregate investment-to-capital ratios for structures, machinery, vehicles, and total. Capital is constructed using the PIM.