Dynamic Investment and Product Market Rivalry: The Network Q Model

Maria Cecilia Bustamante[*]
University of Maryland

Bruno Pellegrino[†]
Columbia University, NBER

August 2026

Abstract

We present a new dynamic model of corporate investment in imperfectly competitive product markets that extends the neoclassical (Q) theory of capital to accommodate heterogeneous, multi-product firms and a rich hedonic demand system. Our model 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 equilibrium existence and uniqueness results along with global analytical solutions for the firms' investment policies. We then take our model to the data for the universe of US public companies and obtain four novel insights: 1) product market competition is a key driver of aggregate investment and capital allocation; 2) shocks to firms' cost of capital generate highly heterogeneous investment and markup responses across firms, and thus impact industry concentration; 3) monopoly rents account for a large, rising share of firms' value; 4) mergers consummated since 1995 have led to a modest decline in the aggregate capital formation of merging firms, yet firm-level markup increases have been highly heterogeneous.

JEL classification: C7, D2, E2, G3
Keywords: heterogeneous firms, investment, networks, product markets

[*]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, Janice Eberly, Joel Flynn, Xavier Gabaix, Ben Golub, Pete Kyle, Hugo Hopenhayn, Ernest Liu, Yueran Ma, Pascal Noel, Koki Okumura, Ludwig Straub, Luke Taylor, Neng Wang, Olivier Wang, Ivan Werning, and workshop participants at NBER Corporate Finance Spring Meetings, the Minnesota Macro Theory Workshop, the New York Macro Theory Group, the Network Sciences in Economics Conference 2026, the BSE Forum 2026, as well as seminar participants at NWU Kellogg, Minnesota Carlson, UMaryland Smith, Columbia Finance & Macro, and UToronto Rotman for helpful feedback. We thank Lorenzo Calvetti for outstanding research assistance.

1 Introduction

How does product market competition shape firms' investment decisions — and, through them, the accumulation and allocation of capital in the aggregate? This question is of central importance to researchers, policymakers and practitioners alike.

Billionaire investor Warren Buffett, for instance, has repeatedly emphasized product market competition as a key determinant of firms' investment intensity and rates of return. In several of his letters to Berkshire Hathaway shareholders,[1] he pointed to airlines as an example of an industry populated by many competitors producing near-perfect substitutes (airline seats): in his view, such intense competition made the airline industry a "bottomless pit" for capital, plagued by investment arms races and chronically low rates of return.

When it comes to modeling corporate investment, the neoclassical (Q) model, developed sequentially by Jorgenson (1963), Tobin (1969), Abel (1980) and Hayashi (1982), stands out as the workhorse, owing to both its influence and its wide range of applications. Yet Q theory describes a firm in isolation. In reality — as Buffett's example illustrates — firms operate within a competitive ecosystem and their investment decisions are interdependent: a firm's marginal return on capital depends on the demand for its products, which in turn depends on its rivals' capacity.

The interplay between investment and product market competition remains hard to study at scale, both theoretically and empirically. On the one hand, theoretical papers typically focus on dynamic investment models of duopoly or symmetric oligopoly, in order to preserve tractability. These however lack generality as well as broad empirical relevance.[2] On the other hand, the applied empirical literature has documented the difficulty of identifying product market investment spillovers in reduced form;[3] structural models that are estimated econometrically also tend to abstract away from product market interactions, again for tractability reasons.

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 (GHL) demand system of Pellegrino (2025), we model firms' product market interactions as a network: in this (weighted) network firms are nodes, and the edge connecting any two firms measures the substitutability of the goods and services they supply. The intensity of competition faced by each individual firm thus depends on how similar or dissimilar its product is to those produced by its rivals; through the lens of the GHL framework, this in turn depends on the attributes of all of the products that the consumer can purchase. We call this the Network Q model.

The Network Q model is simultaneously a neoclassical model of investment and a dynamic game of oligopoly. We restrict attention to Markov strategies that are linear in the state and generate stable capital dynamics, so that in equilibrium firms' capital stocks follow a vector autoregressive (VAR) process. We then derive necessary and sufficient conditions that characterize the Markov Perfect Equilibria (MPE) of the game. However, as is well understood in the literature on dynamic games (Engwerda, 2005; Monti et al., 2024; Nortmann et al., 2025), the MPE concept is ill-suited to quantitative work: the characterizing equations have an unknown, potentially infinite number of

[1] See in particular the 1991, 2004, and 2007 letters at www.berkshirehathaway.com/letters/letters.html

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

[3] For studies on product market peer effects and investment spillovers, see for instance Bloom, Schankerman, and Van Reenen (2013) and Bustamante and Frésard (2021).

solutions; they admit no closed form and are computationally intractable beyond a handful of firms. In addition the MPE concept requires a far-fetched degree of strategic sophistication on the part of managers, and is therefore subject to the critique of Moll (2025). All three problems stem from "strategic feedback" effects: in an MPE, each manager computes every rival's dynamic response to every possible deviation. Standard remedies that approximate the MPE neutralize feedback effects at the cost of abandoning rational expectations.

We make progress by introducing a minimally invasive approximation of the MPE concept that preserves rational expectations, which we call Proximate Markov Equilibrium (PME). The PME is characterized by the same equations as the MPE minus the strategic-feedback term. We prove that it exists, that it is unique, and that it admits a global analytical solution that is also a first-order approximation to all MPEs in a neighborhood of a disconnected economy. We further show that the behavioral assumption underpinning the PME admits a microfoundation that holds when the number of firms is large — precisely our setting.

We obtain closed-form expressions for firms' investment rates, enterprise values and marginal Qs in terms of primitives (discount rates, adjustment-cost parameters, demand and productivity) — something that is unusual for both neoclassical investment models and dynamic oligopoly games, and that allows us to derive new theoretical insights on how product markets shape firm behavior.

In line with previous neoclassical models that feature a single monopolistic firm (e.g., Hayashi, 1982 and Crouzet and Eberly, 2023), 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 endogenously determined alongside 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 the user cost of capital varies with firms' discount rates, we further show that firm-level and aggregate fluctuations in discount rates affect the dynamics of markups and industry concentration.

Our analytical framework 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. 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), converge more slowly to their steady-state level of capital. Conversely, firms that face intense product market competition converge faster. As a result, in the aggregate economy, the density 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.

Another interesting feature of our model is that the equilibrium 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. Each diagonal entry measures the marginal effect of a firm's capital on its own equilibrium value. Meanwhile, the off-diagonal terms capture the impact of firms' investment on their rival firms' values in equilibrium. We also 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 of these measures of spillovers, we show how they relate to the degree of competitive pressure that firms are subject to in equilibrium.

Our model also speaks to an emerging literature that investigates firms' heterogeneous responses to changing investment opportunities (Martin-Baillon, 2021; Winberry, Auclert, Rognlie, and Straub, 2025) and financing conditions (Ottonello and Winberry, 2020; Best, Born, and Menkhoff, 2026). It produces closed-form expressions for the cross-sectional distributions of two key statistics: 1) the Investment Sensitivity (IS), which measures how individual firms' investment responds to changes in the discount rate; 2) the Marginal Propensity to Invest out of Growth Opportunities (MPI^G), which measures how many cents of capital expenditure follow a one-dollar increase in firm value due to improved growth opportunities. Our detailed model of product markets translates into rich firm-level heterogeneity in both these statistics.

After completing the exposition of our theoretical model, we take it to the data for the universe of US public companies from 1995 to 2021. Following Pellegrino (2025), we use the text-based product similarity dataset of Hoberg and Phillips (2016), which covers the universe of US publicly traded firms, to parameterize our hedonic demand system. The resulting matrix of demand elasticities is massive: it contains roughly 10 to 30 million elasticities per year.

Our empirical implementation departs radically from the previous corporate finance literature, which is dominated by reduced-form predictive regressions and structural estimation exercises. By contrast, our model is mapped to the data firm-by-firm (each firm in the model corresponds to a real firm in the data) and is implemented fully structurally: that is, we recover all the model primitives at the firm level. In this respect our empirical approach is close to the empirical industrial organization literature. Unlike IO oligopoly models, however, which are typically solved numerically and can only be mapped to a few firms, our model is solved analytically and can be seamlessly scaled up to the entire Compustat universe.

