Aggregate Dynamics in Lumpy Economies*

Isaac Baley†
UPF, CREi, and BGSE

Andrés Blanco ‡
University of Michigan

April 29, 2020

Abstract

How does an economy's capital respond to aggregate productivity shocks when firms make lumpy investments? We show that capital's transitional dynamics are structurally linked to two steady-state moments: the dispersion of capital-to-productivity ratios—an indicator of capital misallocation—and the covariance of capital-to-productivity ratios with the time elapsed since their last adjustment—an indicator of capital irreversibility. We compute these two sufficient statistics using data on the size and frequency of investment of Chilean plants. Their empirical values indicate significant persistent effects of aggregate productivity shocks, and favor investment models with asymmetric fixed costs and random opportunities for free adjustments.

JEL: D30, D80, E20, E30

Keywords: inaction, lumpiness, transitional dynamics, sufficient statistics, non-convex adjustment costs, firm investment, state-dependent models, time-dependent models.

*We thank our discussants, Fernando Álvarez, Andrés Drenik, and Francesco Lippi, as well as four anonymous referees for valuable feedback. We also benefited from exchanges with Adrien Auclert, Rudi Bachmann, Ricardo Caballero, Andrea Caggese, Jeff Campbell, Javier Cravino, Michael Elsby, Eduardo Engel, Jordi Galí, Manuel García-Santana, Matthias Kehrig, John Leahy, Andrei Levchenko, Matthias Meier, Claudio Michelacci, Kurt Mitman, Pablo Ottonello, Luigi Paciello, Jaume Ventura, Edouard Schaal, Laura Veldkamp, Venky Venkateswaran, and seminar participants at CREi, UPF, Michigan, Banque de France, Bocconi, Cleveland FED, PSE, EIEF, Chicago FED, Duke, Federal Reserve Board, Universitat Autònoma de Barcelona, Banco do Portugal, Nova, EM^{3}C Conference, Transpyrenean Macro Workshop 2019, XXII Workshop in International Economics and Finance, CREI-Bank of Canada Workshop, 2nd Meeting of the Catalan Economic Association, Santiago Macro Workshop, and NBER Monetary Economics. Lauri Esala and Erin Markiewitz provided outstanding research assistance. Isaac Baley gratefully acknowledges the support of the Spanish Ministry of the Economy and Competitiveness, through a Seed Grant (Aggregate Dynamics in Lumpy Economies) from the Barcelona Graduate School of Economics Severo Ochoa Programme for Centres of Excellence in R&D (SEV-2015-0563).

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

‡University of Michigan, jablanco@umich.edu.

1 Introduction

Economies are exposed to productivity, monetary, and many other aggregate shocks. In a frictionless world, agents immediately respond to these shocks and bring the economy back to normal times without any delay. In contrast, in the presence of microeconomic adjustment frictions, agents gradually respond to these shocks slowing the economy's transition.

Lumpiness—periods of inaction followed by bursts of activity—is one of the most pervasive manifestations of microeconomic adjustment frictions. Capital investment, price and wage setting, labor hiring and firing, inventory management, consumption of durable goods, portfolio choice, and many other economic decisions made by firms and households exhibit lumpy adjustment. How persistent are the effects of aggregate shocks in lumpy environments? Understanding the role of lumpy adjustment for the propagation of aggregate shocks is crucial for the design and implementation of policies aimed at stabilizing the business cycle and promoting growth.

We propose a new "sufficient statistics" approach to quantitatively assess the role of lumpiness for transitional dynamics. This approach consists of two steps. First, we represent the speed of convergence of the cross-sectional moments after an aggregate shock as a function of two steady-state moments. The premise is that observing agents' responses to idiosyncratic shocks in normal times conveys information about their responses to an aggregate shock. Second, we compute these steady-state moments using microdata on adjustments. The premise is that the size and timing of the actions taken by adjusters inform us about the behavior of non-adjusters during their period of inaction.

We apply the sufficient statistics approach to investigate the propagation of productivity shocks when firms make lumpy investments. In the first step, we link the speed of convergence of average capital following an aggregate productivity shock to (i) the steady-state *variance* of log capital-to-productivity ratios, and (ii) the *covariance* of log capital-to-productivity ratios with the time elapsed since their last adjustment. These sufficient statistics have a meaningful economic interpretation. The variance of log capital-to-productivity ratios reflects the degree of capital misallocation. In turn, the covariance of log capital-to-productivity ratios with the time elapsed since their last adjustment reflects firms' response to depreciation and the relative costs of shrinking and expanding the capital stock. Thus, our theory indicates that matching the degree of capital misallocation and the irreversibility of investment is critical for understanding the transitional dynamics of aggregate capital.[1]

In the second step, we recover these sufficient statistics using data on the size and timing of investment from manufacturing plants in Chile. We discover that these two steady-state moments are large, implying that micro frictions in investment significantly slow down the propagation of productivity shocks. Because different types of adjustment frictions give rise to different values for these moments, the sufficient statistics also serve as model discrimination devices when compared with their empirical values. As a case in point, we find that the investment data discriminates in favor of lumpy models with random fixed costs, and within this subclass, models that feature large downward irreversibility.

[1] Lanteri, Medina and Tan (2019) make a similar point by showing that the transitional dynamics of domestic production following an import-competition shock depend importantly on the size of frictions in capital reallocation. Likewise, Moll (2014) shows in a model with financial frictions that the speed of transitions and steady-state level of capital misallocation jointly depend on the persistence of idiosyncratic productivity shocks.

In summary, when applying our methodology to the study of lumpy investment, we establish structural links between the transitional dynamics of aggregate capital that follow an aggregate productivity shock, steady-state moments such as capital misallocation and the prevalence of irreversibility, and the nature of capital adjustment costs. More generally, our sufficient statistics capture the economic forces that shape aggregate dynamics, serve as model discrimination devices, provide researchers with a unique set of moments to be targeted by lumpy models, and guide empirical efforts in collecting the statistics that are most informative for the theory.

Next, we explain the theory in more detail and provide economic intuitions for the results.

Sufficient statistics for aggregate dynamics. Consider the following economic environment. There is a continuum of agents. Each agent's uncontrolled state x follows a diffusion, $dx_t = -\nu dt + \sigma dB_t$, where the trend is common and the Brownian shocks are idiosyncratic. Payoffs depend on the state x. To control their state, agents pay an adjustment cost. This adjustment cost is different for upward and downward adjustments, and there are random opportunities for free adjustments that arrive at a constant rate. The decision rule consists of (1) a constant reset point x^*, to which agents set their state when they decide to adjust, and (2) the timing of adjustments, which occur when the state reaches one of two thresholds $\{x^-, x^+\}$ or there is a free opportunity to adjust. The economy features a steady-state distribution of idiosyncratic states $F(x)$. We conceptualize aggregate variables as functions of cross-sectional moments of the state (e.g., the mean, the variance, or other higher-order moments).

In this environment, we characterize analytically the transitional dynamics of the cross-sectional distribution's moments after a common exogenous disturbance. Consider the following hypothetical experiment. Initially, the economy is at its steady-state distribution $F(x)$. At time zero, there is a small, identical, and once-and-for-all change in the state of all agents—an aggregate shock—that displaces the distribution away from the stationary one. As agents gradually respond to this shock by actively changing their state, the distribution follows a deterministic transition to its steady state. Assuming that agents follow their steady-state decision rules $\{x^-, x^*, x^+\}$ along the transition, what can we say about the speed of convergence to steady state?

As a first step, we define our notion of the speed of convergence as the area under the impulse-response function of any moment of the state relative to its steady-state value. We label this object the cumulative impulse response (CIR). The CIR is a useful metric of convergence: It summarizes in one scalar both the impact and persistence of the economy's response, eases comparison across models, and represents a"multiplier" of aggregate shocks. A larger CIR implies slower convergence and more persistent effects of the aggregate shock.[2]

Our first theoretical result proves analytically that the CIR can be expressed, up to first order, as a linear combination of two steady-state cross-sectional moments. In particular, the CIR of the average of the distribution depends on (i) the *steady-state variance of the state*, $\mathbb{V}ar[x]$, and (ii) the *covariance of the state with its age* a, $\mathbb{C}ov[x, a]$, where age is the time elapsed since the last adjustment.

[2] Álvarez, Le Bihan and Lippi (2016) and Baley and Blanco (2019) use the CIR to compare the real output effect of monetary shocks across different price-setting models.

A major challenge to applying our sufficient statistics approach arises if $F(x)$ is unobservable, as in the majority of applications. Thus the steady-state moments cannot be computed directly from the data. Fortunately, as economists, we observe detailed panel data $\Omega = \{\Delta x, \tau\}$ with information on the size of adjustments Δx and the duration of completed inaction spells τ. Our second theoretical result provides analytic mappings from the data Ω to moments of the invariant distribution $F(x)$ and parameters of the stochastic process (ν, σ^2, x^*). Importantly, to obtain these mappings, we exploit, exclusively, the properties of Markov processes and the fact that the reset state x^* is constant across agents and time.

Taken together, our two theoretical results provide researchers with the sufficient statistics that characterize the transitional dynamics of aggregate variables in lumpy environments, together with mappings that allow us to infer the sufficient statistics and parameters of the stochastic process using microdata.

Capital dynamics with lumpy investment. To investigate the propagation of aggregate productivity shocks, we set up a parsimonious investment model with adjustment costs, in the spirit of Caballero and Engel (1999) and related literature.[3] Firms produce output with capital. They are subject to depreciation, technological growth, and idiosyncratic productivity shocks. To change their capital, firms pay a fixed cost that scales with firms' size and could be different for upward and downward adjustments. Also, firms face random opportunities for free adjustments. Defining the state x as the log capital-to-productivity ratio, this model falls into the basic environment described above. Moreover, the presence of drift (due to depreciation and growth) and policy asymmetry (due to asymmetric costs) renders this environment an ideal laboratory to show the power of sufficient statistics to discriminate between models of inaction.

Using plant-level investment data from Chile, we recover the two sufficient statistics: The steady-state variance of log capital-to-productivity ratios, $\mathbb{V}ar[x]$, and the covariance of log capital-to-productivity ratios with their age, $\mathbb{C}ov[x, a]$. Concretely, we recover the sufficient statistics using the following empirical moments: the average and dispersion in duration of inaction, the average and dispersion in investment size, and the covariance between duration and size. What do the data tell us about the role of micro lumpiness for capital dynamics? Which type of investment rules match the data?

Our quantitative analysis of the investment model yields three main results. First, the inferred sufficient statistics in the data imply that the speed of convergence is slow. We obtain a CIR of 3.6: A 1% decrease in aggregate productivity generates a total deviation in capital-to-productivity ratios of 3.6% above steady state along the transition path. The implied half-life of the response is approximately 2.5 years. Second, we find that allowing for *asymmetric policies* through large downward irreversibility, and *randomness* in the adjustment process through random opportunities for free adjustments, is key for correctly matching the inferred sufficient statistics that shape aggregate capital dynamics. Lastly, the inferred parameters of the stochastic process are in line with estimates obtained using alternative methods, such as simulated moments.

[3] Similar environments have been studied by Dixit and Pindyck (1994); Bertola and Caballero (1994); and Caballero, Engel and Haltiwanger (1995), and more recently in quantitative general equilibrium by Veracierto (2002); Khan and Thomas (2008); Bachmann, Caballero and Engel (2013); and Winberry (2019).

Contributions to the literature. We highlight three contributions to previous work.

First, we provide sufficient statistics that capture the role of micro lumpiness for aggregate dynamics. Álvarez, Le Bihan and Lippi (2016) provide the first step in this direction by studying the transitions of the first moment of the distribution in economies with zero drift and symmetric policies. They show that in a large class of price-setting models, the CIR of real output following a monetary shock is proportional to the kurtosis of price changes times the average price duration.[4] Their theoretical strategy links the CIR directly to the observables in the data. Our strategy is different, because we split this challenging problem into two simpler subproblems: From the CIR to steady-state moments and from steady-state moments to the data. In our view, this approach has various advantages. It improves our understanding of the economic forces behind these links. It eases the analysis of sufficient statistics in richer economic environments than previously studied, including drift and asymmetric policies. And finally, it allows us to characterize the transitional dynamics of higher-order moments beyond the average.[5]

Second, we strengthen the bridge between two branches of the literature that study lumpy economies with different objectives and methodologies. The first aims to understand the role of lumpiness for the propagation of aggregate shocks; see Caplin and Spulber (1987); Caplin and Leahy (1991, 1997); and Caballero and Engel (1991) for early work. The second aims to quantify the role of lumpiness for productivity losses in steady state. For example, Álvarez, Beraja, Gonzalez-Rozada and Neumeyer (2018) analyze inefficient price dispersion and Asker, Collard-Wexler and De Loecker (2014) analyze capital misallocation due to adjustment costs. To our best knowledge, we are the first to show theoretically that structural links exist between transitional dynamics and the steady state in lumpy economies. We believe our approach may engage researchers in exploiting the connections between these two dimensions of the same environment.

Third, in the context of the investment, we identify the relevant moments that lumpy models must target to explain aggregate capital dynamics correctly. In this way, our work speaks to the debate about the nature of adjustment frictions in the quantitative investment literature. This literature targets moments that appear, ex ante, to be sensible choices, e.g., the frequency of inaction and investment spikes. These moments are important, but do not fully capture the role of lumpiness for aggregate dynamics, which the sufficient statistics do.[6] Our theory reveals that lumpy investment models should match the first two moments of the joint distribution of duration of inaction and size of adjustment.

Structure of the paper. Section 2 presents a parsimonious model of lumpy investment that serves as a framework to develop the theory. Section 3 establishes the link between the CIR and steady-state moments. Section 4 establishes links between steady-state moments and microdata. Section 5 applies the theory to Chilean data. Section 6 discusses the scope and potential extensions of the theory.

[4] This sufficient statistic is valid for the price-setting models in Taylor (1980); Calvo (1983); Reis (2006); Golosov and Lucas (2007); Nakamura and Steinsson (2010); Midrigan (2011); and Álvarez and Lippi (2014), among other.

[5] In a frictionless environment, Gabaix, Lasry, Lions and Moll (2016) characterize the speed of convergence of higher-order moments of the cross-sectional distribution to study the dynamics of inequality.

[6] See Thomas (2002); Cooper and Haltiwanger (2006); Gourio and Kashyap (2007); Khan and Thomas (2008); Bachmann, Caballero and Engel (2013); and Winberry (2019) for various calibration strategies that deliver different conclusions about the role of lumpiness for aggregate dynamics.

2 A Parsimonious Model of Lumpy Investment

How does an economy's capital respond to aggregate productivity shocks when firms face capital adjustment frictions? In this section we present a parsimonious model of lumpy investment to derive sufficient statistics that characterize the relevance of micro lumpiness for aggregate dynamics. We first study the problem of an individual firm and characterize its optimal investment policy in terms of capital-to-productivity ratios. Then we consider the steady state of an economy with a continuum of ex ante identical firms and perturb it with an aggregate productivity shock. Finally, we define the cumulative impulse response, or CIR, of aggregate capital, which measures the persistence of transitional dynamics.

2.1 The problem of an individual firm

Time is continuous and extends forever. Consider an individual firm that produces output using capital. This firm faces capital adjustment frictions and a constant real interest rate r.

Technology and shocks. The firm produces output y_t using capital k_t according to a production function with decreasing returns to scale

$$y_t \;=\; (z_t e_t)^{1-\alpha} k_t^{\alpha}, \quad \alpha < 1. \tag{1}$$

The firm's total productivity is driven by aggregate z_t and idiosyncratic e_t components. Aggregate productivity z_t grows deterministically at a rate $\mu_z > 0$,

$$\mathrm{d}\log(z_t) \;=\; \mu_z \,\mathrm{d}t. \tag{2}$$

Idiosyncratic productivity shocks e_t follow a geometric Brownian motion with zero drift and volatility σ,

$$\mathrm{d}\log(e_t) \;=\; \sigma \,\mathrm{d}W_t, \tag{3}$$

where W_t is a Wiener process. The capital stock, if uncontrolled, depreciates at a constant rate $\zeta > 0$.

