Dynamic Investment and Product Market Rivalry: The Network Q Model

Maria Cecilia Bustamante[*]
University of Maryland

Bruno Pellegrino[†]
Columbia University, NBER

February 2026

Abstract

We present a new dynamic model of corporate investment in imperfectly-competitive product markets, extending the neoclassical (Q) theory of capital to a multi-firm, multi-product, fully-structural model. Our model embeds a state-of-the-art hedonic demand system, endogenizes firms' markups and generalizes Tobin's Q to a matrix (or network) of product market spillovers, which captures how each firm's investment affects that of its rivals. We provide existence and uniqueness results along with exact, global analytical solutions for the Markov Perfect Equilibrium investment policies. We then take our model to the data for the universe of U.S. public companies and obtain five novel insights: 1) product market competition is a key force driving aggregate investment and capital allocation; 2) the persistence of firm's capital stocks increased over the past 25 years (i.e. capital became "stickier"); 3) monopoly rents account for a large, rising share of firms' value; 4) positive shocks to firms' cost of capital increase markups and concentration; 5) mergers consummated since 1995 have led to a modest decline in aggregate capital formation; at the firm-level the resulting increases in markups are highly heterogeneous.

[*]email: mcbustam@umd.edu - address: Van Munching Hall (4423), 7699 Mowatt Ln, College Park, MD 20742

[†]email: bp2713@columbia.edu - address: Kravis Hall (768) 665 W 130th St, New York, NY 10027

We thank Nico Crouzet, Joel Flynn, Loukas Karabarbounis, Pete Kyle, Ernest Liu, Pascal Noel, Olivier Wang, and workshop participants at Columbia New Empirical Finance Conference, UMD Brown Bag and the New York Macro Theory Group for helpful feedback. We thank Lorenzo Calvetti for excellent research assistance.

1 Introduction

How do firms make forward-looking strategic investment decisions? This is a central question of economics, with important implications for a wide range of real-world problems (e.g. forecasting, optimal taxation, antitrust). It has been extensively studied in various sub-fields of economics, including corporate finance, industrial organization and macroeconomics.

When we consider the behavior of a single firm in isolation, the neoclassical (Q) model developed sequentially by Jorgenson (1963), Tobin (1969), Abel (1980) and Hayashi (1982) stands out within this wide literature, both in terms of influence and wide range of applications. In practice, however, firms do not operate in isolation as in these canonical neoclassical models. Firms operate in a competitive environment, and thus their investment decisions are interrelated. To date, the equilibrium effect of product-market linkages on investment remain hard to analyze at scale, both theoretically and empirically. On one end, due to strategic complexity, theoretical papers typically focus on dynamic investment models of duopoly or symmetric oligopoly, which lack broad empirical relevance.[1] On the other end, the applied empirical literature faces the difficulty of properly identifying investment spillovers in reduced form, and structurally estimating complex models of strategic behavior.[2]

In this paper, we provide a highly tractable and scalable Q model of investment in which firms that produce differentiated products compete in oligopolistic product markets. Leveraging the Generalized Hedonic-Linear demand system (GHL) of Pellegrino (2025), we model the entirety of the product market interactions across all firms in the model as a network: in this (weighted) network firms are nodes, and the pairwise degree of substitution between the goods and services they provide are the edges. We call this the Network-Q model.

Our model is thus simultaneously a neoclassical model *and* a dynamic game of differentiated goods oligopoly. Our equilibrium concept is Markov Perfect Equilibrium (MPE), appropriately refined by the standard restrictions that best response functions are linear and that the system is dynamically stable. These restrictions imply that along the MPE path, the firms' capital stocks follow a Vector AutoRegressive (VAR) process. In our theoretical exposition, we prove existence and uniqueness results for this equilibrium. Furthermore, we are able to obtain a global, analytical characterization of the MPE equilibrium path. Specifically, we obtain closed-form expressions for firms' investment rates, enterprise values and marginal Q's in terms of the underlying primitives (discount rates, adjustment cost parameters, etc...). This is unusual for both neoclassical investment models as well as dynamic oligopoly games, and underscores the tractability of our framework.

In line with previous neoclassical studies considering imperfect competition (e.g Hayashi, 1982 and Crouzet and Eberly, 2023), we show that the wedge between each firm's average Q and marginal Q is associated with monopoly rents. However, our model advances this literature in one important respect: markups in our model are endogenous, and co-determined with investment in equilibrium. A firm's investment is thus affected by its market power; vice-versa a firm's investment affects its market power and that of its peers. Because capital expenses vary with firms' discount rates, we further show that firm-level and aggregate fluctuations in discount rates affect the dynamics of markups and industry concentration.

[1] See the seminal papers Maskin and Tirole (1987), Maskin and Tirole (1988a) and Maskin and Tirole (1988b), as well as the neoclassical symmetric duopoly model with adjustment costs by Reynolds (1991).

[2] See, for instance, Bloom, Schankerman, and Van Reenen (2013) and Bustamante and Fresard (2021) for studies on product market peer effects and spillovers.

Our analytical characterization yields a fresh insight into how firms make investment decisions: namely, product market competition plays a key role in determining the rate of capital accumulation at the firm level. This is not only true in the steady state, but also along the transition dynamics. This is because, in determining the speed of capital adjustments, firms trade off capital adjustment costs against the profit loss that would be incurred by *not* adjusting capacity in response to rivals' investment. Firms that face mild competitive pressures (as captured by the concavity of their profit function), relative to adjustment costs (captured by the convexity of the capital installation cost function) will converge more slowly to their steady state level of capital. Vice-versa, firms that face intense product market competition will converge faster. As a result, in the aggregate economy, the sparsity of the product market network and the initial allocation of capital among firms jointly determine how fast the economy grows its capital stock over time.

The intensity of competition faced by each individual firm in turn depends on how similar or dissimilar their product is to those produced by its rivals. Through the lenses of the generalized hedonic-linear demand framework, this in turn depends on the attributes of all of the products that the consumer can purchase.

A second interesting theoretical consequence of modeling the investment of a multiplicity of firms that compete with each other is that Tobin's Q becomes a matrix (or, equivalently, a network): specifically, it is the Jacobian of the firms' values with respect to the vector of capital stocks. The diagonal entries of this matrix correspond to marginal Q, as defined by Hayashi (1982). Meanwhile, the off-diagonal terms capture the spillover of firms' investment their rival firms' values. We are also able to derive a closed form expression for an alternative metric of investment spillovers, namely the effect of a unit increase in a firm's capital on a rival's investment (Bloom, Schankerman, and Van Reenen, 2013). For both these spillovers metric, we show how they relate to the degree of competitive pressure that firms are subject to in equilibrium.

After completing the theoretical exposition of our model, we take it to the data for the universe of U.S. public companies from 1995 to 2021. Our empirical implementation departs in a very significant way from the previous literature in empirical and structural corporate finance, where empirical applications of the Q model are dominated by reduced-form predictive regressions and structural estimation of key model parameters. By contrast, our model is mapped to the data firm-by-firm (each firm in the model corresponds to an actual real firm in the data) and is implemented fully structurally (we recover all model primitives). Following Pellegrino (2025), we use the text-based product similarities dataset created by Hoberg and Phillips (2016), which covers the universe of publicly-traded firms, to parameterize our massive hedonic demand system (for an average year, it contains several millions of elasticities). In this respect our empirical approach is much closer to empirical dynamic oligopoly models from the industrial organization literature. This novel empirical framework opens up a much wider range of applications for the Q model, such as predictive simulations and counterfactual analyses.

Because our model is significantly over-identified, we begin our quantitative analysis by assessing the model's goodness of fit. In particular, we show that the model performs well in matching the cross-section of variables that can't be perfectly fitted - namely: 1) firm revenues; 2) investment rates; 3) firm valuations. We then precede to utilize our model to deliver five key novel insights.

First, we show that the transition dynamics of capital accumulation have significantly slowed down during the past decades: in other words, capital became stickier, primarily due to technological factors (production became more capital dependent, increasing the relative importance of adjustment costs). We demonstrate this by performing a spectral decomposition of the transition matrix that governs the capital accumulation process. Its eigenvalues (which by construction are comprised in the unit interval) have systematically shifted towards 1 between 1995 and 2021, indicating a higher persistence.

As our second quantitative exercise, we project all firms' steady-state valuations, and (through the lens of our model) break them down into two fundamental additive components. The first captures the replacement value of capital. The second is the present value of all future monopoly rents. We find that aggregate firm value has risen between 1995 and 2021, and monopoly rents have contributed significantly the rise of firm valuations over time. The model decomposition also suggests that monopoly rents have risen significantly as a proportion of aggregate firm value: monopoly rents make up to 46.35% of aggregate firm value in 1995, and increase to 52.28% of aggregate firm value in 2021. Our results confirm the importance of market power in determining firms' valuations, as noted by Covarrubias, Gutiérrez, and Philippon (2020) and Crouzet and Eberly (2023), among others.

We then consider multiple counterfactuals, which highlight the ability of our model to address relevant questions for both researchers and policymakers. By means of a counterfactual on firms' competitive conduct, we confirm product market competition is a key force driving aggregate investment. Our baseline estimation focuses on the canonical Cournot case, in which firms internalize the downward sloping residual demand, but do not internalize their own product market spillovers on peer firms. We compare this case to a cartel scenario in which all product market spillovers are internalized. Moving from Cournot competition to full cartelization reduces aggregate steady-state capital by 30.31%. The result confirms that investment spillovers significantly contribute to aggregate investment.

Our next quantitative exercise is motivated by the fact that, in our model, firms pass their costs of capital onto their customers through their markups. In this exercise, we show that shocks to discount rate affect markups and concentration, and alter the inter-temporal allocation of capital. We show this by simulating a 1pp unanticipated, persistent increase in the discount rate: we find that the model responds by moving to an alternative MPE path characterized by higher industry concentration and markups. While our model does not feature monetary policy in the strict sense, this finding suggests that persistent shocks to the interest rate may have the potential to alter industry structure and product market competition among firms; directionally, this effect would dampen the effectiveness of monetary interventions with respect to a Neo-Keynesian benchmark.

Because our model also allows us to simulate mergers, in our last quantitative exercise we assess the impact of takeover activity on aggregate investment between 1995 and 2021. Compared to the counterfactual case of no takeover activity after 1995, we infer that the documented takeover activity has led to a modest reduction of 1.35% of all corporate investment on aggregate. Simultaneously, however, we observe significant variation in the impact of mergers in the cross section: while the median merging firm barely alters its price, mergers in the 99th percentile range of the distribution generate an increase in markups of approximately 0.73% within five years' time of the merger.

The paper concludes with a discussion of extensions and alternative applications of our network Q model. The model can be extended to incorporate more complex forms of strategic interaction beyond simple product market competition, including input-output relations and common ownership. While we focus on the Cournot outcome in which firms compete in capacity, the model can be reinterpreted to allow for competition in prices and sticky adjustment menu costs. While our baseline model focuses on a single adjustment cost function for all firms, we show in an extension how to incorporate heterogeneous adjustment costs and discount rates. Lastly, although our core empirical application focuses on firm level data, our formulation allows for multi-product firms and can be readily applied to study investment data at the product or establishment level.

Previous Literature. Our study relates to a broad range of contributions in finance, macroeconomics and industrial organization. The most closely-related work is two recent contributions by Okumura (2025) and Hopenhayn and Okumura (2025), who develop a model of endogenous growth that is based on the same demand system (GHL) and mode of competition used in our own model. Our contributions are complementary: they both make important progress in developing dynamic extensions of the static network oligopoly model of Pellegrino (2025), but focus on different applications and research questions. While Okumura (2025) and Hopenhayn and Okumura (2025) focus on growth dynamics in a continuous-time model with R&D spillovers, we develop instead a discrete-time neoclassical investment application. More broadly, our model contributes to the rapidly growing literature on network oligopoly and GHL demand (e.g. Eeckhout and Veldkamp, 2023, Bizzarri and Vega-Redondo, 2024, Voelkening, 2024, Ederer and Pellegrino, 2025, Mukherjee et al., 2025, Miyashita, 2026).

As mentioned above, our model falls squarely in the tradition of neoclassical models, including the seminal contributions of Jorgenson (1963), Tobin (1969), Abel (1980), Hayashi (1982) and Abel and Eberly (1994). This literature has been recently revived by the recent debates around the rise of on market power (Syverson, 2019, Syverson, 2024) and the role of intangible capital (Eisfeldt, Kim, and Papanikolaou, 2022; Eisfeldt and Papanikolaou, 2013), with key contributions by Peters and Taylor (2017), and Crouzet and Eberly (2021, 2023). Eberly and Wang (2009) study capital reallocation in a neoclassical setting, whereas two recent contributions by Chodorow-Reich, Smith, Zidar, and Zwick (2024) and Chodorow-Reich (2025) examine the role of tax policy.

Our model also has its roots in a long-standing theoretical literature on dynamic duopoly and oligopoly games with linear demand, dating back to Maskin and Tirole (1987). Reynolds (1987, 1991) pioneered early models of dynamic duopoly and symmetric oligopoly with capacity adjustment costs.

We also contribute to a literature at the boundary between finance, macroeconomics and industrial organization, including studies on the secular change in industry concentration and markups (De Loecker, Eeckhout, and Unger, 2020, Covarrubias, Gutiérrez, and Philippon, 2020, Kwon, Ma, and Zimmermann (2024), Hoberg and Phillips, 2025), the relationship between market power and stock market valuations (Bustamante and Donangelo, 2017,Corhay, Kung, and Schmid, 2025, Boppart, Klenow, Laski, and Li (2025); Cho, Grotteria, Kremens, and Kung, 2023).

Because the model we describe is a type of network game, our paper also connects to a booming literature on networks, which spans both microeconomic theory and macroeconomics. On the micro side, we note the work of Ushchev and Zenou (2018), Galeotti, Golub, and Goyal (2020) and Galeotti, Golub, Goyal, Talamas, and Tamuz (2024). The literature on networks in macroeconomics includes (among many others): Acemoglu et al. (2012), Bloom, Schankerman, and Van Reenen (2013), Carvalho, Nirei, Saito, and Tahbaz-Salehi (2021) Baqaee and Farhi (2020, 2024), and Baqaee (2018). Liu and Tsyvinski (2024) develop a dynamic model that incorporates adjustment costs of changing inputs. Vom Lehn and Winberry (2022) describe the "investment network" and analyze its role in business cycle fluctuations.

Last, we add to a recent literature that enriches macroeconomic models with oligopolistic competition, which includes (among others) Atkeson and Burstein (2008), Grassi (2017), Wang and Werning (2022) and Burstein, Carvalho, and Grassi (2025).

2 The Network Q Model

2.1 Hedonic Consumer Demand

In this section, we present our dynamic model of corporate investment in imperfectly-competitive product markets. Our model is populated by generically-defined *business units*, which can be interpreted as establishments, plants, product lines or single-product firms, depending on the application. We model the supply and investment decision at the level of business units, which may be either independent from each other or instead owned by the same firm, as we elaborate below.

Time is discrete and indexed by t. We index each business unit with the subscript $n \in \{1,2,...,N\}$. Consumer demand follows the GHL demand framework of Pellegrino (2025). We model each business unit's product n as a bundle of characteristics, sorted into two distinct types. Characteristics that are common across all products are indexed by $o = 1,2,...,O$, whereas product-specific characteristics are directly indexed by their product label n. For each product n, we define a_{on} as the quantity of common characteristic o provided by that product. Every product is represented by an $\mathscr{O}$-dimensional, strictly-positive column vector $\mathbf{a}_i$ with unit length, expressed formally as:

$$\mathbf{a}_n = \begin{bmatrix} a_{1n} & a_{2n} & \dots & a_{On} \end{bmatrix}' \tag{1}$$

$$\text{with} \qquad \sum_{o=1}^{O} a_{on}^2 = 1 \quad \forall\, n \in \{1,2,...,N\} \tag{2}$$

This vector $\mathbf{a}_n$ locates product n within the common characteristic space. Considering all products at the same time, these location vectors can be organized into a single $O \times N$ matrix $\mathbf{A}$:

$$\mathbf{A} \equiv \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \cdots & \mathbf{a}_n \end{bmatrix} \equiv \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1N} \\ a_{21} & a_{22} & \cdots & a_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ a_{O1} & a_{O2} & \cdots & a_{ON} \end{bmatrix} \tag{3}$$

Let y_{nt} represent the volume of product n that is produced and consumed at time t, and denote by $\mathbf{y}_t$ the consumption bundle demanded by the representative consumer at time t. In accordance with hedonic demand theory, consumers aggregate characteristics from various products through linear combination, with utility derived from these characteristics rather than the products themselves. If we denote the total consumption of common characteristic o at time t as x_{ot}, then:

$$x_{ot} = \sum_{n=1}^{N} a_{on}\, y_{nt} \qquad \text{or (in linear algebra notation)} \qquad \mathbf{x}_t = \mathbf{A}\mathbf{y}_t \tag{4}$$

For product-specific characteristics, we assume that each unit of product n delivers exactly one unit of its corresponding unique characteristic. Hence, y_n is both the output of product line n as well as the units of product-specific characteristics embedded in one unit of product n.

Our model employs a representative consumer framework, with utility that is quadratic in both the common characteristics (x_{ot}) and the idiosyncratic characteristics (y_{nt}) associated with each of

$n = 1,...,N$ products in any period t. In matrix notation, this corresponds to the function:

$$U(\mathbf{x}_t, \mathbf{y}_t, L_t) \stackrel{\text{def}}{=} \alpha\left(\mathbf{x}_t^{'}\mathbf{b}^x - \frac{1}{2}\mathbf{x}_t^{'}\mathbf{x}_t\right) + (1-\alpha)\left(\mathbf{y}_t^{'}\mathbf{b}^y - \frac{1}{2}\mathbf{y}_t^{'}\mathbf{y}_t\right) - L_t \tag{5}$$

where $\alpha \in (0,1)$ is the utility weight assigned to common characteristics and the vectors $\mathbf{b}^x$ and $\mathbf{b}^y$ represent characteristic-specific preference shifters and L_t is the total number of work hours supplied by the representative consumer at time t.

The representative consumer takes the vector of product prices by $\mathbf{p}_t$ as given and is endowed with shares of all the firms in the economy. The budget constraint of the consumer is then

$$\sum_{n=1}^{N} p_{nt} y_{nt} \leq L_t + \Pi_t \tag{6}$$

where Π_t is aggregate profits and the price of labor is normalized to 1 (labor is the numéraire good). The consumer's demand for each product n maximizes (5) subject to (6). In matrix form, the inverse demand function $\mathbf{p}_t$ resulting from the consumer's optimization problem is given by:

$$\mathbf{p}_t = \mathbf{b} - (\mathbf{I} + \mathbf{\Sigma})\mathbf{y}_t \tag{7}$$

where $\mathbf{b}$ is defined as a time-invariant $N \times 1$ vector of product-specific demand shifters

$$\mathbf{b} \stackrel{\text{def}}{=} \alpha \mathbf{A}'\mathbf{b}^x + (1-\alpha)\mathbf{b}^y \tag{8}$$

$\mathbf{I}$ is an $N \times N$ identity matrix and the matrix $\mathbf{\Sigma}$ is symmetric and equal to:

$$\mathbf{\Sigma} \stackrel{\text{def}}{=} \alpha\left(\mathbf{A}'\mathbf{A} - \mathbf{I}\right) \tag{9}$$

$\alpha \in (0,1)$ is the same preference parameter from equation (5). The (n,m) entry of $\mathbf{\Sigma}$ measures the degree of substitutability between products n and m, denoted σ_{nm}. When products n and m are more similar in their characteristic bundles, σ_{nm} is larger, implying that an increase in the supply of product n leads to a larger decline in the price of product m.

The key input to construct $\mathbf{\Sigma}$ is the matrix $\mathbf{A}'\mathbf{A}$, whose (n,m) element equals the inner product $\mathbf{a}_n'\mathbf{a}_m$. This inner product has a natural geometric interpretation: it equals the cosine of the angle between the characteristic vectors $\mathbf{a}_n$ and $\mathbf{a}_m$ in the O-dimensional space of common characteristics. For this reason, $\mathbf{a}_n'\mathbf{a}_m$ is called the *cosine similarity* between products n and m. The cosine similarity ranges from zero (when products share no common characteristics and are thus orthogonal in characteristic space) to one (when products have identical characteristic bundles). Products with higher cosine similarity provide more overlapping mixes of characteristics to consumers, making them closer substitutes in demand.

The parameter α serves two distinct roles in our framework. First, as shown in equation (5), it determines the relative weight consumers place on common versus product-specific characteristics when evaluating utility. Second, through equation (9), it scales the strength of competitive interactions in the product market network. A higher value of α implies that cross-price effects are stronger for any given level of similarity in the space of product characteristics.

This network approach to modeling product market competition offers a critical advantage over traditional methods: it allows the data to define product markets endogenously, rather than imposing market boundaries through administrative classifications.

2.2 Production and Investment

Each product in the economy is produced by a single business unit n, which in turn it is administered by a manager whom we also index by n. The optimization problem of the manager builds on the neoclassical literature: each business unit uses capital and a composite variable input which we interpret as labor to produce a perishable amount y_{nt} of product n. Production exhibits constant returns to scale. Consistent with recent quasi-experimental evidence on firm-level capital-labor substitution (Moreau, 2019; Benzarti and Harju, 2021) we assume a Leontief production function:[3] each unit of product$n$ requires the supply of $1/z_n$ units of capital and c_n units of labor, so that:

$$y_{nt} = z_n k_{nt}; \qquad \ell_{nt} = c_n y_{nt} \tag{10}$$

where ℓ_{nt} captures the amounts of hours allocated to business unit n by the representative consumer, z_n represents the productivity associated of capital of business unit n, and c_n represents the marginal cost of operating business unit n. We interpret k_{nt} as a composite of heterogeneous types of capital held by business unit n (i.e., including both tangible and intangible types of capital).

In each period t, manager n starts with capital stock k_{nt} and decides how much to invest and thus what output will be at time $t+1$. We define i_{nt} as the level of investment in capital of business unit n in period t. The law of motion for capital is then given by:

$$k_{nt+1} = i_{nt} + (1-\delta_n) k_{nt} \tag{11}$$

where δ_n is the rate at which the capital stock k_{nt} depreciates. Investment entails costs $\psi_n(k_{nt+1}; k_{nt})$ which drain the business unit's available cash flows each period. Consistent with Abel and Eberly (1994, among others), we assume that the function $\psi_n(k_{nt+1}; k_{nt})$ is convex. Furthermore, we assume that it takes the functional form:

$$\psi_n(k_{nt+1}; k_{nt}) \stackrel{\text{def}}{=} p_n^k i_{nt} + \frac{\theta}{2}(k_{nt+1} - k_{nt})^2 \tag{12}$$

The first term implies that the purchasing cost of capital equals p_n^k, whereas the second term incorporates quadratic adjustment costs of investment with curvature $\theta > 0$. We interpret the function $\psi_n(k_{nt}, k_{nt+1})$ as the business unit's total costs associated with capital expenditures in period t.

Notice that our specification of equation (12) assumes a common adjustment cost parameter θ for all business units. While this may seem restrictive at first sight, in practice our setting allows for plenty of heterogeneity in the effective burden of adjustment costs, which vary substantially across business units through differences in the productivity of capital z_n. All else equal, business units with low productivity z_n find it relatively costlier to adjust their stock of capital, making them sluggish in responding to changing competitive conditions. Conversely, business units with high productivity z_n attain the same level of output with less capital, allowing them to adjust their capital stock more nimbly over time. Furthermore, in Section 5 we show how the model can be extended to allow for firm-specific adjustment cost parameters θ_n.

[3] Aggregate substitution between capital and labor arises instead from between-firm reallocation, which our model captures through the equilibrium transition dynamics.