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 cannot be perfectly fitted — namely: 1) firm revenues; 2) investment rates; 3) firm valuations. We then use our model to deliver four new insights.

First, we use the model to compute all firms' steady-state enterprise values and break them down into two 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 to 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 47.33% of aggregate firm value in 1995, and increase to 73.93% of aggregate firm value in 2021. Our results confirm the importance of market power in determining firms' valuations, as noted by Crouzet and Eberly (2023) and Eeckhout (2025), among others.

Relatedly, monopoly rents also shape the cross-section of firms' MPIs. For the typical firm, a demand shock that raises steady-state firm value by one dollar increases contemporaneous capital expenditure by roughly ten cents. This is due to a combination of adjustment costs (which cause firms to postpone investment) and market power (which leads firms to increase their markups in the face of a favorable demand or cost shock). However, this figure masks considerable cross-sectional variation: the MPI approaches zero for firms with significant monopoly rents and high capital intensity, and rises above 0.56 for firms subject to significant competitive pressure and low capital intensity. The remaining three results come from counterfactual exercises, which showcase our model's ability to answer policy-relevant questions.

Second, by counterfactually changing firms' conduct, we show that product market competition is a key force driving aggregate investment. We do so by comparing our baseline (Cournot) oligopoly model against an anticompetitive counterfactual: in the former, firms internalize the effect of their investment on their own markups but not the corresponding spillovers to their rivals; in the latter scenario spillovers are internalized, as if firms colluded. Moving from Cournot competition to full cartelization reduces aggregate steady-state capital by 30.88%. This result confirms that investment spillovers significantly contribute to aggregate capital formation.

Third, motivated by the fact that in our model firms pass their costs of capital on to their customers through their markups, we shock the model with an unanticipated, persistent one-percentage-point drop in the discount rate. We show that firms' investment and markups display highly heterogeneous responses to the shock: in general, the economy transitions to a new equilibrium path characterized by higher investment, lower industry concentration, and lower markups. While our model does not feature monetary policy in the strict sense, this finding suggests that monetary policy loosening (implemented as a persistent reduction in the nominal interest rate) may result in higher product market competition among firms.

Fourth, because our model also allows us to simulate mergers, we assess the impact of mergers and acquisitions on aggregate investment between 1995 and 2021. Compared to the counterfactual case of no M&A activity after 1995, we find that realized M&A activity reduced the aggregate corporate investment of merging firms by a modest 1.47%. Simultaneously, however, we observe significant variation in the impact of mergers in the cross-section: while the median merging firm barely alters its price, the merging firm at the 99th percentile increases markups by approximately 0.53% within eight years of the merger.

The paper concludes with a sensitivity analysis of our quantitative results (where we study alternative parameterizations of the model) and with a discussion of extensions and alternative applications of our Network Q model. We show that 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. We also show that our model can be modified to allow for Bertrand competition (as opposed to quantity competition) with sticky prices. 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 connects several literatures across finance, macroeconomics and industrial organization. The most closely related papers are two recent contributions by Okumura (2025) and Hopenhayn and Okumura (2025), who develop models of endogenous growth that are 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 where firms' R&D expenditures enhance productivity, we develop instead a discrete-time neoclassical investment application, where firms invest to change their scale. 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, 2026; 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 revived by the recent debate over the rise of market power (Syverson, 2019; Syverson, 2025) 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 in investment. Our quantitative analyses also complement existing methodological studies on the estimation of neoclassical investment models, including Erickson and Whited (2000), Erickson and Whited (2002) and Cooper and Haltiwanger (2006), among others.

Our model 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 (Autor, Dorn, Katz, Patterson, and Van Reenen, 2020; De Loecker, Eeckhout, and Unger, 2020; Kwon, Ma, and Zimmermann, 2024; Hoberg and Phillips, 2025), and the relationship between market power and stock market valuations (Bustamante and Donangelo, 2017; Cho, Grotteria, Kremens, and Kung, 2023; Eeckhout, 2025; Corhay, Kung, and Schmid, 2025; Boppart, Klenow, Laski, and Li, 2025), as well as the corporate finance and macro-finance literatures on how product market competition shapes firms' investment and financial policies (Hoberg and Phillips, 2010b; Hoberg and Phillips, 2010a; Dou, Ji, and Wu, 2021).

Because the model we describe is a type of network game, our paper connects to a burgeoning 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), Carvalho and Gabaix (2013), Bloom, Schankerman, and Van Reenen (2013), Carvalho, Nirei, Saito, and Tahbaz-Salehi (2021), Baqaee and Farhi (2020, 2024), Baqaee (2018) and Flynn, Patterson, and Sturm (2022). 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. In our quantitative analysis, we exploit the regularity that productivity distributions exhibit heavy right tails which are well-approximated by Pareto distributions, consistent with Gabaix (2009).

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

2 The Network Q Model

In this section, we present our theory of corporate investment in oligopolistic product markets. We begin by describing the hedonic demand system that underpins product market competition between the firms in our model.

2.1 Hedonic Consumer Demand

Products and Characteristics. 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 supply and investment decisions at the level of business units, which may be either independent from each other or owned by the same firm, as we elaborate below.

Each business unit produces a distinct good. We index each business unit/good combination with the subscript $n \in \{1,2,...,N\}$. Consumer demand follows the GHL demand framework of Pellegrino (2025). We model each product n as a bundle of characteristics, sorted into two distinct types. Characteristics that are common across all products are indexed by $j = 1,2,...,J$, whereas idiosyncratic characteristics are directly indexed by their product subscript n. We define a_{jn} as the quantity of common characteristic j provided by product n. Every product is represented by a J-dimensional vector $\mathbf{a}_n$ whose length is normalized to one:

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

$$\text{with} \qquad \sum_{j=1}^{J} a_{jn}^2 = 1 \quad \forall\, n \in \{1,2,...,N\} \tag{2}$$

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

$$\mathbf{A} \equiv \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \cdots & \mathbf{a}_N \end{bmatrix} \equiv \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1N} \\ a_{21} & a_{22} & \cdots & a_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ a_{J1} & a_{J2} & \cdots & a_{JN} \end{bmatrix} \tag{3}$$

Time is discrete and indexed by t. 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 j at time t as x_{jt}, then:

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

For idiosyncratic 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 and the total quantity of the idiosyncratic characteristic delivered by product line n.

Consumer Utility. Our model employs a representative consumer framework, with utility that is quadratic in both the common characteristics (x_{jt}) and the idiosyncratic characteristics (equal to output y_{nt}) associated with each of $n = 1,...,N$ products in any period t. The representative consumer's utility function is:

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

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 faces a vector of product prices $\mathbf{p}_t = [p_{1t}, p_{2t}, ..., p_{Nt}]'$ and takes it as given. The consumer is also endowed with the shares of all the firms in the economy and thus he/she faces the following budget constraint:

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

where Π_t is aggregate profits and the price of labor is normalized to 1 (labor is the numéraire good). Because utility is linear in L_t and labor is the numéraire, the Lagrange multiplier on the budget constraint equals one, so the consumer's problem reduces to the maximization of consumer surplus. The consumer's demand for each product n maximizes (5) subject to (6). As shown by Pellegrino (2025), in matrix form, the inverse demand function 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

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

$\mathbf{I}$ is an $N \times N$ identity matrix and the matrix $\mathbf{\Sigma}$ is symmetric and by definition 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.

Cosine Similarity. The key input to construct $\mathbf{\Sigma}$ is the matrix $\mathbf{A}'\mathbf{A}$, whose (n,m) entry 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 J-dimensional space of common characteristics. For this reason, $\mathbf{a}_n'\mathbf{a}_m$ is called the *cosine similarity* between products n and m. Because the quantities a_{jn} are non-negative, 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 idiosyncratic characteristics when

evaluating utility. Second, through equation (9), it controls the overall degree of substitution across all products. 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

Production Technology. Each product in the economy is produced by a single business unit n, that is administered by a manager. Each business unit uses capital and a composite variable input which we label "labor" for convenience 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:[4] in order to produce a unit of product n, its producer 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 number of hours allocated to business unit n by the representative consumer, z_n represents the productivity of capital of business unit n, and c_n represents the unit labor 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., it includes both tangible and intangible types of capital).