The firm can control its capital stock through purchasing or selling capital. For every change in its capital stock (investment) $i_t \equiv \Delta k_t$, the firm must pay an adjustment cost θ_t that is proportional to its total productivity.[7] The adjustment cost is different for positive and negative investments, and there exist random opportunities for free adjustments. Concretely, the adjustment cost takes the form

$$\theta_t \;\equiv\; \Theta\left(i_t, \Delta N_t\right) z_t e_t, \tag{4}$$

where N_t is a Poisson counter with arrival rate λ. The function $\Theta\left(i_t, \Delta N_t\right)$ takes the following values:

$$\Theta\left(i_t, \Delta N_t\right) \;\equiv\; \begin{cases} 0 & \text{if } i_t = 0 \text{ or } \Delta N_t = 1, \\ \theta^- & \text{if } i_t > 0 \text{ and } \Delta N_t = 0, \\ \theta^+ & \text{if } i_t < 0 \text{ and } \Delta N_t = 0. \end{cases} \tag{5}$$

[7]For any stochastic process q_t, we use the notation $\Delta q_t = q_t - q_{t^-}$ to denote the limit from the left, where $q_{t^-} \equiv \lim_{s \uparrow t} q_s$.

We label this type of adjustment friction—i.e., asymmetric fixed costs with random free adjustments—*Bernoulli fixed costs*. We consider different costs for downsizing and upsizing the capital stock to reflect, in a parsimonious way, several asymmetric frictions in capital adjustment. In turn, we consider random free adjustments as a proxy for frictions that contain a stochastic element in the time dimension, e.g., information or search frictions.[8] Our analysis shows that both frictions are relevant to match the data.

An advantage of this formulation is that it nests two benchmark cases of strict state- and time-dependence within a more general framework. Setting $\lambda = 0$ shuts down the random free adjustments and collapses the model into a standard state-dependent fixed cost problem, whereas in the limiting case of infinite fixed costs, i.e., $\{\theta^-, \theta^+\} \to \{\infty, \infty\}$, the model collapses into a standard time-dependent problem that allows adjustment only at random dates that arrive at a rate $\lambda > 0$.[9]

Investment problem. Let $V(k, z, e)$ be the value of the firm. Given the initial conditions (k_0, z_0, e_0), the firm chooses a sequence of capital adjustment dates $\{T_h\}_{h=1}^{\infty}$ and investments $\{i_{T_h}\}_{h=1}^{\infty}$, where h counts the number of adjustments, to maximize its expected discounted stream of profits. The sequential problem of the firm is described by

$$V(k_0, z_0, e_0) \;=\; \max_{\{T_h, i_{T_h}\}_{h=1}^{\infty}} \mathbb{E}\left[\int_0^{\infty} e^{-rt} y_t \,\mathrm{d}t \;-\; \sum_{h=1}^{\infty} e^{-rT_h}\left(\theta_{T_h} + i_{T_h}\right)\right], \tag{6}$$

subject to the production function (1), aggregate productivity (2), idiosyncratic productivity (3), adjustment costs (4 and 5), and the law of motion for its capital stock

$$\log(k_t) \;=\; \log(k_0) \;-\; \zeta t \;+\; \sum_{h:\, T_h \le t} \log\left(1 + \frac{i_{T_h}}{k_{T_h^-}}\right), \tag{7}$$

which describes a period's capital stock as a function of the firm's initial stock k_0, the depreciation rate ζ, and the sum of the investments made at prior adjustment dates.

2.2 Optimal Policy

We solve the sequential problem in (6) recursively as a stopping-time problem using the Principle of Optimality. The resulting investment policy is characterized by an asymmetric inaction region $\mathcal{R} \equiv \{(k, z, e) : k^-(z, e) \le k \le k^+(z, e)\}$, where $k^-(z, e)$ and $k^+(z, e)$ are the lower and upper inaction thresholds, together with a reset value $k^*(z, e)$ to which capital is set upon every adjustment. Given these three functions, $\{k^-, k^*, k^+\}$, adjustment happens at every date T_h when the capital stock falls outside the inaction region $\mathcal{R}$ or there is an opportunity of free adjustment:

$$T_h \;=\; \inf\left\{t \ge T_{h-1} :\; k_t \notin \mathcal{R} \;\text{ or }\; \Delta N_t = 1\right\}. \tag{8}$$

[8] Investment with asymmetric adjustment frictions, e.g., partial irreversibility, is studied by Abel and Eberly (1994, 1996); Bertola and Caballero (1994); Dixit and Pindyck (1994); and Lanteri (2018); investment with information frictions is studied by Verona (2014), using the inattentiveness model of limited information in Reis (2006); and investment with search frictions is studied by Kurmann and Petrosky-Nadeau (2007) and Ottonello (2018), among others.

[9] The Bernoulli fixed-cost formulation originated in the pricing literature to match the empirical distribution of price changes. See, for example, Nakamura and Steinsson (2010) and Álvarez and Lippi (2014).

Investment i_{T_h} is the difference between the reset value and the capital immediately before adjustment:

$$i_{T_h} \;=\; k^*(z,e) - k_{T_h^-}. \tag{9}$$

Given the optimal adjustment dates in (8), we define two useful concepts of duration of inaction: the *duration of completed spells*, denoted by τ, equal to the difference of two consecutive adjustment dates

$$\tau_h \;\equiv\; T_h - T_{h-1}, \quad \text{with} \quad T_0 = 0, \tag{10}$$

and the *duration of uncompleted spells, or capital age*, denoted by a, equal to the time elapsed since the last adjustment

$$a_t \;\equiv\; t - \max\{T_h : T_h \leq t\}. \tag{11}$$

Observe that after each adjustment, the capital age is reset to zero, i.e., $a_{T_h} = 0$.

Log capital-to-productivity ratio. To characterize the policy, it is convenient to reduce the state space and recast the problem in terms of a new variable, the log of the capital-to-productivity ratio:

$$\hat{x}_t \;\equiv\; \log\left(\frac{k_t}{z_t e_t}\right). \tag{12}$$

The problem admits this reformulation because of the homothetic production function and the adjustment costs proportional to productivity. Lemma 1 characterizes the firm value and the optimal investment policy in terms of the log capital-to-productivity ratio through the standard sufficient optimality conditions. The firm value and the policy must satisfy: (i) the Hamilton-Jacobi-Bellman equation, which describes the evolution of the firm's value during periods of inaction, (ii) the value-matching conditions, which set the value of adjusting equal to the value of not adjusting at the borders of the inaction region, and (iii) the smooth-pasting and optimality conditions, which ensure differentiability at the borders of inaction and the reset point. To simplify notation, we define $\nu \equiv \zeta + \mu_z$, which reflects the drift affecting the uncontrolled $\hat{x}$'s, and $\rho \equiv r + \lambda - \mu_z - \sigma^2/2$. All proofs appear in Online Appendix A.

Lemma 1. *Let $\mathcal{V}(\hat{x}) : \mathbb{R} \to \mathbb{R}$ be a function of the log capital-to-productivity ratio. If $\mathcal{V}(\hat{x})$ and the values $\{\hat{x}^-, \hat{x}^*, \hat{x}^+\}$ satisfy the following three conditions, then $\mathcal{V}(\hat{x}) = V(k,z,e)/(ze)$ and the optimal policy is $\{k^-, k^*, k^+\} = ze \times \{\exp(\hat{x}^-), \exp(\hat{x}^*), \exp(\hat{x}^+)\}$.*

1. *In the interior of the inaction region, i.e., $\hat{x} \in (\hat{x}^-, \hat{x}^+)$, $\mathcal{V}(\hat{x})$ solves the HJB equation:*

$$\rho\mathcal{V}(\hat{x}) \;=\; \exp(\alpha\hat{x}) \;-\; \nu\frac{\mathrm{d}\mathcal{V}(\hat{x})}{\mathrm{d}\hat{x}} \;+\; \frac{\sigma^2}{2}\frac{\mathrm{d}^2\mathcal{V}(\hat{x})}{\mathrm{d}\hat{x}^2} \;+\; \lambda\left[\mathcal{V}(\hat{x}^*) - (\exp(\hat{x}^*) - \exp(\hat{x}))\right]. \tag{13}$$

2. *At the borders of the inaction region $\{\hat{x}^-, \hat{x}^+\}$, $\mathcal{V}(\hat{x})$ satisfies the value-matching conditions:*

$$\mathcal{V}(\hat{x}^-) \;=\; \mathcal{V}(\hat{x}^*) \;-\; \theta^- \;-\; \left(\exp(\hat{x}^*) - \exp(\hat{x}^-)\right), \tag{14}$$

$$\mathcal{V}(\hat{x}^+) \;=\; \mathcal{V}(\hat{x}^*) \;-\; \theta^+ \;-\; \left(\exp(\hat{x}^*) - \exp(\hat{x}^+)\right). \tag{15}$$

3. *At the borders of the inaction region and the reset state,* $\{\hat{x}^-, \hat{x}^*, \hat{x}^+\}$*,* $\mathcal{V}(\hat{x})$ *satisfies the smooth-pasting and the optimality conditions, respectively:*

$$\mathcal{V}'(\hat{x}) \;=\; \exp(\hat{x}). \tag{16}$$

Online Appendix A.4 presents a system of equations to solve for all of the objects in Lemma 1 in terms of structural parameters $\{r, \alpha, \mu_z, \sigma, \zeta, \theta^+, \theta^-, \lambda\}$. Notice that when expressed in terms of log capital-to-productivity ratios, the inaction region and the reset state are constant and thus memoryless. This implies that each adjustment completely erases the history of idiosyncratic shocks. Also notice that adjustment dates in (8), duration of inaction in (10), and age in (11) can be written as functions of log capital-to-productivity ratios (just exchanging k for $\hat{x}$ in their expressions) and their distributions remain unchanged. In the case of investment, the continuity of the productivity process allows us to recover the investment rate in (7) from the change in the capital-to-productivity ratio:

$$1 + \frac{i_{T_h}}{k_{T_h^-}} \;=\; \frac{k^*(z_{T_h}, e_{T_h})}{k_{T_h^-}} \;=\; \frac{k^*(z_{T_h}, e_{T_h})/(z_{T_h} e_{T_h})}{k_{T_h^-}/(z_{T_h^-} e_{T_h^-})} \;=\; \exp(\hat{x}^* - \hat{x}_{T_h^-}) \;=\; \exp(\Delta \hat{x}_{T_h}). \tag{17}$$

In (17), the first equality applies the definition of investment. In the second equality, we multiply and divide by total productivity at the moment of adjustment and use the continuity of the stochastic process in the denominator to exchange (z_{T_h}, e_{T_h}) for $(z_{T_h^-}, e_{T_h^-})$. In the third and fourth equalities, we substitute the definition of capital-to-productivity ratios and their change.

With the problem of an individual firm fully characterized, we turn to the analysis of an economy with a continuum of firms.

2.3 Economy with a continuum of firms

Consider a continuum of ex ante identical firms that face the problem described in the previous section. We assume the stochastic processes of idiosyncratic productivity W_t and the arrival of free adjustments N_t are independent across firms. This economy features a steady-state distribution $G(\hat{x})$, with density $g(\hat{x})$, that solves the following Kolmogorov forward equation with its boundary conditions:

$$\nu \frac{\mathrm{d}g(\hat{x})}{\mathrm{d}\hat{x}} \;+\; \frac{\sigma^2}{2}\frac{\mathrm{d}^2 g(\hat{x})}{\mathrm{d}\hat{x}^2} \;-\; \lambda g(\hat{x}) \;=\; 0 \;\; \forall \hat{x} \neq \hat{x}^*, \quad \int_{\hat{x}^-}^{\hat{x}^+} g(\hat{x})\,\mathrm{d}\hat{x} \;=\; 1, \quad g(\hat{x}^-) \;=\; g(\hat{x}^+) \;=\; 0. \tag{18}$$

We denote by $\mathbb{E}_g[\cdot]$ the expectations computed with the steady-state distribution g. In particular, the average capital-to-productivity ratio in steady state is $\hat{K} \equiv \int_{\hat{x}^-}^{\hat{x}^+} \exp(\hat{x})\,\mathrm{d}g(\hat{x})\,\mathrm{d}\hat{x} = \mathbb{E}_g[\exp(\hat{x})]$.

Capital gaps. Using the steady-state distribution, for every firm we define the *capital gap* as its log capital-to-productivity ratio relative to the steady-state average:

$$x_t \;\equiv\; \hat{x}_t \;-\; \mathbb{E}_g[\hat{x}], \tag{19}$$

where $\mathbb{E}_g[\hat{x}] \equiv \int_{\hat{x}^-}^{\hat{x}^+} \hat{x}\, \mathrm{d}g(\hat{x})\, \mathrm{d}\hat{x}$. Notice that in the absence of adjustment frictions, capital gaps are always equal to zero. Similarly, we redefine the investment policy by centralizing the borders of the inaction region and the reset state:

$$(x^-,\ x^*,\ x^+) \ = \ (\hat{x}^- - \mathbb{E}_g[\hat{x}],\ \hat{x}^* - \mathbb{E}_g[\hat{x}],\ \hat{x}^+ - \mathbb{E}_g[\hat{x}]). \tag{20}$$

From now on, we will work with capital gaps x. We use $F(x)$ and $f(x)$ to denote the distribution and density of capital gaps, respectively. We will also denote by $\mathbb{E}[\cdot]$ the expectations computed with their steady-state distribution F. At the end of this section, we discuss why this is appropriate for answering the questions we are interested in and compare our definition with the standard approach in the literature—which defines gaps relative to a micro-optimum target. Importantly, given the centralization, the reset gap x^* is understood as the gap of adjusting firms *relative* to the average gap in the steady state.

2.4 Aggregate productivity shock

How does aggregate capital respond to an aggregate productivity shock? Starting from the steady state, we introduce a small and unanticipated decrease in the (log) level of aggregate productivity of size $\delta > 0$, which we label as δ-perturbation. We normalize the arrival date of the aggregate shock to $t = 0$, so the aggregate shock is given by

$$\ln z_0 \ = \ \ln z_{0^-} - \delta. \tag{21}$$

This negative aggregate productivity shock generates a homogeneous increase in the capital gap of all firms, as they now have too much capital relative to their productivity:

$$x_0 \ = \ x_{0^-} + \delta. \tag{22}$$

Let $F_t(x)$ be the distribution t periods after the aggregate shock and $\mathbb{E}_t[\cdot]$ denote expectations computed with F_t. The distribution of capital gaps displaces horizontally to the right relative to the steady-state distribution, i.e., the gap distribution immediately after the aggregate shock is $F_0(x) = F(x - \delta)$.[10] After the initial displacement, the gap distribution evolves according to the firms' policies, and eventually, it converges back to its steady state. Recall that the interest rate r is constant by assumption, which implies that individual investment policies are not a function of the distribution. For this reason, the steady-state policies continue to hold along the transition path.

Panel A in Figure I plots the steady-state density $f(x)$ and the initial density $f_0(x)$ following the δ-perturbation; it also shows an arbitrary m-th cross-sectional moment before and after the shock. Our exercise consists of tracking these moments as they make their way back to their steady-state value.