The problem of the manager of each business unit n is to choose both ℓ_{nt} and k_{nt+1} optimally. The optimal labor policy is intra-period and such that $\ell_{nt} = c_n k_{nt}$. Operating profits maximized with respect to labor are equal to:

$$\pi_n(\mathbf{k}_t) \stackrel{\text{def}}{=} (p_{nt} - c_n) z_n k_{nt} \tag{13}$$

Plugging in the expression for the output price, we obtain:

$$\pi_n(\mathbf{k}_t) = (b_n - c_n) z_n k_{nt} - z_n^2 k_{nt}^2 - \sum_{m \neq n} \sigma_{nm} z_n z_m k_{nt} k_{mt} \tag{14}$$

where $\mathbf{k}_t$ represents the vector of capital choices associated to all business units in the economy, and σ_{nm} is the (n,m) entry of the adjacency matrix $\mathbf{\Sigma}$ in (7). Alternatively, in matrix form, operating profits are expressed as:

$$\boldsymbol{\pi}(\mathbf{k}_t) = \mathbf{K}_t \mathbf{Z}(\mathbf{b} - \mathbf{c}) - \mathbf{K}_t \mathbf{Z}(\mathbf{I} + \mathbf{\Sigma}) \mathbf{Z} \mathbf{k}_t \tag{15}$$

where $\mathbf{Z}$ and $\mathbf{K}_t$ are simply the vectors $\mathbf{z}$ and $\mathbf{k}_t$ rearranged as a diagonal matrices:

$$\mathbf{Z} \stackrel{\text{def}}{=} \begin{bmatrix} z_1, & 0, & \dots & 0 \\ 0 & z_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & z_N \end{bmatrix}; \qquad \mathbf{K}_t \stackrel{\text{def}}{=} \begin{bmatrix} k_{1t}, & 0, & \dots & 0 \\ 0 & k_{2t} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & k_{Nt} \end{bmatrix} \tag{16}$$

Investors discount cash-flows in period $t+1$ to period t with the discount rate r. We define a business unit's cash flows as the operating profits in equation (13) net of investment costs in equation (12). We next proceed to define the manager's objective. In doing so, we allow the manager of business unit n to internalize the spillovers of her supply choices on some other business unit $m \neq n$ if such alternative business unit is also owned by the same firm. This flexibility allows to model complex ownership structures (such as multi-product firms) and study competitive arrangements across business units or firms (such as cartels). Two business units that operate under the same umbrella firm or which are part of the same cartel are assumed to fully internalize each other's product market spillovers (Baker and Bresnahan, 1985), so that there is no cannibalization among co-owned business units.

For every business unit n, let $\mathcal{M}_n$ be the set of business units $m \neq n$ that are owned by the same firm as n. Then, the optimization problem of the manager of business unit n at time t is to maximize the sum from time t onward of the internalized discounted cash flows of all business units associated with the firm owning business unit n, which we index hereafter by $m \in \mathcal{M}_n$. The formulation implies firm value at time t is expressed *cum-dividend*. Specifically, the manager maximizes the joint shareholder value of all business units $m \in \mathcal{M}_n$ belonging to the same firm, by choosing the optimal investment policy of business unit n, taking as given the capital choices of all business units different from n:

$$v_n(\mathbf{k}_t) \stackrel{\text{def}}{=} \max_{(k_{nt+1}, k_{nt+2}, \dots)} \sum_{T=t}^{\infty} \sum_{m \in \mathcal{M}_n} \frac{\pi_m(\mathbf{k}_T) - \psi_m(k_{mT+1}; k_{mT})}{(1+r)^{T-t}} \tag{17}$$

where $v_n(\mathbf{k}_t)$ represents the total firm value relevant for the manager of business unit n, as a function of the vector of capital stock $\mathbf{k}_t$. Each manager then solves equation (17) subject to equations (11), (13), (12), and a transversality condition which ensures the finiteness of the value function as time t goes to infinity. By construction, managers of business units within the same firm maximize identical objective functions, and yet the control variable of each manager within the same firm is the investment decision associated with their specific business unit.

The dynamic oligopoly game that we have set up is a type of *dynamic network game*. A network is defined as a set of nodes and relationships between the nodes, defined by an *adjacency matrix*. In this setting, business units are nodes, and the (scaled) product similarity matrix $\mathbf{\Sigma}$ can be thought of as the adjacency matrix. We also define a network of ownership between business units. Let γ_{mn} be an indicator/dummy variable that equals 1 if $m \in \mathscr{M}_n$ and $n \in \mathscr{M}_m$. Then the symmetric matrix $\mathbf{\Gamma}$, whose (m,n) entry is γ_{mn}, is an adjacency matrix that describes shared ownership between two alternative business units.

2.3 Equilibrium Concept and Dynamic Stability

In what follows, we derive the optimal investment decision of all managers n at time t by maximizing the N value functions in (17) in stacked form. Throughout our analysis, the manager n optimizes on k_{nt+1} while taking as given the decision of all other business units, which we denote collectively with the subscript $-n$. The equilibrium concept that we use is *Markov Perfect Equilibrium* (MPE). This refinement of the Nash Equilibrium requires that the optimal investment policy of business unit n takes the form:

$$k_{nt+1} = f_n(\mathbf{k}_t) \tag{18}$$

It follows that k_{nt+1} is only a function of the vector of contemporaneous capital stocks: the previous history of play is irrelevant for the future evolution of the system. Given this formulation, we can re-write the maximization problem of manager n as the following dynamic program:

$$v_n(\mathbf{k}) = \max_{f_n(\mathbf{k})} \sum_{m \in \mathscr{M}_n} \left[\pi_m(\mathbf{k}) - \psi_m(f_m(\mathbf{k}); k_m)\right] + \frac{v_n(f_n(\mathbf{k}); \mathbf{f}_{-n}(\mathbf{k}))}{1+r} \tag{19}$$

subject to $\mathbf{f}_{-n}$, the optimal investment policy of all other business units. We obtain the equilibrium vector $\mathbf{k}_{t+1}$ by solving simultaneously a system of n stacked Bellman equations, where each equation represents a business unit facing the optimization problem in equation (19).

Definition 1. *A Markov Perfect Equilibrium (MPE) is a vector of value functions*

$$\mathbf{v}(\mathbf{k}) \equiv [v_1(\mathbf{k}), v_2(\mathbf{k}), ..., v_N(\mathbf{k})]^\top \tag{20}$$

and a vector of investment policies

$$\mathbf{f}(\mathbf{k}) \equiv [f_1(\mathbf{k}), f_2(\mathbf{k}), ..., f_N(\mathbf{k})]^\top \tag{21}$$

such that for all business units $n = 1,2,...,N$: (i) given $\mathbf{f}_{-n}$, $v_n(\mathbf{k})$ satisfies the Bellman equation (19); (ii) the investment policy function $\mathbf{f}_n$ solves the optimization problem in (19).

The vector $\mathbf{v}(\mathbf{k}_t)$ represents a function that maps the vector of physical capital stocks into a vector of firm values. Having formally established the equilibrium concept, we now introduce the notion of steady-state.

Definition 2. *Given a Markov Perfect Equilibrium, we define its steady-state as a vector of capital stocks $\mathbb{k}$ such that $\mathbf{f}(\mathbb{k}) = \mathbb{k}$.*

To keep our analysis tractable, we follow Jun and Vives (2004) and use a refinement of the MPE concept that focuses on linear best responses and the dynamic stability of the process that describes the evolution of the state vector $\mathbf{k}$ along the equilibrium path.

Definition 3. *A Stable Markov Perfect Equilibrium in Linear Strategies is a MPE with linear investment policies – that is, where* $\mathbf{f}(\mathbf{k})$ *takes the form*[4]

$$\mathbf{f}(\mathbf{k}) - \mathbb{k} = \mathbf{\Phi}(\mathbf{k} - \mathbb{k}) \tag{22}$$

where all the eigenvalues of matrix $\mathbf{\Phi}$ *lie in the interval* $(0,1)$*, and the vector* $\mathbb{k}$ *is equal to the unique steady state of the model.*

As we elaborate further below, the matrix $\mathbf{\Phi}$ governs the dynamics of the capital accumulation process for each business unit, as well as the spillover effects of a manager's own investment decision on the investment of its peers over time. By construction, along the equilibrium path, the capital stocks of business units follow a vector auto-regression (VAR) process. The equilibrium is dynamically stable if such VAR process is stable in a time-series sense (technically, the eigenvalues of matrix $\mathbf{\Phi}$ lie in the interval (0,1)). This is an important property of the equilibrium: it ensures that, for every initial vector of capital stocks, operating profits do not become infinitely negative.[5] It follows that the vector of capital stocks gradually converges to its steady state level along the MPE path:

$$\lim_{t\to\infty} \mathbf{k}_t = \mathbb{k} \tag{23}$$

For this reason, with some abuse of notation, we shall use $k_{n\infty}$ to denote the steady-state capital stock of firm n, so that:

$$\mathbb{k} \overset{\text{def}}{=} \begin{bmatrix} k_{1\infty} & k_{2\infty} & \ldots & k_{N\infty} \end{bmatrix} \tag{24}$$

By combining equation (22) with the law of motion of capital (equation 11), we obtain the investment policies of all business units as a system of N linear equations:

$$\boldsymbol{\imath}_{t+1}(\mathbf{k}_t) = \mathbf{\Delta}\mathbf{k}_t + (\mathbf{I} - \mathbf{\Phi})(\mathbb{k} - \mathbf{k}_t) \tag{25}$$

where $\mathbf{\Delta}$ is the vector of depreciation rates $\boldsymbol{\delta}$ rearranged as a diagonal matrix:

$$\mathbf{\Delta} \overset{\text{def}}{=} \begin{bmatrix} \delta_1, & 0, & \ldots & 0 \\ 0 & \delta_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \delta_N \end{bmatrix} \tag{26}$$

In the coming sections, we gradually work our way through to explicit expressions for the steady state vector of capital stocks $\mathbb{k}$ and the transition matrix $\mathbf{\Phi}$, and then provide a full characterization of the (unique) Stable Linear Markov Perfect Equilibrium.

[4]Notice that this definition of "linear" is without loss of generality as long as the matrix $(\mathbf{I} - \Phi)$ is invertible, which we verify later on. Specifically, expressing linear strategies as in (22) is the same as: 1) writing the more general expression $f(\mathbf{k}_t) = \boldsymbol{\iota} + \mathbf{\Phi}\mathbf{k}_t$; 2) noticing that the unique steady state is given by the vector $\mathbb{k} = (\mathbf{I} - \Phi)^{-1}\boldsymbol{\iota}$; 3) substituting vector $\boldsymbol{\iota}$ in terms of $\mathbb{k}$ inside $f(\mathbf{k}_t) = \boldsymbol{\iota} + \mathbf{\Phi}\mathbf{k}_t$ to obtain (22).

[5]If the matrix $\mathbf{\Phi}$ were unstable, then for a suitable choice of $\mathbf{k}_t$ the capital stock would diverge unboundedly from the steady state for any given firm. Because the profit function is quadratic, profits and ultimately expected firm value would become infinitely negative—then violating the transversality condition associated with (17).

2.4 Marginal Q: Definition

We formally introduce the concept of marginal-Q, consistent with the seminal contribution of Hayashi (1982). The marginal contribution of a business unit's own capital to the value of its associated firm at time t is defined as:

$$q_{nn}(\mathbf{k}_t) \stackrel{\text{def}}{=} \frac{\partial v_n(\mathbf{k}_t)}{\partial k_{nt}} \tag{27}$$

The double subscript nn is intentional: because the economic environment is populated by multiple business units we can equally well define the marginal effect on the value of the firm associated with business unit n of an additional unit of capital accumulated by another business unit m:

$$q_{nm}(\mathbf{k}_t) \stackrel{\text{def}}{=} \frac{\partial v_n(\mathbf{k}_t)}{\partial k_{mt}} \tag{28}$$

Our model thus generalizes Hayashi's marginal-Q from the idle firm case to network of business units that interact with each other. In our network game, Q is no longer a scalar, but rather a matrix. Specifically, we define $\mathbf{Q}(\mathbf{k}_t)$ as the Jacobian of the vector of firm values $\mathbf{v}(\mathbf{k}_t)$:

$$\mathbf{Q}(\mathbf{k}_t) \stackrel{\text{def}}{=} \frac{\partial \mathbf{v}(\mathbf{k}_t)}{\partial \mathbf{k}_t} \equiv \begin{bmatrix} q_{11}(\mathbf{k}_t) & q_{12}(\mathbf{k}_t), & \dots & q_{1N}(\mathbf{k}_t) \\ q_{21}(\mathbf{k}_t) & q_{22}(\mathbf{k}_t) & \dots & q_{2N}(\mathbf{k}_t) \\ \vdots & \vdots & \ddots & \vdots \\ q_{N1}(\mathbf{k}_t) & q_{N2}(\mathbf{k}_t) & \dots & q_{NN}(\mathbf{k}_t) \end{bmatrix} \tag{29}$$

Due to the presence of product market externalities, the capital accumulation decision of any business unit within firm n can affect the market value of any firm m. If we think of $\mathbf{Q}(\mathbf{k}_t)$ as an adjacency matrix, it describes a network of product market spillovers on business unit (and firm) values. With some abuse of notation, we can use the vector $\mathbf{q}(\mathbf{k}_t)$ to denote the diagonal entries of matrix $\mathbf{Q}(\mathbf{k}_t)$, which correspond to Hayashi's familiar definition of a firm's own marginal product of capital:

$$\mathbf{q}(\mathbf{k}_t) = \operatorname{diag}[\mathbf{Q}(\mathbf{k}_t)] \equiv \begin{bmatrix} q_{11}(\mathbf{k}_t) & q_{22}(\mathbf{k}_t) & \cdots & q_{NN}(\mathbf{k}_t) \end{bmatrix}^\top \tag{30}$$

2.5 Optimality Conditions and the Competitive Pressure Matrix

Using the definition of marginal Q in equation (27), we solve the optimal investment decision of each manager, taking as given the choices of other managers. The first order condition with respect to k_{nt+1} as expressed in equation (19) yields:

$$\frac{q_{nn}(\mathbf{k}_{t+1})}{1+r} = \frac{\partial \psi_n(k_{nt+1}; k_{nt})}{\partial k_{nt+1}} \tag{31}$$

where the left hand side captures the discounted marginal benefit of investment next period and the right hand side captures the marginal costs of investment today. It is worth noticing the function $q_{nn}(\mathbf{k}_{t+1})$ depends not only on the investment of business unit n itself, but also on the investment choices of all the other business units.

Differentiating the maximand of equation (17) with respect to k_{nt+1} we obtain the following system of stacked first order conditions of all business units:

$$\frac{1}{1+r}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c})-\mathbf{\Omega}\mathbf{k}_{t+1}+(\mathbf{I}-\mathbf{\Delta})\mathbf{p}^{k}+\theta(\mathbf{k}_{t+2}-\mathbf{k}_{t+1})\right]=\mathbf{p}^{k}+\theta(\mathbf{k}_{t+1}-\mathbf{k}_{t}) \tag{32}$$

where the left hand side of the equation captures the expected benefit of having a larger stock of capital tomorrow, the right hand the associated cost of investing today. More importantly, we denote the matrix $\mathbf{\Omega}$ as *Competitive Pressure* matrix, such that:

$$\mathbf{\Omega} \overset{\text{def}}{=} \mathbf{Z}(\mathbf{I}+\mathbf{\Sigma}+\mathbf{D}+\mathbf{\Gamma}\circ\mathbf{\Sigma})\mathbf{Z} \tag{33}$$

The symbol "$\circ$" denotes the Hadamard (entry-by-entry) product, and $\mathbf{D}$ is a diagonal matrix of conduct parameters (Weyl and Fabinger, 2013) with binary diagonal entries which are equal to one if and only if the corresponding business unit acts as an oligopolist (and, thus, internalizes she is facing a downward sloping demand curve):

$$\mathbf{D} \overset{\text{def}}{=} \begin{bmatrix} \mathscr{D}_1 & 0 & \dots & 0 \\ 0 & \mathscr{D}_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \mathscr{D}_N \end{bmatrix}, \qquad \mathscr{D}_n \overset{\text{def}}{=} \begin{cases} 0 & \text{if firm } n \text{ is a price} - \text{taker} \\ 1 & \text{if firm } n \text{ is an oligopolist} \end{cases} \tag{34}$$

The entries γ_{nm} of the adjacency matrix $\mathbf{\Gamma}$ are dummy variables capturing ownership, and hence are equal to one if business units n and m are part of the same firm – formally:

$$\mathbf{\Gamma} \overset{\text{def}}{=} \begin{bmatrix} \gamma_{11} & \gamma_{12} & \cdots & \gamma_{1N} \\ \gamma_{21} & \gamma_{22} & \cdots & \gamma_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{N1} & \gamma_{N2} & \cdots & \gamma_{NN} \end{bmatrix} \qquad \gamma_{nm} \overset{\text{def}}{=} \begin{cases} 0 & \text{if } n \notin \mathscr{M}_m, m \notin \mathscr{M}_n \\ 1 & \text{if } n \in \mathscr{M}_m, m \in \mathscr{M}_n \end{cases} \tag{35}$$

The formulation of matrix $\mathbf{\Omega}$ is therefore flexible enough to allow for a multiplicity of competitive conduct cases and policy-relevant counterfactuals. For instance, if we let $\mathbf{D}=\mathbf{\Gamma}=\mathbf{I}$ we have the canonical Cournot conduct (single-product firms that internalize downward sloping demand curves, but not product market spillovers); if we consider instead $\mathbf{D}=\mathbf{I}$, $\mathbf{\Gamma}=\mathbf{1}\mathbf{1}'$ we are simulating an economy where all spillovers are internalized – a large cartel that encompasses all business units; finally, if $\text{diag}(\mathbf{D})=\mathbf{0}$ and $\mathbf{\Gamma}=\mathbf{I}$ we have a competitive economy where the first welfare theorem holds.

More importantly, the Competitive Pressure matrix $\mathbf{\Omega}$ embeds a critical economic interpretation, which is key to our modeling framework: its (n,m) entry is by definition the second-order derivative

$$\frac{\partial^2\left(\sum_{n'\in\mathscr{M}_n}\pi_{nt+1}\right)}{\partial k_{nt+1}\,\partial k_{mt+1}} = -\omega_{nm} \tag{36}$$

$\mathbf{\Omega}$ is therefore a measure of the *concavity* of the firm's (internalized) profit function: any of its n^{th} diagonal entries equals $2z_n^2$ for the canonical Cournot conduct, and z_n^2 when firms are price takers. In turn, the off-diagonal entries measure the effect of a unit change in a competitor's capital stock on the focal firm's marginal rate of return on capital. By the implicit function theorem, $\frac{1}{4}\omega_{nm}^2$ is also the $t+1$ dollar loss that is incurred (internalized) by business unit n if it fails to adjust its capacity in response to a 1 unit change in the capital of its peer m.

In what follows, we shall assume that $\mathbf{\Omega}$ is symmetric positive definite - all of its eigenvalues are strictly positive. This assumption is always satisfied in our empirical implementation.

2.6 Equilibrium: Existence, Uniqueness and Characterization

In this subsection, we establish the existence and uniqueness of the linear stable MPE. By guessing an equilibrium of the form of equation (22) and substituting $\mathbf{k}_{t+2}$ with the corresponding expression in terms of $\mathbf{k}_{t+1}$, we obtain the following analytical solution.

Proposition 1 (Existence, Uniqueness, Necessary and Sufficient Conditions)**.** *Consider the linear investment policy in equation* (*22*)*. There exists a unique Stable Markov Perfect Equilibrium in linear strategies, characterized by the following expressions for the steady state capital stock vector*

$$\mathbb{k} \;=\; \mathbf{\Omega}^{-1}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c})-(r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^{k}\right] \tag{37}$$

and of the transition matrix $\mathbf{\Phi}$*:*

$$\mathbf{\Phi} \;=\; \frac{1}{2}\left[\frac{1}{\theta}\mathbf{\Omega}+(2+r)\mathbf{I}\right]-\sqrt[+]{\frac{1}{4}\left[\frac{1}{\theta}\mathbf{\Omega}+(2+r)\mathbf{I}\right]^{2}-(1+r)\mathbf{I}} \tag{38}$$

where $\sqrt[+]{\mathbf{M}}$ *denotes, for a symmetric square matrix* $\mathbf{M}$ *with all positive eigenvalues, the "principal" square root matrix – that is, the unique square root matrix with all positive eigenvalues.*[6]

Proof. The steady state vector of capital stocks in equation (37) can be easily found by imposing $\mathbf{k}_{t+2}=\mathbf{k}_{t+1}=\mathbf{k}_{t}=\mathbb{k}$ inside equation (32). See Appendix A for the rest of the proof. □

The equilibrium transition matrix $\mathbf{\Phi}$ in equation (38) is the unique, explicit solution of a Non-symmetric Algebraic Riccati Equation (NARE). Neoclassical models typically solve for a relationship between investment policies as a function of firms' marginal product of capital. They rarely provide explicit solutions for firms' investment policies, value functions and marginal Q. Our model is uniquely tractable, in that we can solve for the steady-state vector $\mathbb{k}$ and the transition matrix $\mathbf{\Phi}$.

The diagonal entries of matrix $\mathbf{\Phi}$ in equation (38) capture the how fast firm n's capital stock gets pulled toward the steady state (depending on how far from its k_n is), whereas the off-diagonal terms capture instead the equilibrium investment spillovers across business units in the product market network.

As in the static game of Pellegrino (2025), the steady state equilibrium capital stock vector $\mathbb{k}$ in equation (37) admits an intuitive interpretation from network theory. The competitive pressure matrix $\mathbf{\Omega}$ can also be used (like $\mathbf{\Sigma}$ and $\mathbf{\Gamma}$) as an adjacency matrix to define a network. Then $\mathbb{k}$ can be interpreted as a metric of network centrality – specifically, the metric developed by Katz (1953) and Bonacich (1987). A business unit that is "peripheral" in the network – that is, a business unit that is shielded from competitive pressures (e.g. because it has few productive competitors) – will attain, all else equal, a higher level of capital in the steady-state.

In the following subsections, we use equations (37) and (38) to characterize explicitly the value function of the firm, investment rates and the associated marginal Q's.

[6] See Higham (2008, Chapter 1, Theorem 1.29).

2.7 Transition Dynamics and Spectral Analysis of $\mathbf{\Phi}$

In this subsection we present and discuss an novel and unique theoretical insight that emerges from our model: as it turns out, product market competition not only affect the equilibrium steady state level of capital (37), but also the speed of convergence to steady state.

As is well known from the macroeconomics literature[7], the MPE transition dynamics are governed by the eigenvalues of the matrix $\mathbf{\Phi}$. The fact that business units are heterogeneous in their capital productivity (z_n) and that the competitive pressure matrix $\mathbf{\Omega}$ defines a highly asymmetric network. As some firms are much more exposed to competitive pressure than others, the transition matrix $\mathbf{\Phi}$ has (by construction) rich spectral properties – most importantly, vast dispersion in its eigenvalues.

Balanced Growth Paths. To analyze the spectral properties of matrix $\mathbf{\Phi}$, we introduce some more notation. Let $\mathbf{e}_n$ be the n^{th} eigenvector of $\mathbf{\Phi}$, and let $\hat{\varphi}_n$ be the corresponding n^{th} eigenvalue. By initializing the capital stock vector at $\mathbf{k}_0$, which is removed from $\mathbb{k}$ by an amount that is proportional to any of the n eigenvectors $\mathbf{e}_n$, we can construct n "balanced growth paths". Along these balanced growth paths, all firms' capital stock converge towards the steady state at a rate that is exactly equal to eigenvalue $\hat{\varphi}_n$:

$$(\mathbf{k}_0 - \mathbb{k}) \propto \mathbf{e}_n \qquad \Longrightarrow \qquad (\mathbf{k}_t - \mathbb{k}) = \hat{\varphi}_n^t \, (\mathbf{k}_0 - \mathbb{k}) \tag{39}$$

Due to rich heterogeneity embedded in $\mathbf{z}$ and $\mathbf{\Omega}$, the transition matrix $\mathbf{\Phi}$ will have N different eigenvalues: this implies that, in our model, the MPE transition speed is *state-dependent* – i.e. it is determined by the allocation of capital (along which direction $\mathbf{k}_t$ deviates from the steady state). Fixing an aggregate level of capital stock K, there are allocations of capital (states) $\mathbf{k}$ that are highly "transitory", in the sense that firms will swiftly transition away from them to migrate to the steady state $\mathbb{k}$. Other capital allocations $\mathbf{k}$ are instead "sticky" or "persistent", in the sense that they will produce a slow transition to the steady state.