Law of Motion of Capital and Installation Costs. In each period t, the manager of business unit 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:

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

The first term on the right hand side of (12) is the cost of purchasing new capital; p_n^k is the price of capital for business unit n, assumed to be exogenous. The second term is a quadratic adjustment cost with curvature parameter $\theta > 0$. We interpret the function $\psi_n(k_{nt+1}; k_{nt})$ 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

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

of heterogeneity in the effective burden of adjustment costs, which vary substantially across business units through differences in the productivity of capital z_n. Since adjusting output by Δy requires adjusting capital by $\Delta y/z_n$, the cost of a given output adjustment is $\frac{\theta}{2}(\Delta y/z_n)^2$: in output units, the effective adjustment cost parameter is θ/z_n^2. All else equal, business units with low productivity z_n therefore find it relatively costlier to adjust their output, making them sluggish in responding to changing competitive conditions. Conversely, business units with high productivity z_n attain the same change in output with a smaller change in capital, allowing them to adjust their capacity more nimbly over time. Furthermore, in Section 5 we show how the model can be extended to allow for unit-specific adjustment cost parameters θ_n.

Operating Profits. Next we define π_{nt}, the operating profit of business unit n at time t, as follows:

$$\pi_{nt} \stackrel{\text{def}}{=} (p_{nt} - c_n)\, y_{nt} \tag{13}$$

Substituting the price p_{nt} in terms of the output y_{nt} (imposing consumer maximization) and y_{nt} in terms of the capital stock k_{nt}, we can re-write π_n as a (linear-quadratic) function of the vector of capital stocks:

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

where $\mathbf{k}_t$ is the vector of capital stocks (one row for each business unit) and σ_{nm} is the (n,m) entry of $\mathbf{\Sigma}$ 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 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).

2.3 The Manager's Objective

Next, we define the manager's objective. In doing so, we allow the manager of business unit n to internalize the spillovers of his/her supply choices on the profits of other business units $m \neq n$. This flexibility allows us to model complex ownership structures (multi-product firms or mergers) or alternative 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 the product market spillovers that they generate on related peers (Baker and Bresnahan, 1985).

The optimization problem of the manager n at time t consists of maximizing the sum of the discounted future cash flows of all relevant business units m that belong to a set that we denote by $\mathcal{M}_n$. This set includes the focal business unit n as well as peer business units that belong to the same firm

or cartel. The manager solves this maximization problem by choosing the capital stock of his/her own business unit, k_{nT}, at every future period $T > t$. In so doing, the manager takes the choice of all rival business units as given:[5]

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

where $v_n(\mathbf{k}_t)$ represents the *value function* of manager n at time t. The formulation in (17) implies the value function at time t is expressed *cum-dividend*. By construction, the managers of business units within the same firm or cartel maximize identical objective functions. Hence, while each chooses independently the investment plan of his/her business unit, cannibalization between related business units is fully internalized. This setup is equivalent to a single agent jointly planning the investment decisions of all business units within the same firm or cartel.

Manager n solves problem (17) subject to equations (11), (12), (13), and a transversality condition:

$$\lim_{T \to \infty} (1+r)^{-(T-t)} v_n(\mathbf{k}_T) = 0. \tag{18}$$

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. The sets $\mathscr{M}_n$ similarly define a network of ownership/collusion relationships between business units that can be summarized in a second adjacency matrix that we call $\mathbf{\Gamma}$: the (n,m) entry γ_{nm} is an indicator/dummy parameter that equals 1 if $m \in \mathscr{M}_n$ and $n \in \mathscr{M}_m$. We refer to the symmetric matrix $\mathbf{\Gamma}$ as the *cross-internalization matrix.* Because the sets $\mathscr{M}_n$ group business units into firms or cartels, $\mathbf{\Gamma}$ is block-diagonal (up to a permutation of the business-unit indices).

2.4 Equilibrium Concepts and Dynamic Stability

In solving our investment game, we focus on Markov strategy profiles, as is standard in the literature on dynamic oligopoly (Reynolds, 1991). Specifically, the investment policy of manager n takes the form:

$$k_{nt+1} = f_n(\mathbf{k}_t) \tag{19}$$

hence, next period capital k_{nt+1} depends only on the contemporaneous vector of capital stocks: the previous history of play is irrelevant for the future evolution of the system. Next, we establish a couple of additional useful definitions.

Definition 1. *For any Markov strategy profile, we define a steady state as a vector of capital stocks* $\mathbb{k}$ *such that* $\mathbf{f}(\mathbb{k}) = \mathbb{k}$.

Next, we follow Jun and Vives (2004) and focus our attention on strategy profiles that feature linear investment policies, and which impose the dynamic stability of the process that describes the evolution of the state vector $\mathbf{k}_t$:

Definition 2. *A Stable Linear Markov Strategy Profile is a Markov strategy profile with linear in-*

[5] The latter statement requires further clarification that we will provide in the next subsection. As we shall see, it will matter for the equilibrium definition.

vestment policies where $\mathbf{f}(\mathbf{k}_t)$ *takes the form*[6]

$$\mathbf{f}(\mathbf{k}_t) - \mathbb{k} \;=\; \mathbf{\Phi}(\mathbf{k}_t - \mathbb{k}) \tag{20}$$

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

By combining equation (20) 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:

$$\boldsymbol{i}_t(\mathbf{k}_t) \;=\; \mathbf{\Delta}\mathbf{k}_t + (\mathbf{I} - \mathbf{\Phi})(\mathbb{k} - \mathbf{k}_t) \tag{21}$$

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

$$\mathbf{\Delta} \stackrel{\text{def}}{=} \begin{bmatrix} \delta_1 & 0 & \dots & 0 \\ 0 & \delta_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \delta_N \end{bmatrix} \tag{22}$$

Dynamic Stability. As we elaborate further below, the matrix $\mathbf{\Phi}$ governs the dynamics of the capital accumulation process for all business units, as well as the spillover effects of a manager's own investment decision on his/her peers' investment. Our restriction on linear, stable, Markov strategy profiles implies that the vector of capital stocks $\mathbf{k}_t$ follows a vector auto-regression (VAR) process. The VAR process is stable in a time-series sense because the eigenvalues of matrix $\mathbf{\Phi}$ lie in the interval (0,1): this property ensures that, for every initial vector of capital stocks, operating profits do not diverge, thus enforcing the transversality condition. This restriction implies that the vector of capital stocks gradually converges to its steady-state level along the equilibrium 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 business unit n (the n^{th} entry of $\mathbb{k}$), so that:

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

It is well known that, even after restricting attention to stable linear Markov strategies, dynamic linear-quadratic games can admit multiple Nash equilibria (see Engwerda, 2005, 2016; Nortmann et al., 2025). In what follows, we introduce two stable linear Markov-Nash equilibrium concepts that we will utilize in the rest of the paper.

Markov Perfect Equilibrium (MPE). An equilibrium concept that is standard in dynamic games is Markov Perfect Equilibrium (MPE). In order to provide a formal equilibrium definition, we re-write the manager's maximization problem in (17) in recursive form as a dynamic program, with the

[6]Notice that this definition of "linear" is without loss of generality. Specifically, expressing linear strategies as in (20) is the same as: 1) writing the more general expression $\mathbf{f}(\mathbf{k}_t) = \boldsymbol{\iota} + \mathbf{\Phi}\mathbf{k}_t$; 2) noticing that the unique steady state is given by the vector $\mathbb{k} = (\mathbf{I} - \mathbf{\Phi})^{-1}\boldsymbol{\iota}$; 3) substituting vector $\boldsymbol{\iota}$ in terms of $\mathbb{k}$ inside $\mathbf{f}(\mathbf{k}_t) = \boldsymbol{\iota} + \mathbf{\Phi}\mathbf{k}_t$ to obtain (20).

vector $\mathbf{k}_t$ as the state variable:

$$v_n(\mathbf{k}_t) \;=\; \max_{f_n(\mathbf{k}_t)} \sum_{m \in \mathscr{M}_n} \left[\pi_m(\mathbf{k}_t) - \psi_m(f_m(\mathbf{k}_t); k_{mt})\right] + \frac{v_n(\mathbf{f}(\mathbf{k}_t))}{1+r} \tag{25}$$

To find an equilibrium of the game means to solve the system of value functions

$$\mathbf{v}(\mathbf{k}_t) \;\equiv\; [v_1(\mathbf{k}_t), v_2(\mathbf{k}_t), ..., v_N(\mathbf{k}_t)]' \tag{26}$$

for the associated investment policies

$$\mathbf{f}(\mathbf{k}_t) \;\equiv\; [f_1(\mathbf{k}_t), f_2(\mathbf{k}_t), \ldots, f_N(\mathbf{k}_t)]' \tag{27}$$

The key defining feature of MPE is that each manager knows and takes as given the competitors' policy functions ($\mathbf{f}_{-n}$). Next, we provide a formal definition.