Figure I – Distribution Dynamics and Cumulative Impulse Response (CIR)

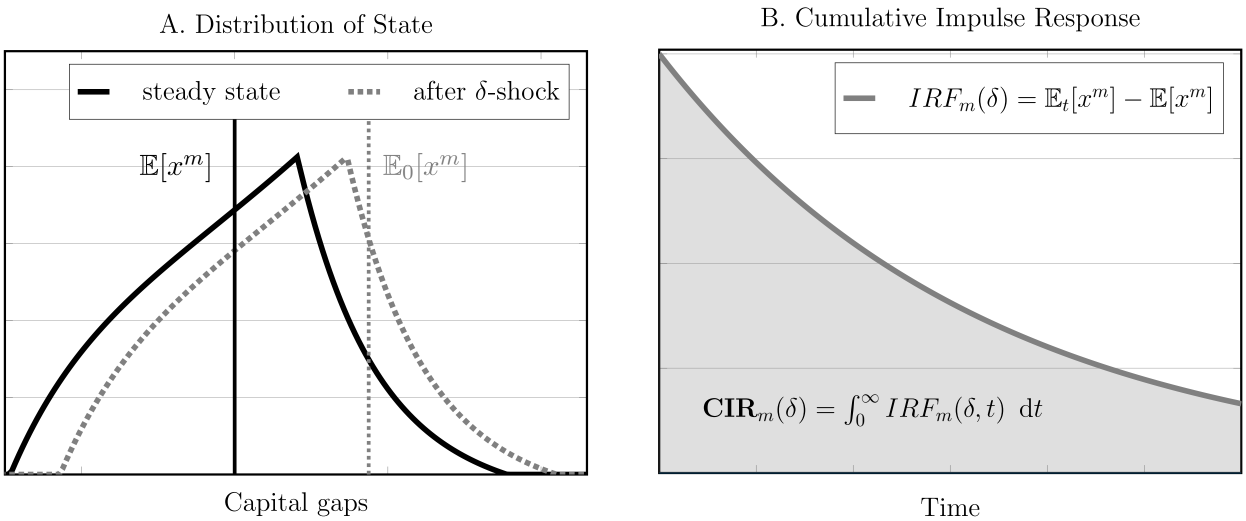

Notes: Panel A shows the steady-state distribution of the idiosyncratic state $f(x)$ and an initial distribution $f_0(x) = f(x-\delta)$ following the δ-shock. It also illustrates an arbitrary m-th cross-sectional moment to be tracked from its initial value $\mathbb{E}_0[x^m]$ toward its steady-state value $\mathbb{E}[x^m]$. Panel B shows the transitional dynamics of the m-th moment in two ways: the impulse response function (IRF, solid line) and the cumulative impulse response (CIR, total area under the IRF).

Aggregate deviations from steady state. We are interested in characterizing the effects of the aggregate productivity shock on the average capital-to-productivity ratio, $\hat{K}_t \equiv \mathbb{E}_{g_t}[\exp(\hat{x})]$, expressed as percent deviations from its steady-state value $\hat{K} \equiv \mathbb{E}_g[\exp(\hat{x})]$. This deviation, up to a first-order approximation, can be expressed as the average capital gap:

$$\frac{\hat{K}_t - \hat{K}}{\hat{K}} = \frac{\mathbb{E}_{g_t}[\exp(\hat{x})]}{\mathbb{E}_g[\exp(\hat{x})]} - 1 = \frac{\mathbb{E}_t[\exp(x)]}{\mathbb{E}[\exp(x)]} - 1 \approx \mathbb{E}_t[x]. \tag{23}$$

To obtain expression (23), the first equality applies the definition of aggregate capital-to-productivity ratios during the transition and in steady state. The second equality is obtained by multiplying and dividing by $\exp(\mathbb{E}_g[\hat{x}])$ and writing in terms of capital gaps. The third step uses a first-order approximation of the exponential function, i.e., $e^x \approx 1 + x$, and applies the definition of capital gaps. In this way, we connect the average capital-to-productivity ratio to the average capital gap.

While the aggregate capital deviations are not exactly equal to the average capital gap, the approximation is quite helpful for exposition. To exactly compute the deviation, one needs all of the moments of the capital gap distribution, as the full expansion is $\mathbb{E}_t[\exp(x)] = \sum_{n=0}^{\infty}(\mathbb{E}_t[x^m]/m!)$. Later in the paper, we characterize the transitions of all moments of x, and Section 6 explains how to compute the CIR of any continuous function of the capital gaps.

[10] We consider an infinitesimal shock δ that displace the cross-sectional distribution away from steady-state. If $f(x-\delta)$ denotes the initial density of capital gaps and $f_x(x)$ denotes its derivative, we can approximate the perturbation as $f(x-\delta) \approx f(x) - \delta f_x(x)$. Since $\delta > 0$, the initial density is equal to a right shift of the steady-state density. The analysis is closely related to the "marginal response function" in Borovička, Hansen and Scheinkman (2014).

2.5 Cumulative Impulse Response (CIR)

To analyze transitional dynamics, we consider the impulse response function (IRF) of the m-th moment of capital gaps following the δ-perturbation. This object is denoted by $\text{IRF}_m(\delta, t)$, it is a function of time, and it is defined as the difference between the moment's value at time t and its steady-state value:

$$\text{IRF}_m(\delta, t) \;\equiv\; \mathbb{E}_t\left[x^m\right] - \mathbb{E}[x^m]. \tag{24}$$

Following Álvarez, Le Bihan and Lippi (2016), we define the cumulative impulse response (CIR), denoted by $\text{CIR}_m(\delta)$, as the area under the $\text{IRF}_m(\delta, t)$ curve across all dates $t \in (0, \infty)$:

$$\text{CIR}_m(\delta) \;\equiv\; \int_0^\infty \text{IRF}_m(\delta, t)\;\mathrm{d}t. \tag{25}$$

Panel B in Figure I plots these two objects. The solid line represents the IRF, and the area underneath it is the CIR. The CIR measures the convergence speed toward steady state. The larger the CIR the longer it takes firms to respond to the aggregate shock through investment, and the slower the convergence. The CIR is a useful metric: It summarizes both the impact and the persistence of the response in one scalar, eases the comparison of models, and represents a "multiplier" of aggregate shocks. It is illustrative to compare the CIR with and without adjustment frictions. In the frictionless environment, individual gaps are always equal to zero. When the aggregate shock hits the economy, all firms respond instantly to keep their gap at zero. Thus the impulse response is a jump with zero area underneath, i.e., $\text{CIR}_m = 0$ for all m.

Remarks on the definition of gaps. Our definition of capital gaps centers log capital-to-productivity ratios around their steady-state average. This contrasts with the standard approach in the literature that defines gaps using a micro target, which is usually the frictionless optimal capital choice (see, for example, Caballero, Engel and Haltiwanger, 1997; Caballero and Engel, 1993; Cooper and Willis, 2004). We base our approach on the fact that specifying a micro target is irrelevant for the study of impulse responses centered around steady-state: The micro target cancels out as it enters symmetrically the impulse response and the steady state. In a similar way, the micro-optimum target does not affect the investment rate distribution. Since investment is all about *changes* in the capital stock, a theory that aims to explain investment only needs to specify the *relative position* of a firm's capital-to-productivity ratio in the overall distribution, and not its absolute level.

Remarks on the framework. In this section, we developed a parsimonious framework to study the propagation of aggregate productivity shocks when firms make lumpy investments. Before we proceed, let us summarize the three key properties required to develop the sufficient statistics approach described in the following sections.

First, our theory requires a constant reset capital-to-productivity ratio across firms and time, which implies that the history of idiosyncratic shocks is fully erased at each capital adjustment. This feature, which results from the Bernoulli fixed-cost formulation, greatly simplifies the analysis of transitional dynamics, as one only needs to keep track of firms until their first adjustment and can be then dropped out (this is formalized in Lemma 2 below). The limitation is that our theory does not cover convex adjustment costs, time-to-build, or other features that may generate correlation between adjustments. In Section 6 we discuss potential avenues to relax this requirement.

Second, our theory requires that steady-state policies hold along the transitional dynamics. To satisfy this requirement, we directly assume constant prices. In Online Appendix C, we develop a general equilibrium model that delivers constant prices as an equilibrium outcome. We extend the simple environment to include labor as a production input, sticky prices, a capital-good producer, and monetary policy. We show that if labor supply has an infinite Frisch elasticity, the capital-good producer can perfectly substitute between labor and output to produce investment goods, and monetary policy systematically responds to aggregate productivity, then prices are constant in equilibrium.

Lastly, our theory only characterizes the dynamics of aggregate variables that are functions of capital gaps—for example, the average capital-to-productivity ratio. Other aggregate variables of interest, such as aggregate output, depend on the joint distribution of capital and productivity. The general equilibrium model in Online Appendix C incorporates the idiosyncratic productivity shocks in such a way that aggregate feasibility only depends on firms' capital-to-productivity ratios instead of capital and productivity separately. To do this, we borrow from the literature and reinterpret idiosyncratic shocks as "capital quality shocks" that simultaneously enter the marginal costs and marginal revenues, counteracting their effect on profits.[11] As a result, capital quality shocks reduce the dimensionality of the aggregate state space from the joint distribution of capital and productivity to the distribution of their ratio.

3 Propagation and steady-state moments

In this section, we establish the theoretical relationships between the dynamics of average capital-to-productivity ratio following an aggregate productivity shock and two steady-state moments of the capital gap distribution. We show analytically that their cumulative impulse response (CIR_1), up to first order, is equal to a linear combination of the cross-sectional variance of capital gaps $\mathbb{V}ar[x]$ and the covariance of capital gaps and their age $\mathbb{C}ov[x,a]$. Then, we generalize this result to the transitional dynamics of any moment m. We also compute the value of the sufficient statistics for some special cases of interest.

3.1 Characterizing the CIR

As the first step in the characterization, Lemma 2 expresses the cumulative impulse response of moment m—the CIR_m defined in (25)—as the solution to a collection of stopping-time problems indexed by the initial capital gap. This Lemma establishes that it is only necessary to keep track of a firm from the arrival of the aggregate shock at $t = 0$ until its *first* adjustment at $t = \tau$. This result is extremely convenient, since it allows us to characterize the propagation of the aggregate productivity shocks without the need to track the evolution of the whole distribution of capital gaps.

[11] In the pricing literature, a similar formulation is used by Woodford (2009); Midrigan (2011); Álvarez and Lippi (2014); Baley and Blanco (2019); and Blanco (2020), among others.

Lemma 2. *The CIR_m can be written recursively as*

$$CIR_m(\delta) \;=\; \int v_m(x)\,\mathrm{d}F(x-\delta), \tag{26}$$

where the function $v_m(x)$ measures the cumulative deviations of the m-th moment from its steady-state value for a firm with initial capital gap x:

$$v_m(x) \;\equiv\; \mathbb{E}\left[\int_0^{\tau} (x_t^m - \mathbb{E}[x^m])\,\mathrm{d}t \Big| x_0 = x\right]. \tag{27}$$

An analogous result to Lemma 2 was first shown by Álvarez, Le Bihan and Lippi (2016) in a driftless and symmetric environment for $m = 1$, noting that after the first adjustment, a firm's expected contribution to the average gap is zero, since positive and negative contributions are equally likely. Thus, in their environment, the average gap conditional on adjustment is equal to zero at every date.

What is surprising about our result is that this property still holds *in the presence of drift and asymmetric policies*. In such an environment, the expected contribution to the average gap by an adjusting firm is clearly no longer equal to zero. In other words, the average gap conditional on adjustment fluctuates. However, once a firm decides it is time to adjust, it incorporates the history of all shocks that have hit it, including the aggregate one. Any subsequent adjustments are purely driven by idiosyncratic conditions that do not contribute to deviations from the steady-state moment $\mathbb{E}[x^m]$—i.e., $v_m(x^*)$ equals zero. Importantly, this property hinges exclusively on properties of Markov processes. The Lemma does not need to assume a specific stochastic process for x, the source of the rigidity, the moment we wish to track, or the type of initial perturbation. The crucial assumption is that an adjustment erases the history of shocks—a property embedded in the constant reset state.

3.2 Sufficient statistics for transitional dynamics

Now we proceed to characterize the CIR as a function of steady-state moments. For expositional purposes, we consider the joint steady-state distribution of capital gaps and age, denoted by $F(x, a)$, and for any two numbers $k, l \in \mathbb{N}$, we define the joint steady-state moments of capital gap and age as

$$\mathbb{E}[x^k a^l] \;\equiv\; \int_x \int_a x^k a^l \;\mathrm{d}F(x,a), \quad \forall k, l \in \mathbb{N}. \tag{28}$$

For the Bernoulli fixed-cost model described in Section 2, Proposition 1 characterizes the $\mathrm{CIR}_m(\delta)$. It considers the first-order Taylor expansion $\mathrm{CIR}_m(\delta) = \delta \mathrm{CIR}'_m(0) + o(\delta^2)$, and expresses the term $\mathrm{CIR}'_m(0)$ as a linear combination of two steady-state moments of the distribution $F(x, a)$.

Proposition 1. *With Bernoulli fixed costs, the CIR_m is given by*

$$\frac{CIR_m(\delta)}{\delta} \;=\; \frac{\mathbb{E}\left[x^{m+1}\right] + \nu \mathbb{C}ov\left[x^m, a\right]}{\sigma^2} \;+\; o(\delta). \tag{29}$$

Equation (29) shows that up to first order, the transitional dynamics of the m–th moment of capital gaps are structurally linked to the $m+1$ steady-state moment plus a covariance term that corrects for the presence of drift. To better understand why these two moments are sufficient statistics for the propagation of aggregate shocks, let us focus on the case $m = 1$, stated in the following Corollary.

Corollary 1. *With Bernoulli fixed costs, the CIR_1 is given by*

$$\frac{CIR_1(\delta)}{\delta} = \frac{\mathbb{V}ar\,[x] + \nu \mathbb{C}ov\,[x, a]}{\sigma^2} + o(\delta). \tag{30}$$

Equation (30) presents the CIR of the mean of the cross-sectional distribution as a linear combination of the steady-state variance of capital gaps, $\mathbb{V}ar\,[x]$, and the steady-state covariance between capital gaps and their age, $\mathbb{C}ov\,[x, a]$. Since aggregate shocks z_t and idiosyncratic shocks e_t enter symmetrically into a firms' capital gaps, how firms respond to idiosyncratic shocks is very informative about the way in which they respond to aggregate shocks. In the rest of this section, we explore this link between firms' response to idiosyncratic shocks (measured via steady-state moments) and aggregate shocks (measured via the CIR) to interrogate how aggregate shocks propagate in this environment.

Insensitivity to productivity shocks. To explain the link between the two sides of Equation (30) heuristically, we propose a notion of "insensitivity to productivity shocks." Let $\tilde{W}_t \equiv (W_t - W_{t-a_t})/\sigma$ be the sum of all idiosyncratic shocks received by a firm since its last adjustment, normalized by idiosyncratic volatility. We define the economy's *insensitivity to idiosyncratic shocks* as the covariance of capital gaps with the process $\tilde{W}_t$ in the population of firms, i.e., $\mathbb{C}ov[x_t, \tilde{W}_t]$. Intuitively, if firms sluggishly incorporate changes in their productivity into their capital stock, then the cross-section features a strong relationship between capital gaps and idiosyncratic shocks. In that case, the covariance is large. In the opposite case without adjustment frictions, firms are extremely sensitive to idiosyncratic shocks, so they continuously adjust their gaps to keep them at zero, yielding a zero covariance.

Let us link our definition of insensitivity to productivity shocks to the CIR_1. The capital gap of any firm at time t can be written as $x_t = x^* - \nu a_t + \sigma^2 \tilde{W}_t$. Multiplying both sides by x_t, taking the cross-sectional average, and using $\mathbb{E}[x] = 0$, we obtain $\mathbb{V}ar[x] = -\nu \mathbb{C}ov[x, a] + \sigma^2 \mathbb{C}ov[x, \tilde{W}]$. Rearranging, it yields $\mathbb{C}ov[x, \tilde{W}] = (\mathbb{V}ar[x] + \nu \mathbb{C}ov[x, a])/\sigma^2$, which is exactly the expression for the CIR_1 in (30).

This analysis reveals two novel insights. First, when the drift is zero, the propagation of aggregate productivity shocks is proportional to the steady-state variance of capital gaps, normalized by idiosyncratic volatility, i.e., $\mathbb{V}ar[x]/\sigma^2$. This ratio of "ex post" to "ex ante" dispersions is a sufficient statistic for aggregate capital's insensitivity to productivity shocks: A large ratio signals that firms are very insensitive, and thus the average capital gap slowly converges back to its steady state. The second insight is that when the drift is different from zero, the covariance term $\mathbb{C}ov[x, a]$ appears in the expression to correct for the additional dispersion generated by the drift, so as to identify correctly the insensitivity to productivity shocks.