Firm-level Tradeoff and Persistence. The rich spectral properties of $\mathbf{\Phi}$ also imply that, outside of the balanced growth path, individual firms will adjust their capital stock at vastly different speeds. It also means that the speed itself at which each firm adjusts its capacity will change as the transition unfolds. The speed at which any given firm n adjusts its capital stock depends on whether firm n is associated with high or low eigenvalues. In turn, we predict that such speed also depends on the configuration of the product market. Specifically, firms that are shielded from competitive pressure (in a certain mathematical sense) load on higher eigenvalues of $\mathbf{\Phi}$, and thus adjust their capacity more slowly. Vice-versa, firms that are *central* in the network of competitive pressures will adjust their capital stock much faster.

We start building intuition for this claim by noting that the analytical solution for $\mathbf{\Phi}$ (equation 38) depends directly on the competitive pressure matrix $\mathbf{\Omega}$. However, equation (38) includes matrix powers and roots that obfuscate somewhat the relation between the two matrices. To make the link between these two matrices more transparent, we can leverage the properties of $\mathbf{\Omega}$ and a key result from matrix theory.

Corollary 1. *Let $\hat{\omega}$ the vector of eigenvalues of $\mathbf{\Omega}$ and $\mathbf{E}$ be the corresponding (orthogonal) matrix of eigenvectors, so that*

$$\mathbf{\Omega} = \mathbf{E} \, \mathrm{diag}(\hat{\omega}) \, \mathbf{E}^* \tag{40}$$

[7] See Ljungqvist and Sargent (2018) for a review.

Figure 1: Eigenvalues of the Transition Matrix

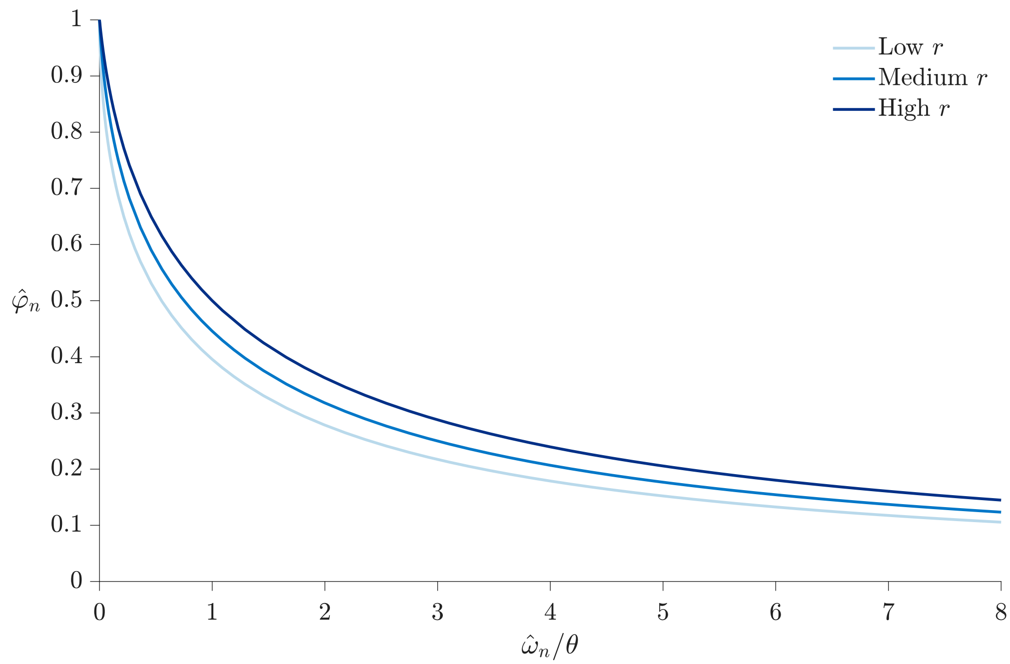

and $\mathbf{E}^* \equiv \mathbf{E}' \equiv \mathbf{E}^{-1}$ *is the transpose or inverse of* $\mathbf{E}$ (*in this case the two coincide due to* $\mathbf{E}$ *being orthogonal*). *Also, let* $\hat{\boldsymbol{\varphi}}$ *be the vector eigenvalues of* $\mathbf{\Phi}$. *Then,* $\mathbf{\Phi}$ *and* $\mathbf{\Omega}$ *are simultaneously diagonalizable:*[8]

$$\mathbf{\Phi} = \mathbf{E}\,\mathrm{diag}\left(\hat{\boldsymbol{\varphi}}\right)\mathbf{E}^* \tag{41}$$

and their eigenvalues are related entry-by-entry by the following the scalar function:

$$\hat{\varphi}_n = \frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right) - \sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right)^2 - (1+r)} \tag{42}$$

which is simply the scalar equivalent of equation (38).

Proof. Embedded in the proof of Proposition 1, in Appendix A. □

This corollary uncovers the economic incentives that underlie the link between $\mathbf{\Omega}$ and $\mathbf{\Phi}$. Figure 1 displays $\hat{\varphi}_n$ as a function of $\hat{\omega}_n/\theta$ and r. First, consistent with previous studies such as Chodorow-Reich (2025), we find that all eigenvalues $\hat{\varphi}_n$ are increasing functions of the curvature of adjustment costs θ and of the discount rate r. That is, higher adjustment costs and higher discount rates slow down the speed of convergence to the steady state. It is immediately obvious why higher adjustment costs translate into higher persistence of capital. The intuition for why higher discount rates also induce more persistence in capital is that, in deciding future capacity, business units evaluate adjustment costs that are paid in the current period against profits that are realized in future periods. A higher discount rate increases the relative importance adjustment costs, leading to higher persistence of capital.

[8] This implies that $\mathbf{\Phi}$ and $\mathbf{\Omega}$ share a basis and thus commute.

As a novel insight, our model shows that $\hat{\varphi}_n$ is a decreasing function of the ratio of $\hat{\omega}_n/\theta$ and is homogeneous of degree zero in $\left(\hat{\omega}_n, \theta^{-1}\right)$. In other words, from a quantitative standpoint, the eigenvalue $\hat{\omega}_n$ stands out as the "antagonist" of the adjustment cost parameter θ in determining the speed of transition. To build economic intuition behind this result, let us recall the economic interpretation of the competitive pressure matrix $\mathbf{\Omega}$: it measures the concavity of the firms' profit functions – or, equivalently the economic loss from failing to adjust capacity in response to a change in a rival's capital stock.

This observation takes us to the heart of the dynamic economic trade-off faced by business units: when deciding how much to invest, managers trade off the direct cost of adjusting the current capital stock (captured by the convexity of the installation function ψ_n) against the future loss in profits that they incur by *not* adjusting the capital stock (captured by the concavity of the profit function). The convexity of the adjustment cost function, modulated by θ, is the economic force that "pulls" capital stocks towards the current level $\mathbf{k}_t$. The concavity of the profit function vector, modulated by the eigenvalues of $\mathbf{\Omega}$, is the economic force that "pushes" capital stocks towards the steady state $\mathbb{k}$. The Markov Perfect Equilibrium presented in Proposition 1 is the outcome of these two economic forces balancing each other out as a result of the firms' forward-looking decision-making.

Special Case: Symmetry. The matrix $\mathbf{\Omega}$ and its eigenvalues are themselves non-trivial functions of N (the number of firms), $\mathbf{z}$ (the capital productivity) and the substitution matrix $\mathbf{\Sigma}$. While it is interesting to ask how the transition matrix $\mathbf{\Phi}$ depends on these individual objects, it is impossible to derive analytical results for the general case. Therefore, to further investigate the relationship between $\mathbf{\Phi}$ and these more fundamental objects, we narrow down our focus to the special (symmetric) case where all firms are identical and engage in standard Cournot competition. For this case, we can obtain analytical results. As we show in Appendix A, even in this special case the competitive pressure matrix $\mathbf{\Omega}$ has multiple eigenvalues: yet, only one of such eigenvalues is associated with the unit eigenvector, and is therefore of interest.

Proposition 2. *In a Cournot oligopoly with* $\mathbf{D} = \mathbf{\Gamma} = \mathbf{I}$ *and* N *identical single product firms, the unique eigenvalue of matrix* $\mathbf{\Omega}$ *associated with the unit eigenvector* $\mathbf{1}$*, which we denote* $\hat{\omega}_{\mathbf{1}}$*, is given by*

$$\hat{\omega}_{\mathbf{1}} = z^2 \left[2 + (N-1)\sigma\right] \tag{43}$$

Denoting $\hat{\varphi}_{\mathbf{1}}$ *the corresponding eigenvalue of matrix* $\mathbf{\Phi}$*, it follows that*

$$\frac{\partial \hat{\varphi}_{\mathbf{1}}}{\partial z} < 0 \qquad \frac{\partial \hat{\varphi}_{\mathbf{1}}}{\partial \sigma} < 0 \qquad \frac{\partial \hat{\varphi}_{\mathbf{1}}}{\partial N} < 0 \tag{44}$$

Proof. See Appendix A. □

The proposition shows that either lower dependence on capital (higher z), stronger substitution among products (σ), or a higher number of firms (N), all result in lower persistence in capital, and a faster convergence to the steady state.

In sum, our model delivers the novel insight that the rate of capital accumulation is not only affected by adjustment costs or the discount rate, as well-known from the previous literature: it is also regulated by the degree of product market competition.

2.8 Equilibrium Enterprise Values

We now utilize the equilibrium investment policy to solve for the vector of firm-level value functions $\mathbf{v}(\cdot)$ associated with each business unit n. This is necessary to validate our solution of the MPE, but also to estimate investment equations empirically. To begin with, we note that the term in square brackets of equation (32) corresponds to the vector of firm-level $\mathbf{q}(\mathbf{k}_t)$. Leveraging the expression for $\mathbb{k}$ that we previously obtained, we can re-write $\mathbf{q}(\mathbf{k}_t)$ as a function of the price of capital and of the deviation of the vector of capital stocks $\mathbf{k}_t$ from the steady state.

Proposition 3. *The equilibrium vector of marginal products of capital (or marginal Qs) associated with each business unit is equal to*:

$$\mathbf{q}(\mathbf{k}_t) \;=\; (1+r)\mathbf{p}^k + \mathbf{\Lambda}(\mathbb{k} - \mathbf{k}_t) \tag{45}$$

where the matrix $\mathbf{\Lambda}$ *is defined as*

$$\mathbf{\Lambda} \stackrel{\text{def}}{=} \theta(\mathbf{I} - \mathbf{\Phi}) + \mathbf{\Omega} \tag{46}$$

implying that the corresponding steady state value is

$$\mathbf{q}(\mathbb{k}) \;=\; (1+r)\mathbf{p}^k \tag{47}$$

Proof. See Appendix A. □

Equation (45) is important for several reasons. First, it tells us that along the equilibrium path the matrix $\mathbf{\Lambda}$ defines a time-invariant linear mapping between marginal Q and the deviation of the state vector $\mathbf{k}_t$ from its steady-state. In other words, if $\mathbf{\Lambda}$ is known, $\mathbf{q}(\mathbf{k}_t)$ reveals each firm's capital stock's deviation from the steady state value. Second, the equation provides a closed-form expression for each firm's own marginal product of capital. Last, because $q_{nn}(k_{nt};\mathbf{k}_{-nt})$ is the derivative of $v_n(k_{nt};\mathbf{k}_{-nt})$ with respect to a firm's own capital stock, (45) is also a system of differential equations, which we can solve to recover enterprise value $\mathbf{v}(\cdot)$.

To obtain an exact solution for the Bellman Equation, however, we also need a boundary condition. We obtain it by evaluating equation (19) at the steady state. First, notice that in the steady state the cum-dividend value of the firm is given by:

$$v_n(\mathbb{k}) \;=\; \frac{1+r}{r} \sum_{m \in \mathcal{M}_n} \left[\pi_m(\mathbb{k}) - \delta_m p_m^k k_{m\infty}\right] \tag{48}$$

The expression for firm value under steady state is straightforward. The value of the firm equals the present value of its operating profits across all business units, less the cost of maintaining its capital stock. By expanding this expression and replacing the previously-derived vector of steady-state capital stocks ($\mathbb{k}$) we obtain an even more compact expression for the steady-state firm value.

Proposition 4. *In the steady-state, the value of the firm* $v_n(\mathbb{k})$ *is equal to the sum of the resale value of the steady-state capital stock and* ρ_n *– the present discounted value of future steady-state monopoly rents*:

$$v_n(\mathbb{k}) \;=\; \sum_{m \in \mathcal{M}_n} \left[(1+r)p_m^k k_{m\infty} + \rho_m\right] \tag{49}$$

where ρ_n is the present value of oligopoly rents in steady state, namely:

$$\rho_n \stackrel{\text{def}}{=} \sum_{T=t}^{\infty} \frac{1}{(1+r)^{(T-t)}} \left[p_n(\mathbb{k}) - \frac{\partial p_n y_n}{\partial y_n}(\mathbb{k}) \right] y_n(\mathbb{k}) \; = \; \frac{1+r}{r} \cdot z_n^2 k_{n\infty}^2 \tag{50}$$

Proof. See Appendix A. □

Armed with the boundary condition in equation (49), we can now proceed to recover the value function vector $\mathbf{v}(\cdot)$ by solving the system of differential equations in (45).

Proposition 5. *The vector of maximized Bellman equations in (19), corresponding to the equilibrium enterprise values, is solved by:*

$$\mathbf{v}(\mathbf{k}_t) \; = \; \mathbf{v}(\mathbb{k}) + (\mathbf{k}_t - \mathbb{k}) \circ \left[(1+r)\mathbf{p}^k + \left(\mathbf{\Lambda} - \frac{1}{2}\mathbf{I} \circ \mathbf{\Lambda} \right)(\mathbf{k}_t - \mathbb{k}) \right] \tag{51}$$

in scalar notation:

$$v_m(\mathbf{k}_t) \; = \; v_m(\mathbb{k}) + (k_{mt} - k_{m\infty}) \left[(1+r) p_m^k - \sum_{n=1}^{N} \frac{\lambda_{mn}}{2^{\mathbb{I}\{m=n\}}} (k_{nt} - k_{n\infty}) \right] \tag{52}$$

Proof. Differentiating each row of (52) with respect to k_i we verify that it satisfies (45). To verify that it satisfies the boundary condition (49) we plug $\mathbf{k}_t = \mathbb{k}$. □

2.9 Product Market Spillovers and the Equilibrium Q Network

There is a rich literature that discusses product market spillovers and peer effects linked to corporate investment. Within this literature, product market spillovers have been defined in a variety of ways. As a first step, aligned with previous work by Bloom, Schankerman, and Van Reenen (2013), we study how a firm's investment policy changes when we supply to that firm or one of its competitors an additional unit of capital. The answer to this question follows from computing the Jacobian the vector of firm's investment policies (equation 25):

$$\frac{\partial i_{nt}}{\partial k_{mt}} \quad = \quad \begin{cases} -(1-\delta_n) + \varphi_{nn} & \text{for } m = n \\ \varphi_{nm} & \text{for } m \neq n \end{cases} \tag{53}$$

We confirm that product market spillovers in investment are captured by the off-diagonal terms of the transition matrix $\mathbf{\Phi}$.

As investment rates are rarely solved explicitly, the neoclassical literature also studies product market spillovers in investment by studying the associated product market spillovers in marginal and average products of capital. As a next step, we specify a closed-form expression for the entire $\mathbf{Q}$ matrix, allowing us to characterize product market spillovers on firms' values.

Proposition 6. *The equilibrium matrix $\mathbf{Q}(\mathbf{k}_t)$ is given by:*

$$\mathbf{Q}(\mathbf{k}_t) \; = \; (1+r)\mathbf{P}^k + \text{diag}[\mathbf{\Lambda}(\mathbb{k} - \mathbf{k}_t)] + (\mathbb{k} - \mathbf{k}_t)\mathbf{1}' \circ \left(\mathbf{1}\mathbf{1}' - \mathbf{I}\right) \circ \mathbf{\Lambda} \tag{54}$$

and thus the corresponding steady-state value is:

$$\mathbf{Q}(\mathbb{k}) \; = \; (1+r)\mathbf{P}^k \tag{55}$$

where $\mathbf{P}^k$ *is simply* $\mathbf{p}^k$ *rearranged as a diagonal matrix:*

$$\mathbf{P}^k \overset{\text{def}}{=} \begin{bmatrix} p_1^k & 0, & \dots & 0 \\ 0 & p_2^k & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & p_N^k \end{bmatrix} \tag{56}$$

Proof. Follows trivially from the definition of $\mathbf{Q}(\mathbf{k}_t)$ in equation (29) and knowledge of the solution to the Bellman equation (52). □

Equation (54) shows that $\mathbf{Q}(\mathbf{k}_t)$ becomes diagonal in the steady state. Specifically, at the steady state, a small increase in the capital stock of a business unit has no effect on another firm's value. By contrast, product market spillovers due to changes in the capital stock of a business unit do affect other firms' values outside steady state so that $\mathbf{k}_t \neq \mathbb{k}$.

In addition, even at the steady-state, there are relevant second-order product market spillovers. By second-order product market spillovers we intend the effect of a change in a business unit's stock of capital on the marginal product of capital of any other firm. These second-order spillovers, which do cause peer business units to change their investment policy, are governed by the matrix $\mathbf{\Lambda}$, which by definition is the Jacobian of the firms' own marginal Q vector:

$$\frac{\partial \mathbf{q}(\mathbf{k}_t)}{\partial \mathbf{k}_t} \equiv -\mathbf{\Lambda} \tag{57}$$

By definition of matrix $\mathbf{\Lambda}$, we infer that second-order product market spillovers marginal Q relate to both for the adjustment costs of capital (captured by $\theta(\mathbf{I}-\mathbf{\Phi})$), as well as the competitive pressure firms are subject to (captured by $\mathbf{\Omega}$).

Put together, the model captures how a change in the capital stock of a rival business unit has a "first-order" effect on its own value with the off-diagonal terms of matrix $\mathbf{Q}$, defined as $q_{mn}(\mathbf{k}_t)$. Furthermore, due to decreasing returns to scale, the size of each business unit (i.e., the magnitude of its accrued stock of capital) is itself relevant in explaining the magnitude of product market spillovers. These "second-order" product market spillovers on firms' marginal products of capital are captured by the matrix $\mathbf{\Lambda}$. The result relates to the size effect also detected in the law of motion of markups in equilibrium.

2.10 Neoclassical Investment Equations

Next, we re-express our optimal investment rule in equation (25) as a vector of stacked neoclassical investment equations (Hayashi, 1982).

Proposition 7. *The vector of business units' investment policies obeys the following neoclassical equation:*

$$\boldsymbol{i}_{t+1}(\mathbf{k}_t) = \mathbf{\Delta}\mathbf{k}_t + \left[\theta\mathbf{I} + \mathbf{\Omega}(\mathbf{I}-\mathbf{\Phi})^{-1}\right]^{-1}\left[\mathbf{q}(\mathbf{k}_t) - \mathbf{q}(\mathbb{k})\right] \tag{58}$$

Proof. See Appendix A. □

The functional form of equation (58) evokes Hayashi (1982) and subsequent studies. For perfectly competitive business units, this literature establishes marginal Q as a "sufficient statistic" to predict investment. By contrast, for business units or firms subject to imperfect product market competition, Hayashi (1982) already observes a single-product firm's own marginal Q is no longer a "sufficient statistic". In our setting, to predict the investment of business unit n, it is not sufficient to know its own marginal Q - we also need to know the marginal Q of every other business unit affecting its residual demand. While there is an elegant linear mapping between the vector of corporate investment $\boldsymbol{i}_{t+1}(\mathbf{k}_t)$ and $\mathbf{q}(\mathbf{k}_t)$, product market spillovers ensure that a business unit's own marginal Q is insufficient to explain its investment rate.

2.11 Equilibrium Markups

Having solved the equilibrium investment policy, we derive an explicit expression for the endogenous markup charged by each business unit any period. Trivially, if $\operatorname{diag}(\mathbf{D}) = \mathbf{0}$ and $\boldsymbol{\Gamma} = \mathbf{I}$ we have a competitive economy where markups are zero for all business units, and thus the analysis is strictly relevant for the case in which managers do not take prices as given so that $\mathbf{D} = \mathbf{I}$.

We define $\mu_n(\mathbf{k}_t)$ as the markup charged by business unit n at state $\mathbf{k}_t$, capturing the percent difference of the output price with respect to the marginal cost:

$$\mu_n(\mathbf{k}_t) \stackrel{\text{def}}{=} \frac{p_n(\mathbf{k}_t)}{c_n} \tag{59}$$

Using the solution for the steady state vector $\mathbb{k}$ derived in Proposition 1, we derive an explicit expression for the equilibrium vector of markups or any state $\mathbf{k}_t$.

Proposition 8. *The equilibrium vector of markups associated with the unique Stable Markov Perfect Equilibrium in Proposition 1, for any state $\mathbf{k}_t$, is equal to*

$$\boldsymbol{\mu}(\mathbf{k}_t) = \boldsymbol{\mu}(\mathbb{k}) + \mathbf{C}^{-1}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}(\mathbb{k}-\mathbf{k}_t) \tag{60}$$

where the vector of equilibrium markups at the steady state is equal to:

$$\boldsymbol{\mu}(\mathbb{k}) = \mathbf{C}^{-1}\left\{\mathbf{b} + (\mathbf{I}+\boldsymbol{\Sigma})\,\mathbf{Z}\boldsymbol{\Omega}^{-1}\left[(r\,\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^k - \mathbf{Z}(\mathbf{b}-\mathbf{c})\right]\right\} \tag{61}$$

and matrix $\mathbf{C}$ is simply the vector of marginal costs rearranged as a diagonal matrix:

$$\mathbf{C} \stackrel{\text{def}}{=} \begin{bmatrix} c_1 & 0, & \dots & 0 \\ 0 & c_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & c_N \end{bmatrix} \tag{62}$$

Proof. See Appendix A. □

The proposition above provides two important insights. First, equation (60) shows that some of the cross-sectional variation in markups is transitory: a business unit that is operating below steady-state capacity will optimally reduce its markup over time as its stock of capital increases.

Second, equation (61) shows that the markup charged over variable costs will embed, over the long run, a compensation for capital costs. The matrix $\mathbf{C}^{-1}(\mathbf{I}+\mathbf{\Sigma})\mathbf{Z}\mathbf{\Omega}^{-1}$ in equation (61) captures the extent to which business units pass the cost of capital through to their customers via the markup. This result contributes to existing studies in showing that not only investment is affected by economic rents (as in Crouzet and Eberly, 2023 and others), but also how these economic rents are, in turn, affected by capital expenses in equilibrium.

2.12 Single Product Cournot Case: The Wedge between Marginal and Average Q

As a final step in our theoretical analysis, before moving onto the quantification of the model, we characterize the relation between marginal Q in equation (58) (which is empirically unobservable) and the associated market to book asset ratio (or average Q). The result wedge between both quantities has been thoroughly studied in the neoclassical literature, for practical reasons (see, e.g., Hayashi, 1982 and Hennessy, 2004). In our setting, the wedge reveals the extent to which firms internalize the impact of their investment on rival behavior and, thus, on their own continuation value.