Definition 3. *A Markov Perfect Equilibrium in stable linear strategies consists of a vector of value functions* $\mathbf{v}(\mathbf{k}_t)$ *and a vector of investment policies* $\mathbf{f}(\mathbf{k}_t)$ *from Definition 2 that solve the Bellman equation (25) maximization problem for all business units n, subject to*

$$k_{mt+1} \;=\; \begin{cases} i_{mt} + (1-\delta_m) k_{mt} & \text{for} \quad m = n \\ k_{m\infty} + \boldsymbol{\varphi}'_m (\mathbf{k}_t - \mathbb{k}) & \text{for} \quad m \neq n \end{cases} \tag{28}$$

where $\boldsymbol{\varphi}'_m$ *denotes the* m^{th} *row of* $\boldsymbol{\Phi}$*, and taking the competitors' policy functions* $\mathbf{f}_{-n}(\mathbf{k}_t)$ *as given. Moreover we require that for all business units the matrix* $\boldsymbol{\Phi}_{-n,-n}$ (*obtained by removing row and column n from* $\boldsymbol{\Phi}$) *has a spectral radius lower than* $1+r$*.*

In an MPE, managers internalize how rivals' future actions respond to every counterfactual deviation. On the one hand, this is desirable because it makes MPE subgame perfect. On the other hand, this leads to a loss of tractability that is well known and documented in the literature on dynamic linear-quadratic games (see, e.g., Engwerda, 2005; Monti et al., 2024; Nortmann et al., 2025). It leads to equilibrium multiplicity, and specifically to the existence of an unknown, potentially infinite number of equilibria. Moreover we shall see shortly that, while equations that characterize the set of MPE can be derived, they lack analytical solutions and cannot be solved computationally either when the number of players becomes large.

Proximate Markov Equilibrium (PME). To overcome these challenges, the previous literature has utilized MPE approximations that depart from rational expectations. We propose here instead an alternative equilibrium concept that approximates MPE, preserves rational expectations and can be constructed from established equilibrium concepts: the Proximate Markov Equilibrium (PME).

The PME is a Nash equilibrium in stable linear strategies that admit a Markov representation, where managers solve their dynamic program by taking as given their rivals' chosen sequence of capital stocks, denoted by $\mathbf{k}_{-n}^{(t)}$

$$\mathbf{k}_{-n}^{(t)} \;\stackrel{\text{def}}{=}\; \{\mathbf{k}_{-nt}, \mathbf{k}_{-nt+1}, \mathbf{k}_{-nt+2}, ...\} \tag{29}$$

Compare this with MPE, where managers take as given instead their rivals' policy functions $\mathbf{f}_{-n}(\cdot)$.

Our primary justification for the PME is an approximation result. Informally speaking, we show that

the PME equilibrium path first-order approximates the locally unique stable MPE branch around a disconnected product network ($\mathbf{\Sigma} = \mathbf{0}$), where all MPEs coincide with the PME. Proposition 3 in Section 2.6 will present this result formally.

The assumption that manager n takes $\mathbf{k}_{-n}^{(t)}$ as given makes the PME a type of *Open-Loop Equilibrium* (Reynolds, 1991). The key difference between the PME and a standard Open-Loop equilibrium lies in how we micro-found this assumption. The early dynamic duopoly literature micro-founds this assumption in one of two ways: 1) the managers pre-commit to their respective investment plans; 2) managers observe their rivals' planned investment, but not the investment rule $\mathbf{f}_{-n}(\mathbf{k}_t)$ that determines it. We set aside the first microfoundation, as it is clearly unpalatable in our setting.

The second microfoundation has been criticized, in the dynamic duopoly literature, on the grounds that it makes Open-Loop equilibria self-disproving: if indeed all firms follow the Open-Loop strategy profile, then the future path of capital stocks (known in equilibrium by all players) reveals the rivals' policy functions. Any manager n is able (and has the incentive) to reverse-engineer the transition matrix $\mathbf{\Phi}$, and then use it to devise a profitable unilateral deviation.

PME equips the (unique) Open-Loop equilibrium with a technical condition (called *unrecoverability*) that shores up its second microfoundation. It is based on the simple observation that, in the presence of a large number of players and an asymmetric product network, managers cannot reverse-engineer their rivals' policy functions. Technically, we say that $\mathbf{\Phi}$ is *ε-unrecoverable*: for a given computation precision ε, the transition matrix $\mathbf{\Phi}$ cannot be econometrically recovered from the sequence of future equilibrium capital stocks. We prove that this condition is satisfied when the number of business units N exceeds a certain threshold. This requirement is met by a wide margin in our quantitative implementation.[7] For brevity, we provide the formal definition of unrecoverability in the Appendix.

Intuitively, in a PME profitable deviations do exist off the equilibrium path, but they cannot be exploited: players take their rivals' investment plans as given because the size and complexity of the game make it impossible to recover their rivals' investment policy. The idea that players' beliefs cannot be disproven by equilibrium play makes our PME concept closely related to *Self-Confirming Equilibria* (Fudenberg and Levine, 1993).

Under these assumptions, the manager's maximization problem in (17) can be re-written in recursive form as:

$$v_n\left(k_{nt};\mathbf{k}_{-n}^{(t)}\right) = \max_{f_n(\mathbf{k}_t)} \sum_{m\in\mathcal{M}_n} \left[\pi_m(\mathbf{k}_t) - \psi_m(f_m(\mathbf{k}_t);k_{mt})\right] + \frac{v_n\left(f_n(\mathbf{k}_t);\mathbf{k}_{-n}^{(t+1)}\right)}{1+r} \tag{30}$$

the value function of each manager n has the stock of current capital k_{nt} as the state variable, and is conditioned by the entire path of rivals' capital stocks.[8] As before, finding an equilibrium of the game means finding a vector of policy functions $\mathbf{f}(\mathbf{k}_t)$ that represents the equilibrium path in Markov form and solves the system of N stacked Bellman equations in (30). We next provide a formal definition.

Definition 4. *A Proximate Markov Equilibrium* (*PME*) *with precision ε consists of a vector of value functions $\mathbf{v}(\mathbf{k}_t)$ and a vector of investment policies $\mathbf{f}(\mathbf{k}_t)$ from Definition 2 that solve the manager's*

[7] At double-precision accuracy ($\varepsilon \approx 2.2\times10^{-16}$), the threshold is 35 firms in 1995 and 31 firms in 2021, against cross sections of 5,707 and 3,157 firms respectively.

[8] For convenience, we use the compact notation $\mathbf{v}(\mathbf{k}_t)$ for the vector of value functions of both the PME and the MPE, suppressing the additional conditioning arguments in the PME. This is without loss of generality: the path $\mathbf{k}_{-n}^{(t)}$ is correctly forecasted in equilibrium, and differentiation of $v_n(\mathbf{k}_t)$ with respect to k_{nt+1} in the PME takes that sequence as given.

maximization problem in (17) for all business units $n = 1,2,...,N$, subject to

$$\mathbf{k}_{t+1} = \boldsymbol{i}_t + (\mathbf{I} - \boldsymbol{\Delta})\mathbf{k}_t \tag{31}$$

taking the sequence of competitors' capital stocks $\{\mathbf{k}_{-nt}, \mathbf{k}_{-n,t+1}, \ldots\}$ as given, and where the transition matrix $\boldsymbol{\Phi}$ is ε-unrecoverable.[9]

The PME has several desirable properties: we can prove its existence and uniqueness, and the unique solution is available in closed form. Unlike similar tractable approximations in the literature, the PME can be micro-founded without invoking commitment or forcing players to hold inconsistent beliefs along the equilibrium path. Finally, it is far more robust than the MPE to the critique of Moll (2025), who highlighted how a strict enforcement of rational expectations in models characterized by significant heterogeneity can imply that model agents engage in implausibly complex computations.