Extreme sensitivity. Besides the frictionless case, an environment with a very large drift parameter ($\nu \to \infty$) also features extreme sensitivity to idiosyncratic shocks. In this limiting case, the joint steady-state distribution of gaps and age $F(x,a)$ weakly converges to the distribution of an economy without idiosyncratic shocks $\sigma^2 = \lambda = 0$. Capital gaps in the limit are generated by the process $dx_t = -\nu dt$ so that they become an affine function of age, i.e., $x_t = x^* - \nu a_t$. Multiplying both sides by x_t and taking expectations, we obtain $\mathbb{V}ar[x] = -\nu \mathbb{C}ov[x,a]$ and thus $\text{CIR}_1 = 0$. Corollary 2 formalizes this argument.[12]

Corollary 2. *With Bernoulli fixed costs, when the drift goes to infinity, then*

$$\lim_{\nu\to\infty,\sigma^2>0} \frac{\mathbb{V}ar[x]}{\nu \mathbb{C}ov[x,a]} = -1, \tag{31}$$

which implies that $\lim_{\nu\to\infty,\sigma^2>0} CIR_1(\delta)/\delta = 0$.

This limiting case speaks to a classic money-neutrality result in the pricing literature developed by Caplin and Spulber (1987).[13] That paper considers an environment with nonzero drift (due to inflation) and zero idiosyncratic risk. The authors show that aggregate money shocks have no effect on real output: Money is neutral. In our jargon, the CIR_1 equals zero when $\lim_{\sigma^2\to 0,\nu>0} (\nu/\sigma^2) = \infty$. Corollary 2 replicates an analogous neutrality result by taking an equivalent limit: $\lim_{\nu\to\infty,\sigma^2>0} (\nu/\sigma^2) = \infty$. In the investment environment, this result implies that technological innovations working through capital depreciation ζ or economic growth μ_z—the two components of the drift—directly affect the propagation of aggregate productivity shocks.

In summary, the linear combination of the variance of capital gaps $\mathbb{V}ar[x]$ and the covariance between capital gaps and their age $\mathbb{C}ov[x,a]$ encodes firms' insensitivity to the idiosyncratic shocks, acting as a sufficient statistic for the speed at which the mean of the cross-sectional distribution converges back to the steady state following an aggregate productivity shock.

The CIR of higher-order moments. The macro literature mainly focuses on the dynamics of cross-sectional averages—e.g, capital, output, and inflation. However, there is increasing interest in the dynamics of higher-order moments. The sufficient statistics for the CIR_m in Proposition 1, for $m > 1$, could in principle be helpful to researchers interested in measuring and characterizing higher-order dynamics using models of lumpy adjustment. For example, the CIR_2—which measures the dynamics of the second moment—relates to the steady-state third moment. This relationship could be used to connect the cyclical fluctuations in the *dispersion* of investment rates (Bachmann, Caballero and Engel, 2013; Bachmann and Bayer, 2014), marginal products of capital (Oberfield, 2013; Sandleris and Wright, 2014), or prices (Vavra, 2014; Nakamura, Steinsson, Sun and Villar, 2018) to the steady-state *skewness* of those distributions. In turn, the CIR_3—which measures the dynamics of the third moment—relates to the steady-state fourth moment. This relationship could be used to connect the cyclical fluctuations in the *skewness* of sales' growth (Salgado, Guvenen and Bloom, 2019) to their steady-state *kurtosis*.

[12]Clearly, as the drift goes to infinity, the duration of inaction goes to zero as well, $\mathbb{E}[\tau] \to 0$. This mechanically makes the CIR_1 equal to zero, as firms are constantly adjusting. However, the result in Corollary 2 is stronger: It says that even after escaling the fixed cost such that $\mathbb{E}[\tau]$ is constant for every drift level, the CIR_1 is also zero in the limit.

[13]We thank Fernando Álvarez for suggesting to explore the connection to this limiting case.

3.3 Sufficient statistics as model discrimination devices

Our analysis shows that the CIR_1 with Bernoulli fixed costs is structurally linked to two steady-state cross-sectional moments of the capital-to-productivity distribution. This means that any configuration of this model that generates the empirical values of the sufficient statistics is relevant to the study of aggregate dynamics. Nevertheless, the theory imposes restrictions that can systematically rule in some model configurations and rule out others. To illustrate the power of sufficient statistics in distinguishing across model configurations, we use the two benchmark cases nested in our framework.

Fully state-dependent adjustments. The Bernoulli fixed-cost model nests the widely used state-dependent model of investment (Caballero and Engel, 1999) by shutting down free adjustments ($\lambda = 0$). This model does not impose any restriction on the sign of the covariance $\mathbb{C}ov[x, a]$. On the one hand, the drift reduces capital gaps as they get older, which pushes the covariance to be negative. On the other hand, the combination of idiosyncratic shocks with downward irreversibility—due to higher costs associated with downward adjustments than with upward adjustments—increases capital gaps as they get older, which pushes the covariance to be positive. Therefore, the size of the drift relative to the strength of downward irreversibility determines the value of the covariance of gaps and age.

However, this configuration restricts the variance of gaps $\mathbb{V}ar[x]$ to be small, because the distribution of adjustment size gets concentrated at the borders of the inaction region (we discuss this in detail in the following section). Therefore, if the empirical variance of capital gaps is large, as we infer from the Chilean plant-level data, this configuration will fall short of explaining the data.

Fully time-dependent adjustments. Another model nested in our framework consists of fully time-dependent adjustments, and it is obtained when both fixed costs go to infinity $\lim\{\theta^-, \theta^+\} \to \{\infty, \infty\}$. In the limit, the inaction region disappears and adjustments occur at a constant rate λ—i.e., they are exponentially distributed. In contrast to the fully state-dependent model, this model generates a large variance of gaps due to the free adjustments.

However, this configuration restricts the covariance between gaps and age to be negative $\mathbb{C}ov[x, a] < 0$. This is because firms cannot decide *when* to adjust, and thus the drift renders capital gaps negative as time goes by.[14] Therefore, if the empirical covariance is positive, as we infer from the Chilean data, this configuration will fall short of explaining the data.

Taken together, a positive covariance of gaps an age, $\mathbb{C}ov\,[x, a]$, is a tell-tale sign of state dependence, while a large variance of gaps, $\mathbb{V}ar[x]$, is a tell-tale sign of time dependence. We use these facts when applying the theory to the data in Section 5.

3.4 The CIR in the time-dependent limit

We conclude with an additional result for the time-dependent limit of the Bernoulli fixed-cost model. Corollary 3 shows that the sufficient statistic for the CIR_1 equals the *average age* of the capital gaps

[14]The covariance is $\mathbb{C}ov[x, a] = -\nu\overline{\mathbb{E}}[\tau]^2 < 0$. See Online Appendix B.1 for the proof.

$\mathbb{E}[a]$.[15] This result could be useful for researchers working with pure time-dependent models of inaction.

Corollary 3. *Consider the limit of the Bernoulli fixed-cost model with* $\lim(\theta^-, \theta^+) \to (\infty, \infty)$*. For any drift* $\nu \in \mathbb{R}$*, the* CIR_1 *is given by the average age of the capital gaps:*

$$\frac{CIR_1(\delta)}{\delta} \;=\; \mathbb{E}\left[a\right] \;+\; o(\delta). \tag{32}$$

In the time-dependent limit, the age of capital gaps informs about the speed at which firms adjust to the aggregate productivity shock. If firms adjusted continuously, average age would be zero and the economy would reach its steady-state immediately. In contrast, with more inaction, firms allow their capital gaps to get older, and the CIR_1 increases.

It is straightforward to show that in symmetric and driftless environments, average age equals the variance of gaps normalized by idiosyncratic volatility, i.e., $\mathbb{E}[a] = \mathbb{V}ar[x]/\sigma^2$. This means that the sufficient statistic in the time-dependent limit coincides with that in the Bernoulli fixed-cost model for finite λ in equation (30).[16] This result is in line with Álvarez, Lippi and Passadore (2017), who prove that in symmetric and driftless environments, propagation dynamics are very similar for state-dependent and time-dependent models as long as δ-perturbations are not very large. However, this equivalence of sufficient statistics breaks with nonzero drift and asymmetric policies, suggesting that propagation dynamics across alternative configurations of adjustment costs can be very different, even for small perturbations.

While Corollary 3 entertains a time-dependent model derived as a limiting case of the Bernoulli fixed-cost model, the fact that average age $\mathbb{E}[a]$ is a sufficient statistic for the CIR_1 is valid more generally for *any* fully time-dependent model—i.e., models in which adjustment dates τ_i are *iid* draws from an arbitrary distribution (not necessarily exponential) and independent of the capital gap $T_{h+1}|x_t = T_{h+1}$. Two well-known examples of this family, common in the pricing literature, are fixed adjustment dates, as in Taylor (1980), and sticky information, as in Reis (2006) and Álvarez, Lippi and Paciello (2016). Our proof covers this general class of models.

4 Steady-state moments and microdata

We have established a structural link between the CIR_m and the steady-state moments of capital gaps. The challenge ahead lies in computing these moments, as capital gaps are difficult to observe. Fortunately, the actions of adjusters—investment rates and adjustment dates—are readily available in the microdata. Let $\Omega \equiv (\Delta x, \tau)$ denote a panel of observations of capital gap changes (or adjustment size) and the duration of completed inaction spells, and $R(\Omega)$ denote their distribution. This section shows how to use the distribution of observable actions $R(\Omega)$ to infer the behavior of non-adjusters, and reverse engineer the steady-state cross-sectional moments of $F(x, a)$ and the parameters of the stochastic process.[17]

To establish the inverse mapping from data Ω to steady-state moments and parameters, we need two

[15]Online Appendix A.5 characterizes the CIR_m in the time-dependent limit for any m, with and without a drift.
[16]See Online Appendix B.2 for the proof.
[17]Clearly, the duration of uncompleted spells or *age* a can be measured in the data, but not the joint distribution of (x, a).

inputs: A parametric stochastic process for the uncontrolled gaps (in our case, a Brownian motion with drift ν and volatility σ), and a constant reset gap x^*. These inputs, together with the properties of Markov processes, are enough to pin down these mappings. We emphasize that we *do not need to assume a specific model of inaction* (e.g., the Bernoulli fixed-cost model) as long as it delivers a constant reset state. For this reason, in this part of the theory, we consider the reset state x^* a parameter.

Notation. To express our results succinctly, we use the following notation. We denote with bars the cross-sectional moments computed with the distribution of adjusters $R(\Omega)$ (e.g., $\overline{\mathbb{E}}[\cdot]$ and $\overline{\mathbb{C}ov}[\cdot,\cdot]$). We denote with tildes the variables that are expressed relative to their mean (e.g., $\tilde{\tau} \equiv \tau/\overline{\mathbb{E}}[\tau]$). Lastly, for any random variable $y \in \mathbb{R}$ and $\psi > 0$, we define the *generalized coefficient of variation* as $\overline{\mathbb{CV}}^{\psi}[\cdot] \equiv (\overline{\mathbb{E}}[y^{\psi}] - \overline{\mathbb{E}}[y]^{\psi})/\overline{\mathbb{E}}[y]^{\psi}$.[18] Note that for $\psi > 1$, $\overline{\mathbb{CV}}^{\psi}[\cdot]$ is increasing in y's dispersion due to Jensen's inequality for convex functions.

4.1 Recovering parameters from microdata

Proposition 2 provides mappings that allow an economist with observables Ω to make inferences about the parameters (ν, σ^2, x^*).

Proposition 2. *Let $\Omega \equiv (\Delta x, \tau)$ be a panel of observations of adjustment size and duration of inaction. Then, the following relationships hold:*

1. *The drift ν and volatility σ^2 of the stochastic process for capital gaps are recovered as*

$$\nu = \frac{\overline{\mathbb{E}}[\Delta x]}{\overline{\mathbb{E}}[\tau]}, \tag{33}$$

$$\sigma^2 = \frac{\overline{\mathbb{E}}[\Delta x^2]}{\overline{\mathbb{E}}[\tau]} - 2\nu x^*. \tag{34}$$

2. *The reset capital gap x^* is given by*

$$x^* = \nu(\overline{\mathbb{E}}[\tau] - \mathbb{E}[a]) + \overline{\mathbb{C}ov}\left[\tilde{\tau}, \Delta x\right]. \tag{35}$$

Drift and volatility. Expressions in (33) and (34) provide a mapping to infer the parameters of the stochastic process. The first expression shows that in a stationary environment, the average adjustment size $\overline{\mathbb{E}}[\Delta x]$ must compensate for the average drift between two adjustments $\nu\overline{\mathbb{E}}[\tau]$. Similarly, the dispersion in adjustment size $\overline{\mathbb{E}}[\Delta x^2]$ reflects the cumulative shocks received during the inaction period: A high dispersion in adjustment size must mean that either idiosyncratic volatility is high or the time between two adjustments is high. Álvarez, Le Bihan and Lippi (2016) obtain an analogous expression for idiosyncratic volatility σ^2 in the symmetric and driftless environment, i.e. $x^* = \nu = 0$, given by $\sigma^2 = \overline{\mathbb{E}}[\Delta x^2]/\overline{\mathbb{E}}[\tau]$. The capital gaps' negative drift increases $\overline{\mathbb{E}}[\Delta x^2]$ and decreases $\overline{\mathbb{E}}[\tau]$; thus the new term $2\nu x^*$ in our expression appears to correctly identify idiosyncratic volatility from the ratio of these two statistics.

[18] For $\psi = 2$, the standard definition of the coefficient of variation squared arises: $\overline{\mathbb{CV}}^2[\cdot] \equiv \overline{\mathbb{V}ar}[y]/\overline{\mathbb{E}}[y]^2$.

Reset capital gap. Equation (35) shows how to recover the reset gap x^* from the microdata. This object carries important information about optimal behavior in environments with nonzero drift and asymmetric policies.

The value of the reset gap is derived from the restriction that the distribution of gaps has mean zero—i.e., $\mathbb{E}[x] = 0$. To see this, let us write the capital gap of an inactive firm as $x_t = x^* + d_t$, where $d_t \equiv x_t - x^*$ equals the deviation from the reset point. Since gaps have a zero mean, the reset gap equals the negative of the average deviation of inactive firms: $0 = \mathbb{E}[x] = x^* + \mathbb{E}[d_t]$ or simply $x^* = -\mathbb{E}[d_t]$.[19] Notice that in the case with zero drift and a symmetric policy, the average deviation of inactive firms (and consequently the reset state) is always zero, because positive and negative deviations cancel each other out. In contrast, with nonzero drift or an asymmetric policy, the average deviation is likely to be different from zero, so the reset state must compensate accordingly to ensure that the mean of the stationary distribution remains at zero.

Importantly, a nonzero drift and an asymmetric policy may push average deviations in opposite directions. For instance, a negative drift pushes average deviations down, while an asymmetric policy arising from very costly downward adjustment pushes average deviations up (i.e, if $\theta^+ >> \theta^-$, units that have positive gaps tend to get stuck and do not want to adjust down).[20] The reset state summarizes how firms optimally balance these opposing forces. Both terms in expression (35) are relevant for identifying the roles of the drift and asymmetric policies, so we cannot attribute a precise mechanism to each term. However, we can focus on three special cases that showcase each mechanism separately.

Role of policy asymmetry. We start by discussing how the reset state reflects policy asymmetry. To isolate this effect, we assume away the drift and free adjustments, i.e., $\nu = \lambda = 0$. In this case, the reset gap is a weighted average of gap changes, with weights equal to relative durations of inaction, i.e., $x^* = \overline{\mathbb{E}}[\tilde{\tau}\Delta x]$. The analysis makes use of Figure II. Panel A plots three distributions of capital gaps (symmetric, left-skewed, and right-skewed) that correspond to alternative assumptions about fixed costs, which generate different policies. Panel B plots the corresponding distributions of capital gap changes. The idea is to use Panel B (observables) to infer Panel A (unobservables).

If the distribution of gaps is symmetric around zero (black solid line), positive and negative adjustments are equality likely ($\Pr(\Delta x > 0) = \Pr(\Delta x < 0) = 1/2$), implying a zero reset gap $x^* = \overline{\mathbb{E}}[\Delta x] = 0$. If the distribution of gaps is right-skewed (gray solid line), the upper border is further from x^* than the lower border; thus we observe very few negative adjustments ($\Pr(\Delta x > 0) = 4/5 > \Pr(\Delta x < 0) = 1/5$). In the weighted average, negative adjustments get overweighted by their longer relative duration, generating a negative reset gap $x^* = \overline{\mathbb{E}}[\tilde{\tau}\Delta x] = -u < 0$. Thus a negative reset gap suggests a right-skewed distribution generated by relatively more costly downward adjustments ($\theta^+ > \theta^-$). An analogous argument applies for positive reset gaps that suggest a left-skewed distribution. Overall, the sign of the reset gap is indicative of the shape of the gap distribution and the adjustment costs that generate it.