For analytical convenience and in alignment with our quantitative application, we focus on the single-product Cournot case ($\mathbf{D}=\mathbf{\Gamma}=\mathbf{I}$), where each business unit represents a single-product firm. In this case, we define the average Q of firm n as follows:

$$\bar{q}_n(\mathbf{k}_t) \overset{\text{def}}{=} \frac{v_n(\mathbf{k}_t)}{k_{nt}} \tag{63}$$

The wedge between average and marginal Q is then given by:

$$\bar{q}_n(\mathbf{k}_t) - q_n(\mathbf{k}_t) = \frac{\rho_n}{k_{nt}} + \lambda_{nn} \cdot \frac{k_{nt}+k_{n\infty}}{k_{nt}} \cdot \frac{k_{nt}-k_{n\infty}}{k_{nt}} \tag{64}$$

Using a first-order Taylor approximation around the steady state for the second term on the right hand side, the expression then simplifies to:

$$\bar{q}_n(\mathbf{k}_t) - q_n(\mathbf{k}_t) \approx \frac{\rho_n}{k_{nt}} + \lambda_{nn}(k_{nt}-k_{n\infty}) \tag{65}$$

At any point in time, as long as firms have market power, the first term in equation (64) is positive and implies that average Q overstates marginal Q, where ρ_n captures the present value of oligopoly rents for firm n in the (long-run) steady state. In turn, the second term captures changes in oligopoly rents due to short-term adjustments in the firm's stock of capital.

More simply, at the steady state, the wedge between average and marginal Q is strictly equal to the present value of future oligopoly rents per unit of capital:

$$\bar{q}_n(\mathbb{k}) - q_n(\mathbb{k}) = \frac{\rho_n}{k_{n\infty}} \tag{66}$$

The result is consistent with the argument that average Q overstates marginal Q when firms have monopolistic power in (Hayashi, 1982). We extend the implications of this early study on monopoly firms to the general case of oligopoly, while solving explicitly for firms' endogenous markups in equilibrium.

3 Data and Calibration

In this section, we outline the data used to take the model to the data and show how to recover the model primitives. The dataset that we use spans the near-universe of US publicly-traded corporations from 1995 to 2021.

3.1 Data Sources

Compustat. We rely on two key data sources. The first is balance sheet data from Compustat, which we download from WRDS. From that database we use sales/revenues, which are mapped to revenues firm in our model ($p_i y_i$). We also use Costs of Goods Sold, which are mapped to labor costs ($c_i y_i$). Finally, we also use Capex, PPE, R&D Expenditures and Selling, General and Administrative Costs (SGA), which we use to construct a panel data of total capital stocks, as explained in the next subsection. Under the assumption that COGS per unit of output equals marginal costs of production times production, we later compute markups empirically as revenues divided by COGS, in line with previous work by De Loecker, Eeckhout, and Unger (2020).

Hoberg and Phillips (2016) Product Similarity. We follow Pellegrino (2025) and rely on the text-based product cosine similarity data of Hoberg and Phillips (2016, henceforth HP) to measure the matrix of product cosine similarities $\mathbf{A}'\mathbf{A}$. The HP dataset applies text analysis methods to the business description section (Item 1) of annual 10-K regulatory filings. For each firm-year, HP extract a vector indicating the frequency with which different words appear in the product description. Their word selection focuses on substantive terms—specifically, nouns and proper nouns that are not ubiquitous across filings—yielding a vocabulary of approximately 61,000 unique words. The similarity score for any firm pair equals the cosine of the angle between their respective word vectors, bounded between zero and one. Unlike traditional classifications such as NAICS or SIC that group firms based on common production technologies, the HP approach captures product-level competition by identifying firms that describe their offerings using similar language. Additionally, the HP data offer several practical advantages: annual updates that reflect evolving product portfolios, graduated similarity measures rather than binary industry membership, and direct derivation from legally binding disclosures that firms have strong incentives to report accurately. The dataset spans nearly all Compustat firms and is available from 1995 onward.

Validation of the HP Cosine Similarities. Pellegrino (2025) provides multiple validation tests supporting the use of HP similarity scores within the GHL demand framework. A first test examines whether firms naturally group into coherent product markets. When the HP network is visualized using standard algorithms, the resulting clusters align closely with independently defined product-market categories from S&P's Global Industry Classification system, suggesting the text data successfully capture competitive relationships at the market level. A second set of tests, originally conducted by HP and confirmed by Pellegrino (2025), compare the HP-derived industry classifications (TNIC) against conventional approaches. When benchmarked against firms' own disclosures of their main competitors in Capital IQ filings, the text-based classifications prove more accurate than traditional SIC or NAICS groupings. Most importantly for our empirical strategy, Pellegrino (2025) shows that demand elasticities derived from the GHL system—when parameterized using HP data—align well with structural estimates from prior industrial organization studies. Drawing on established work in automobiles Berry, Levinsohn, and Pakes (1995), breakfast cereals Nevo (2001), and personal computers Goeree (2008), he finds that model-implied elasticities track the patterns in these

TABLE 1: SUMMARY STATISTICS (1995-2021)

Description	Notation	Mean	Std. Dev.	p.25	p.75
Revenues (2021 US$mln)	$p_{nt}y_{nt}$	4,548	19,628	55	1,740
Costs of Goods Sold (2021 US$mln)	$c_{nt}y_{nt}$	3,016	14,658	23	1,032
Market Value of Assets (2021 US$mln)	v_{nt}	18,835	124,043	152	4,722
Total Capital Stock (2021 US$mln)	k_{nt}	4,928	22,034	68	1,689
Price of Capital	p_n^k	1.177	0.100	1.125	1.238
Depreciation Rate	δ_n	0.154	0.048	0.130	0.188
Investment Rate	i_{nt}/k_{nt}	0.370	33.398	0.149	0.319
Average Q (log)	$\log \bar{q}_{nt}$	0.745	1.469	-0.201	1.497
Capital Requirement (log)	$\log z_n$	-5.282	1.587	-6.328	-4.596

microeconometric benchmarks. Collectively, these validation exercises justify our interpretation of HP cosine scores as empirical analogs to the theoretical product similarity matrix $\mathbf{A}'\mathbf{A}$.

In every year, our sample is composed of firms that appear in both Compustat and HP's database and, either have no missing data from the original Compustat dataset (including capital for the previous and the next year) or variables have been reconciled from accounting equations and linear interpolation. Moreover, there are no outliers (defined as firms that experience a change in the capital stock by a factor of over 2). This results in a sample size of 5,719 in 1995 and 3,202 in 2021.

3.2 Construction of the Total Capital Panel Dataset

Our measured state variable, the stock of capital k_{nt}, combines both tangible and intangible capital. To implement our model empirically, we build a new database of total capital stocks for US publicly-traded firms, building on the foundational methodologies of Peters and Taylor (2017, henceforth PT) and Eisfeldt, Kim, and Papanikolaou (2022, henceforth EKP). Specifically, we measure total capital as the sum of three components: knowledge capital (accumulated R&D spending), organization capital (accumulated SG&A spending), and physical capital (property, plant, and equipment). For each component, we apply the perpetual inventory method using industry-specific depreciation rates from the Bureau of Economic Analysis, combined with component-specific price deflators from multiple sources including the BEA, FRED, and EUKLEMS & INTANProd databases. For knowledge capital, we use industry-specific BEA depreciation rates, with a default depreciation rate of 15% for industries without coverage, while for organization capital, we use a depreciation rate of 20% for all firms. For physical capital, we use industry-specific BEA depreciation rates that vary across equipment and structures.

A key innovation in our approach is the construction of firm-specific chain price indexes for total capital. Rather than deflating all capital components with a single economy-wide deflator—as is common in the literature—we construct separate price indexes for each capital component and then aggregate them using time-varying, firm-specific weights based on each component's nominal share in the firm's total capital stock. For knowledge capital, we use the BEA's R&D price deflator. For organization capital, we use price indexes from EUKLEMS & INTANProd, spliced with the BEA's intellectual property products (IPP) deflator to extend coverage back to 1950 and forward to 2021.

For physical capital, we construct a chain-weighted index from separate BEA price deflators for equipment and structures, using the relative investment shares as weights. This approach allows us to account for the fact that firms with different capital compositions face different effective price changes over time, which is particularly important given the divergent price trends across tangible and intangible assets over our sample period.

Our database builds on PT and EKP with several refinements suited to our application. Relative to PT, we include 100% of SG&A instead of 30%. As explained by EKP, this choice is motivated by the lack of a reliable calibration for this parameter. Treating it in this way also allows the measure to capture other important forms of intangible capital, such as organizational or brand capital. Additionally, while PT use a single Tobin's Q deflator for all types of capital, our firm-specific chain price indexes capture heterogeneity in price exposure across firms. Relative to EKP, our primary refinement lies in the aggregation approach: whereas they apply component-specific deflators and then aggregate in real terms, our chain-weighting method more closely follows national accounting conventions and allows for changing composition effects over time. We provide additional details on data sources and construction in Appendix B.

Table 1 reports the summary statistics of our working sample before merging with the HP data. The distributions of firm's stock of capital, investment and depreciation rates are consistent with those reported in the alternative working samples studied by PT and EKP. The log values of firms' productivities of capital ($\log z_n$) are inferred from our own calibration exercise, on which we elaborate next.

3.3 Calibration

We now describe our calibration strategy for the model parameters. We begin by considering externally calibrated parameters that are either standard in the literature or directly observable from the data, then turn to the calibration of productivity ($\mathbf{z}$), adjustment costs (θ), and the steady-state capital distribution ($\mathbb{k}$), which require careful treatment due to the high degree of dimensionality of our parameter space is ($\mathbf{z}$ has thousands of entries) and the even larger number of equality and inequality restrictions imposed by the model, which imply our model is heavily over-identified.

Externally calibrated parameters. We set the utility weight on common characteristics to $\alpha = 0.12$, following Pellegrino (2025), who calibrates this parameter by matching model-implied demand elasticities to micro econometric estimates from the industrial organization literature. Firm-specific depreciation rates δ_n are computed as the time-varying weighted average of component-specific BEA depreciation rates, as described in Section 3.2. The discount rate r is year-specific and set equal to the estimate by Gormsen and Huber (2025).

Set-identification and point-calibration of the adjustment cost parameter θ. The curvature parameter θ governs the convexity of adjustment costs. We identify an upper bound for θ by imposing two economically motivated restrictions. First, the implied distribution of $z_{nt}^2 k_{nt}$ must have a positive mean. This yields an upper bound:

$$\bar{\theta}_1 = \frac{\sum_{n=1}^{N} \sum_{t \in \mathcal{T}_n} \pi_{nt}/k_{nt} - (r + \delta_{nt}) p_{nt}^{k}}{\sum_{n=1}^{N} \sum_{t \in \mathcal{T}_n} (1+r)(k_{nt} - k_{nt-1}) - (k_{nt+1} - k_{nt})} \tag{67}$$

where $\mathcal{T}_n$ is the set of years where data for firm n is available. Second, we require that adjustment costs cannot exceed operating profits net of the user cost of capital, as this would imply firms have no economic reason to exist, which yields a second upper bound:

$$\bar{\theta}_2 = \frac{\sum_{n=1}^{N} \sum_{t \in \mathcal{T}_n} \pi_{nt} - (r + \delta_{nt}) p^k_{nt} k_{nt}}{\sum_{n=1}^{N} \sum_{t \in \mathcal{T}_n} \frac{1}{2} (k_{nt} - k_{nt-1})^2} \tag{68}$$

In practice, the second constraint is the binding one in our data, therefore it must be that $\theta \in \left(0, \overline{\theta}_2\right)$. In our setting, $\overline{\theta}_2 = 1.145 \times 10^{-4}$.

Within this set-identified interval, we specify a point estimate of θ by matching the aggregate magnitude of adjustment costs to the empirical benchmark of Cooper and Haltiwanger (2006), who structurally estimate adjustment costs using plant-level manufacturing data. The authors' preferred specification estimates adjustment costs to be approximately 3.1% of operating profits. We calibrate θ such that aggregate adjustment costs in our model equal 3.1% of aggregate operating profits. The calibration yields $\theta = 5.808 \times 10^{-5}$, which falls comfortably within the set-identified bounds. The approach ensures that our adjustment cost parameter both satisfies economically motivated restrictions and aligns with credible estimates from the neoclassical investment literature.

Estimation of capital productivity z_n. The key identification challenge that we face is recovering the high (N)-dimensional capital productivity parameter vector $\mathbf{z}$ (where N is the number of firm-level observations in each cross section of the working sample, 5,719 in 1995 and 3,202 in 2021). For each firm in the working sample, we notice that the productivity parameter z_n appears squared in the first-order condition (32). Rearranging this condition, we can express:

$$z_n^2 = \frac{\pi_{nt}}{k_{nt}^2} - (r + \delta_n) \frac{p_n^k}{k_{nt}} - \theta \left[(1 + r) \frac{k_{nt} - k_{nt-1}}{k_{nt}} - \frac{k_{nt+1} - k_{nt}}{k_{nt}} \right] \tag{69}$$

In principle, equation (32) would allow us, conditional on a calibrated value of θ, to infer z_n^2 from observable data. However, because our model is over-identified (it has more equality and inequality restrictions than parameters), it cannot perfectly fit equation (69). As a result, its right-hand side can take negative values in the data. These negative values are inadmissible since z_n^2 must be positive by construction.

To extract a robust cross-section of firm-level productivities, we employ a statistical approach, and exploit the well-established regularity that productivity distributions exhibit heavy right tails that are well-approximated by Pareto distributions (Gabaix, 2009; Luttmer, 2007). This empirical fact holds robustly for the noisy estimate that we can recover from equation (69). We therefore assume that the underlying cross-section of true capital productivity z_n follows a Pareto distribution. Hence, z_n^2 also follows a Pareto distribution: we obtain the corresponding empirical estimate of z_n by fitting a Pareto distribution to the empirical noisy estimate of z_n^2 obtained from equation (69). After fitting the Pareto distribution, we rescale the entire productivity vector by a scalar factor chosen to match the aggregate sum of model-predicted firm values to the aggregate sum of observed market values. We provide additional details of these steps in Appendix (B).

Enforcement of positive steady-state capital. While $\mathbb{k}$ is an equilibrium object rather than a primitive, we must ensure it satisfies $\mathbb{k} > 0$: this condition is both intuitive (we expect firms that we observe in the data to remain active in the long-term) as well as a requirement for tractability (taking the model to the data requires that the data generating process be an interior equilibrium). Yet, the fact that our model is over-identified, and that (despite our best efforts) capital stocks will

Figure 2: Empirical Investment Spillover Network

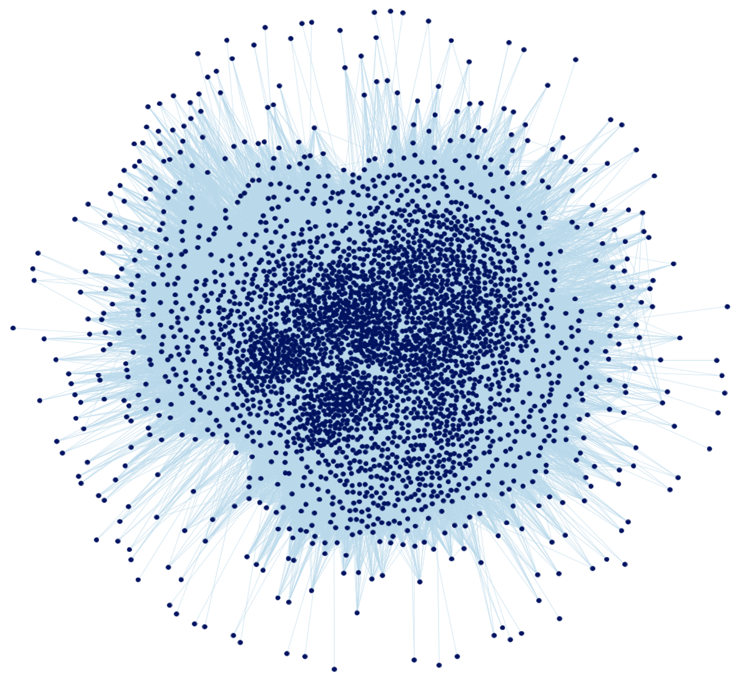

always be measured imperfectly, leads to the direct estimate of $\Bbbk$ from the transition equation (22) to be negative for a small subset of firms (<10%). To enforce positivity while preserving the cross-sectional distribution, we fit a Gompertz distribution[9] to the predicted ratio $k_{n\infty}/k_{nt+1}$ and use it to estimate positive steady-state values. We provide additional details in Appendix B.

Recovery of demand and cost parameters. Given $\mathbf{z}$, θ, and $\Bbbk$, we recover the remaining parameters directly from the data and equilibrium conditions. Marginal costs are set to $c_n = \text{COGS}_n/(y_{nt})$, assuming all reported operating costs are variable. The demand intercept vector $\mathbf{b}$ is recovered by inverting the steady-state equation (37).

3.4 Estimated Transition Matrix

Figure 2 displays the matrix of investment spillovers (equation (53)) or, equivalently, the off-diagonal entries of the transition matrix $\mathbf{\Phi}$, as a weighted network. Each edge represents the magnitude of investment spillovers between firm pairs. Specifically, we plot edges corresponding to the top 320,000 entries of $-\mathbf{\Phi}$ (multiplied by negative one so that competitive spillovers—which dominate empirically—appear as positive weights), with edge thickness proportional to spillover strength. The network exhibits a characteristic core-periphery structure. Central firms, positioned densely in the interior, experience strong bilateral interactions with numerous other firms, reflecting either intense product market competition with many rivals or high capital productivity z_n that reduces capital adjustment frictions and amplifies responsiveness to competitive conditions. Peripheral firms, scat-

[9] We use a Gompertz distribution because it matches the empirical properties of the distribution of $\Bbbk$, which is negatively skewed and in logs has a fat left tail, while maintaining non-negative support. As a robustness check, we also fit a Weibull distribution using the same moment-matching procedure; results are virtually identical.

FIGURE 3: EMPIRICAL TRANSITION MATRIX: EIGENVALUES

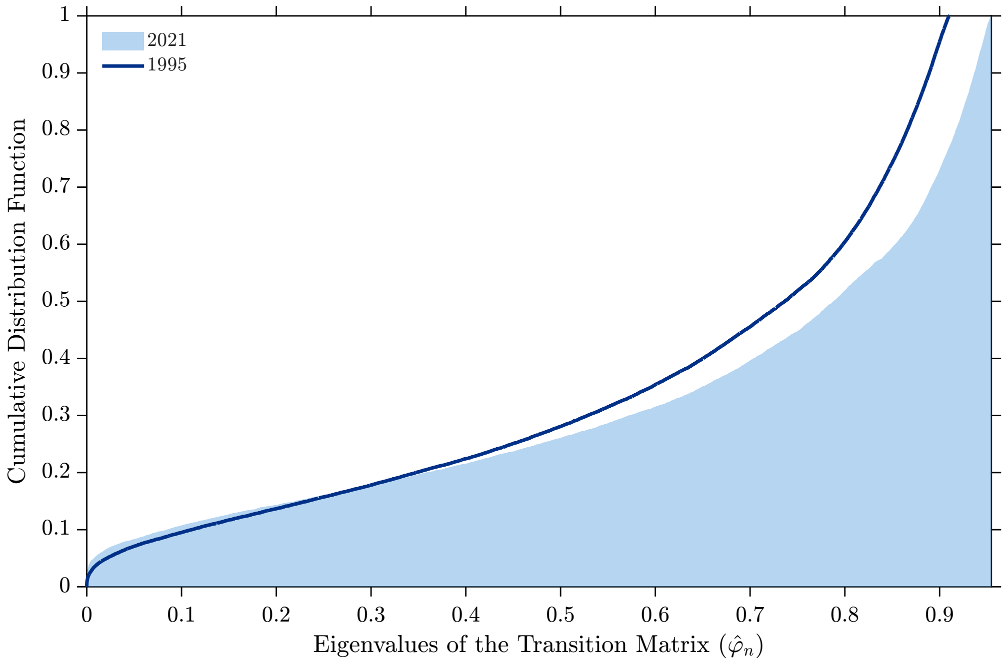

tered along the exterior with sparse connections, exhibit weak investment spillovers either because they occupy isolated niches in the product market network, or because low capital productivity generates substantial adjustment costs that dampen their investment response to competitive pressures.

The network displays a moderate degree of clustering, suggesting that investment spillovers organize firms into loosely defined groups sharing similar competitive dynamics. This cluster structure proves less pronounced than the one observed in the underlying Hoberg-Phillips product similarity network, as documented in Pellegrino (2025). This attenuation reflects the fact that $\mathbf{\Phi}$ incorporates not only product market similarity—captured by the substitution matrix $\mathbf{\Sigma}$—but also firm-level heterogeneity in capital productivities z_n, which introduces cross-sectional variation in adjustment costs and thereby modulates the mapping from product similarity to investment spillovers. Firms with low z_n exhibit weak investment responses even to closely positioned competitors, whereas firms with high z_n respond vigorously to product market conditions. This heterogeneity obscures the clean industry boundaries visible in pure product similarity networks, yielding a transition matrix that blends product market structure with technological constraints on capital adjustment.

Figure 3 presents the cumulative distribution of eigenvalues $\hat{\varphi}_n$ of the estimated transition matrix $\mathbf{\Phi}$ for 1995 and 2021. These eigenvalues characterize the persistence of firm's capital stock along the balanced growth path: values near zero indicate rapid convergence to the steady state, whereas values approaching unity imply instead slow adjustment. We observe substantial mass at both extremes of the cumulative distribution: many eigenvalues cluster near zero (indicating rapid convergence along those modes) while others concentrate near 70% (indicating much slower adjustment).

The large variation in eigenvalues implies that predicting the speed of convergence for capital on aggregate is inherently difficult. The actual adjustment speed depends on the current capital distribution ($\mathbf{k}_t - \mathbb{k}$): if the economy happens to be displaced primarily along fast-adjusting eigenvector directions (associated with firms with high productivity or high product market centrality), convergence will be rapid; if the displacement aligns with slow-adjusting modes (associated with firms with low productivity or high market power), aggregate capital will be stickier. This state-dependence distinguishes our network model from simpler frameworks where convergence rates are uniform across all initial conditions, and highlights that both productivity heterogeneity and network positioning jointly determine the economy's response to shocks.

Another fact that stands out from Figure 3 is that the 2021 distribution (approximately) first-order stochastically dominates the 1995 distribution. The rightward shift indicates that capital stocks have become more persistent over time. The increased persistence reflects rising capital dependence across firms: firms in 2021 operate capital more intensively than their 1995 counterparts (due to a reduction in z_n), generating larger eigenvalues ω_n in the matrix $\mathbf{\Omega}$, and consequently higher persistence in firms' capital stocks $\hat{\varphi}_n$.

3.5 Model Fit

Next, we assess the model's performance by evaluating its fit against the firm-level data used in calibration. In order to do so, we must start by noting that our structural model is over-identified: it has more structural equation restrictions (equalities and inequalities) than degrees of freedom. Given observed revenues, investment, and capital for each firm-year, perfectly matching both observed revenues and investment rates would require inferring, for some firms, negative productivity ($z_n < 0$) or steady-state capital ($k_{n\infty} < 0$). Such values would be recovered from the structural equations. After we impose the economically necessary restrictions $z_n > 0$ and $\mathbb{k} > 0$, it becomes clear that (by construction) the model cannot simultaneously fit all observed data perfectly. This is a feature, not a bug. Specifically, the variables that the model won't match perfectly are revenues and investment rates.