2.5 Optimality Conditions and the Competitive Pressure Matrix $\boldsymbol{\Omega}$

Marginal Q: Definition. Ahead of solving the optimal investment decision of each manager, we define the *decision marginal* Q for each manager in the network, a concept which builds on the seminal contribution of Hayashi (1982). The *decision marginal* Q is the marginal continuation value of an additional unit of capital, as perceived by the manager of business unit n — formally:

$$q_n(\mathbf{k}_t) \overset{\text{def}}{=} \frac{\partial v_n(\mathbf{k}_t)}{\partial k_{nt}}. \tag{32}$$

We then use the vector $\mathbf{q}(\mathbf{k}_t)$ to denote the vector of decision marginal Qs (one entry for each business unit):

$$\mathbf{q}(\mathbf{k}_t) \equiv \begin{bmatrix} q_1(\mathbf{k}_t) & q_2(\mathbf{k}_t) & \cdots & q_N(\mathbf{k}_t) \end{bmatrix}' \tag{33}$$

First Order Condition. Next, using the definition of decision marginal Q in equation (32) we write the (necessary) first-order condition for a solution to the unit n's optimization problem:

$$\frac{q_n(\mathbf{k}_{t+1})}{1+r} = p_n^k + \theta(k_{nt+1} - k_{nt}) \tag{34}$$

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 paid today. The function $q_n(\mathbf{k}_{t+1})$ depends not only on the capital stock of business unit n itself, but also on the capital stocks of all other business units.

Notice that the exact expression for $q_n(\mathbf{k}_{t+1})$ depends on the applicable equilibrium concept. In order to write a generic expression and highlight the differences between MPE and PME, we use a dummy parameter — $\mathbb{1}_{\text{MP}}$ — which equals one in MPE, and zero in PME. The left hand side of the first-order condition in (34) is then equal to:

$$q_n(\mathbf{k}_{t+1}) = \frac{\partial \sum_{m \in \mathcal{M}_n} \pi_m(\mathbf{k}_{t+1})}{\partial k_{n,t+1}} + (1-\delta_n)p_n^k + \theta\left(k_{n,t+2} - k_{n,t+1}\right) + \frac{\mathbb{1}_{\text{MP}}}{1+r} \cdot \underbrace{\frac{\partial v_n(\mathbf{k}_{t+2})}{\partial \mathbf{k}_{-n,t+2}} \cdot \frac{\partial \mathbf{k}_{-n,t+2}}{\partial k_{n,t+1}}}_{\text{Markov-Perfect Feedback}} \tag{35}$$

[9] Without the unrecoverability assumption, the resulting equilibrium is simply an Open-Loop equilibrium in linear stable Markov strategies.

The first additive term of this expression is the internalized profit generated by an additional unit of capital at time $t+1$. The second is equal to the resale value at $t+1$ of the undepreciated portion of the same unit of capital. The third term is the marginal change in adjustment costs at $t+1$. The last term encapsulates the economic distinction between the MPEs and the PME in the model. The term represents the indirect (feedback) effect of the investment decision of business unit n in period t on its continuation value in period $t+2$, mediated by the reactions of all its rivals. In the PME ($\mathbb{1}_{\text{MP}} = 0$) the manager is unable to internalize these feedback effects. By contrast, in the case of MPEs ($\mathbb{1}_{\text{MP}} = 1$), managers know their rivals' complete policy functions and internalize how an investment deviation changes rivals' actions from period $t+1$ onward. Both concepts use Markov policies; they differ only in whether firms are able to internalize feedback effects when the number of firms in the network is sufficiently large.

The Competitive Pressure Matrix. Stacking all firms in vector notation, we can then express the first term of equation (35) as follows:

$$\left[\frac{\partial \sum_{m\in\mathscr{M}_1}\pi_m(\mathbf{k}_{t+1})}{\partial k_{1t+1}},\ldots,\frac{\partial \sum_{m\in\mathscr{M}_N}\pi_m(\mathbf{k}_{t+1})}{\partial k_{Nt+1}}\right]' = \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\,\mathbf{k}_{t+1} \tag{36}$$

The matrix $\mathbf{\Omega}$, which we call hereafter the *Competitive Pressure* matrix, is defined as:

$$\mathbf{\Omega} \overset{\text{def}}{=} \mathbf{Z}(\mathbf{I}+\mathbf{\Sigma}+\mathbf{D}+\mathbf{\Gamma}\circ\mathbf{\Sigma})\mathbf{Z} \tag{37}$$

where 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 the corresponding business unit acts as an oligopolist (and, thus, internalizes that it is facing a downward sloping demand curve):

$$\mathbf{D} \overset{\text{def}}{=} \begin{bmatrix} \mathscr{D}_1 & 0 & \ldots & 0 \\ 0 & \mathscr{D}_2 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & \mathscr{D}_N \end{bmatrix}, \qquad \mathscr{D}_n \overset{\text{def}}{=} \begin{cases} 0 & \text{if firm } n \text{ is a price-taker} \\ 1 & \text{if firm } n \text{ is an oligopolist} \end{cases} \tag{38}$$

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

$$\mathbf{\Gamma} \overset{\text{def}}{=} \begin{bmatrix} \gamma_{11} & \gamma_{12} & \cdots & \gamma_{1N} \\ \gamma_{21} & \gamma_{22} & \cdots & \gamma_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{N1} & \gamma_{N2} & \cdots & \gamma_{NN} \end{bmatrix} \qquad \gamma_{nm} \overset{\text{def}}{=} \begin{cases} 0 & \text{if } n \notin \mathscr{M}_m,\ m \notin \mathscr{M}_n \\ 1 & \text{if } n \in \mathscr{M}_m,\ m \in \mathscr{M}_n \end{cases} \tag{39}$$

The formulation of matrix $\mathbf{\Omega}$ is 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 quantity oligopoly 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.

The Competitive Pressure matrix $\mathbf{\Omega}$ is key to our modeling framework: its (n,m) entry is, under oligopolistic conduct, the negative of the second-order derivative

$$\frac{\partial^2\left(\sum_{n'\in\mathcal{M}_n}\pi_{n't+1}\right)}{\partial k_{nt+1}\partial k_{mt+1}} = -\omega_{nm} \tag{40}$$

$\mathbf{\Omega}$ is therefore a measure of the *concavity* of the firm's (internalized) profit functions: its n^{th} diagonal entry equals $2z_n^2$ for the simple oligopolistic conduct, and z_n^2 when firms are price takers. 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, $\omega_{nm}^2/(2\omega_{nn})$ — which equals $\omega_{nm}^2/\left(4z_n^2\right)$ under oligopolistic conduct — 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 positive definite: All of its eigenvalues are strictly positive. This assumption is always satisfied in our empirical implementation.

2.6 Markov Equilibria: Characterization, Existence, Uniqueness

In this subsection, we first characterize all linear stable MPEs and the PME of the model using a common analytical framework. We then proceed to establish the existence and uniqueness of the linear stable PME, for which we obtain a convenient analytical solution.