[19] By definition, deviations of adjusting firms are equal to zero. Note the difference between a deviation $d_t = x_t - x^*$ and a gap change $\Delta x_t = x^* - x_\tau$, where the latter is an actual change in the gap by a firm that has taken action at $t = \tau$.

[20] We emphasize this possibility because with the Chilean data, we recover a large negative drift and a reset gap that is very close to zero; taken together, these numbers suggests a large role for asymmetric adjustment costs with $\theta^+ >> \theta^-$.

Figure II – Reset State, Capital Gaps, and Capital Gap Changes for $\nu = \lambda = 0$

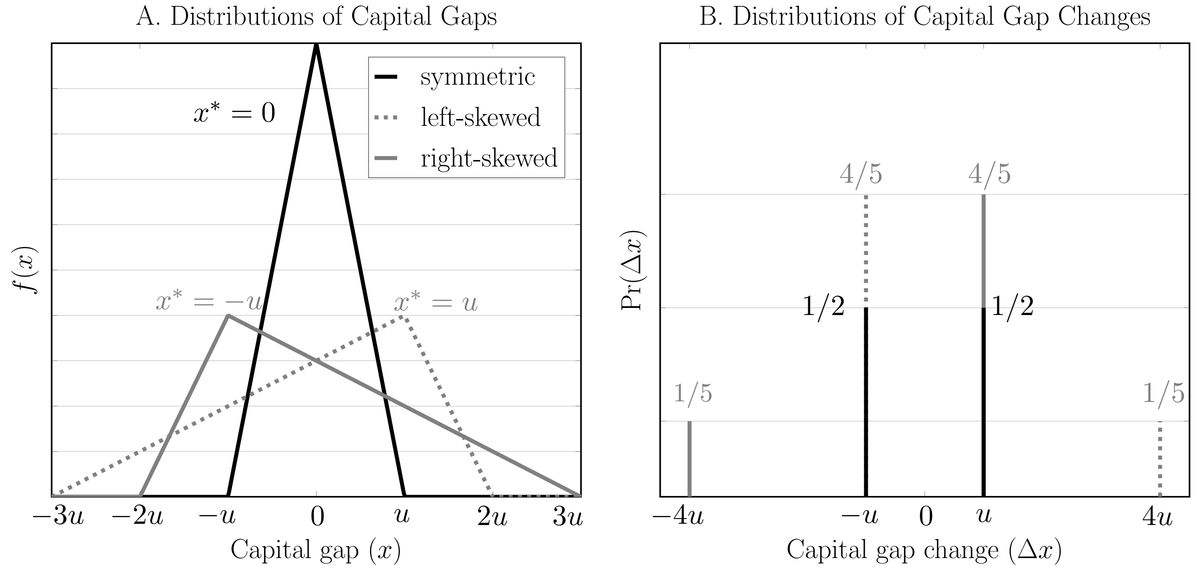

Notes: Panel A plots three distributions of capital gaps x in the Bernoulli fixed-cost model without drift and free adjustments. The symmetric distribution (black solid line) is generated by the policy $(x^-, x^*, x^+) = (-u, 0, u)$; the left-skewed distribution (gray dotted line) is generated by $(x^-, x^*, x^+) = (-3u, u, 2u)$; and the right-skewed distribution (gray solid line) is generated by $(x^-, x^*, x^+) = (-2u, -u, 3u)$. Panel B shows the corresponding three distributions of capital gap changes.

One may naively conjecture that the simple average gap change $\overline{\mathbb{E}}[\Delta x]$ provides information about asymmetries. However, this intuition is flawed: Larger adjustments happen with lower occurrence, so that $\overline{\mathbb{E}}[\Delta x] = 0$ regardless of the policy, eliminating any possibility to learn from the average gap change (see Panel B of Figure II). Our analysis shows that appropriately reweighting the distribution of gap changes using relative durations $\tilde{\tau}$ circumvents this identification challenge.[21]

Role of the drift. Now we discuss the role of the drift by considering two polar cases of the Bernoulli fixed-cost model with symmetric policies. In the limiting case without idiosyncratic shocks ($\sigma^2 \to 0$), the duration of all inaction spells is identical for all firms—say $\overline{\mathbb{E}}[\tau] = \mathcal{T}$—and evidently the covariance between duration and size in (35) disappears. The surviving term, $\nu(\overline{\mathbb{E}}[\tau] - \mathbb{E}[a])$, reflects how the reset gap compensates for the expected erosion caused by the drift that accumulates between adjustments. The expected erosion is proportional to the average length of completed spells ($\overline{\mathbb{E}}[\tau]$, adjusting firms) minus the average length of uncompleted spells ($\mathbb{E}[a]$, inactive firms). The reset gap becomes $x^* = \nu\mathcal{T} - \nu\mathbb{E}[a] = \nu\mathcal{T}/2 > 0$, where $\mathbb{E}[a] = \mathcal{T}/2$ is the average length of uncompleted spells.

Now consider the limiting case with infinite adjustment costs, $\lim(\theta^-, \theta^+) \to (\infty, \infty)$, so that investments occur at a constant rate $\lambda > 0$. Due to the *iid* nature of adjustment dates, the expected duration for adjusting and non adjusting firms is identical $\overline{\mathbb{E}}[\tau] = \mathbb{E}[a]$.[22] The first term in equation (35) disappears and the reset gap is identified by the second term: $x^* = \overline{\mathbb{C}ov}[\tilde{\tau}, \Delta x]$. It is easy to show that

[21] Online Appendix A.8 presents an instructive example, in which we compute explicitly the steady-state moments and other objects of interest in an environment with zero drift and a right-skewed distribution, as presented in Figure II.

[22] More generally, the relationship between $\mathbb{E}[\tau]$ and $\mathbb{E}[a]$ is determined by the dispersion and the tails of the duration distribution, as explained in the next subsection.

$\overline{\mathbb{C}ov}[\tilde{\tau}, \Delta x] = \nu \overline{\mathbb{E}}[\tau]$, corroborating that this covariance correctly identifies the role of the drift on the average capital gap.[23]

We have discussed at length the forces that shape the reset state for various reasons. It enters directly into the formulas for volatility (see equation 34) and steady-state moments (as explained in the next subsection). It helps us assess the shape of distribution of x. But most importantly, the economic forces that shape the reset point—and the way they manifest in the data—are also responsible for determining the value of $\mathbb{C}ov[x, a]$, one of the sufficient statistics for the CIR_1.

4.2 Recovering steady-state moments from microdata

Proposition 3 provides mappings from observables Ω to steady-state moments of capital gaps.

Proposition 3. *Let $\Omega \equiv (\Delta x, \tau)$ be a panel of observations of adjustment size and duration of inaction. Construct the gaps immediately before adjustment as $x_\tau = x^* - \Delta x$ using (35). Then the following relationships hold:*

1. *Average age relates to the average and the dispersion of duration of inaction as*

$$\mathbb{E}[a] \;=\; \frac{\overline{\mathbb{E}}[\tau]}{2}\left(1 + \overline{\mathbb{CV}}^2[\tau]\right). \tag{36}$$

2. *With zero drift ($\nu = 0$), the steady-state moments for any $m \geq 1$ are given by*

$$\mathbb{E}[x^m] \;=\; \frac{2}{(m+1)(m+2)}\left(\frac{\overline{\mathbb{E}}\left[x_\tau^{m+2}\right] - x^{*m+2}}{\overline{\mathbb{E}}\left[\Delta x^2\right]}\right), \tag{37}$$

$$\mathbb{E}[x^m a] \;=\; \frac{2\overline{\mathbb{E}}[\tau]}{(m+1)(m+2)}\left(\frac{\overline{\mathbb{E}}\left[\tilde{\tau} x_\tau^{m+2}\right] - \mathbb{E}[x^{m+2}]}{\overline{\mathbb{E}}[\Delta x^2]}\right). \tag{38}$$

3. *With nonzero drift ($\nu \neq 0$), the steady-state moments for any $m \geq 1$ are given by*

$$\mathbb{E}[x^m] \;=\; \frac{\Sigma^2}{(m+1)}\left(\frac{x^{*m+1} - \overline{\mathbb{E}}\left[x_\tau^{m+1}\right]}{\overline{\mathbb{E}}\left[\Delta x\right]}\right) + \frac{m\sigma^2}{2\nu}\mathbb{E}[x^{m-1}], \tag{39}$$

$$\mathbb{E}[x^m a] \;=\; \frac{\overline{\mathbb{E}}[\tau]}{(m+1)}\left(\frac{\mathbb{E}[x^{m+1}] - \Sigma^2\overline{\mathbb{E}}\left[\tilde{\tau} x_\tau^{m+1}\right]}{\overline{\mathbb{E}}[\Delta x]}\right) + \frac{m\sigma^2}{2\nu}\mathbb{E}[x^{m-1}a] \tag{40}$$

   with $\Sigma^2 \equiv \sigma^2/(\sigma^2 + 2\nu x^*)$.

Average age. Equation (36) relates the average age (the average length of uncompleted spells in the whole population) to the average and the dispersion in duration (the length of completed spells for adjusters), where the dispersion is measured using the coefficient of variation squared. The relationship between average age and average duration is straightforward: If adjusters take longer to adjust on average,

[23] See Online Appendix B.3 for the proof.

then the average capital gap in the cross-section will be older. Why does the dispersion in duration also increase age? The reason is what we call the *fundamental renewal property*: The probability that a random firm has an expected duration of inaction of τ is increasing in τ—i.e., many inaction spells are short, but the average spell is attributable to firms with long duration.[24] Dispersion in duration implies that some firms take a longer time to adjust, and those firms are more representative of the economy. This naturally raises the average age of capital gaps in the economy.

To analyze the relationships between steady-state moments and the microdata, the following Corollary presents simplified expressions for the case $m = 1$ and zero reset gap. These are derived by evaluating (37) to (40) at $m = 1$ and using (33) and (34).

Corollary 4. *Assume the reset point is zero, i.e., $x^* = 0$, so that $x_\tau = -\Delta x$. Then, we recover the steady-state moments $\mathbb{V}ar[x]$ and $\mathbb{C}ov[x, a]$ as follows:*

(i) For zero drift ($\nu = 0$),

$$\mathbb{V}ar[x] = \frac{1}{6} \underbrace{\overline{\mathbb{E}}[\Delta x^2]}_{avg.\ size} \Big(\underbrace{1 + \overline{\mathbb{CV}}^2 [\Delta x^2]}_{size\ dispersion} \Big), \tag{41}$$

$$\mathbb{C}ov[x, a] = \frac{\overline{\mathbb{E}}[\tau]}{3} \underbrace{\left(\frac{\overline{\mathbb{E}}\left[\tilde{\tau} x_\tau^3\right] - \mathbb{E}[x^3]}{\overline{\mathbb{E}}[\Delta x^2]} \right)}_{excess\ asymmetry\ of\ adjusters}, \tag{42}$$

(ii) For nonzero drift ($\nu \neq 0$),

$$\mathbb{V}ar[x] = \frac{1}{3} \underbrace{\overline{\mathbb{E}}[\Delta x]^2}_{avg.\ size} \Big(\underbrace{1 + \overline{\mathbb{CV}}^3 [\Delta x]}_{size\ dispersion} \Big), \tag{43}$$

$$\mathbb{C}ov[x, a] = \frac{\overline{\mathbb{E}}[\tau]}{2} \underbrace{\left(\frac{\mathbb{E}[x^2] - \overline{\mathbb{E}}\left[\tilde{\tau} x_\tau^2\right]}{\overline{\mathbb{E}}[\Delta x]} \right)}_{excess\ asymmetry\ of\ adjusters} + \frac{\sigma^2}{2\nu} \mathbb{E}[a]. \tag{44}$$

Variance of capital gaps. The drivers behind the cross-sectional variance of capital gaps $\mathbb{V}ar[x]$ are described in equations (41) and (43) for the cases without and with drift, respectively. The first term in these expressions relates to average adjustment size (measured via squared gap changes or gap changes squared, respectively) and the second term relates to dispersion of adjustment size (measured through generalized coefficients of variation).[25] Clearly, large average adjustments signal more dispersed gaps. But what is the connection between the dispersion in adjustment size and the dispersion of capital gaps? It is the *fundamental renewal property* again: The average behavior in the economy is attributable to firms with longer periods of inaction, which coincidentally are firms that make larger adjustments. Accordingly,

[24] This property has been widely studied in labor economics. For example, Mankiw (2014) states: "Many spells of unemployment are short, but most weeks of unemployment are attributable to long-term unemployment."

[25] With zero drift, the dispersion in adjustment size is measured by the generalized coefficient of variation of Δx^2 with $\psi = 2$ (the standard definition of the coefficient of variation); with nonzero drift, it is measured by the generalized coefficient of variation with $\psi = 3$ (this is closer to measuring skewness, as the presence of the drift alters the notion of dispersion).

higher dispersion in x_τ^2 (squared gaps of adjusters) increases $\mathbb{E}[x^2]$ (squared gaps of non-adjusters, which equal $\mathbb{V}ar[x]$).

Covariance of capital gaps with their age. The drivers behind the covariance between capital gaps and their age $\mathbb{C}ov[x, a]$ are described in (42) and (44). As with the reset gap, this covariance can be positive or negative, depending on the relative importance of the drift and policy asymmetry.

When the drift is equal to zero, the covariance in (42) is proportional to the excess asymmetry in the capital gaps of adjusters relative to non-adjusters—namely, the difference in the third moments of their respective distributions, i.e., $\overline{\mathbb{E}}\left[\tilde{\tau} x_\tau^3\right] - \mathbb{E}[x^3]$. A positive difference reflects a right-skewed distribution of gaps and vice versa. Note that the distribution of adjusters is weighted by the relative duration $\tilde{\tau}$, as with the reset state. The two additional terms in (42) are rescaling factors: The term $\overline{\mathbb{E}}[\Delta x^2]$ in the denominator ensures that the covariance is of order 1 (cancelling the cubic powers in the numerator) and the term $\overline{\mathbb{E}}[\tau]$ accounts for age's dependence on σ^2.

When the drift is different from zero, the covariance in (44) is obtained from different moments, but the economic interpretation is the same. In this case, the excess asymmetry of adjusters relative to non-adjusters is measured through the second moments of the respective distributions, i.e., $\mathbb{E}[x^2] - \overline{\mathbb{E}}\left[\tilde{\tau} x_\tau^2\right]$. Again, the distribution of adjusters is reweighted by relative duration $\tilde{\tau}$ and there is a rescaling factor, $\overline{\mathbb{E}}[\Delta x]$, to ensure that the covariance remains linear. Lastly, the term $\sigma^2 \mathbb{E}[a]/\nu$ compensates for the direct effect of idiosyncratic volatility on second moments (see equation 34).

4.3 The CIR in terms of microdata

Now we combine our two main theory results—the mapping from CIR_1 to steady-state moments in Corollary 1 and the mapping from steady-state moments to microdata in Corollary 4— in order to express concisely the propagation of an aggregate productivity shock in terms of microdata Ω. Corollary 5 presents the results.