Figure 4 displays three key validation exercises, plotting model-predicted values against their empirical counterparts for all firm-year observations in our sample. The first two panels examine the model's ability to match targeted moments. The top panel shows the revenues-to-capital ratio ($p_{nt}y_{nt}/k_{nt}$), achieving a correlation of ρ =0.92 between model and data. The middle panel displays the investment rate (i_{nt}/k_{nt}), with a correlation of ρ =0.52.

The bottom panel provides an important validation using the cross-sectional distribution of average Q, defined as the market value of the firm's assets divided by its capital stock (v_{nt}/k_{nt}). Average Q is an untargeted quantity in our calibration—the model does use this moment in recovering any parameter. Nevertheless, the model achieves a correlation of ρ =0.63 with the empirical average Q, demonstrating that the calibrated network structure and productivity distribution successfully capture firms' market valuations. The result indicates that the model's structural predictions extend beyond the specific moments used in calibration, capturing fundamental aspects of firm valuation that emerge from the underlying economic forces in the model.

For reference, the goodness of fit captured in the correlations associated with both investment and average Q in Figure 4 also indicate high explanatory power compared to that of estimated investment equations in the existing neoclassical literature. For instance, the linear regression in the seminal paper by Hayashi (1982) fits the investment to capital ratio up to an R-square of 0.46, whereas more recent work by Peters and Taylor (2017) incorporating intangible capital and correcting for measurement error report R-squares ranging from 0.233 to 0.374.

Figure 4: Model Fit

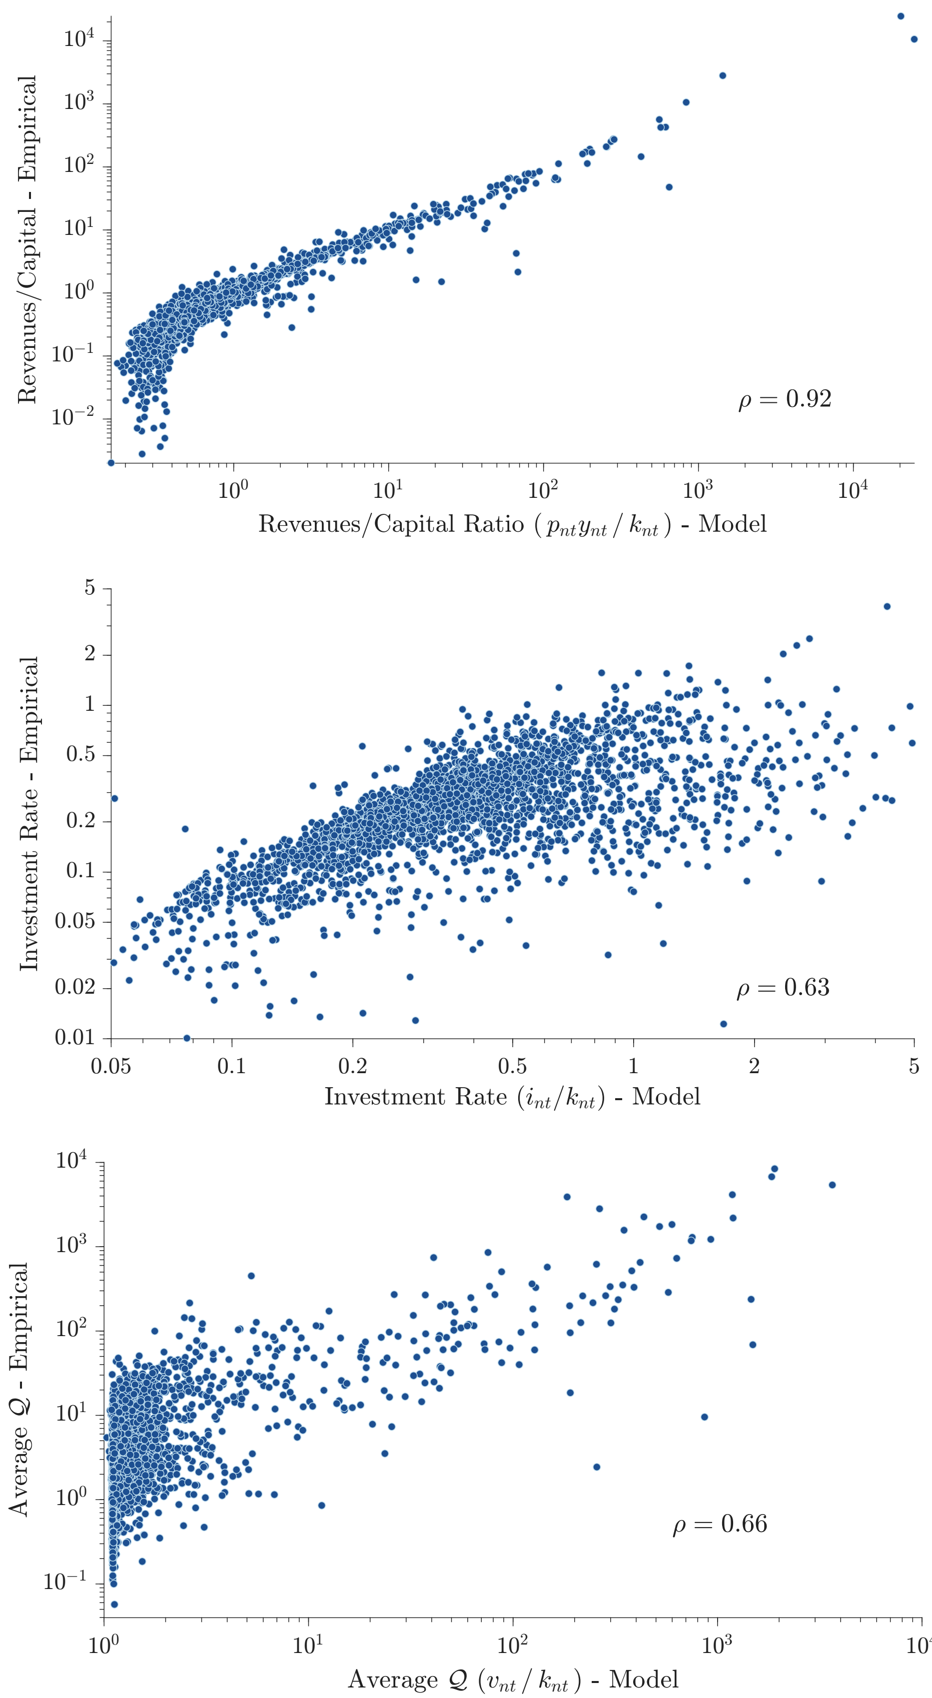

4 Model Quantification

In this section, we bring our calibrated model to bear on a series of quantitative exercises. Our baseline implementation assumes simple Cournot competition, where each firm in our Compustat sample corresponds to an independent business unit that internalizes its own downward-sloping residual demand but not its product market spillovers on rivals. We begin by analyzing the properties of this factual equilibrium, and then turn to counterfactual exercises that illustrate the model's capacity for policy-relevant analysis.

4.1 Analysis of the Factual Equilibrium

We begin our quantitative analysis by taking the model to the data in 2021 and examining the properties of the calibrated factual Cournot equilibrium. We focus on three sets of results: the aggregate and firm-level transition dynamics of capital, the relationship between average and marginal Q, and the decomposition of steady-state enterprise values.

Transition Dynamics. We begin our quantitative analysis by examining the equilibrium dynamics of the calibrated model. Figure 5 displays the evolution of aggregate capital stock in the US economy over a 15-year forecast horizon, starting from the observed 2021 cross-section. The model predicts a smooth increase from approximately \$20.82 trillion in 2021 to \$29.82 trillion by 2035, representing aggregate capital growth of roughly \$8.99 trillion over the period. This trajectory reflects the collective adjustment of thousands of individual firms as they converge toward their steady-state capital levels, each responding to their own capital intensity, adjustment costs, and competitive environment.

The slightly concave shape of the aggregate capital path—with faster growth in earlier years that gradually decelerates—is characteristic of convergence dynamics in our model. Starting from 2021, when many firms are displaced from their long-run equilibrium positions, the economy experiences relatively rapid capital accumulation as firms close the gap to steady state. As firms approach their optimal capital stocks $\mathbb{k}$, the investment rate naturally slows, even though adjustment continues.

The annual aggregate growth rate implied by the model—approximately 2.88%—reflects both the fundamental forces driving capital accumulation (capital productivity, demand conditions, and competitive interactions) and the frictions that slow adjustment (convex adjustment costs parameterized by θ).

Figure 6 reveals the firm-level heterogeneity masked by the smooth aggregate trajectory shown in Figure 5. The fan chart displays the cross-sectional distribution of capital deviations from steady state, measured as percentage gaps. In 2021, firms exhibit substantial dispersion, with the 5th percentile approximately -73.18% below steady state and the 95th percentile roughly 35.36% above. The distribution is notably asymmetric: more firms operate below their optimal capital levels than above, indicating that under-capitalization is more prevalent than over-capitalization in the initial cross-section. The median firm itself starts approximately -37.65% below its long-run equilibrium.

Over the forecast horizon, the model predicts gradual convergence towards the steady state for the entire distribution. The percentile bands systematically narrow and drift toward zero, reflecting firms' heterogeneous investment responses to their initial capital gaps. Firms substantially below steady state invest heavily to close the distance, whereas those above their steady state allow capital to depreciate faster than they replace it. By 2030, the interquartile range (40th-60th percentile band) has tightened considerably, spanning only about $\pm$13.68% around steady state.

Figure 5: Transition Dynamics of the Aggregate Capital Stock

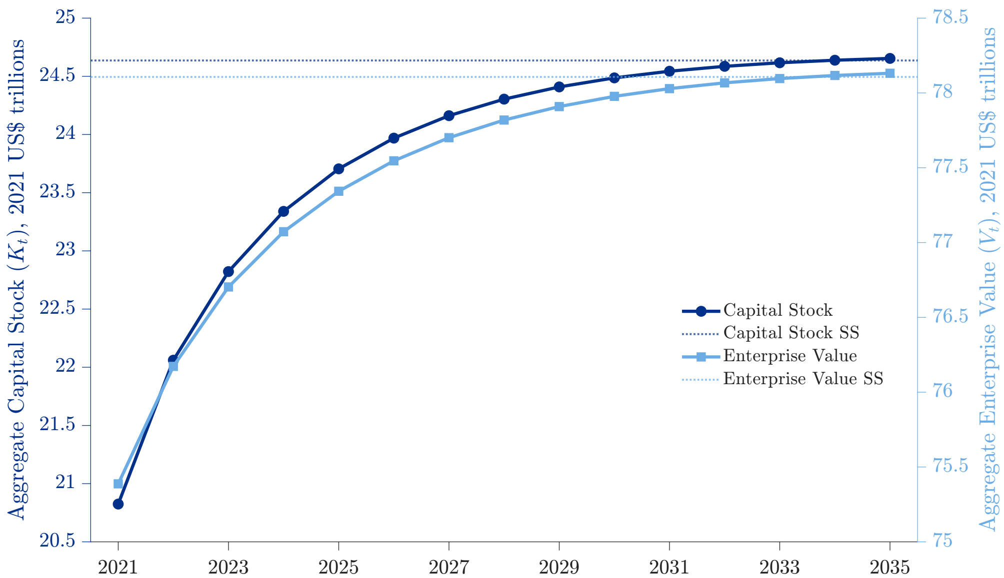

The heterogeneous adjustment speeds visible in Figure 6 directly reflect the eigenvalue distribution analyzed in Figure 3. Firms whose current deviations align with fast-adjusting eigenvector directions (corresponding to small eigenvalues of $\mathbf{\Phi}$) converge rapidly, while those aligned with slow-adjusting modes take much longer to reach steady state. Overall, the graph reveals an economy characterized by rich transitional dynamics, where firms that start with a stock of capital above their steady state tend to adjust their capital more slowly.

Average and Marginal Q. Figure 7 plots the model-implied firm-level estimates of average Q as of 2021 against the corresponding estimates for marginal Q. In spite the relatively high correlation of 0.89, the fit between both quantities masks substantial heterogeneity in the cross section. The wedge between average and marginal Q is such that most observations lie strictly below the 45-degree line, confirming that monopoly rents significantly contribute to enterprise values for most firms as of 2021. Further validating the workings of the model, firms associated with relatively larger dots in the plot (and, thus, with higher market capitalization) are typically further away from the 45-degree line: a higher wedge between average and marginal Q is naturally associated with higher enterprise value, driven by with higher markups and, therefore, higher market power.

Firm Value Decomposition. A recent literature has explored the link between the stock markets and market power, suggesting a link between the rising market power of US publicly traded firms and their rising valuations (see Eggertsson, Robbins, and Wold, 2021; Cho, Grotteria, Kremens, and Kung, 2023). Because our model features endogenous markups and enterprise values, it can help inform this debate. In particular, Equation (48) provides a useful decomposition of firm n's steady state market value in terms of two components: the replacement value of capital ($p_n^k k_{n\infty}$) and the present value of monopoly rents (ρ_n). We implement this decomposition in Figure 8, which depicts the aggregate market value of all firms in our sample for 1995 and 2021 as two pie charts.

The first thing that we can observe from the figure is that the absolute scale of firm value in steady state grew dramatically, from \$33.82 trillion in 1995 to \$69.10 trillion in 2021. This increase reflects

FIGURE 6: FIRM-LEVEL TRANSITION DYNAMICS

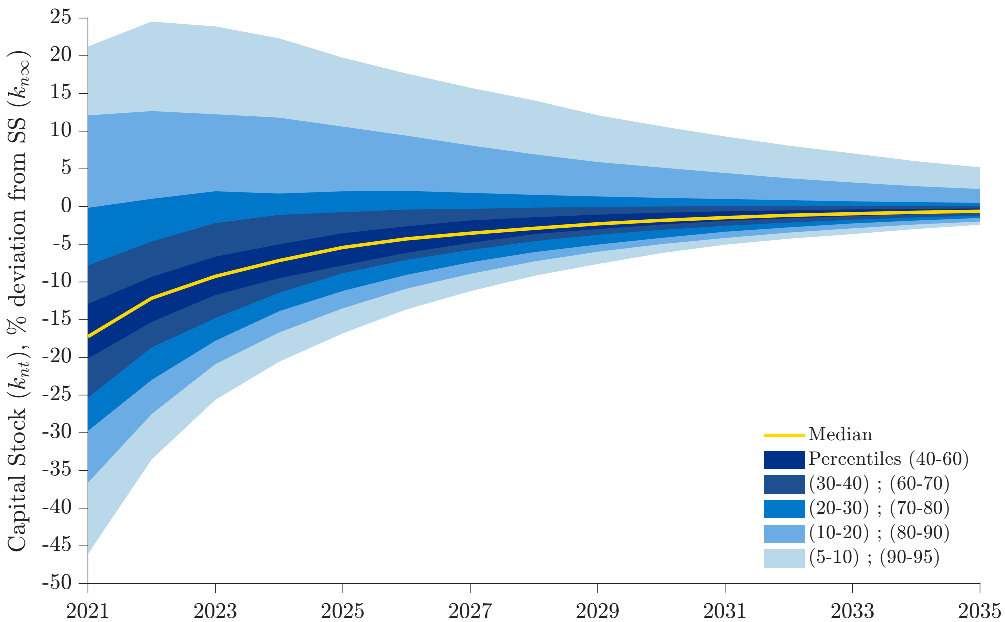

both capital accumulation as well as the increase in monopoly rents associated with rising market concentration.

Next, we consider the shares of the two component. In 1995, the replacement value of capital accounted for 53.65% of total enterprise value in steady state, whereas monopoly rents comprised 46.35%. By 2021, this composition changed drastically: the replacement value of capital decreased to just 47.72% of firm value in steady state, whereas monopoly rents have risen to 52.28%.

In sum, three interesting facts emerge from this analysis: 1) monopoly rents account for a significant fraction of firms' market values; 2) monopoly rents are key to explaining the rising dollar valuations of US firms; 3) the relative shares of these two components have increased significantly over time.

Our findings from Figure 8 are therefore consistent with the recent findings of Cho et al. (2023), who find that markups account for a large share of rising US stock market valuations. As a remark, however, the market power measure used by Cho et al. (2023) is the (short-term) markup over variable costs, which excludes compensation for fixed/dynamic inputs such as capital. By contrast, our measure of markups accounts for firms' long term capital expenses as well. Because they find expected markups to correlate with intangible investments (which are included in our definition of total capital), both of our findings can be rationalized by rising markups providing compensation for accumulated past investments, particularly intangible capital, which during our sample period accounts for an increasing share of the aggregate capital stock (Eisfeldt and Papanikolaou, 2014; Crouzet and Eberly, 2021).

4.2 Counterfactual Analysis

In the next subsection, we exploit the counterfactual-building capabilities of our model to provide additional insights on the role of product market competition, discount rate shocks, and M&A activities. Before we proceed to do so, we provide a technical clarification on the practical implementation

FIGURE 7: AVERAGE Q AND MARGINAL Q

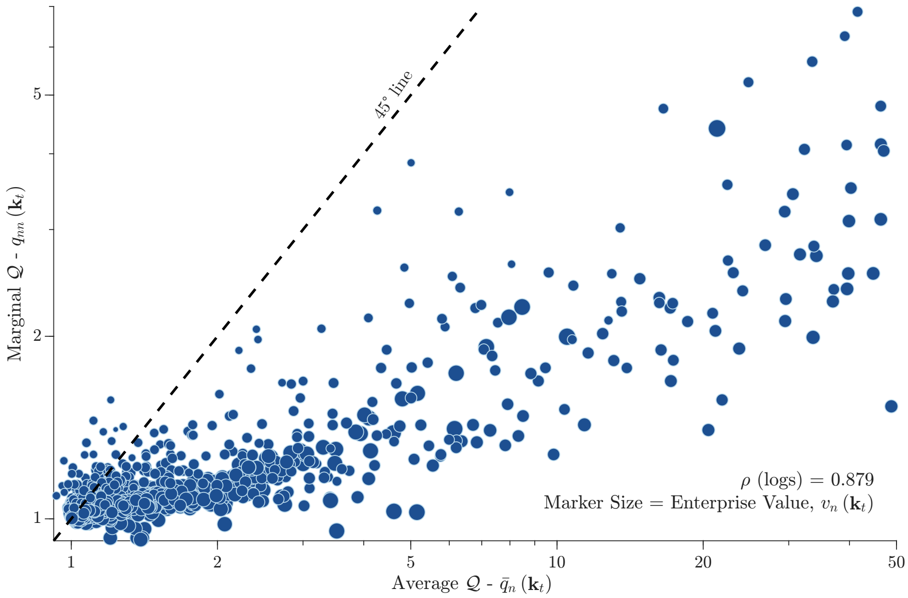

of the counterfactual.

How to Implement Counterfactuals with Non-negativity Constraints. One minor challenge that must be dealt with when performing counterfactual exercises (mergers, changes in r or marginal costs, etc.), is that the interior steady state implied by Proposition 1 does not preclude capital stocks from being negative for some business units. To perform counterfactuals, we develop a method to enforce a non-negativity constraint on $\mathbb{k}$. To build intuition on such method, we notice that the steady state presented in equation (37) corresponds to the Nash equilibrium of the corresponding *static* game obtained by shutting down adjustment costs (i.e., setting $\theta = 0$) - specifically, it solves:

$$k_{n\infty} \equiv \underset{k_n}{\text{argmax}}\ z_n k_n (b_n - c_n) - (r + \delta_n) p_n^k - \frac{1}{2} \sum_{m \neq n} z_m z_n \sigma_{mn} k_m k_n \tag{70}$$

As discussed in Pellegrino (2025), this static game is a type of *Potential* game, which can be represented by a scalar function $\mathscr{P}(\mathbf{k})$ or *Potential*, defined as follows:

$$\mathscr{P}(\mathbf{k}) \stackrel{\text{def}}{=} \mathbf{k}' \left[\mathbf{Z}(\mathbf{b} - \mathbf{c}) - (r\mathbf{I} + \mathbf{\Delta}) \mathbf{p}^k \right] - \frac{1}{2} \mathbf{k}' \mathbf{\Omega} \mathbf{k} \tag{71}$$

Regardless of whether a non-negativity constraint applies to $\mathbf{k}$, the solution of the static game (70) is also the maximizer of the potential function (71), as the Kuhn-Tucker Conditions are the same.

To compute counterfactuals, we can therefore find the steady state capital stock by maximizing the potential function $\mathscr{P}(\mathbf{k})$, with all primitives ($\mathbf{\Omega}$, $\mathbf{b}$, $\mathbf{c}$, etc...) evaluated according to the counterfactual, and subject to the non-negativity constraint $\mathbf{k} \geq \mathbf{0}$. While this constrained steady state does not generally admit a closed-form expression, it is a convex quadratic program with linear constraints; in our implementation we solve it using `quadprog` in MATLAB: it completes the optimization in a

FIGURE 8: DECOMPOSITION OF STEADY STATE AGGREGATE FIRM VALUE

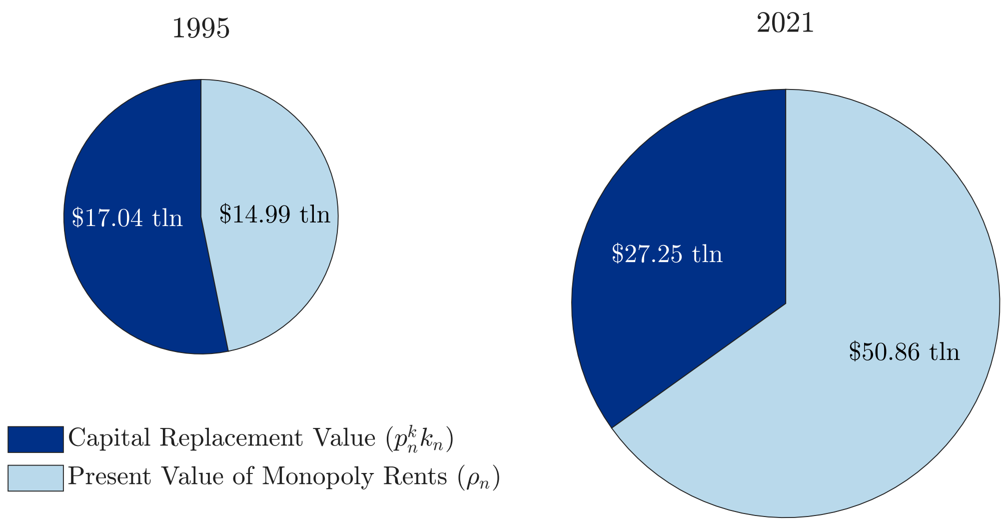

few seconds on an ordinary laptop even when the number of firms is in the thousands. This makes it straightforward to compute the counterfactual $\Bbbk$, while ensuring its non-negativity. A separate discussion is required for the computation of the counterfactual transition dynamics: here, we present two possible ways to proceed.

Active Set implementation. The first method relies on restricting the oligopoly game to subset of firms that are guaranteed to be active in the steady state counterfactual. Specifically, let $\mathscr{A} \equiv \{n : k_{n,\infty} > 0\}$ be the set of firms that operate with positive capital stock in the steady state counterfactual that solves (70). We restrict the counterfactual economy to $\mathscr{A}$ (retain the $\mathscr{A} \times \mathscr{A}$ sub-matrices of the primitives and set $k_{nt} = 0$ for all $n \notin \mathscr{A}$ and all t), and compute the transition matrix on the reduced system using the analytical interior formula (38). The advantage of this approach is internal consistency: conditional on the active set $\mathscr{A}$, the interior solution in equation (37) holds exactly and non-negativity is enforced by construction. The drawback is that extensive-margin adjustment is "front-loaded": firms with $k_{n,\infty} = 0$ are treated as inactive from $t = 0$, ruling out gradual exit dynamics.