Proposition 1 (Stable Linear Markov Equilibria: Necessary and Sufficient Conditions)**.** *Consider the linear investment policy in equation (20). All MPEs and PMEs are characterized by the following necessary and sufficient conditions, which include an expression for* $\mathbb{k}$

$$\mathbb{k} = \left[\mathbf{\Omega} + \mathbb{1}_{\text{MP}}\cdot\mathbf{S}_1(\mathbf{\Phi},\mathbf{\Omega})\right]^{-1}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^k\right] \tag{41}$$

and the following matrix equation, which pins down $\mathbf{\Phi}$ *as an implicit function of* $(\mathbf{\Omega},\theta,r)$

$$\left[\mathbf{I} - \frac{\mathbb{1}_{\text{MP}}}{\theta}\cdot\mathbf{S}_2(\mathbf{\Phi},\mathbf{\Omega})\right]\mathbf{\Phi}^2 - \left[\frac{1}{\theta}\mathbf{\Omega} + (2+r)\mathbf{I}\right]\mathbf{\Phi} + (1+r)\mathbf{I} = \mathbf{0} \tag{42}$$

where the matrices $\mathbf{S}_1(\mathbf{\Phi},\mathbf{\Omega})$ *and* $\mathbf{S}_2(\mathbf{\Phi},\mathbf{\Omega})$ *are non-linear functions of* $(\mathbf{\Phi},\mathbf{\Omega})$ *that capture Markov-perfect strategic feedback incentives, for which we provide formulas in Appendix A for the case of single-unit managers* $(\mathbf{\Gamma}=\mathbf{I})$*; the conditions for the PME, obtained by setting* $\mathbb{1}_{\text{MP}}=0$*, hold for a general cross-internalization matrix* $\mathbf{\Gamma}$*. If* $\mathbb{1}_{\text{MP}}=1$*, these conditions further include the requirement that* $\mathbf{\Phi}^{-1}$ *has strictly positive diagonal entries.*[10]

Proof. See Appendix A. □

Hence, the MPEs and the PME are characterized by identical equations for the transition matrix and steady-state capital, except for the feedback effects embedded in the matrices $\mathbf{S}_1(\mathbf{\Phi},\mathbf{\Omega})$ and $\mathbf{S}_2(\mathbf{\Phi},\mathbf{\Omega})$, which only appear in the equations that characterize the MPEs (the dummy variable $\mathbb{1}_{\text{MP}}$ "turns on/off" such feedback effects).

[10] If $\mathbb{1}_{\text{MP}}=0$ this requirement is automatically implied by equation (42) alongside dynamic stability, and thus need not be imposed separately (see the proof of Proposition 2).

The terms $\mathbf{S}_1(\mathbf{\Phi},\mathbf{\Omega})$ and $\mathbf{S}_2(\mathbf{\Phi},\mathbf{\Omega})$ are the source of equilibrium multiplicity and numerical infeasibility in the MPE case ($\mathbb{1}_{\text{MP}} = 1$). They are infinite converging sums that depend on an $N \times N \times N$ array, where $N = 4{,}000$ in a typical year of our working sample. Computing the investment policy requires solving a system of N coupled $N \times N$ Algebraic Riccati Equations. The undesirable multiplicity of the MPEs is a standard result in dynamic linear-quadratic games (Engwerda, 2005; Nortmann et al., 2025) and in option-exercise games such as Dai, Jiang, and Wang (2024).

Equations (41) and (42) clearly illustrate why we previously referred to PME as a "minimally invasive" approximation of the MPE. Standard approximate equilibrium concepts restore tractability by imposing assumptions that remove Markov Perfect feedback effects, while also introducing additional deformations in the underlying economics. Our Proximate Markov Equilibrium accomplishes tractability by removing selectively the feedback effects ($\mathbb{1}_{\text{MP}} = 0$) while leaving all of the other underlying MPE economic forces unaffected.

We proceed to establish the existence and uniqueness of the linear stable PME of our model, and then characterize it analytically.

Proposition 2 (Existence and Uniqueness of the Analytical PME)**.** *Consider the linear investment policy in equation (20). There exists a unique PME, characterized by the following explicit expressions for the steady-state capital stock vector*

$$\mathbb{k} \;=\; \mathbf{\Omega}^{-1}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^k\right] \tag{43}$$

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

$$\mathbf{\Phi} \;=\; \frac{1}{2}\left[\frac{1}{\theta}\mathbf{\Omega} + (2+r)\mathbf{I}\right] - \sqrt[+]{\frac{1}{4}\left[\frac{1}{\theta}\mathbf{\Omega} + (2+r)\mathbf{I}\right]^2 - (1+r)\mathbf{I}} \tag{44}$$

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.*[11]

Proof. The steady-state vector of capital stocks is trivially found by imposing $\mathbb{1}_{\text{MP}} = 0$ and $\mathbf{k}_{t+2} = \mathbf{k}_{t+1} = \mathbf{k}_t = \mathbb{k}$ inside equations (34), (35) and (36). See Appendix A for the rest of the proof. □

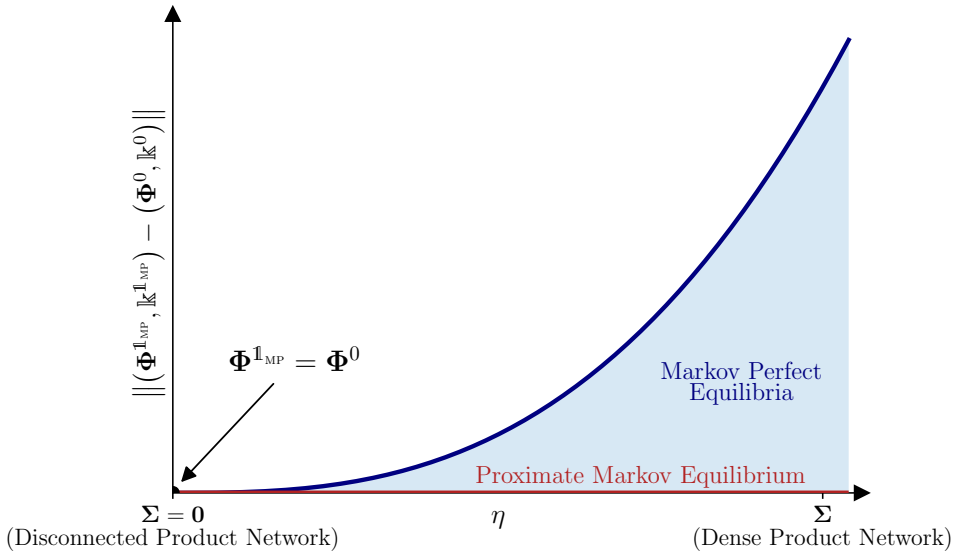
The equilibrium transition matrix $\mathbf{\Phi}$ in equation (44) is the unique, explicit solution of a single Nonsymmetric Algebraic Riccati Equation (NARE). Neoclassical models typically solve for investment policies as a function of firms' marginal Q. They rarely provide explicit solutions for firms' investment policies, value functions and marginal Q. The PME of 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 (44) capture how fast firm n's capital stock gets pulled toward its steady state, whereas the off-diagonal terms capture instead the equilibrium investment spillovers across business units.

Notice that the PME steady state in equation (43) does not depend on the adjustment cost parameter θ: it coincides exactly with the static network Cournot equilibrium of Pellegrino (2025), once production costs are augmented by the user cost of capital $(r+\delta_n)p_n^k$. Adjustment costs thus shape

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

FIGURE 1: ILLUSTRATION — FIRST-ORDER EQUIVALENCE OF PME AND MPE



the transition path, but not the long-run allocation of capital. As in that static game, the steady-state equilibrium capital stock vector $\mathbb{k}$ admits an intuitive interpretation from network theory. The competitive pressure matrix $\boldsymbol{\Omega}$ can also be used (like $\boldsymbol{\Sigma}$ and $\boldsymbol{\Gamma}$) 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.

The analytical PME also provides a local approximation to MPE. To formalize this result, we consider a perturbation $\eta\boldsymbol{\Sigma}$ of the matrix $\boldsymbol{\Sigma}$ around $\boldsymbol{\Sigma} = \mathbf{0}$. At $\eta = 0$, the product network is disconnected, the terms $\mathbf{S}_1(\boldsymbol{\Phi}, \boldsymbol{\Omega})$ and $\mathbf{S}_2(\boldsymbol{\Phi}, \boldsymbol{\Omega})$ collapse to zero and all MPEs coincide with the PME. The following proposition shows that their transition matrices and steady states continue to agree to first order as product network links are introduced: the strategic-feedback terms are second order in their strength.