Corollary 5. *The CIR_1 can be expressed as a function of microdata as follows:*

(i) With zero drift and symmetric policy, i.e., $\nu = x^* = 0$:

$$\frac{CIR_1(\delta)}{\delta} = \underbrace{\frac{\overline{\mathbb{E}}[\tau]}{2}}_{avg.\ duration} \underbrace{\frac{\overline{\mathbb{K}ur}[\Delta x]}{3}}_{size\ dispersion}. \tag{45}$$

(ii) With nonzero drift and asymmetric policy:

$$\frac{CIR_1(\delta)}{\delta} = \underbrace{\frac{\overline{\mathbb{E}}[\tau]}{2}}_{avg.\ duration} \left[\underbrace{\frac{\overline{\mathbb{CV}}^2[\tau] - 1}{2}}_{duration\ dispersion} + \underbrace{\frac{\overline{\mathbb{E}}\left[\widetilde{\Delta x}^3\right]}{\overline{\mathbb{E}}\left[\widetilde{\Delta x}^2\right]}}_{size\ dispersion} - \underbrace{\overline{\mathbb{C}ov}[\tilde{\tau}, \widetilde{\Delta x}^2]}_{covariance\ btw\ duration\ and\ size} \right]. \tag{46}$$

The previous expressions summarize many economic forces that shape the propagation of aggregate shocks in lumpy economies and show how these forces are reflected in the data. We use these expressions to organize a short literature review.

Zero drift and symmetric policy. Álvarez, Le Bihan and Lippi (2016) characterize the CIR_1 for zero drift and a symmetric policy ($\nu = x^* = 0$), and obtain their well-known "kurtosis" formula. Their result is presented in (45). In the price-setting context of their paper, x represents markup gaps and Δx represents price changes, and the formula expresses the response of real output to a one-time monetary policy shock as the product of the kurtosis of price changes and the expected time between adjustments.[26] The kurtosis of adjustment size—a measure of dispersion—has proven to be extremely useful sufficient statistics in evaluating the empirical relevance of various models of price adjustment, as long as the drift (inflation) is not too large.

Three amplification channels with nonzero drift and asymmetric policies. It is well known that average duration matters for propagation because it reflects the average speed at which agents adjust to the aggregate shock. However, according to expression (46), there are three additional channels that shape the persistence of aggregate dynamics: (i) dispersion in duration, (ii) dispersion in adjustment size, and (iii) the covariance between duration and size. The first two statistics have been analyzed in the context of the pricing literature, so we refer to this literature for a discussion. The third statistic is novel and is a contribution of our analysis. We will discuss each in turn.

Dispersion in the duration of inaction amplifies the CIR_1 because it reflects the coexistence of fast and slow adjusters, and this latter group is responsible for slowing the response to the shock (it is also more representative of the economy, due to the renewal property). This insight is formalized by Carvalho and Schwartzman (2015) and Álvarez, Lippi and Paciello (2016) in fully time-dependent models with zero drift. Our formula formally demonstrates that this insight extends beyond fully time-dependent models and also applies to environments with nonzero drift and asymmetric policies.

Dispersion in adjustment size amplifies the CIR_1 because it weakens the "selection effect"— namely, that adjusting firms are those with the largest need for adjustment. When the measure of marginal firms—firms whose gaps lie in the neighborhood of the adjustment thresholds—is small, most firms are dispersed away from their adjustment threshold, and hence the distribution of the adjustment size shows large dispersion. This insight is explored quantitatively in Midrigan (2011) and demonstrated theoretically by Álvarez and Lippi (2014) in multi-product menu-cost models. This channel is now well understood within the context of price-setting, but it is not the case for investment or other fields. We formally demonstrate that accounting for the dispersion of adjustment size is key to studying the aggregate dynamics in lumpy economies more generally, beyond menu-cost models of price-setting.

Lastly, the covariance between duration of inaction and adjustment size also shapes the CIR_1, as it reflects the presence of asymmetric policies. Identifying and quantifying this channel is one of the key payoffs from our theory. A naive approach is to identify an asymmetric policy through an asymmetric dis-

[26] In a model with monopolistic price-setters, the average markup gap is equal, up to a first order, to the aggregate real output. Therefore, the CIR_1 tracks the deviations of real output relative to steady state following a monetary shock δ.

tribution of adjustments. However, we have already shown that this is incorrect, because time-dependent models—which are inherently symmetric—also generate asymmetric adjustments in the presence of drift. Our analysis shows that the correct way to identify asymmetric policies from the drift effect is through the excess asymmetry of the distribution of adjusters relative to non-adjusters, as measured by the covariance between the relative duration of inaction and relative adjustment size squared, $\overline{\mathbb{C}ov}[\tilde{\tau}, \widetilde{\Delta x}^2]$. This statistic complements alternative methodologies that aim to diagnose whether frictions in capital allocation mainly affect upsizing firms or downsizing firms (Caballero and Engel, 2007; Lanteri, Medina and Tan, 2019).

5 Empirical Application

This section revisits the lumpy investment model described in Section 2 and applies the theoretical results obtained in Sections 3 and 4 using establishment-level data from Chile. First, we construct the distributions of capital gap changes Δx and duration of inaction τ from the data. Second, we use these empirical distributions as inputs to our formulas and obtain parameters, sufficient statistics, and the CIR_1 as outputs. Third, we use the sufficient statistics to discriminate across configurations of the Bernoulli cost model and settle on the best calibration to explain the data. Fourth, we connect the steady-state dispersion of capital gaps—one of the sufficient statistics—to the notion of capital misallocation and compare our estimate with the literature. Lastly, we discuss how to deal with firm heterogeneity.

5.1 Data description

Sources. 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 and structures, a capital category that represents 30% of all investment in the manufacturing sector and features the strongest lumpy behavior. 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. Online Appendix I describes the sample selection, variable construction, and analysis of each capital category separately: structures, vehicles, machinery, and equipment.

Capital stock and investment rates. We construct the capital stock series using the perpetual inventory method (PIM). Given an initial capital stock K_0, a plant's capital stock in year t is

$$K_t = (1-\zeta)K_{t-1} + I_t/D_t, \tag{47}$$

where ζ is the depreciation rate, D_t is the gross fixed capital formation deflator, and initial capital K_0 is a plant's self-reported nominal capital stock at current prices for the first year in which it is nonnegative. Gross nominal investment I_t is based on information on purchases, reforms, improvements, and sales of fixed assets. We define the investment rate ι_t as the ratio of real gross investment to the capital stock [27]

$$\iota_t \;\equiv\; \frac{I_t/D_t}{K_{t-1}}. \tag{48}$$

5.2 Construction of capital gap changes and duration of inaction

To apply the theory, for each plant and each inaction spell h, we record the capital gap change upon action Δx_h and the spell's duration τ_h. We construct capital gap changes with investment rates in (48):

$$\Delta x_h \;=\; \begin{cases} \log\left(1+\iota_h\right) & \text{if} \quad |\iota_h| > \underline{\iota}, \\ 0 & \text{if} \quad |\iota_h| < \underline{\iota}. \end{cases} \tag{49}$$

The parameter $\underline{\iota} > 0$ captures the idea that small maintenance investments do not incur the fixed cost. Following Cooper and Haltiwanger (2006), we set $\underline{\iota} = 0.01$, such that all investment rates below 1% in absolute value are excluded and considered to be part of the inaction spell. Given the capital gap changes, we define an adjustment date T_h from $\Delta x_{T_h} \neq 0$ and compute a spell's duration as the difference between two adjacent adjustment dates: $\tau_h = T_h - T_{h-1}$. Finally, we truncate the distribution at the 2nd and 98th percentiles of the investment distribution to eliminate outliers.[28]

Figure III – Empirical Distribution of Capital Gap Changes

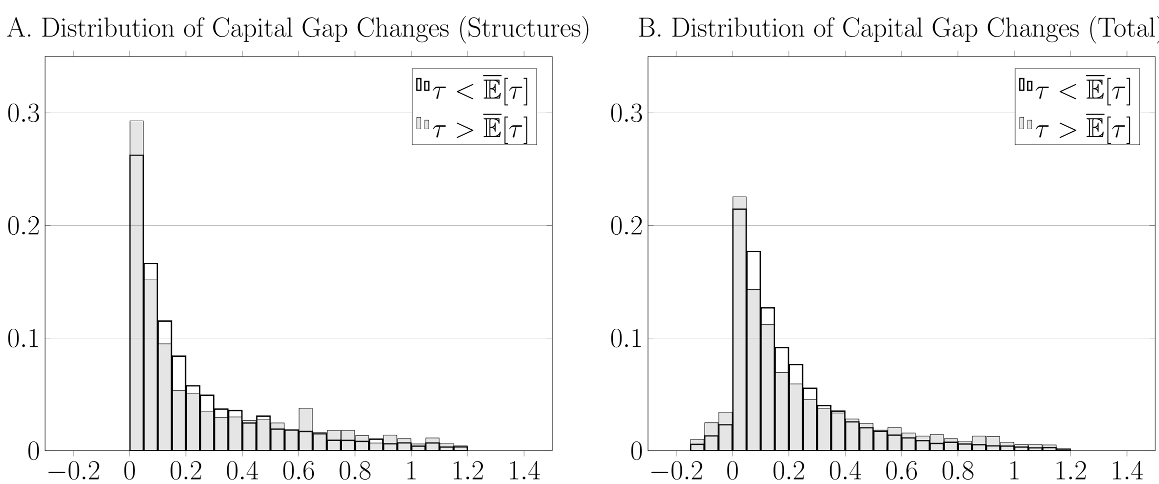

Notes: Own calculations using establishment-level data from Chile. Sample: Firms with at least 10 years of data, truncation at 2nd and 98th percentiles, and an inaction threshold of $\underline{\iota} = 0.01$. Panel A plots the distribution of capital gap changes Δx for structures, while Panel B plots the distribution for total capital. Solid bars = inaction spells with duration below average; white bars = inaction spells with duration above average.

[27] Note that the investment rate equals the ratio in the last term of equation (7): $\iota_{T_h} \equiv i_{T_h}/k_{T_h^-} = (k_{T_h} - k_{T_h^-})/k_{T_h^-}$, where $k_{T_h^-} = \lim_{t\uparrow T_h} k_t$. In contrast to the continuous-time model, in which investment is computed as the difference in the capital stock between two consecutive instants, in the data we compute it as the difference between two consecutive years.

[28] Table X in the Online Appendix L presents descriptive statistics on investment rates. For comparison, the table includes the numbers reported by Cooper and Haltiwanger (2006) for US manufacturing plants and by Zwick and Mahon (2017) for US firms from tax records.

Figure III plots the resulting cross-sectional distribution of capital gap changes for structures (Panel A) and total capital (Panel B). Both histograms show sizable asymmetry, inaction, and positive skewness. In each figure, we plot the distribution for two subsamples: observations with duration of inaction above the average duration (gray bars) and below the average duration (white bars). Notice that capital gap changes in both subsamples lie on top of each other, which is a sign of lack of covariance between adjustment size and duration. Below, we interpret this fact through the lens of the theory.

5.3 Putting the theory to work

We put the theory to work by computing the cross-sectional statistics of capital gap changes and duration to infer the parameters and sufficient statistics related to the CIR_1. We apply the formulas in Propositions 2 and 3. Table I summarizes the results. The left side of the table shows the inputs from the data: cross-sectional statistics of duration and capital gaps changes. The right side shows the outputs from the theory: parameters (ν, σ^2, x^*), sufficient statistics $(\mathbb{V}ar[x], \mathbb{C}ov[x, a])$, and the CIR_1.

Table I – Inputs from Micro Data and Outputs from the Theory

	Inputs from Data			**Outputs from Theory**	
	Structures	**Total**		**Structures**	**Total**
Duration			**Parameters**		
$\overline{\mathbb{E}}[\tau]$	2.510	1.749	ν	0.095	0.119
$\overline{\mathbb{CV}}^2[\tau]$	1.107	0.872	σ^2	0.049	0.049
			x^*	0.006	0.028
Capital Gaps Changes			**Sufficient Statistics**		
$\overline{\mathbb{E}}[\Delta x]$	0.239	0.207	$\mathbb{V}ar[x]$	0.124	0.092
$\overline{\mathbb{E}}[\Delta x^2]$	0.126	0.098	$\mathbb{C}ov[a, x]$	0.592	0.293
$\overline{\mathbb{E}}[x_\tau^3]$	-0.089	-0.057	$\mathbb{E}[a]$	2.644	1.637
$\overline{\mathbb{K}ur}[\Delta x]$	4.635	5.683			
Covariances			**CIR_1**		
$\overline{\mathbb{C}ov}[\tilde{\tau}, \Delta x]$	0.019	0.015	Drift + Asymmetric	3.661	2.562
$\overline{\mathbb{E}}[\tilde{\tau} x_\tau^2]$	0.141	0.103	Driftless + Symmetric	1.939	1.657

Notes: Own calculations using establishment-level data from Chile. Sample: Firms with at least 10 years of data, truncation at 2nd and 98th percentiles, and inaction threshold of $\underline{i} = 0.01$. Capital category: Structures and total. Gap of adjusters: $x_\tau = x^* - \Delta x$. Normalized duration: $\tilde{\tau} \equiv \tau / \overline{\mathbb{E}}[\tau]$.

Inputs from microdata. We focus our discussion on the values obtained for structures. The duration of inaction has an average of $\overline{\mathbb{E}}[\tau] = 2.51$ years and a coefficient of variation squared of $\overline{\mathbb{CV}}^2[\tau] = 1.11$, suggesting substantial heterogeneity in the adjustment frequency. Adjustment size has an average of $\overline{\mathbb{E}}[\Delta x] = 0.24$; a second moment of $\overline{\mathbb{E}}[\Delta x^2] = 0.13$; a third moment of $\overline{\mathbb{E}}[x_\tau^3] = -0.09$ (the distribution is right-skewed); and a kurtosis of $\overline{\mathbb{K}ur}[\Delta x] = 4.64$ (the distribution is leptokurtic). The covariance between adjustment size and relative duration is almost zero $\overline{\mathbb{C}ov}[\tau, \Delta x] \;=\; 0.02$, as suggested by Figure III. The duration-weighted second moment of gap changes is $\overline{\mathbb{E}}[\tilde{\tau} x_\tau^2] = 0.14$.

We compute average age in two ways: directly from the data and using the formula in (36), which connects it to duration of inaction $\mathbb{E}[a] = \overline{\mathbb{E}}[\tau](1 + \overline{\mathbb{CV}}^2[\tau])/2 = 2.64$ (for this reason, we show it in the second column with other outputs from the theory). We obtain very similar numbers using both methods, which is a good robustness check.[29] Now we input these statistics into the formulas we derived in the previous section.

Output from theory: Parameters. From (33), the inferred drift is 9.5%:

$$\nu \;=\; \frac{\overline{\mathbb{E}}[\Delta x]}{\overline{\mathbb{E}}[\tau]} \;=\; \frac{0.240}{2.519} \;=\; 0.095. \tag{50}$$

The drift reflects the depreciation rate ζ, productivity growth μ_z, and changes in relative prices between consumption and capital goods (ignored in the model). Since the inferred drift is different from zero, we apply (34) to estimate the volatility of idiosyncratic shocks:

$$\sigma^2 \;=\; \underbrace{\frac{\overline{\mathbb{E}}[\Delta x^2]}{\overline{\mathbb{E}}[\tau]}}_{0.05} \;-\; \underbrace{2\nu x^*}_{0.00} \;=\; 0.05. \tag{51}$$

The volatility estimate is in line with the value used by Khan and Thomas (2008). Lastly, using (35), we estimate the reset capital gap to be 0.5% above the average capital gap in steady state:

$$x^* \;=\; \underbrace{\nu(\overline{\mathbb{E}}[\tau] - \mathbb{E}[a])}_{-0.01} \;+\; \underbrace{\overline{\mathbb{C}ov}\,[\tilde{\tau}, \Delta x]}_{0.02} \;=\; 0.01. \tag{52}$$

Although small, the positive reset capital gaps suggests that the distribution of capital gaps is right-skewed and that the drift effect is overcome by policy asymmetry (recall the discussion in Section 4.1). In particular, we infer that downward adjustment is costlier than upward adjustment.

Output from theory: Sufficient statistics. Using (39), we infer a steady-state variance of capital gaps of

$$\mathbb{V}ar[x] \;=\; \underbrace{\frac{\Sigma^2}{3}}_{0.38} \left(\underbrace{\frac{x^{*3} - \overline{\mathbb{E}}[x_\tau^3]}{\overline{\mathbb{E}}[\Delta x]}}_{0.33} \right) \;=\; 0.12, \tag{53}$$

with $\Sigma^2 = \sigma^2/(\sigma^2 + 2\nu x^*) = 0.98$. The estimated variance of gaps is large and suggests an important role for free adjustments.[30]

[29] We thank Francesco Lippi for suggesting this robustness exercise.