Full-set (interior-$\mathbf{\Phi}$) implementation. An alternative solution is to keep all firms in the state vector, use the constrained $\Bbbk$ from the static quadratic programming (allowing some components to be zero), and compute $\mathbf{\Phi}$ from the interior formula (38) on the full system. The main advantage of this solution is that it preserves the full network of interactions and yields smooth dynamics; the obvious drawback is that it is best interpreted as a heuristic approximation: the transition dynamics so computed do not correspond exactly to the genuine MPE of the counterfactual because the law of motion consistent with equations (22) and (38) needs not satisfy the Kuhn–Tucker conditions of the constrained dynamic game.

In order to make the counterfactuals as comparable as possible to the factual equilibrium, we implement the second method (the "full set" implementation) in computing our counterfactuals. This specific choice does not materially affect any of our conclusions.

Counterfactual 1: Competition and Investment. A central question in finance and macroeconomics is how important is product market competition for capital accumulation. While the liter-

FIGURE 9: ANTICOMPETITIVE COUNTERFACTUAL: AGGREGATE CAPITAL STOCK

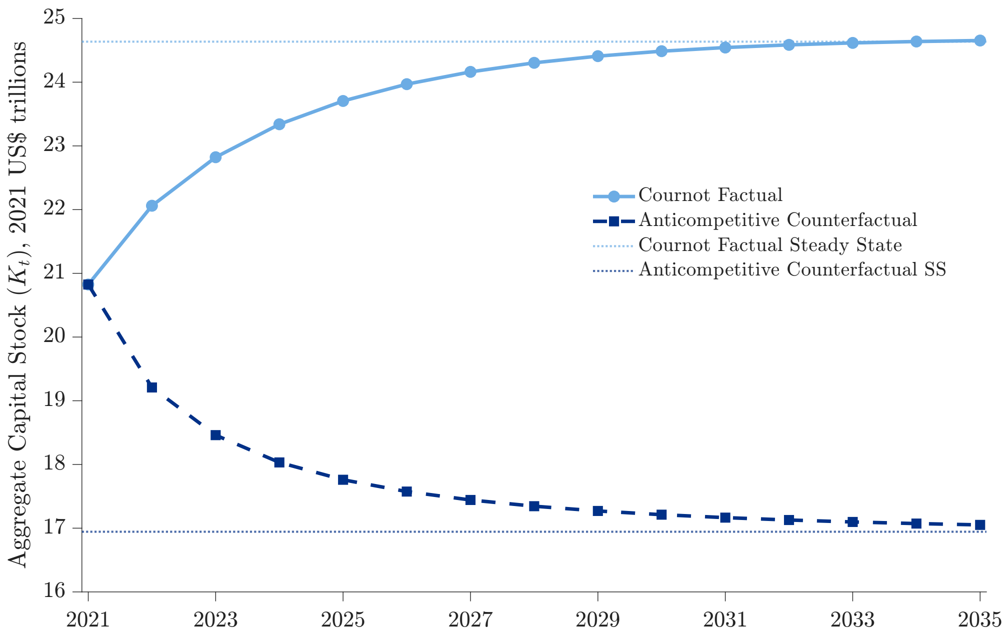

ature has long recognized that market structure affects static allocative efficiency, quantifying the dynamic effects of competition on aggregate investment has proven challenging. In canonical macroeconomic models featuring monopolistic competition or perfect competition, the concept of "shutting down competition" lacks a clear operational definition for counterfactual analysis. Our framework's explicit parameterization of competitive conduct—through the matrices $\mathbf{D}$ and $\mathbf{\Gamma}$, which govern how much firms internalize product market spillovers—renders this question tractable. We can precisely compare alternative competitive regimes while holding constant all other economic primitives.

We quantify the impact of product market competition on corporate investment by comparing the factual Cournot equilibrium to an anti-competitive counterfactual in which all firms fully internalize product market spillovers. Under the Cournot factual equilibrium, firms internalize their downward-sloping residual demand curves but do not account for strategic spillovers on rival firms. In our anti-competitive counterfactual all business units operate as a single cartel: we set $\mathbf{\Gamma} = \mathbf{1}\mathbf{1}'$ such that the network effects matrix becomes $\mathbf{\Omega} = 2\mathbf{Z}(\mathbf{I}+\mathbf{\Sigma})\mathbf{Z}$. Under this regime, firms coordinate their investment decisions to maximize joint profits, fully internalizing the negative externalities their capital accumulation imposes on competitors through increased output and lower equilibrium prices. The answer to our central question is economically significant: moving from Cournot competition to full cartelization reduces aggregate steady-state capital by 30.31%.

Figure 9 displays the aggregate implications. Commencing from the observed 2021 capital stock of \$20.82 trillion, the factual Cournot economy converges monotonically to its steady state of \$29.82 trillion over approximately 15 years. The anti-competitive counterfactual, in contrast, exhibits a pronounced contraction, with aggregate capital declining by 30.31% to \$20.78 trillion in steady state. This substantial reduction reflects the coordinated restriction of capacity to sustain elevated markups—the canonical cartel incentive to limit supply extends naturally to dynamic investment decisions. The transition path reveals that this adjustment occurs relatively rapidly, with the bulk of the conver-

FIGURE 10: ANTICOMPETITIVE COUNTERFACTUAL: CROSS-SECTION

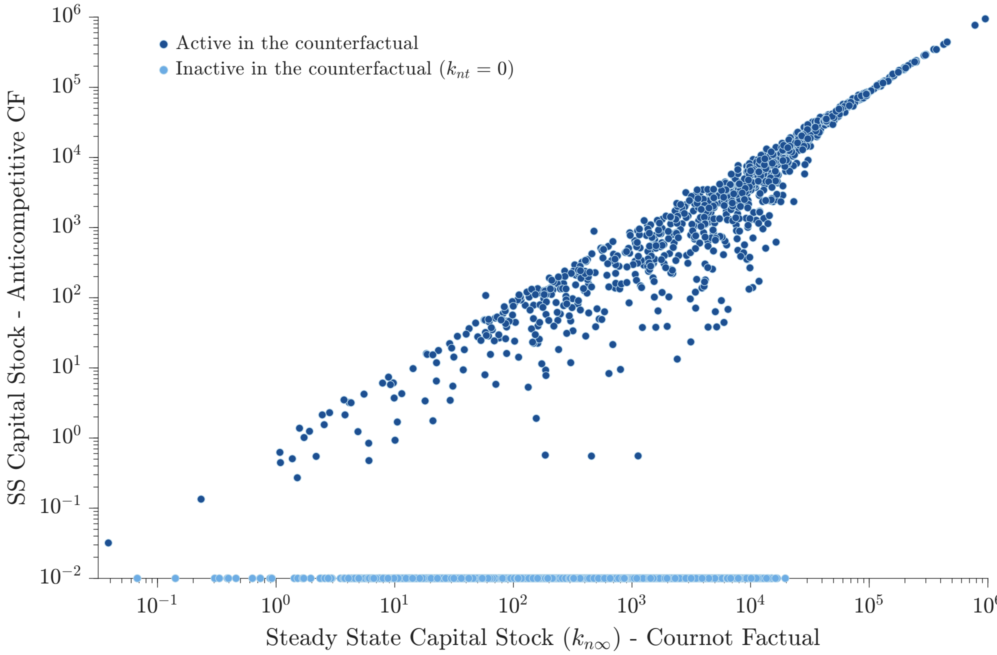

gence completed within 5 years.

The cross-sectional heterogeneity in firm-level responses, illustrated in Figure 10, reveals important distributional implications. The figure plots each firm's steady-state capital in the counterfactual against its factual steady state, with observations below the 45-degree line indicating investment reductions. Nearly all active firms reduce their capital stocks substantially under cartel coordination, though the magnitude of adjustment varies considerably across the size distribution. Notably, a non-trivial fraction of firms—depicted in light blue—exit entirely in the counterfactual steady state, with their optimal capital stocks falling to zero. These exits occur predominantly among smaller firms with high product similarity to larger competitors, for whom the cartel finds it optimal to consolidate capacity into fewer, larger producers. This pattern indicates that anti-competitive coordination not only suppresses aggregate investment but also fundamentally reshapes market structure toward greater concentration.

Figure 11 illustrates the markup consequences of this investment reallocation. The distribution of markups in 2021 (light blue shaded region) exhibits moderate central tendency, with a median markup of approximately 1.98. The factual steady state (dashed line) displays modest compression towards a median markup of 1.89 as the economy converges to its long-run equilibrium. The anti-competitive counterfactual (solid line), however, generates a pronounced shift in the entire distribution. Markups increase substantially for the majority of firms, with the median rising to approximately 2.01—a 6.43% increase—and the upper tail extending considerably beyond 100.81. This reflects both the direct effect of coordinated output restriction and the indirect effect of reduced capital stocks, which further constrain aggregate supply. The marked contrast between the factual and counterfactual markup distributions underscores how competitive conduct shapes not only the level but also the cross-sectional dispersion of market power.

FIGURE 11: DISTRIBUTION OF MARKUPS

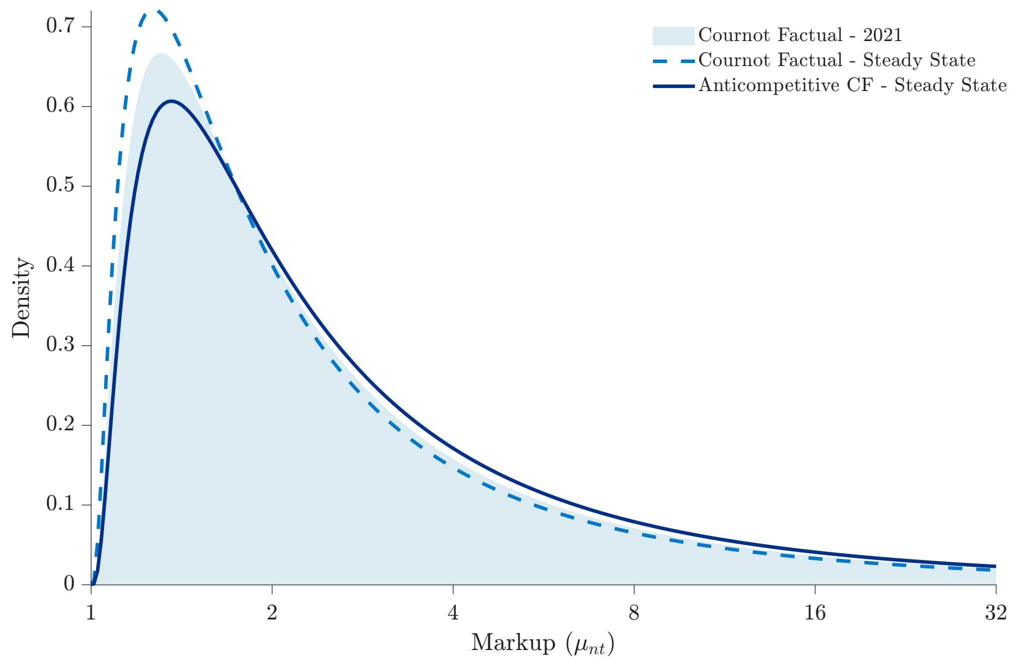

These results indicate that imperfect competition in product markets has quantitatively substantial implications for aggregate capital accumulation. This finding complements existing work on the welfare costs of oligopoly, by highlighting a distinct dynamic channel: beyond the static allocative distortions from markup pricing, imperfect competition distorts the inter-temporal allocation of resources by attenuating incentives for capital formation.

Counterfactual 2: Discount Rates and Market Structure. Our second counterfactual examines the market structure implications of changes in discount rates. Specifically, we study the case of a one percentage point increase in the aggregate discount rate ($r \to r + 0.01$), holding all other primitives constant, including the competitive conduct regime. This experiment isolates the effect of aggregate financial conditions on market dynamics, abstracting from changes in competitive behavior. The counterfactual is particularly relevant for understanding how cyclical fluctuations in aggregate discount rates—associated to either counter-cyclical fluctuations in the aggregate risk premium, or pro-cyclical fluctuations in the risk free rate—interact with an oligopolistic market structure.

Figure 12 displays the evolution of market concentration under both scenarios. The factual Cournot economy exhibits an initial decline in the Herfindahl index from 56.67 in 2021 to approximately 53.97 by 2022, followed by a gradual increase toward 60.83 in steady state. This non-monotonic pattern reflects the interaction between adjustment dynamics and the cross-sectional distribution of initial capital stocks. The discount rate shock counterfactual, in contrast, generates a substantially different trajectory. Concentration rises immediately and monotonically, reaching 78.19 in steady state—significantly higher than the factual outcome. This pronounced increase indicates that elevated aggregate discount rates fundamentally alter the competitive landscape by differentially affecting firms across the size distribution.

The mechanism underlying this concentration effect operates through heterogeneous responses to

FIGURE 12: EVOLUTION OF CONCENTRATION

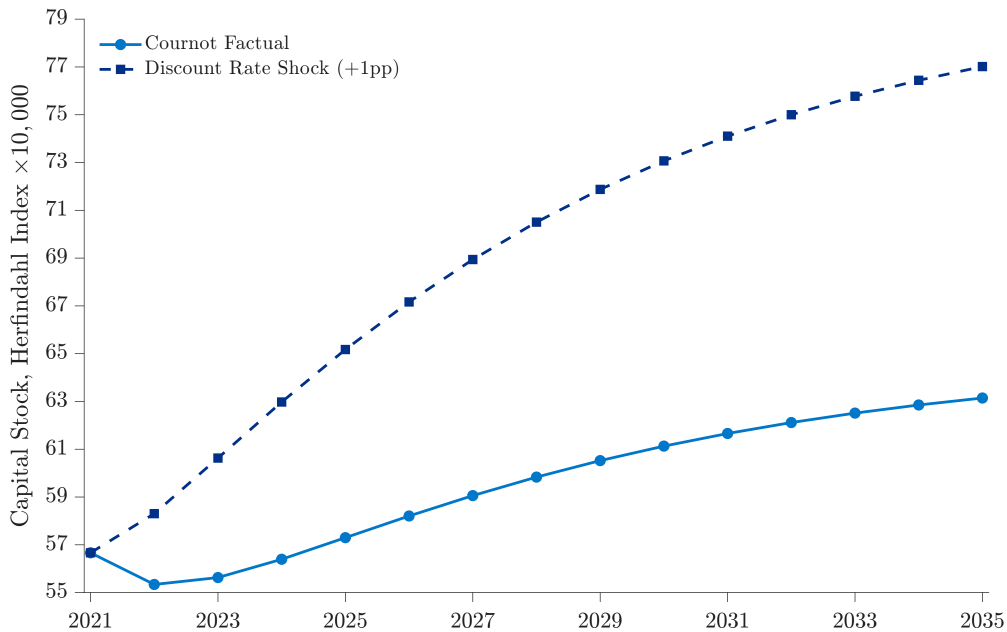

the aggregate discount rate shock. An aggregate discount rate reduces the present value of future profits, attenuating investment incentives for all firms. However, this effect is not uniform: smaller firms, which face steeper marginal adjustment costs relative to their capital bases and operate closer to their exit thresholds, reduce investment more sharply than larger incumbents. The resulting reallocation of market share toward larger firms manifests as increased concentration. Importantly, this channel operates even holding constant the competitive conduct regime—the increase in concentration strictly relates to the differential impact of financial conditions across the firm size distribution, rather than changes in strategic behavior.

Figure 13 illustrates the markup consequences of this reallocation effect. The fan chart plots the ratio of markups in the discount rate counterfactual to markups in the factual equilibrium, with values above unity indicating higher markups under the shock. The median firm experiences a modest markup increase, rising gradually to 0.62% by 2027. However, this aggregate pattern masks substantial heterogeneity. The upper percentiles of the distribution exhibit considerably larger markup increases, with firms in the 90th-95th percentile experiencing increases exceeding 1.24%. This dispersion reflects the oligopolistic nature of competition: as smaller firms reduce their capital stocks more aggressively in response to the discount rate shock, their output falls disproportionately, reducing competitive pressure for larger firms which then increase their own markups. The mechanism is micro-founded by a fundamental asymmetry associated with decreasing returns to scale: when the aggregate discount rate increases, smaller or less productive firms experience larger proportional changes in their operating costs, in comparison to their larger, more productive competitors. The widening fan over time indicates that these markup effects amplify along the transition path, as the cumulative impact of differential investment decisions compounds.

These findings demonstrate that discount rate changes have asymmetric effects across the firm distribution in oligopolistic markets, with implications extending beyond aggregate investment to mar-

Figure 13: Response of Markups to Discount Rate Shock

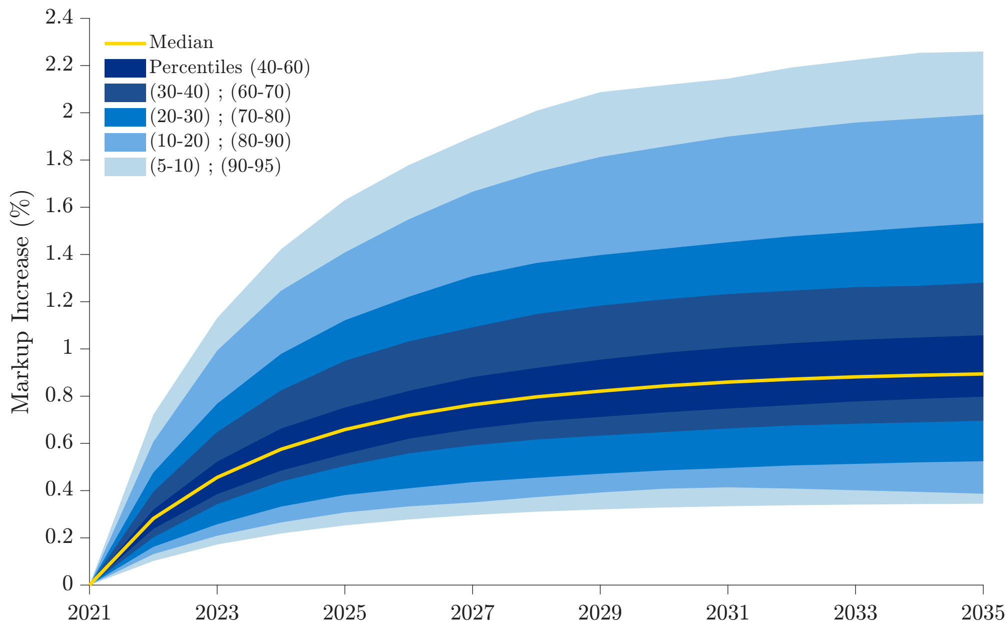

ket structure and pricing behavior. The concentration channel we identify suggests that business cycle fluctuations in aggregate discount rates may inadvertently affect product market competition, with persistently elevated aggregate discount rates potentially contributing to increased market concentration and higher markups. This mechanism is absent in models with monopolistic competition or perfect competition, where firm-level heterogeneity does not interact with strategic product market spillovers.

Counterfactual 3: Mergers. In a recent monograph, Philippon (2019) identifies rising market power as a potential explanatory factor for sluggish aggregate investment in the United States, and suggests that intense M&A activity could be one of the driver of rising market power. Gutierrez and Philippon (2023) further examine diverging trends in antitrust policy between EU and the US, and suggest that timid merger enforcement in the United States is a contributing factor to rising concentration and market power.

An interesting question that arises from the existing literature is therefore what has been the real impact of M&A activity on aggregate investment and prices over the last 25 years. This is the objective of our third counterfactual exercise. We take the model to the data in 1995 and simulate the impact of all mergers between firms in our sample observed in the SDC Platinum database through 2021. This counterfactual isolates the effect of common ownership on investment decisions by comparing an economy where merged entities internalize product market spillovers (setting $\mathbf{\Gamma}$ according to observed merger activity) against a baseline where all firms remain independent. The analysis encompasses 4,455 mergers involving 3,289 firms—representing 57.51% of our 1995 sample—providing substantial variation in merger exposure across the firm distribution.

Two important remarks merit emphasis on how we perform the exercise. First, we assume mergers generate no marginal cost reductions or product quality improvements: these could potentially render some mergers welfare-enhancing. Second, we treat all observed mergers as horizontal trans-

Figure 14: Mergers Simulation: Investment

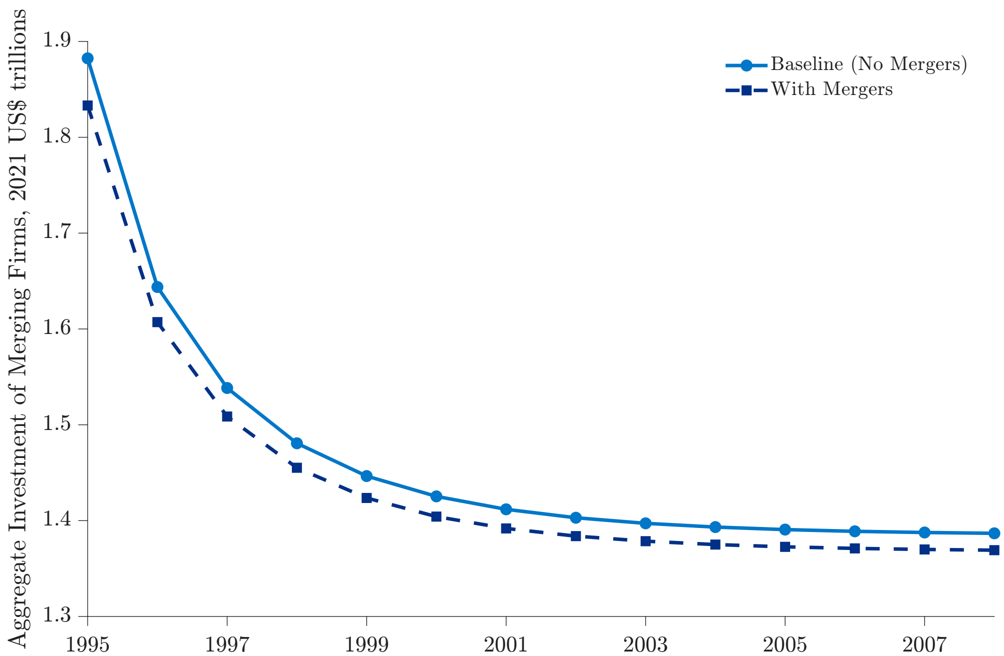

actions between competing firms. In practice, some transactions may represent vertical integration or conglomerate mergers that could benefit consumers through the reduction of double marginalization. Hence, our results should be interpreted as an upper bound on the anticompetitive harm caused by merger activity.

Figure 14 displays the aggregate investment dynamics of merging firms under both scenarios. In the baseline economy without mergers, aggregate investment among these firms exhibits a pronounced decline from \$1.98 trillion in 1995 to approximately \$1.50 trillion by 2003, reflecting the natural adjustment dynamics as firms converge toward their steady-state capital stocks. The merger counterfactual, however, generates a different trajectory: every year along the transition path (and all the way to steady state) the aggregate investment of merging firms is lower, reaching \$1.48 trillion by 2003.

The associated investment suppression observed in Figure 14 thus reflects the canonical mechanism identified in our model: when firms internalize product market spillovers through common ownership, they account for the negative externality their capital accumulation imposes on affiliated entities through increased output and lower equilibrium prices. Interestingly, however, while the observed magnitude of the decline in aggregate investment is roughly \$20.30 billion and therefore highly significant in absolute values, it represents a modest decline in relative terms, as it captures a 1.35% reduction with respect to the no-merger baseline. This finding suggests that the increased merger activity in recent decades in the US is not the likely driver of the sharp reduction in aggregate investment.