Proposition 3 (First-Order Equivalence of PME and MPE)**.** *Consider a perturbation $\eta\boldsymbol{\Sigma}$ of the matrix $\boldsymbol{\Sigma}$ around the disconnected product network $\boldsymbol{\Sigma} = \mathbf{0}$.*[12] *Assume that the eigenvalues of $\boldsymbol{\Omega}(0)$ are distinct.*[13] *Let $\left(\boldsymbol{\Phi}^0(\eta), \mathbb{k}^0(\eta)\right)$ denote the unique stable PME pair at η, and let $\mathscr{K}^1(\eta)$ denote the corresponding set of stable $(\boldsymbol{\Phi}, \mathbb{k})$ pairs satisfying the MPE equations in Proposition 1. There exist $\bar{\eta} > 0$ and a neighborhood $\mathscr{N}$ of $\left(\boldsymbol{\Phi}^0(0), \mathbb{k}^0(0)\right)$ such that, for every $\eta \in [0, \bar{\eta})$, there is a unique stable MPE solution inside $\mathscr{N}$*

$$\mathscr{K}^1(\eta) \cap \mathscr{N} = \left\{\left(\boldsymbol{\Phi}^1(\eta), \mathbb{k}^1(\eta)\right)\right\} \tag{45}$$

where the locally unique MPE pair is twice continuously differentiable in η. Around the disconnected

[12] We choose $\eta \geq 0$ sufficiently small that the corresponding competitive pressure matrix $\boldsymbol{\Omega}(\eta)$ remains positive definite.

[13] Since $\boldsymbol{\Sigma} = \mathbf{0}$ at $\eta = 0$, the matrix $\boldsymbol{\Omega}(0) = \mathbf{Z}(\mathbf{I} + \mathbf{D})\mathbf{Z}$ is diagonal, so this amounts to requiring that the scalars $z_n^2(1 + \mathscr{D}_n)$ be pairwise distinct — a condition that fails only on a set of measure zero. It is not needed for the PME, whose transition matrix is symmetric and therefore has real eigenvalues by construction.

product network ($\boldsymbol{\Sigma} = \mathbf{0}$)*, the PME solution approximates the MPE to first order — formally:*

$$\boldsymbol{\Phi}^1(0) = \boldsymbol{\Phi}^0(0), \qquad \mathbb{k}^1(0) = \mathbb{k}^0(0), \tag{46}$$

$$\boldsymbol{\Phi}^1(\eta) - \boldsymbol{\Phi}^0(\eta) = O\left(\eta^2\right), \qquad \mathbb{k}^1(\eta) - \mathbb{k}^0(\eta) = O\left(\eta^2\right). \tag{47}$$

Proof. See Appendix A. □

For all of the reasons discussed above, throughout the rest of the paper we focus on the Proximate Markov Equilibrium of the model. In the next subsections, we use equations (43) and (44) associated with the PME of the dynamic game to explicitly characterize the value function of the firm, investment rates and the associated marginal Q's.

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

We move next to discuss a novel and unique theoretical insight that emerges from our model: it turns out that the intensity of product market competition affects not only the equilibrium steady-state level of capital (43), but also the speed of convergence to the steady state.

As is well known from the macroeconomics literature[14], the speed of the transition along the equilibrium path is governed by the eigenvalues of the matrix $\boldsymbol{\Phi}$. Because business units differ by their capital productivity (z_n) and by the substitutability of the goods they produce (as measured by $\boldsymbol{\Sigma}$), the competitive pressure matrix $\boldsymbol{\Omega}$ represents a highly asymmetric network: some firms are significantly more exposed to competitive pressure than others. It follows that the transition matrix $\boldsymbol{\Phi}$ inherits rich spectral properties from $\boldsymbol{\Omega}$ — most importantly, vast dispersion in its eigenvalues.

Balanced Growth Paths. To analyze the spectral properties of matrix $\boldsymbol{\Phi}$, we introduce some more notation. Let $\mathbf{e}_n$ be the n^{th} eigenvector of $\boldsymbol{\Phi}$, 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 stocks converge towards the steady state at a rate that is exactly equal to eigenvalue $\hat{\varphi}_n$:

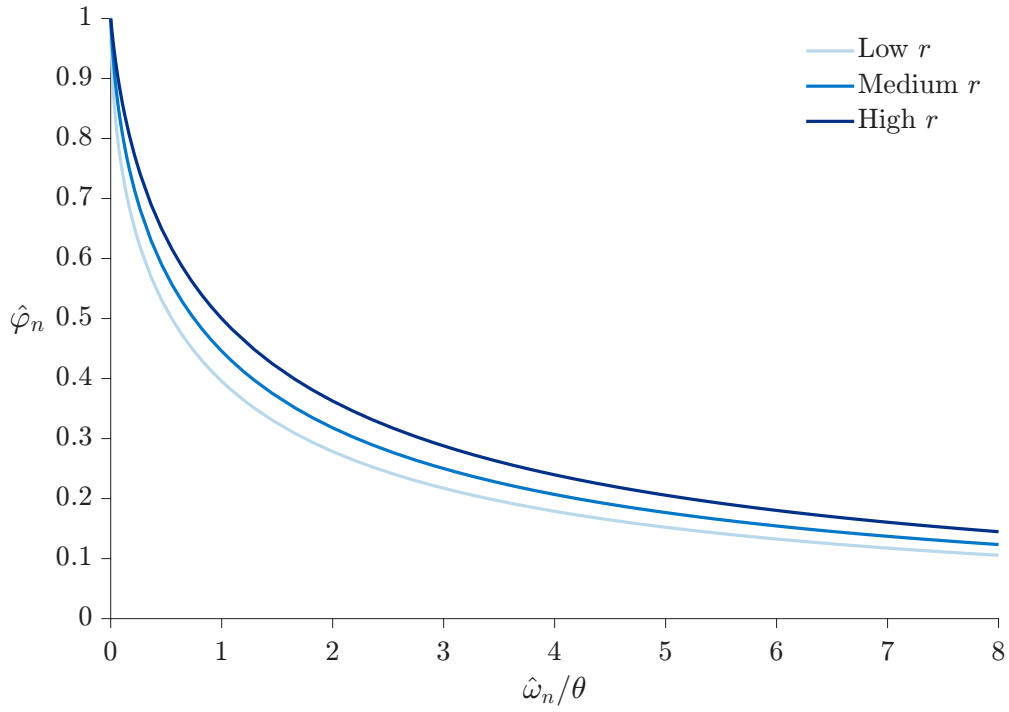
$$(\mathbf{k}_0 - \mathbb{k}) \propto \mathbf{e}_n \qquad \Longrightarrow \qquad (\mathbf{k}_t - \mathbb{k}) = \hat{\varphi}_n^t \, (\mathbf{k}_0 - \mathbb{k}) \tag{48}$$

Due to the rich structure of $\boldsymbol{\Omega}$, the transition matrix $\boldsymbol{\Phi}$ will generally have N different eigenvalues: this implies that, in our model, the PME 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 capital allocations (states) 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 are instead sticky (i.e., persistent), in the sense that firms will transition slowly to the steady state.

Firm-level Trade-off and Persistence. The rich spectral properties of $\boldsymbol{\Phi}$ also imply that, outside of the balanced growth path, individual firms will adjust their capital stock at markedly 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 firm n adjusts its capital stock depends on whether (at any

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

FIGURE 2: EIGENVALUES OF THE TRANSITION MATRIX



point in time) 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 44) depends directly on the competitive pressure matrix $\mathbf{\Omega}$. However, equation (44) 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 (to Proposition 2)**.** *Let $\hat{\boldsymbol{\omega}}$ be the vector of eigenvalues of $\mathbf{\Omega}$ and $\mathbf{E}$ be the corresponding (orthogonal) matrix of eigenvectors, so that*

$$\mathbf{\Omega} \;=\; \mathbf{E}\,\mathrm{diag}(\hat{\boldsymbol{\omega}})\,\mathbf{E}^{*} \tag{49}$$

and $\mathbf{E}^{*} \equiv \mathbf{E}' \equiv \mathbf{E}^{-1}$ is the transpose or inverse[15] of $\mathbf{E}$. Also, let $\hat{\boldsymbol{\varphi}}$ be the vector of eigenvalues of the PME transition matrix $\mathbf{\Phi}$. Then, $\mathbf{\Phi}$ and $\mathbf{\Omega}$ are simultaneously diagonalizable:[16]

$$\mathbf{\Phi} \;=\; \mathbf{E}\,\mathrm{diag}\left(\hat{\boldsymbol{\varphi}}\right)\mathbf{E}^{*} \tag{50}$$

[15] In this case the two coincide due to $\mathbf{E}$ being orthogonal.