[30] A purely state-dependent model without free adjustments ($\lambda = 0$) that matches the drift, idiosyncratic volatility, and average duration implies a variance of $\mathbb{V}ar[x] = 0.072$, 40% lower than what is observed in the data. See Table II.

Using (40), we infer a covariance between capital gaps and age of

$$\mathbb{C}ov[x,a] \;=\; \underbrace{\frac{\overline{\mathbb{E}}[\tau]}{2}}_{1.26} \left(\underbrace{\frac{\mathbb{E}[x^2] - \Sigma^2 \overline{\mathbb{E}}\left[\tau x_\tau^2\right]}{\overline{\mathbb{E}}[\Delta x]}}_{-0.08} \right) \;+\; \underbrace{\frac{\sigma^2}{2\nu}\mathbb{E}[a]}_{0.70} \;=\; 0.60. \tag{54}$$

The positive covariance between capital gaps and age means that plants that have not adjusted for a long time—their capital is old—have larger capital-to-productivity ratios than those that have recently adjusted. In other words, it is harder for firms to downsize in response to negative productivity shocks than to upsize in response to positive ones. There is strong downward irreversibility.[31]

To summarize, the large variance of gaps $\mathbb{V}ar[x]$ and the positive covariance with age $\mathbb{C}ov[x,a]$ strongly suggest that firms follow a hybrid investment policy with both time- and state-dependent components. Time-dependent adjustments increase the dispersion of gaps, while state-dependent asymmetric adjustments generate a positive covariance. Since the Bernoulli fixed-cost model nests these two alternatives, it serves as an adequate laboratory to study the relative importance of these components. In the following section, we search for a configuration of the Bernoulli model that best explains the data.

Output from theory: CIR_1. Assuming that the Bernoulli fixed-cost model is a good description of the data, the CIR_1 implied by the sufficient statistics is

$$\frac{\text{CIR}_1(\delta)}{\delta} \;=\; \underbrace{\frac{\mathbb{V}ar[x]}{\sigma^2}}_{2.53} \;+\; \underbrace{\frac{\nu\;\mathbb{C}ov[x,a]}{\sigma^2}}_{1.13} \;=\; 3.66. \tag{55}$$

This number says that a negative aggregate productivity shock of 1% generates a cumulative deviation of 3.66% in the average capital-to-productivity ratio (and in aggregate capital, up to first order) above its steady-state value along the transition path. As firms gradually scale down to accommodate the fall in aggregate productivity, capital is adjusted downward by selling it or letting it depreciate. In other words, aggregate productivity shocks have a long-run multiplier effect of approximately 3.6. We approximate the half-life of the response assuming exponential decay, obtaining 2.5 years ($\ln(2) \times \text{CIR}_1$).

Notice that naively applying the kurtosis formula in (45), which is invalid for the investment environment with significant negative drift and policy asymmetry, implies a speed of convergence of $\text{CIR}_1(\delta)/\delta = 1.94$. This underestimates the persistence of an aggregate productivity shock by 50%.

5.4 Parametrization of the Bernoulli fixed-cost model

With the estimated parameters of the stochastic process and the steady-state moments at hand, a natural question arises: Which configuration of the Bernoulli fixed-cost model generates the Chilean data? How important are the fixed costs relative to free adjustment opportunities?

To answer these questions, we explore the benchmark configurations nested within the Bernoulli fixed cost model to assess their ability to generate the data. These special cases clearly illustrate the

[31] A purely time-dependent model that matches the drift, idiosyncratic volatility, and average duration implies a covariance of $\mathbb{C}ov[x,a] = -0.602$, which has the same magnitude but the opposite sign of the one observed in the data. See Table II.

relationship between the structure of adjustment frictions and the sufficient statistics. In all exercises, we take as given the estimated parameters of the stochastic process and match the average duration of inaction spells. Notice that, given the parameters, matching average duration imposes additional constraints—e.g., the average adjustment size is also matched by equation (50). Table II summarizes the calibrated parameters.

Table II – Configurations of the Bernoulli Fixed-cost Model

Parameters	**Chilean Data**	**(I) Bernoulli** $\lambda = 0$	**(II) Bernoulli** $\{\theta^-, \theta^+\} \to \{\infty, \infty\}$	**(III) Extended Bernoulli** $\lambda^- \neq \lambda^+$
θ^- (for $i_t > 0$)	—	0.025	∞	3.203
θ^+ (for $i_t < 0$)	—	∞	∞	∞
λ^- (for $i_t > 0$)	—	0	0.397	0.800
λ^+ (for $i_t < 0$)	—	0	0.397	0
Moments				
$\overline{\mathbb{E}}[\tau]$	2.519	2.519*	2.519*	2.519*
$\mathbb{V}ar[x]$	0.124	0.072	0.182	0.112
$\mathbb{C}ov[x, a]$	0.600	0.654	−0.602	0.551
x^*	0.005	−0.141	0.239	−0.030
CIR_1	3.663	2.736	2.519	3.318

Notes: Data from Chilean plants. Configurations: (I) State-dependent Bernoulli with $\lambda = \lambda^+ = \lambda^- = 0$. (II) Time-dependent Bernoulli with $\lim\{\theta^-, \theta^+\} \to \{\infty, \infty\}$ and $\lambda = \lambda^+ = \lambda^-$. (III) Extended Bernoulli with $\lambda^+ \neq \lambda^-$. Parameters for the stochastic process: $\nu = 0.095$ and $\sigma^2 = 0.050$. * = targeted moment.

Column I considers a purely state-dependent model by shutting down the free adjustments, $\lambda = 0$. Anticipating that this configuration generates a tiny variance of gaps, and to give it the largest possibility of matching the positive covariance of capital gaps and age, we set an inaction threshold for negative investments of $\theta^+ = \infty$ (this is effectively a one-sided inaction region that imposes a strong downward irreversibility). To match average duration, the inaction threshold for positive investments is $\theta^- = 0.025$. The physical adjustment costs represent 0.02% of yearly revenues.[32]The implied sufficient statistics are $\mathbb{V}ar[x] = 0.072$ and $\mathbb{C}ov[x, a] = 0.654$. The CIR_1 equals 2.736, which is 25% below the data.

Column II considers the limiting case with infinite fixed costs, $\{\theta^-, \theta^+\} \to \{\infty, \infty\}$, which produces a purely time-dependent model. We calibrate the arrival rate of free adjustments $\lambda = 0.397$ to match average duration $\overline{\mathbb{E}}[\tau]$. As expected, this model produces a significant variance of gaps of $\mathbb{V}ar[x] = 0.182$, larger than in the data. However, it produces a negative covariance with age of $\mathbb{C}ov[x, a] = -0.602$, which we do not observe in the data. Surprisingly, this configuration implies a CIR_1 of 2.519, which is very similar to the one obtained in the state-dependent model of Column (1).

The previous analysis illustrates how two extreme calibrations can generate the same CIR_1 by matching one of the two sufficient statistics. The state-dependent model correctly captures the covariance of gaps with age—but misses the variance of gaps—whereas the time-dependent model does the opposite. Calibrations that lie between these two extremes only decrease the CIR_1. We conclude that the Bernoulli fixed-cost model falls short of generating the critical moments in the data.

[32]Assuming an output-capital elasticity of $\alpha = 0.3$, the average yearly payment of adjustment costs relative to revenue is $(\theta^- \Pr[x_\tau = x^-] + \theta^+ \Pr[x_\tau = x^+])/(\overline{\mathbb{E}}[\tau]\mathbb{E}_g[\exp(\alpha\hat{x})])$.

Extended Bernoulli fixed-cost model Can a simple modification of the Bernoulli model enable it to explain the data? The answer is yes. The extension considers different rates of free adjustments for positive and negative investments, λ^- and λ^+, respectively. Importantly, we verify numerically that the sufficient statistics for the CIR$_1$ remain valid under this extension.[33]

Column III shows the calibration. The extended model breaks the trade-off between asymmetry and randomness embedded in the original model, and it does an excellent job of matching the data. The best match is for fixed costs of $(\theta^-, \theta^+) = (3.203, \infty)$ and arrival rates of free adjustments of $(\lambda^-, \lambda^+) = (0.800, 0)$. The average physical adjustment costs represent 0.07% of yearly revenues. The implied values for the sufficient statistics are $\mathbb{V}ar[x] = 0.112$ and $\mathbb{C}ov[x, a] = 0.551$, and the CIR$_1$ equals 3.318, which is 90% of its empirical value.

As in the pure state-dependent case, we obtain a one-sided inaction region that matches the positive covariance of gaps and age. Additionally, the free adjustments introduce a random element in the policy that increases the variance of gaps, but it applies exclusively to upward adjustments. Finally, notice that the reset state x^* implied by the extended Bernoulli model is closer to the data than in the other two configurations.

5.5 On the steady-state dispersion of capital gaps

Our theory highlights the importance of matching the variance of capital gaps $\mathbb{V}ar[x]$ to correctly assess the propagation of aggregate productivity shocks, and estimates this sufficient statistic using data on investment rates exclusively. Here, we connect this sufficient statistic to the idea of *capital misallocation*, defined as the cross-firm dispersion in the log marginal revenue product of capital or MRPK (Restuccia and Rogerson, 2013), and compare our estimation strategy with the standard approach that requires additional data on value added or sales.[34]

First, we show that capital gaps are collinear to log MPRK deviations from steady-state, and therefore capital misallocation equals the dispersion of capital gaps. We start from the production technology in (1). Since all firms produce the same good, the output price is normalized to $p = 1$ for all firms. Then, the marginal revenue product of capital $\text{MRPK}_t \equiv \partial(py_t)/\partial k_t$ equals

$$\text{MRPK}_t \;=\; \alpha(z_t e_t)^{1-\alpha}\,(k_t)^{\alpha-1}\,. \tag{56}$$

Taking logs, using the definition of capital-to-productivity ratios $\hat{k} = \log(k/(ze))$, and subtracting the steady-state cross-sectional average $\mathbb{E}_g[\hat{k}]$, we obtain an expression for the log MRPK deviations from steady state in terms of capital gaps $x_t = \hat{k}_t - \mathbb{E}_g[\hat{k}]$:

$$\log \text{MRPK}_t - \mathbb{E}[\log \text{MRPK}] \;=\; (\alpha - 1)x_t. \tag{57}$$

[33] Online Appendix F explains how to compute the CIR$_1$ numerically for the extended Bernoulli fixed-cost model, and verifies that it is perfectly approximated by the sufficient statistics in expression (30).

[34] We refer to the dispersion of log MPRK as capital misallocation, but recognize that this dispersion could be efficient if it is generated by factors such as technological constraints to adjustment, as in our framework.

Squaring and taking the expectation with respect to the steady-state distribution of gaps $F(x)$, we obtain an expression for capital misallocation in terms of the variance of capital gaps:

$$\mathbb{V}ar[\log \mathrm{MRPK}] \;=\; (\alpha-1)^2 \mathbb{V}ar[x]. \tag{58}$$

Now, using our estimate from the Chilean data of $\mathbb{V}ar[x] = 0.124$ in (53), together with a standard value for the output-capital elasticity of $\alpha = 0.3$, we provide an estimate of economy-wide capital misallocation:

$$Std[\log \mathrm{MRPK}] \;=\; |\alpha - 1| \times \sqrt{\frac{\Sigma^2}{3} \frac{x^{*3} - \overline{\mathbb{E}}[x_\tau^3]}{\overline{\mathbb{E}}[\Delta x]}} \;=\; 0.7 \times 0.35 = 0.24. \tag{59}$$

To put this number in context, we conduct a measurement exercise in the spirit of Hsieh and Klenow (2009), using data on value added and capital to estimate misallocation, and obtain

$$Std[\log \mathrm{MRPK}] \;=\; Std\left[\log\left(\frac{\text{value added}}{\text{capital}}\right)\right] \;=\; 1.4. \tag{60}$$

Our estimate in (59) represents 18% of the estimate obtained with value added in (60). We attribute the difference to the fact that our calculation abstracts from labor misallocation and other "wedges" that generate additional dispersion, which may be due to markups, technology, information frictions, and measurement error. While incorporating these features into the model is outside the scope of this paper, our estimate is useful for gauging the relative importance of adjustment costs vis-à-vis other mechanisms, which lie between modest (David and Venkateswaran, 2019) and exceptional (Asker, Collard-Wexler and De Loecker, 2014).

The main advantage of our methodology is that it only requires investment data to estimate capital misallocation. To the extent that revenues are noisier than investment—e.g., due to measurement error or transitory shocks—our approach provides researchers with a suitable alternative. This comes at the cost of assuming that firms completely close their gaps with every investment, which is captured with a constant reset point x^* across firms and time. Heterogeneity in reset points across firms—due, for instance, to differences in technology or markups—can be accommodated easily in our theory, as we show in the following section. However, heterogeneity in reset points across time does pose an important challenge to our theory, precluding the analysis of convex adjusted costs, information and financial frictions, and any other sources of serial correlation in the reset point.

5.6 Heterogeneity and aggregation

One of the key assumptions of our theory is that one can exploit the cross-section to learn about the behavior of individual firms over time. In practice, fixed heterogeneity may affect the computation and interpretation of the cross-sectional statistics and the CIR. Here we extend the model in Section 2 to allow for fixed heterogeneity and discuss how to address it.[35]

[35] Online Appendix Section D presents the details of this extension and the proofs for the results.

Multisector economy. We divide the continuum of firms into N sectors denoted by $n \in \{1, 2, \ldots, N\}$, with sectoral masses $\{\gamma_n\}$ such that $\sum_{n=1}^{N} \gamma_n = 1$. Sectors differ in output-capital elasticity (α_n), adjustment frictions ($\theta_n^-, \theta_n^+, \lambda_n$), depreciation rate ζ_n, and idiosyncratic volatility σ_n. There is a common aggregate productivity component z_t with drift μ_z. We define the sectorial drift as $\nu_n \equiv \zeta_n + \mu_z$.

The crucial departure from the one-sector model is that now capital gaps refer to the capital-productivity ratio *relative to the sectorial average* instead of the economy's average. Let $\hat{x}_t \equiv \log\left(k_t/(z_t e_t)\right)$ be the log capital-to-productivity ratio of a firm. Investment policies consist of sectorial inaction regions $(\hat{x}_n^-, \hat{x}_n^+)$ and sectorial reset points x_n^*, which generate sectorial steady-state distributions $g_n(\hat{x})$. Let $x_t \equiv \hat{x}_t - \mathbb{E}_{g_n}[\hat{x}]$ denote the capital gap of a firm in sector n, where $\mathbb{E}_{g_n}[\hat{x}]$ is the steady-state cross-sectional average of $\hat{x}$ in sector n. Firms' policy and implied steady-state distributions in the multisector model satisfy the optimality conditions (13 to 16) and KFE (18) with parameters indexed by n. We denote with subscript n the moments computed with the sectorial steady-state distribution of gaps $F_n(x)$ (e.g., $\mathbb{V}ar_n[x]$ and $\mathbb{C}ov_n[x, a]$). Given the laws of motion for sectorial capital gaps, Lemma 3 characterizes the economy-wide CIR_1 after a common unanticipated decrease in aggregate productivity of size $\delta > 0$.

Lemma 3. *In a multisector economy with Bernoulli fixed costs, the CIR_1 can be expressed as a function of sectorial steady-state moments:*

$$\frac{CIR_1(\delta)}{\delta} = \mathbb{E}_{\tilde{\gamma}}\left[\underbrace{\frac{\mathbb{V}ar_n[x] + \nu_n \mathbb{C}ov_n[x,a]}{\sigma_n^2}}_{CIR_{1,n}(\delta)}\right] + o(\delta), \tag{61}$$

where $\mathbb{E}_{\tilde{\gamma}}[\cdot]$ is the average using capital-adjusted sectorial weights $\tilde{\gamma}_n \equiv \frac{\gamma_n \exp(\mathbb{E}_{g_n}[\hat{x}])}{\sum_{l=1}^{N} \gamma_l \exp(\mathbb{E}_{g_l}[\hat{x}])}$.