Analyzing merger activity more granularly at the firm level, we next observe striking heterogeneity on the impact of mergers on markups. Figure 15 illustrates this dispersion through a percentile fan chart focused on the upper half of the distribution, plotting the ratio of markups in the case

FIGURE 15: MERGERS SIMULATION: MARKUPS

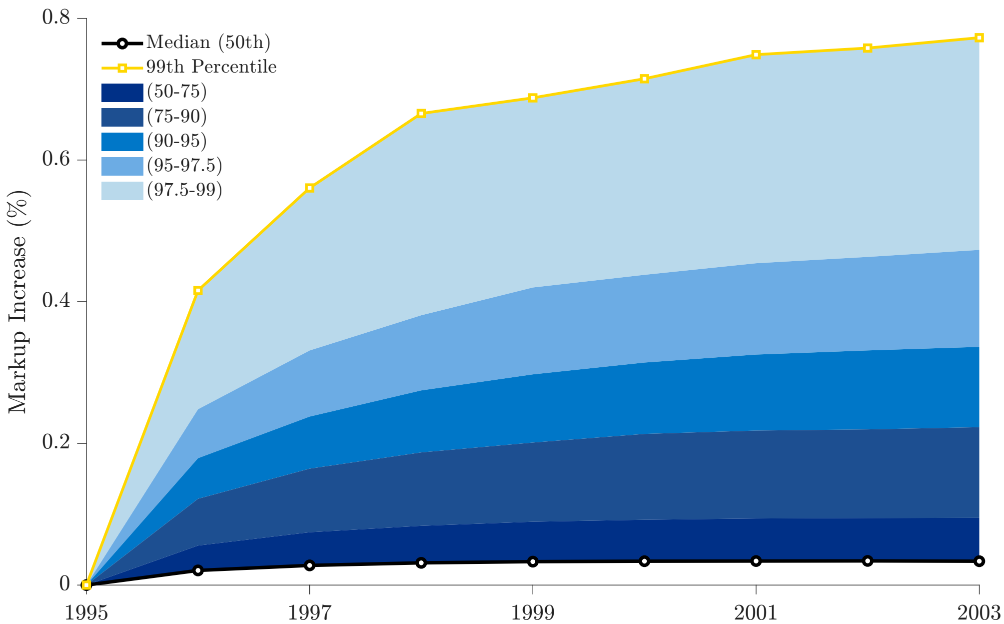

with mergers to markups in the alternative counterfactual scenario without mergers. The median merging firm experiences a modest initial markup increase, rising gradually to 0.03% by 2003. However, this central tendency masks dramatic variation in the upper tail. Firms between the 90th and 95th percentiles exhibit markup increases above 0.21% in 2003, while those in the 95th-99th percentile range experience increases exceeding 0.32% in 2003. Most remarkably, mergers in the 99th percentile range generate a markup increase of approximately 0.73% in 2003—an economically dramatic effect indicating near-monopolistic pricing power in the relevant product market.

Put together, our findings suggest that even if mergers on aggregate have had a modest distortion on aggregate investment, a few mergers may have had very effects on markups.

5 Sensitivity Analysis and Model Extensions

5.1 Model Parameters: Sensitivity Analysis

The adjustment cost parameter θ and the discount rate r play a central role in our model: the former governs the speed at which firms reallocate capital in response to shocks, while the latter determines how forward-looking firms are and, consequently, how strongly future competitive conditions shape current investment decisions. Given this, one should expect that the quantitative magnitudes of our main results and counterfactual exercises will be somewhat sensitive to the calibration of these parameters. It is therefore natural to ask how the results vary under alternative calibration strategies.

We address this question systematically in Appendix C, where we reproduce all the figures from Sections 3 and 4 under three alternative specifications. In Appendix C.1, we relax the assumption that θ is constant over time and instead allow it to be calibrated independently for each year, accommodating potential time variation in the frictions governing capital reallocation. In Appendix C.2, we replace our baseline discount rate with the 10-year US Treasury rate plus 5 percentage points. In

Appendix C.3, we use analyst-based estimates of the implied cost of capital from Lee, So, and Wang (2021) as an alternative measure of$r$. In Appendix C.4, we multiply the baseline estimate of θ by 1.5. Across all four specifications, the qualitative patterns documented in the main text remain robust.

5.2 Heterogeneous Adjustment Costs and Discount Rates

Our baseline model assumes homogeneous adjustment costs and discount rates across all business units, with θ and r being constant parameters. This subsection demonstrates how the model can be generalized to accommodate firm-specific heterogeneity in these parameters, where each business unit n faces its own adjustment cost parameter θ_n and discount rate r_n.

Under this generalization, the investment cost function in equation (12) becomes:

$$\psi_n(k_{nt+1};k_{nt}) \stackrel{\text{def}}{=} p_n^k i_{nt} + \frac{\theta_n}{2}(k_{nt+1} - k_{nt})^2 \tag{72}$$

The main cost of this added heterogeneity is that we lose the elegant closed-form solution for the transition matrix $\mathbf{\Phi}$ presented in equation (46). As mentioned earlier, that expression is the obtained as the to a Non-Symmetric Algebraic Riccati Equation (NARE), which we derive in Appendix A, and in the general case with heterogeneous discount rates and adjustment costs is given by:

$$\mathbf{\Phi}^2 - \left[\mathbf{\Theta}^{-1}\mathbf{\Omega} + (2\mathbf{I} + \mathbf{R})\right]\mathbf{\Phi} + (\mathbf{I} + \mathbf{R}) = \mathbf{0} \tag{73}$$

where $\mathbf{\Theta}$ and $\mathbf{R}$ are diagonal matrices containing the firm-specific adjustment costs and discount rates. This NARE generally lacks a closed-form solution and must be solved numerically through iterative procedures. An added disadvantage of this extension is that $\mathbf{\Phi}$ and $\mathbf{\Omega}$ are no longer simultaneously diagonalizable.

Despite this added computational complexity, the model still retains a remarkable degree of tractability, as well as its fundamental economic insights: product market spillovers continue to generate rich network effects in investment decisions, and the steady-state capital allocation still reflects firms' positions in the product similarity network.

An important empirical advantage of this extension is that it allows for sector-specific estimation of adjustment cost parameters. By applying the identification procedure at the sector level rather than across the entire sample, we can obtain estimates of θ_n that vary across industries, then capturing important differences in the nature of capital adjustment across sectors, consistent with the findings in Bustamante (2016).

Concerning heterogeneous discount rates, in turn, we stress that our model is deterministic and hence absent further frictions managers all managers should discount cashflows as the same risk free rate r. Even incorporating uncertainty into the model, the risk free rate would be the appropriate discount rate to use under the risk neutral measure. This extension can be used, however, as a tool to quantify the effective variation in discount rates across managers, and then explore the associated economic drivers.

5.3 More Complex Firm Interactions and Ownership Structures

Our baseline model can be extended to incorporate more complex forms of strategic interaction beyond simple product market competition. The key requirement for tractability is that the first-order conditions remain linear in quantities and capital stocks.

One important generalization involves expanding the interpretation of the adjacency matrix $\mathbf{\Sigma}$ beyond pure product substitution. The matrix can incorporate input-output linkages between business units, where changes in one firm's capital stock affect others through supply chain relationships rather than just demand-side competition. This allows the model to capture vertical relationships and complementarities alongside the horizontal substitution effects emphasized in our baseline specification.

A particularly relevant extension involves relaxing the binary nature of the ownership matrix $\mathbf{\Gamma}$. In the baseline model, the elements γ_{mn} of this matrix take values of either 0 or 1, indicating whether business units m and n are owned by the same firm. By allowing γ_{mn} to take continuous values between 0 and 1, the model can capture more nuanced ownership and control arrangements.

For instance, partial cross-holdings between firms can be modeled by setting γ_{mn} to the fraction of firm n owned by the entity controlling firm m. This extension naturally accommodates the analysis of cross-shareholdings, joint ventures, and strategic alliances where firms partially internalize spillovers without full common ownership.

The model also lends itself to studying common ownership structures, where institutional investors hold significant stakes in multiple competing firms. Following Ederer and Pellegrino (2025), the degree of common ownership can be captured through appropriate calibration of the $\mathbf{\Gamma}$ matrix, allowing the analysis of how institutional ownership concentration affects investment incentives and market outcomes.

These extensions preserve the model's tractability while significantly expanding its scope for analyzing complex organizational forms and market structures. The key insight remains that the pattern of strategic interactions—whether through product markets, ownership structures, or input-output relationships—determines both the steady-state allocation of capital and the dynamic adjustment process following shocks to the economic environment.

5.4 Bertrand Competition with Sticky Prices

While our baseline model is based on quantity competition (Cournot) and capital adjustment costs, the modeling framework can be adapted to work with price competition (Bertrand) and price adjustment (menu) costs. It is well known that these two classes of model bear many similarities. Here we briefly review a variant of our model based on Bertrand competition with sticky prices: We start by re-writing the vector of profits in terms of prices

$$\boldsymbol{\pi}_t(\mathbf{p}_t) = \operatorname{diag}(\mathbf{p}_t - \mathbf{c})(\mathbf{I} + \mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{p}_t) \tag{74}$$

Next, we re-define the adjustment cost function ψ in terms of prices, as opposed to capital stock:

$$\psi(p_{nt+1}, p_{nt}) = \frac{\theta}{2}(p_{nt+1} - p_{nt})^2 \tag{75}$$

As in the Cournot case, business units maximize the present discounted value of of profits, net of adjustment costs. Their optimization leads the following system of stacked first-order conditions:

$$\frac{1}{1+r}\left[(\mathbf{I} + \mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{p}_{t+1}) - \mathbf{\Upsilon}\mathbf{p}_{t+1} + \mathbf{\Upsilon}\mathbf{c} + \theta(\mathbf{p}_{t+2} - \mathbf{p}_{t+1})\right] = \theta(\mathbf{p}_{t+1} - \mathbf{p}_t) \tag{76}$$

where $\boldsymbol{\Upsilon}$ is a diagonal matrix that shares its diagonal entries with the inverse of $(\mathbf{I}+\boldsymbol{\Sigma})$. As for the Cournot case, we shall posit a linear Stable Markov Perfect equilibrium. Here the policy functions are next-period prices that are linear in current prices (as opposed to next period capital stocks that are linear in current capital stocks):

$$\mathbf{p}_{t+1} - \mathbb{p} = \boldsymbol{\Phi}(\mathbf{p}_t - \mathbb{p}) \tag{77}$$

By setting $\mathbf{p}_{t+2} = \mathbf{p}_{t+1} = \mathbb{p}$ we find the MPE steady-state price vector. As in Section 2, the equilibrium transition matrix is obtained as the solution of a Non-symmetric Algebraic Riccati Equation (NARE):

$$\mathbb{p} = \mathbf{c} + [\mathbf{I} + (\mathbf{I}+\boldsymbol{\Sigma})\boldsymbol{\Upsilon}]^{-1}(\mathbf{b}-\mathbf{c}) \tag{78}$$

$$\boldsymbol{\Phi} = \frac{1}{2}\Xi - \sqrt[+]{\frac{1}{4}\Xi^2 - (1+r)\mathbf{I}} \qquad \text{where} \qquad \Xi \overset{\text{def}}{=} \frac{1}{\theta}\left[\left(\mathbf{1}\mathbf{1}' + \mathbf{I}\right) \circ (\mathbf{I}+\Sigma)^{-1}\right] + \frac{2+r}{1+r}\mathbf{I} \tag{79}$$

This Bertrand variant demonstrates the versatility of our Network Q framework: by substituting price adjustment costs for capital adjustment costs and strategic price-setting for quantity competition, we obtain an analogous dynamic equilibrium with the same mathematical structure, suggesting that our core insights about network effects and strategic spillovers extend beyond the specific modeling choices of our baseline Cournot specification. Consistent with Wang and Werning (2022), this price-based formulation opens the door to additional applications in monetary economics, where the interaction between nominal rigidities and strategic firm behavior could provide new insights into the transmission of monetary policy shocks.

6 Conclusions

This paper provides a new dynamic model of corporate investment in oligopolistic product markets, extending the neoclassical Q theory of capital to a multi-firm, multi-product, fully-structural model. By embedding an hedonic demand system into a neoclassical investment setting, we endogenize markups and characterize in matrix form the network of investment spillovers. Moreover, as we obtain an explicit solution for firms' investment policies in matrix form, we then apply the model directly to US firm-level data, considering a product market network of thousands of firms.

The quantitative analyses in the paper illustrate the multiple practical applications of the model, as a tool for both researchers and policy makers alike. Specifically, we make inferences on the aggregate effects on investment of specific corporate events—such as cartels or mergers—and run arguably relevant exercises to policy-makers—such as studying the impact of an increase in aggregate discount rates on markups and concentration, when firms compete in an oligopoly.

Combined together, the theoretical predictions and the quantitative findings of our paper confirm that product market competition is a key force driving investment and capital allocation, both at the micro and macro levels. Arguably the more convenient property of our product market networks' approach to Q-theory is that it shows how firms' strategic behavior (at the micro level) ultimately affects investment on aggregate (or at the macro level). For this reason, future possible applications of this model span to multiple fields in economics, including macroeconomics, corporate finance, and international macro-finance.

References

Abel, Andrew B. 1980. "Empirical investment equations: An integrative framework." In *Carnegie-Rochester conference series on public policy*, Volume 12. 39–91, Elsevier.

Abel, Andrew, and Janice Eberly. 1994. "A Unified Model of Investment Under Uncertainty." *American Economic Review*.

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The network origins of aggregate fluctuations." *Econometrica* 80 (5): 1977–2016.

Atkeson, Andrew, and Ariel Burstein. 2008. "Pricing-to-market, trade costs, and international relative prices." *American Economic Review* 98 (5): 1998–2031.

Baker, Jonathan B, and Timothy F Bresnahan. 1985. "The gains from merger or collusion in product-differentiated industries." *The Journal of Industrial Economics* 33 (4): 427–444.

Baqaee, David Rezza. 2018. "Cascading failures in production networks." *Econometrica* 86 (5): 1819–1838.

Baqaee, David Rezza, and Emmanuel Farhi. 2020. "Productivity and Missalocation in General Equilibrium." *The Quarterly Journal of Economics* 135 (1): 105–163.

Baqaee, David Rezza, and Emmanuel Farhi. 2024. "Networks, Barriers, and Trade." *Econometrica* 92 (2): 505–541.

Benzarti, Youssef, and Jarkko Harju. 2021. "Using Payroll Tax Variation to Unpack the Black Box of Firm-Level Production." *Journal of the European Economic Association* 19 (5): 2737–2764.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica* 63 (4): 841–890.

Bizzarri, Matteo, and Fernando Vega-Redondo. 2024. "Common Ownership in Production Networks."

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying technology spillovers and product market rivalry." *Econometrica* 81 (4): 1347–1393.

Bonacich, Phillip. 1987. "Power and centrality: A family of measures." *American journal of sociology* 92 (5): 1170–1182.

Boppart, Timo, Peter J Klenow, Reiko Laski, and Huiyu Li. 2025. "Idea Rents and Firm Growth."Technical report, Working paper.

Burstein, Ariel T, Vasco M Carvalho, and Basile Grassi. 2025. "Bottom-up markup fluctuations." *The Quarterly Journal of Economics* 140 (4): 2619–2684.

Bustamante, M. Cecilia. 2016. "How do frictions affect corporate investment? A structural approach." *Journal of Financial and Quantitative Analysis* 51 1863–1895.

Bustamante, M Cecilia, and Andres Donangelo. 2017. "Product market competition and industry returns." *The Review of Financial Studies* 30 (12): 4216–4266.

Bustamante, M. Cecilia, and Laurent Fresard. 2021. "Does firm investment respond to peers' investment?" *Management Science* 67 (8): 4703–4724.

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi. 2021. "Supply chain disruptions: Evidence from the great east japan earthquake." *The Quarterly Journal of Economics* 136 (2): 1255–1321.

Cho, Thummim, Marco Grotteria, Lukas Kremens, and Howard Kung. 2023. "The present value of future market power." *Available at SSRN 3921171*.

Chodorow-Reich, Gabriel. 2025. "The Neoclassical Theory of Firm Investment and Taxes: A Reassessment." *NBER Working Paper*.

Chodorow-Reich, Gabriel, Matthew Smith, Owen M Zidar, and Eric Zwick. 2024. "Tax policy and investment in a global economy."Technical report, National Bureau of Economic Research.

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

Corhay, Alexandre, Howard Kung, and Lukas Schmid. 2025. "Q: Risk, rents, or growth?" *Journal of Financial Economics* 165 (C): .

Covarrubias, Matias, Germán Gutiérrez, and Thomas Philippon. 2020. "From good to bad concentration? US industries over the past 30 years." *NBER Macroeconomics Annual* 34 (1): 1–46.

Crouzet, Nicolas, and Janice Eberly. 2021. "Intangibles, markups, and the measurement of productivity growth." *Journal of Monetary Economics* 124 S92–S109.

Crouzet, Nicolas, and Janice Eberly. 2023. "Rents and intangible capital: A q+ framework." *The Journal of Finance* 78 (4): 1873–1916.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The rise of market power and the macroeconomic implications." *The Quarterly journal of economics* 135 (2): 561–644.

Eberly, Janice, and Neng Wang. 2009. "Capital Reallocation and Growth." *American Economic Review* 99 (2): 560–566.

Ederer, Florian, and Bruno Pellegrino. 2025. "A Tale of Two Networks: Common Ownership and Product Market Rivalry." *Review of Economic Studies*.

Eeckhout, Jan, and Laura Veldkamp. 2023. "Data and markups: A macro-finance perspective."Technical report, Working paper.

Eggertsson, Gauti B., Jacob A. Robbins, and Ella Getz Wold. 2021. "Kaldor and Piketty's facts: The rise of monopoly power in the United States." *Journal of Monetary Economics* 124 S19–S38.

Eisfeldt, Andrea, Edward Kim, and Dimitris Papanikolaou. 2022. "Intangible value." *Critical Finance Review* 11 (2): 299–332.

Eisfeldt, Andrea L, and Dimitris Papanikolaou. 2013. "Organization capital and the cross-section of expected returns." *The Journal of Finance* 68 (4): 1365–1406.

Eisfeldt, Andrea L, and Dimitris Papanikolaou. 2014. "The value and ownership of intangible capital." *American Economic Review* 104 (5): 189–194.

Gabaix, Xavier. 2009. "Power laws in economics and finance." *Annu. Rev. Econ.* 1 (1): 255–294.

Galeotti, Andrea, Benjamin Golub, and Sanjeev Goyal. 2020. "Targeting interventions in networks." *Econometrica* 88 (6): 2445–2471.

Galeotti, Andrea, Benjamin Golub, Sanjeev Goyal, Eduard Talamas, and Omer Tamuz. 2024. "Robust market interventions." *arXiv preprint arXiv:2411.03026*.

Goeree, Michelle Sovinsky. 2008. "Limited information and advertising in the US personal computer industry." *Econometrica* 76 (5): 1017–1074.

Gormsen, Niels Joachim, and Kilian Huber. 2025. "Corporate discount rates." *American Economic Review* 115 (6): 2001–2049.

Grassi, Basile. 2017. "IO in I-O: Size, Industrial Organization, and the Input-Output Network Make a Firm Structurally Important." *Work. Pap., Bocconi Univ., Milan, Italy*.

Gutierrez, German, and Thomas Philippon. 2023. "How European markets became free: A study of institutional drift." *Journal of the European Economic Association* 21 (1): 251–292.

Hayashi, Fumio. 1982. "Tobin's marginal q and average q: A neoclassical interpretation." *Econometrica: Journal of the Econometric Society* 213–224.

Hennessy, Christopher. 2004. "Tobin's Q, Debt Overhang, and Investment." *Journal of Finance*.

Higham, Nicholas J. 2008. "Functions of Matrices: Theory and Computation."

Higham, Nicholas J. 2014. "Sylvester's influence on applied mathematics." *Mathematics Today* 50 (4): 202–206.

Hoberg, Gerard, and Gordon Phillips. 2016. "Text-Based Network Industries and Endogenous Product Differentiation." *Journal of Political Economy* 124 (5): 1423–1465.

Hoberg, Gerard, and Gordon M Phillips. 2025. "Scope, Scale, and Concentration: The 21st-Century Firm." *The Journal of Finance* 80 (1): 415–466.

Hopenhayn, Hugo, and Koki Okumura. 2025. "Dynamic Oligopoly and Innovation: A Quantitative Analysis of Technology Spillovers and Product Market Competition."

Jorgenson, Dale W. 1963. "Capital theory and investment behavior." *The American economic review* 53 (2): 247–259.

Jun, Byoung, and Xavier Vives. 2004. "Strategic incentives in dynamic duopoly." *Journal of Economic Theory* 116 (2): 249–281.

Katz, Leo. 1953. "A new status index derived from sociometric analysis." *Psychometrika* 18 (1): 39–43.

Kwon, Spencer Y, Yueran Ma, and Kaspar Zimmermann. 2024. "100 years of rising corporate concentration." *American Economic Review* 114 (7): 2111–2140.

Lee, Charles MC, Eric C So, and Charles CY Wang. 2021. "Evaluating firm-level expected-return proxies: implications for estimating treatment effects." *The Review of Financial Studies* 34 (4): 1907–1951.

Liu, Ernest, and Aleh Tsyvinski. 2024. "A Dynamic Model of Input-Output Networks." *Review of Economic Studies* 91 (6): 3608–3644.

Ljungqvist, Lars, and Thomas J Sargent. 2018. *Recursive Macroeconomic Theory (4th ed.)*. MIT press.

Luttmer, Erzo GJ. 2007. "Selection, growth, and the size distribution of firms." *The Quarterly Journal of Economics* 122 (3): 1103–1144.

Maskin, Eric, and Jean Tirole. 1987. "A theory of dynamic oligopoly, III: Cournot competition." *European economic review* 31 (4): 947–968.

Maskin, Eric, and Jean Tirole. 1988a. "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs." *Econometrica* 56 (3): 549–569.

Maskin, Eric, and Jean Tirole. 1988b. “A Theory of Dynamic Oligopoly, II: Price Competition, Kinked Demand Curves, and Edgeworth Cycles.” *Econometrica* 56 (3): 571–599.

Miyashita, Masaki. 2026. “Characteristics Design: A Hedonic Approach to Optimal Product Differentiation.” *Working Paper*.

Moreau, Flavien. 2019. “Inferring Capital-Labor Substitution from Firm-level Distortions.”Technical report, Working paper.

Mukherjee, Abhiroop, Bruno Pellegrino, Alminas Zaldokas, Yiman Ren, and Tomas Thornquist. 2025. “New Products.” Working Paper.

Nevo, Aviv. 2001. “Measuring market power in the ready-to-eat cereal industry.” *Econometrica* 69 (2): 307–342.

Okumura, Koki. 2025. “Ownership Structure and Economic Growth.”

Pellegrino, Bruno. 2025. “Product differentiation and oligopoly: A network approach.” *American Economic Review* 115 (4): 1170–1225.

Peters, Ryan, and Luke Taylor. 2017. “Intangible capital and the investment- q relation.” *Journal of Financial Economics* 123 251–272.

Philippon, Thomas. 2019. *The great reversal: How America gave up on free markets*. Harvard University Press.

Reynolds, Stanley S. 1987. “Capacity investment, preemption and commitment in an infinite horizon model.” *International Economic Review* 69–88.

Reynolds, Stanley S. 1991. “Dynamic oligopoly with capacity adjustment costs.” *Journal of Economic Dynamics and Control* 15 (3): 491–514.

Sylvester, James Joseph. 1909. *The Collected Mathematical Papers of James Joseph Sylvester...* Volume 3. University Press.

Syverson, Chad. 2019. “Macroeconomics and Market Power: Context, Implications, and Open Questions.” *Journal of Economic Perspectives* 33 (3): 23–43.

Syverson, Chad. 2024. “Markups and markdowns.” *Annual Review of Economics* 17.

Tobin, James. 1969. “A general equilibrium approach to monetary theory.” *Journal of money, credit and banking* 1 (1): 15–29.

Ushchev, Philip, and Yves Zenou. 2018. “Price competition in product variety networks.” *Games and Economic Behavior* 110 226–247.

Voelkening, Daniel. 2024. “Product Design in an Hedonically Differentiated Duopoly.” *Working Paper*.