A few observations are in place. First, equation (61) shows that the economy-wide CIR_1 equals the weighted average of sectorial CIRs, labeled $\text{CIR}_{1,n}$. The capital-adjusted weights $\tilde{\gamma}_n$ depend on the mass of firms and the relative size of the average capital-to-productivity ratio in each sector. Second, the dynamics in the multisector economy only depend on the average of sectorial statistics. As a consequence, sectorial dispersion is irrelevant up to first order—i.e., one can always construct a "representative" sector that delivers the same aggregate CIR_1, as in Blanco and Cravino (2019).

To assess the importance of heterogeneity in shaping the CIR, we take a stand on the definition of a sector. We consider two alternative definitions: (i) 2-digit industries within manufacturing, and (ii) plant size, where size is proxied by the average number of workers in the sample period. Table III presents the sectoral sufficient statistics and the implied $\text{CIR}_{1,n}$ for structures. The last two columns show the weighted average of sectorial statistics following (61) and the pooled statistics computed in Section 5.3.

Table III shows that the estimated parameters and the sufficient statistics are quite stable across our two definitions of sectors. In particular, the variance of capital gaps $\mathbb{V}ar[x]$ and the covariance of capital gaps and age $\mathbb{C}ov[x, a]$ are positive and significant in each sector, which implies large sectorial CIRs in the range of 3.04 to 5.20. Moreover, when comparing the last two columns, the difference between sector-weighted and pooled statistics is tiny. These results suggest that sectorial heterogeneity in observable characteristics does not alter the conclusions regarding the propagation of aggregate productivity shocks.

Table III – Sufficient statistics for CIR_1 by 2-digit industry and firm size

2-Digit	**Food**	**Text.**	**Wood**	**Paper**	**Chem.**	**Min.**	**Metal**	**Mach.**	**Weighted**	**Pooled**
ν_n	-0.09	-0.09	-0.12	-0.12	-0.11	-0.09	-0.09	-0.10	-0.10	-0.10
σ_n^2	0.04	0.04	0.07	0.07	0.06	0.05	0.05	0.04	0.06	0.05
$\mathbb{V}ar_n[x]$	0.12	0.09	0.16	0.17	0.14	0.18	0.14	0.10	0.14	0.12
$\mathbb{C}ov_n[x, a]$	0.63	0.34	0.45	0.65	0.56	1.16	0.94	0.38	0.66	0.59
$Weight_n$	0.24	0.05	0.04	0.14	0.23	0.07	0.15	0.08	—	—
$\text{CIR}_{1,n}$	3.95	3.12	3.15	3.47	3.27	5.20	4.22	3.04	3.73	3.66

Firm size	**Small**	**Medium**	**Large**	**Extra Large**	**Weighted**	**Pooled**
ν_n	-0.10	-0.08	-0.10	-0.08	-0.09	-0.10
σ_n^2	0.05	0.04	0.04	0.04	0.04	0.05
$\mathbb{V}ar_n[x]$	0.13	0.11	0.10	0.10	0.11	0.12
$\mathbb{C}ov_n[x, a]$	0.60	0.62	0.43	0.80	0.61	0.59
$Weight_n$	0.25	0.25	0.26	0.24	—	—
$\text{CIR}_{1,n}$	3.62	4.04	3.38	4.36	3.83	3.66

Notes: Own calculations using establishment-level data for Chile. Sample: Firms with 10 years of data or more, truncation at 2nd and 98th percentiles, and threshold of $\underline{\iota} = 0.01$. Capital category: Structures. Sectors: (I) 2-digit industries within manufacturing = {Food, Textiles, Wood, Paper, Chemicals, Minerals, Metallurgy, Machinery}. (II) Firm size: Small = 1^{st} quartile, Medium = 2^{nd} quartile, Large = 3^{rd} quartile, Extra Large = 4^{th} quartile. Weighed = Average of sectorial statistics using capital-adjusted weights $\{\tilde{\gamma}_n\}$, approximated with average sectorial capital. Pooled = Economy-wide statistics.

The aggregation results in Lemma 3 and Table III assume that we can perfectly classify firms in their corresponding sector—i.e., the parameters and statistics are equal for all firms within a particular sector. Otherwise, if there exist important dimensions of unobserved heterogeneity, the sectorial steady-state statistics may be biased. Online Appendix D elaborates on this issue, identifies three potential sources of bias, and suggests ways to correct them.

Robustness checks. We conduct a series of robustness checks. Throughout these checks, we consistently obtain similar parameters and sufficient statistics. All results appear in the Online Data Appendix. First, we consider alternative samples of balanced and unbalanced panels (Table XVII). Second, we explore alternative methodologies to construct capital gap changes. We consider different depreciation rates, cut values that define inaction $\underline{\iota}$, and truncations that eliminate outliers (Tables XVIII, XIX and XX). Third, we assess the role of time trends by computing statistics in subperiods of 10 years (Table XXI) and the role of the business cycle conditioning on the cycle phase, where recession periods are identified following the OECD index of economic indicators (Table XXII). Lastly, we repeat the entire analysis for Colombia using the Annual Manufacturers Survey (*Encuesta Anual Manufacturera*) for the period 1995–2016.

6 Conclusion

We propose a sufficient statistics approach as a novel way to take models of lumpy adjustment to the data. Applying this approach to lumpy investment, we discover that the transitional dynamics of capital are structurally linked to the degree of capital misallocation and irreversibility in steady state. Our results indicate that policies that impact the lumpiness of investment—e.g., investment tax credits (Chen, Jiang, Liu, Suárez Serrato and Xu, 2019)—have a direct effect on the speed at which the economy propagates aggregate shocks.

We conclude by briefly discussing extensions to the theory and promising avenues for future research. A first extension considers the transitions of arbitrary smooth functions of the state $\eta(x)$. Using a Taylor approximation around zero, we write the CIR of $\eta(x)$ in terms of a sequence of CIRs, weighted by the Taylor factors.[36] This extension can be applied, for example, to compute the exact CIR of aggregate capital in (23). A second extension consists of incorporating mean reversion in the stochastic process of capital gaps, which is important in several empirical applications. In this case, the uncontrolled state follows an Ornstein–Uhlenbeck process and we recover the parameters (which now include a mean-reversion rate) through a non linear system of equations.[37] These equations could, in principle, be useful for testing for the persistence of productivity shocks directly from investment data.

Looking forward, we foresee four interesting avenues for future developments that would extend the scope of our theory. We focus on a one-dimensional state. Extending the theory to multidimensional states would allow for interactions between heterogeneity and lumpiness in multiplant firms (Kehrig and Vincent, 2019), lumpiness in various inputs (Hawkins, Michaels and Oh, 2015), or the interaction of lumpy behavior across different choices, such as investment and price-setting (Sveen and Weinke, 2007).

Second, we assume full adjustment upon action—i.e., agents completely close their gaps with every adjustment. Extending the theory to accommodate partial adjustments would allow for interactions of lumpiness with convex adjustment costs or imperfect information. This line of work would continue the recent contributions by Baley and Blanco (2019), who provide bounds for the CIR in environments with learning. The innovation lies in carrying an additional state (in the learning case, the aggregate forecast error) to characterize aggregate dynamics.

Third, we characterize the CIR, but not the complete profile of the impulse-response function; moreover, we only consider small perturbations around steady state. Extending the theory to characterize the full IRF and more general perturbations would allow us to study non linearities and different types of aggregate shocks. Contemporaneous work by Álvarez and Lippi (2019) makes progress in this direction by characterizing the impulse response function using eigenvalue-eigenfunction decompositions.

Finally, we assume that steady-state policies hold along the transition path. Thus, the tools we develop are helpful for gauging the strength of partial equilibrium responses, which in turn determine the magnitude of the general equilibrium effect (Auclert and Rognlie, 2018). Incorporating feedback from the distribution to individual policies, e.g., strategic complementaries, is likely the most important extension ahead, but also the most challenging.

[36]See Online Appendix E.3.

[37]See Online Appendix E.4.

References

Abbring, J. H. (2012). Mixed hitting-time models. *Econometrica*, **80** (2), 783–819.

Abel, A. B. and Eberly, J. C. (1994). A unified model of investment under uncertainty. *American Economic Review*, **84** (5), 1369.

— and — (1996). Optimal investment with costly reversibility. *The Review of Economic Studies*, **63** (4), 581–593.

Álvarez, F., Beraja, M., Gonzalez-Rozada, M. and Neumeyer, P. A. (2018). From hyperinflation to stable prices: Argentina's evidence on menu cost models. *The Quarterly Journal of Economics*, **134** (1), 451–505.

—, Borovičková, K. and Shimer, R. (2016). *Decomposing duration dependence in a stopping time model*. Tech. rep., National Bureau of Economic Research.

—, Le Bihan, H. and Lippi, F. (2016). The real effects of monetary shocks in sticky price models: a sufficient statistic approach. *American Economic Review*, **106** (10), 2817–51.

— and Lippi, F. (2014). Price setting with menu cost for multiproduct firms. *Econometrica*, **82** (1), 89–135.

— and Lippi, F. (2019). The analytic theory of a monetary shock. *EIEF Working Paper*.

—, Lippi, F. and Paciello, L. (2011). Optimal price setting with observation and menu costs. *Quarterly Journal of Economics*.

—, — and Paciello, L. (2016). Monetary shocks in models with inattentive producers. *The Review of Economic Studies*, **83** (2), 421–459.

—, — and Passadore, J. (2017). Are state-and time-dependent models really different? *NBER Macroeconomics Annual*, **31** (1), 379–457.

Asker, J., Collard-Wexler, A. and De Loecker, J. (2014). Dynamic inputs and resource (mis) allocation. *Journal of Political Economy*, **122** (5), 1013–1063.

Auclert, A. and Rognlie, M. (2018). *Inequality and aggregate demand*. Tech. rep., National Bureau of Economic Research.

Bachmann, R. and Bayer, C. (2014). Investment dispersion and the business cycle. *American Economic Review*, **104** (4), 1392–1416.

—, Caballero, R. J. and Engel, E. M. (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. (2019). Firm uncertainty cycles and the propagation of nominal shocks. *AEJ: Macro*, **11** (1), 276–337.

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

Blanco, A. (2020). Optimal inflation target in an economy with menu costs and an occasionally binding zero lower bound. *American Economic Journal: Macroeconomics*.

— and Cravino, J. (2019). Price rigidities and the relative ppp. *Journal of Monetary Economics*.

Borovička, J., Hansen, L. P. and Scheinkman, J. A. (2014). Shock elasticities and impulse responses. *Mathematics and Financial Economics*, **8** (4), 333–354.

Caballero, R. J. and Engel, E. (1991). Dynamic (s,s) economies. *Econometrica*, **59** (6), 1659–1686.

— and — (1999). Explaining investment dynamics in u.s. manufacturing: A generalized (s, s) approach. *Econometrica*, **67** (4), 783–826.

—, — and Haltiwanger, J. (1995). Plant-level adjustment and aggregate investment dynamics. *Brookings papers on economic activity*, pp. 1–54.

—, — and — (1997). Aggregate employment dynamics: Building from microeconomic evidence. *The American Economic Review*, **87** (1), 115–137.

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

— and — (2007). Price stickiness in ss models: New interpretations of old results. *Journal of monetary economics*, **54**, 100–121.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics*, **12** (3), 383–398.

Caplin, A. and Leahy, J. (1991). State-Dependent Pricing and the Dynamics of Money and Output. *The Quarterly Journal of Economics*, **106** (3), 683–708.

— and — (1997). Aggregation and optimization with state-dependent pricing. *Econometrica: Journal of the Econometric Society*, pp. 601–625.

Caplin, A. S. and Spulber, D. F. (1987). Menu costs and the neutrality of money. *The Quarterly Journal of Economics*, **113**, 287–303.

Carvalho, C. and Schwartzman, F. (2015). Selection and monetary non-neutrality in time-dependent pricing models. *Journal of Monetary Economics*, **76**, 141–156.

Chen, Z., Jiang, X., Liu, Z., Suárez Serrato, J. C. and Xu, D. (2019). *Tax Policy and Lumpy Investment Behavior: Evidence from China's VAT Reform*. Tech. rep., National Bureau of Economic Research.

Cooper, R. and Willis, J. L. (2004). A comment on the economics of labor adjustment: Mind the gap. *The American Economic Review*, **94** (4), 1223–1237.

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

David, J. M. and Venkateswaran, V. (2019). The sources of capital misallocation. *American Economic Review*, **109** (7), 2531–67.

Dixit, A. K. and Pindyck, R. (1994). *Investment under uncertainty*. Princeton university press.

Gabaix, X., Lasry, J.-M., Lions, P.-L. and Moll, B. (2016). The dynamics of inequality. *Econometrica*, **84** (6), 2071–2111.

Golosov, M. and Lucas, R. E. (2007). Menu costs and phillips curves. *Journal of Political Economy*, pp. 171–199.

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

Hawkins, W., Michaels, R. and Oh, J. (2015). *The joint dynamics oF capital and employment at the plant level*. Tech. rep., Mimeo, Yeshiva University.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, **124** (4), 1403–1448.

Kehrig, M. and Vincent, N. (2019). Good dispersion, bad dispersion. *Duke Working Paper*.

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.

Kurmann, A. and Petrosky-Nadeau, N. (2007). Search frictions in physical capital markets as a propagation mechanism. *CIRPEE Working Paper 07-12*.

Lancaster, T. (1990). *The econometric analysis of transition data*. 17, Cambridge university press.

Lanteri, A. (2018). The market for used capital: Endogenous irreversibility and reallocation over the business cycle. *American Economic Review*, **108** (9), 2383–2419.

—, Medina, P. and Tan, E. (2019). Capital-reallocation frictions and trade shocks. *Economic Research Initiatives at Duke (ERID) Working Paper*, (293).

Mankiw, N. G. (2014). *Principles of macroeconomics*. Cengage Learning.

Midrigan, V. (2011). Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, **79** (4), 1139–1180.

Moll, B. (2014). Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review*, **104** (10), 3186–3221.

Nakamura, E. and Steinsson, J. (2010). Monetary non-neutrality in a multisector menu cost model. *The Quarterly Journal of Economics*, **125** (3), 961–1013.

—, —, Sun, P. and Villar, D. (2018). The elusive costs of inflation: Price dispersion during the us great inflation. *The Quarterly Journal of Economics*, **133** (4), 1933–1980.

Oberfield, E. (2013). Productivity and misallocation during a crisis: Evidence from the chilean crisis of 1982. *Review of Economic Dynamics*, **16** (1), 100–119.

Ottonello, P. (2018). Capital unemployment, financial shocks, and investment slumps. *Working Paper*.

Reis, R. (2006). Inattentive Producers. *The Review of Economic Studies*, **73** (3), 793–821.

Restuccia, D. and Rogerson, R. (2013). Misallocation and productivity. *Review of Economic Dynamics*, **1** (16), 1–10.

Salgado, S., Guvenen, F. and Bloom, N. (2019). *Skewed business cycles*. Tech. rep., National Bureau of Economic Research.

Sandleris, G. and Wright, M. L. (2014). The costs of financial crises: Resource misallocation, productivity, and welfare in the 2001 argentine crisis. *The Scandinavian Journal of Economics*, **116** (1), 87–127.

Stokey, N. (2009). *The Economics of Inaction*. Princeton University Press.

Sveen, T. and Weinke, L. (2007). Lumpy investment, sticky prices, and the monetary transmission mechanism. *Journal of Monetary Economics*, **54**, 23–36.

Taylor, J. B. (1980). Aggregate dynamics and staggered contracts. *The Journal of Political Economy*, pp. 1–23.

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

Vavra, J. (2014). Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics*, **129** (1), 215–258.

Veracierto, M. L. (2002). Plant-level irreversible investment and equilibrium business cycles. *American Economic Review*, **92** (1), 181–197.

Verona, F. (2014). Investment dynamics with information costs. *Journal of Money, Credit and Banking*, **46** (8), 1627–1656.

Winberry, T. (2019). Lumpy investment, business cycle, and stimulus policy. *American Economic Review*, **Forthcoming**.

Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics*, **56**, S100–S124.

Zwick, E. and Mahon, J. (2017). Tax policy and heterogeneous investment behavior. *American Economic Review*, **107** (1), 217–48.