Vom Lehn, Christian, and Thomas Winberry. 2022. “The investment network, sectoral comovement, and the changing US business cycle.” *The Quarterly Journal of Economics* 137 (1): 387–433.

Wang, Olivier, and Ivan Werning. 2022. “Dynamic Oligopoly and Price Stickiness.” *American Economic Review* 112 (8): 2018–2049.

Weyl, E Glen, and Michal Fabinger. 2013. “Pass-through as an economic tool: Principles of incidence under imperfect competition.” *Journal of political economy* 121 (3): 528–583.

Online Appendix:

Dynamic Investment and Product Market Rivalry: The Network Q Model

Maria Cecilia Bustamante
University of Maryland

Bruno Pellegrino
Columbia University

A Additional Proofs

Proof of Proposition 1 (continued). We derive most of the proof using matrices $\mathbf{\Theta}$ and $\mathbf{R}$ in general form, to accommodate for the cases of heterogeneous adjustment costs and firm-specific discount rates discussed in an extension to the baseline model. To obtain the solution in Proposition 1, we further impose below $\mathbf{\Theta} = \theta \mathbf{I}$ and $\mathbf{R} = r\mathbf{I}$.

As a first step, we divide both side by replace $\mathbf{k}_{t+2}$ using the investment rule (22) and re-express the system of stacked first order conditions (32):

$$\begin{aligned} \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_{t+1} + (\mathbf{I}-\mathbf{\Delta})\mathbf{p}^k + \mathbf{\Theta}[\mathbb{k} + \mathbf{\Phi}(\mathbf{k}_{t+1}-\mathbb{k}) - \mathbf{k}_{t+1}] = \\ = (\mathbf{I}+\mathbf{R})\left[\mathbf{p}^k + \mathbf{\Theta}(\mathbf{k}_{t+1}-\mathbf{k}_t)\right] \end{aligned} \tag{80}$$

and distribute:

$$\begin{aligned} \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_{t+1} + (\mathbf{I}-\mathbf{\Delta})\mathbf{p}^k + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbb{k} - \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbf{k}_{t+1} &= \\ = \quad (\mathbf{I}+\mathbf{R})\mathbf{p}^k + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_{t+1} - (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_t & \end{aligned} \tag{81}$$

We now rearrange the equation so that all the multiples of $\mathbf{k}_{t+1}$ are on the left:

$$\begin{aligned} (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_{t+1} + \mathbf{\Omega}\mathbf{k}_{t+1} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbf{k}_{t+1} &= \\ = \quad \mathbf{Z}(\mathbf{b}-\mathbf{c}) - (\mathbf{R}+\mathbf{\Delta})\mathbf{p}^k + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbb{k} + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_t & \end{aligned} \tag{82}$$

next, notice that the sum of the first three terms of the right hand side is equal to $\mathbf{\Omega}\mathbb{k}$

$$\begin{aligned} (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_{t+1} + \mathbf{\Omega}\mathbf{k}_{t+1} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbf{k}_{t+1} &= \\ = \quad \mathbf{\Omega}\mathbb{k} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbb{k} + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_t & \end{aligned} \tag{83}$$

Subtracting $(\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_t$ from both sides and collecting the terms appropriately we have:

$$[\mathbf{\Omega} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi}) + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}](\mathbf{k}_{t+1}-\mathbb{k}) = (\mathbf{I}+\mathbf{R})\mathbf{\Theta}(\mathbf{k}_t-\mathbb{k}) \tag{84}$$

if we pre-multiply by $\mathbf{\Theta}^{-1}$ and invert the matrix in square brackets this equation can then be re-written as:

$$(\mathbf{k}_{t+1}-\mathbb{k}) = \left[\mathbf{\Theta}^{-1}\mathbf{\Omega} + (2\mathbf{I}+\mathbf{R})\mathbf{I} - \mathbf{\Phi}\right]^{-1}(\mathbf{I}+\mathbf{R})(\mathbf{k}_t-\mathbb{k}) \tag{85}$$

Then, comparing with (22), we see that $\mathbf{\Phi}$ must satisfy:

$$\left[\mathbf{\Theta}^{-1}\mathbf{\Omega}+(2\mathbf{I}+\mathbf{R})\mathbf{I}-\mathbf{\Phi}\right]^{-1}(\mathbf{I}+\mathbf{R}) = \mathbf{\Phi} \tag{86}$$

rearranging we obtain:

$$\mathbf{\Phi}^2-\left[\mathbf{\Theta}^{-1}\mathbf{\Omega}+(2\mathbf{I}+\mathbf{R})\right]\mathbf{\Phi}+(\mathbf{I}+\mathbf{R}) = \mathbf{0} \tag{87}$$

Equation (87) is a Non-symmetric Algebraic Riccati Equation (NARE), a variation of the Algebraic Riccati Equation (ARE) that is commonly encountered in linear-quadratic regulator problems from the optimal control literature. In general, neither the canonical AREs from linear quadratic optimal control, nor the NARE that we derived in (87) admit a closed-form solution. It is also known that they can admit multiple solutions, or none.

Our specific version is a quadratic matrix equation of the type $\mathbf{X}^2+\mathbf{M}_1\mathbf{X}+\mathbf{M}_2=\mathbf{0}$. This equation was first studied by Sylvester (1909)[10] He proved that when $\mathbf{M}_1\mathbf{M}_2=\mathbf{M}_2\mathbf{M}_1$ and $\mathbf{M}_1^2-4\mathbf{M}_2$ has a square root, then all solutions take the form:

$$\mathbf{X}=-\frac{1}{2}\mathbf{M}_1\pm\sqrt{\frac{1}{4}\mathbf{M}_1^2-\mathbf{M}_2} \tag{88}$$

This condition is satisfied in our case if $\mathbf{\Theta}=\theta\mathbf{I}$ and $\mathbf{R}=r\mathbf{I}$. Applying this solution to (87) yields:

$$\mathbf{\Phi} = \frac{1}{2}\left[\frac{1}{\theta}\mathbf{\Omega}+(2+r)\mathbf{I}\right]\pm\sqrt{\frac{1}{4}\left[\frac{1}{\theta}\mathbf{\Omega}+(2+r)\mathbf{I}\right]^2-(1+r)\mathbf{I}} \tag{89}$$

As discussed by Higham (2008), the square root of a matrix with positive eigenvalues may admit up to $2^{\mathscr{E}}$ solutions, where $\mathscr{E}$ is the number of unique eigenvalues of the radicand. This is because the square roots of a matrix can be found by taking the square root of its eigenvalues (which in our setting are guaranteed to be positive). Those square roots in turn each admit both a positive and a negative solution. Considering all sign combinations we obtain $2^{\mathscr{E}}$ potential solutions.

We show next that focusing on the unique principal solution (with all positive eigenvalues) preceded by a leading minus sign is both necessary and sufficient for $\mathbf{\Phi}$ to be stable. To do so, we exploit the fact that all candidate solutions $\mathbf{\Phi}$ are simultaneously diagonalizable with $\mathbf{\Omega}$. That is, let $\hat{\omega}_n$ be the n^{th} eigenvalue of $\mathbf{\Omega}$ and let $\mathbf{E}$ be the matrix of eigenvectors of $\mathbf{\Omega}$, so that its spectral decomposition is

$$\mathbf{\Omega} = \mathbf{E}\,\mathrm{diag}(\hat{\boldsymbol{\omega}})\,\mathbf{E}^{-1} \tag{90}$$

also, let $\hat{\varphi}_n$ be the n^{th} eigenvalue of $\mathbf{\Phi}$ .The matrices $\mathbf{\Omega}$ and $\mathbf{\Phi}$ are simultaneously diagonalizable, that is:

$$\mathbf{\Phi} = \mathbf{E}\,\mathrm{diag}\left(\hat{\boldsymbol{\varphi}}\right)\mathbf{E}^{-1} \tag{91}$$

and their respective n^{th} eigenvalues are related by the following function, the scalar equivalent of equation (38)

$$\hat{\varphi}_n = \frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)\pm\sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)^2-(1+r)} \tag{92}$$

This is a known result in matrix theory that applies when $\mathbf{\Omega}$ is symmetric (it is by construction)

[10]See also the review of Higham (2014).

and when $\mathbf{\Phi}$ is a function of $\mathbf{\Omega}$ that results from the concatenation of 1) summation to multiples of identity matrices; 2) scalar multiplications and 3) symmetric matrix power functions. This latter condition can be verified by examining Equation (89).

Because for all $\hat{\omega}_n > 0$, $\theta > 0$ and $r > 0$ we have

$$\frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)+\sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)^2-(1+r)} \geq 1 \tag{93}$$

$$\frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)-\sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)^2-(1+r)} \leq 1 \tag{94}$$

it follows that the unique principal square root matrix, preceded by the minus sign, yields the only stable solution $\mathbf{\Phi}$.

□

Proof of Proposition (2). Consider first the single-product, symmetric Cournot oligopoly case. Matrix $\mathbf{\Omega}$ is then equal to a $N \times N$ real symmetric matrix whose diagonal entries are all equal to $2z^2$ and whose off-diagonal entries are all equal to $z^2\sigma$, namely:

$$\mathbf{\Omega} = z^2\sigma\mathbf{1}\mathbf{1}' + z^2(2-\sigma)\mathbf{I}. \tag{95}$$

Consider the two terms on the right hand side of the equation. In the first term, matrix $\mathbf{1}\mathbf{1}'$ has rank 1, so its eigenvalues are either equal to N (for the eigenvector $\mathbf{1}$) or 0 (for the $N-1$ eigenvectors orthogonal to vector $\mathbf{1}$). The second term shifts every eigenvalue by $z^2(2-\sigma)$. The symmetric case in which all firms are endowed with the same level of capital is associated with the eigenvector $\mathbf{1}$. Hence, the unique eigenvalue of matrix $\mathbf{\Omega}$ for the symmetric, single product Cournot oligopoly case is equal to $z^2[2+(N-1)\sigma]$, as stated in equation (43). □

Proof of Proposition 3. Rearrange equation (22) as follows

$$\mathbf{k}_{t+1} - \mathbf{k}_t = (\mathbf{\Phi}-\mathbf{I})(\mathbf{k}_t - \mathbb{k}) \tag{96}$$

and plug it in the definition of $\mathbf{q}(\mathbf{k}_t)$

$$\mathbf{q}(\mathbf{k}_t) = \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_t + (\mathbf{I}-\mathbf{\Delta})\mathbf{p}^k + \theta(\mathbf{k}_{t+1} - \mathbb{k}) \tag{97}$$

to obtain an expression in terms of the contemporaneous state vector:

$$\mathbf{q}(\mathbf{k}_t) = \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_t - \mathbf{\Delta}\mathbf{p}^k + \mathbf{p}^k + \theta(\mathbf{\Phi}-\mathbf{I})(\mathbf{k}_t - \mathbb{k}) \tag{98}$$

Add and subtract $r\mathbf{p}^k$ to obtain:

$$\mathbf{q}(\mathbf{k}_t) = \left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^k\right] - \mathbf{\Omega}\mathbf{k}_t + (1+r)\mathbf{p}^k + \theta(\mathbf{\Phi}-\mathbf{I})(\mathbf{k}_t - \mathbb{k}) \tag{99}$$

rearrange equation (37) as follows:

$$\mathbf{Z}(\mathbf{b}-\mathbf{c})-(r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^{k} \;=\; \boldsymbol{\Omega}\Bbbk \tag{100}$$

and substitute inside equation (99) to obtain:

$$\mathbf{q}(\mathbf{k}_t) \;=\; \boldsymbol{\Omega}(\Bbbk-\mathbf{k}_t)+(1+r)\mathbf{p}^{k}+\theta(\mathbf{I}-\boldsymbol{\Phi})(\Bbbk-\mathbf{k}_t) \tag{101}$$

finally, collect terms to obtain the following version of equation (45):

$$\mathbf{q}(\mathbf{k}_t) \;=\; (1+r)\mathbf{p}^{k}+[\theta(\mathbf{I}-\boldsymbol{\Phi})+\boldsymbol{\Omega}](\Bbbk-\mathbf{k}_t) \tag{102}$$

□

Proof of Proposition 4. expand the expression for $\boldsymbol{\pi}(\Bbbk)$ and collect $\mathbb{K}$:

$$\frac{r}{1+r}\mathbf{v}(\Bbbk) \;=\; \mathbb{K}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c})-\mathbf{Z}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\Bbbk-\boldsymbol{\Delta}\mathbf{p}^{k}\right] \tag{103}$$

Then rearrange equation (37)

$$\mathbf{Z}(2\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\Bbbk \;=\; \mathbf{Z}(\mathbf{b}-\mathbf{c})-(r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^{k} \tag{104}$$

to obtain the term in square brackets in equation

$$\mathbf{Z}(\mathbf{b}-\mathbf{c})-\mathbf{Z}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\Bbbk-\boldsymbol{\Delta}\mathbf{p}^{k} \;=\; r\mathbf{p}^{k}+\mathbf{Z}^{2}\Bbbk \tag{105}$$

And replace it inside equation (103) to obtain:

$$\frac{r}{1+r}\mathbf{v}(\Bbbk) \quad = \quad r\,\mathbf{P}^{k}\,\Bbbk+\mathbb{K}\,\mathbf{Z}^{2}\Bbbk \tag{106}$$

Then divide both sides by $r/(1+r)$. □

Proof of Proposition 7. Rewrite equation (45) as

$$\Bbbk-\mathbf{k}_t \;=\; \boldsymbol{\Lambda}^{-1}\left[\mathbf{q}(\mathbf{k}_t)-\mathbf{q}(\Bbbk)\right] \tag{107}$$

substitute it inside equation (25) to obtain

$$\boldsymbol{i}_{t+1}(\mathbf{k}_t) \;=\; \boldsymbol{\Delta}\mathbf{k}_t+(\mathbf{I}-\boldsymbol{\Phi})[\boldsymbol{\Theta}(\mathbf{I}-\boldsymbol{\Phi})+\boldsymbol{\Omega}]^{-1}\left[\mathbf{q}(\mathbf{k}_t)-\mathbf{q}(\Bbbk)\right] \tag{108}$$

and use the fact that, for symmetric matrices $\mathbf{M}_1(\mathbf{M}_2)^{-1}=\left(\mathbf{M}_2\mathbf{M}_1^{-1}\right)^{-1}$ to simplify to (58). □

Proof of Proposition 8. Consider the derivation of steady state markups in Equation (61). Multiply both sides of the steady state level of capital by $\mathbf{Z}$

$$\mathbf{y}(\mathbb{k}) \;=\; \mathbf{Z}\boldsymbol{\Omega}^{-1}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c})-(r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^{k}\right] \tag{109}$$

substituting the expression in the inverse demand function $\mathbf{p}(\mathbb{k})$:

$$\mathbf{b}-(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{y}(\mathbb{k}) \;=\; \mathbf{b}-(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\boldsymbol{\Omega}^{-1}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c})-(r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^{k}\right] \tag{110}$$

finally we multiply both sides by $\mathbf{C}^{-1}$ and rearrange to obtain Equation (61):

$$\boldsymbol{\mu}(\mathbb{k}) \;=\; \mathbf{C}^{-\mathbf{1}}\mathbf{b}+\mathbf{C}^{-1}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\boldsymbol{\Omega}^{-1}\left[(r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^{k}-\mathbf{Z}(\mathbf{b}-\mathbf{c})\right] \tag{111}$$

The solution outside steady state in (60) follows from combining the definition of markup with the law of motion of investment in equilibrium, and rearranging terms. □

B Additional Details on Data and Calibration

Calibration of productivity vector z. Our empirical estimate of z_n^2 is obtained in four steps. First, we obtain a noisy estimate of z_n^2 as the residual from the first order condition (equation 69).

Second, we smooth transitory noise in firm-level productivity estimates by applying a centered 7-period moving average to the raw estimate of z_n^2 computed from equation (69). Third, for each year t in our sample, we compute each firm's percentile rank and estimate the Pareto scale and shape parameters by regressing the noisy estimate of $\log z_{nt}$ on the corresponding log of 1 minus the corresponding quantile, using only firms in the top 5% of the distribution: the year-specific intercept identifies the scale parameter, while the slope identifies the shape parameter.

Finally, we use the estimated coefficients to predict $\log z_n$ for all firms in year t, including those below the top 5%. This effectively extrapolates the Pareto tail across the entire distribution. The logic behind this approach is that while measurement error may severely distort z_n estimates for low- and medium-productivity firms, the highest-z_n firms are more reliably identified.

By allowing the regression coefficients to vary by year, we accommodate potential time trends in productivity dispersion over our 1995–2021 sample period. The validity of assumed Pareto distribution is reflected in the R^2 of the regression, which is typically 98 to 99%, with a typical sample size of about 200 observations.

Rescaling productivity to match aggregate enterprise values. The Pareto-fitted productivity estimates yield a distribution of z_n that captures the cross-sectional shape but may not match the aggregate level of firm values in the data. We therefore apply a scalar rescaling factor $\mathcal{Z}$ to the entire productivity vector, choosingλ_z to satisfy the aggregate moment condition:

$$\sum_{n=1}^{n} v_n(\mathbf{k}_t;\mathcal{Z}\mathbf{z}) \;=\; \sum_{n}(1+r)\hat{v}_{nt} \tag{112}$$

FIGURE 16: DISTRIBUTION OF STEADY-STATE CAPITAL STOCK

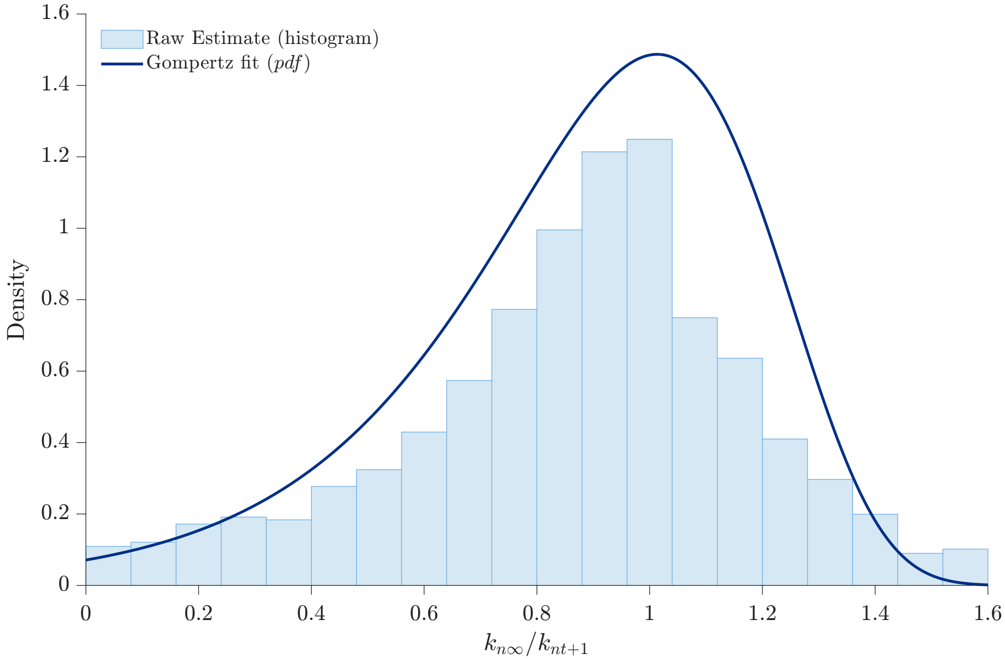

where $v_n(\mathbf{k}_t; \lambda_z \mathbf{z})$ denotes the model-predicted enterprise value for firm$n$ as a function of the rescaled productivity vector, and $\hat{v}_{nt}$ is the observed market value of assets. The factor $(1+r)$ accounts for the cum-dividend adjustment in the model's enterprise value definition. We solve for $\mathcal{Z}$ numerically using a root-finding algorithm (`fzero` in MATLAB). The final calibrated productivity vector is used in all subsequent model calculations and counterfactual simulations.

Calibration of steady-state capital $\Bbbk$. The steady-state capital distribution $\Bbbk$ is an equilibrium object rather than a primitive parameter, but we must ensure it satisfies the positivity constraint $\Bbbk > 0$. From the transition dynamics in Proposition 1, we can express the steady state as:

$$\Bbbk = (\mathbf{I} - \mathbf{\Phi})^{-1}(\mathbf{k}_{t+1} - \mathbf{\Phi}\mathbf{k}_t) \tag{113}$$

As for $\mathbf{z}$, this "raw" calculation can yield negative values due to measurement error in observed capital stocks. To ensure positivity while preserving the cross-sectional distribution, we fit a Gompertz distribution to the ratio$k_{n\infty}/k_{nt+1}$ (matching its median and interquartile range) and then use the fitted distribution to assign strictly positive steady-state values to each firm. Results are robust to replacing the Gompertz with a Weibull distribution fitted using the same procedure.

Figure 16 shows the Gompertz fit of the projected raw steady state capital stocks.

Recovery of demand and cost parameters. Given the calibrated values of $\mathbf{z}$, θ, and $\Bbbk$, we recover the remaining structural parameters. Output y_{nt} is simply obtained as $z_n k_{nt}$. The marginal cost c_n is equal to the ratio of costs of goods sold (as reported in Compustat) to output y_{nt} in each firm-year. This assumes that all reported operating costs are variable, consistent with our Leontief production technology. The demand intercept vector $\mathbf{b}$ is then recovered by inverting the steady-state condition from equation (37):

$$\mathbf{b} \ = \ \mathbf{c} + \mathbf{Z}^{-1}\left[\mathbf{\Omega}\Bbbk + (r\mathbf{I} + \mathbf{\Delta})\mathbf{p}^k\right] \tag{114}$$

where $\mathbf{\Omega} = \mathbf{Z}(2\mathbf{I} + \mathbf{\Sigma})\mathbf{Z}$ is the equilibrium network effects matrix. The price vector $\mathbf{p}$ is then obtained from the inverse demand function (7).

Notice that at this point, if we compute revenues as the newly obtained vector of prices ($\mathbf{p}$) times output ($\mathbf{y}$) we will find it to be different from the vector of revenues observed in the Compustat data. This is to be expected: because our model is significantly over-identified (as many equality restrictions as parameters, plus additional inequalities) it cannot perfectly fit the observed revenues (model revenues absorb the residual variation). The same goes for the vector of next period capital $\mathbf{k}_{t+1}$, which also cannot be perfectly fit by the model. The last step therefore consists of updating the vector of revenues $p_{nt}y_{nt}$ with the model-consistent estimate.

This calibration strategy balances several competing objectives: (1) using observable data where possible, (2) imposing economically interpretative restrictions to handle measurement error, (3) ensuring all parameters satisfy their theoretical constraints (particularly positivity), and (4) maintaining internal consistency with the model's equilibrium conditions. The distributional approach for $\mathbf{z}$ and $\mathbb{k}$ is particularly important, as it allows us to extract robust cross-sectional patterns from noisy firm-level data while ensuring the model's predictions remain economically sensible.

C Online Appendix: Sensitivity Analysis

This appendix evaluates the sensitivity of our main quantitative results to alternative calibration choices. We consider three variations. First, in Appendix C.1, we allow the adjustment cost parameter θ to be year-specific rather than uniform across all years, accommodating potential time variation in the frictions governing capital reallocation.

Second, in Appendix C.2, we replace our baseline discount rate r with the 10-year US Treasury rate plus 5 percentage points, providing a robustness check on the role of the cost of capital in driving the model's dynamics.

Third, in Appendix C.3 we use analyst-based estimates of the implied cost of capital from Lee, So, and Wang (2021) as an alternative measure of the discount rate r.

Fourth, in AppendixC.4, we multiply the baseline estimate of the adjustment cost parameter θ by 1.5, testing the sensitivity of our results to higher capital reallocation frictions. In each case, we reproduce the full set of figures from the main text to facilitate direct comparison with the baseline results.