[16] This is equivalent to saying that $\mathbf{\Phi}$ and $\mathbf{\Omega}$ share a basis or that they commute.

and their respective n^{th} eigenvalues are related by the following scalar function:

$$\hat{\varphi}_n \;=\; \frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)-\sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta}+2+r\right)^2-(1+r)} \tag{51}$$

which is simply the scalar equivalent of equation (44)*.*

Proof. Embedded in the proof of Proposition 2, in Appendix A. □

This corollary allows us to unpack the economic link between the competitive pressure matrix $\mathbf{\Omega}$ and the PME transition matrix $\mathbf{\Phi}$. Figure 2 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. 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 of adjustment costs, leading to higher persistence of capital.

As a novel insight, our model shows that $\hat{\varphi}_n$ is a decreasing function of the ratio $\hat{\omega}_n/\theta$. 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 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 PME presented in Proposition 2 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 single-product oligopolistic 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 4. *In a single-product oligopoly with* $\mathbf{D} = \mathbf{\Gamma} = \mathbf{I}$ *and N identical single-product firms,*

the eigenvalue of matrix $\mathbf{\Omega}$ *associated with the unit eigenvector* $\mathbf{1}$*, which we denote* $\hat{\omega}_{\mathbf{1}}$*, is given by*

$$\hat{\omega}_{\mathbf{1}} = z^2[2+(N-1)\sigma] \tag{52}$$

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

$$\frac{\partial\hat{\varphi}_{\mathbf{1}}}{\partial z}<0 \qquad \frac{\partial\hat{\varphi}_{\mathbf{1}}}{\partial\sigma}<0 \qquad \frac{\partial\hat{\varphi}_{\mathbf{1}}}{\partial N}<0 \tag{53}$$

Proof. See Appendix A. □

The proposition shows that lower dependence on capital (higher z), stronger substitution among products (σ), and 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 is well known from the previous literature: it is also regulated by the degree of product market competition.

2.8 Enterprise Value, Marginal Q and Value Function

Next, we compute two objects related to firm value. We previously called $v_n(\cdot)$ the *value function*: this is the objective function that the manager of business unit n maximizes when choosing investment. We use instead the notation $v_n^e(\cdot)$ to distinguish the *enterprise value* of business unit n, as a function of the current state: this is the present discounted value of operating profits net of investment costs, evaluated using the equilibrium strategy profile. While in the canonical neoclassical model these two objects coincide by construction, in our multi-firm setting we need to differentiate them. This is perhaps the most counterintuitive aspect of the equilibrium.

While it is still the case that, for a single-unit firm (as in the one-firm Q model), both quantities evaluate to the same dollar amount in equilibrium, in our model they are distinct *functions* of the state vector $\mathbf{k}_t$. The reason is that, in determining the optimal investment policy, the manager of business unit n differentiates $v_n(\cdot)$ with respect to $k_n^{(t)}$ while taking as given the competitors' future sequence of capital stocks $\mathbf{k}_{-n}^{(t)}$. Therefore, $v_n(\cdot)$ is a function of k_{nt} parameterized by $\mathbf{k}_{-n}^{(t)}$.

By contrast, the enterprise value $v_n^e(\cdot)$ returns the value of business unit n as a function of the current state $\mathbf{k}_t$ alone. It is obtained by writing the present discounted value of all future profits of n and substituting $\mathbf{k}_{t+1},\mathbf{k}_{t+2},...$ in terms of $\mathbf{k}_t$ using the equilibrium policy function. It therefore incorporates all feedback effects.

As a stepping stone to deriving closed-form expressions for both $v_n(\cdot)$ and $v_n^e(\cdot)$, we first compute their steady-state values. We start by noticing that, in the steady-state, the enterprise value of each business unit must be equal to its cum-dividend value, and therefore equal to the present discounted value of its future cash flows:

$$v_n^e(\mathbb{k}) = \frac{1+r}{r}\left[\pi_n(\mathbb{k}) - p_n^k\,\delta_n\,k_{n\infty}\right] \tag{54}$$

The yearly cash flow (in square brackets) is the unit's operating profit, less the cost of maintaining its capital stock (price of capital × depreciated capital). The steady-state value function of manager

n, in turn, is simply the sum of the enterprise values of all internalized business units:

$$v_n(\mathbb{k}) = \sum_{m \in \mathscr{M}_n} v_m^e(\mathbb{k}) \tag{55}$$

Both expressions hold for any conduct matrix $\mathbf{D}$ and cross-internalization matrix $\mathbf{\Gamma}$. By expanding the cash flow and substituting the previously-derived vector of steady-state capital stocks ($\mathbb{k}$) we obtain an even more compact expression for the steady-state firm values.

Proposition 5. *In the single-product Cournot case, the PME steady-state value function $v_n(\mathbb{k})$ is equal to the sum of the cum-dividend resale value of the steady-state capital stock and ρ_n — the present discounted value of future monopoly rents*:

$$v_n(\mathbb{k}) = v_n^e(\mathbb{k}) = (1+r)\, p_n^k\, k_{n\infty} + \rho_n \tag{56}$$

where ρ_n is the present value of monopoly rents in steady state, defined as follows:

$$\rho_n \overset{\text{def}}{=} \sum_{T=t}^{\infty} \frac{1}{(1+r)^{(T-t)}} \left[p_n(\mathbb{k}) - \frac{\partial p_n y_n}{\partial y_n}(\mathbb{k}) \right] y_n(\mathbb{k}) = \frac{1+r}{r} \cdot z_n^2 k_{n\infty}^2 \tag{57}$$

Proof. See Appendix A. □

Having solved for the value function $v_n(\mathbb{k})$ in steady state, we can now solve for $v_n(\cdot)$ and $v_n^e(\cdot)$ at a generic state. We first recover the vector of equilibrium firm values $\mathbf{v}^e(\mathbf{k}_t)$ by re-computing the present discounted value of the cash flows for each firm, and substituting future state vectors $\mathbf{k}_{T>t}$ in terms of today's state vector $\mathbf{k}_t$.

Proposition 6. *The vector of PME enterprise values is*:

$$\mathbf{v}^e(\mathbf{k}_t) = \mathbf{v}^e(\mathbb{k}) + \mathbb{Q}(\mathbf{k}_t - \mathbb{k}) - \operatorname{diag}\left[\mathbf{W}(\mathbf{k}_t)\, \mathbf{Z}(\mathbf{I}+\mathbf{\Sigma})\mathbf{Z} + \frac{\theta}{2}(\mathbf{I}-\mathbf{\Phi})\, \mathbf{W}(\mathbf{k}_t)\, (\mathbf{I}-\mathbf{\Phi}) \right] \tag{58}$$

where $\mathbb{Q}$ is the Jacobian of the enterprise value vector, evaluated at the steady state

$$\mathbb{Q} \overset{\text{def}}{=} \left. \frac{\partial \mathbf{v}^e(\mathbf{k})}{\partial \mathbf{k}} \right|_{\mathbf{k}=\mathbb{k}} = (1+r)\mathbf{P}^k - (\mathbb{K}\mathbf{Z}\mathbf{\Sigma}\mathbf{Z} - \operatorname{diag}[\mathbf{Z}(\mathbf{D}+\mathbf{\Gamma}\circ\mathbf{\Sigma}-\mathbf{I})\mathbf{Z}\mathbb{k}]) \left(\mathbf{I} - \frac{\mathbf{\Phi}}{1+r} \right)^{-1} \tag{59}$$

where $\mathbb{K} \overset{\text{def}}{=} \operatorname{diag}(\mathbb{k})$ and $\mathbf{P}^k \overset{\text{def}}{=} \operatorname{diag}(\mathbf{p}^k)$ are the steady-state capital vector and the vector of capital prices rearranged as diagonal matrices, and the (state-dependent) matrix $\mathbf{W}(\mathbf{k}_t)$ is defined as follows:

$$\mathbf{W}(\mathbf{k}_t) \overset{\text{def}}{=} \mathbf{E} \left\{ \left[\mathbf{E}^* (\mathbf{k}_t - \mathbb{k})(\mathbf{k}_t - \mathbb{k})' \mathbf{E} \right] \circ \left(\mathbf{1}\mathbf{1}' - \frac{\hat{\boldsymbol{\varphi}}\hat{\boldsymbol{\varphi}}'}{1+r} \right)^{\circ(-1)} \right\} \mathbf{E}^* \tag{60}$$

The superscript operator $\circ(-1)$ denotes the entry-by-entry (Hadamard) reciprocal matrix.

Proof. See Appendix A. □

We now move on to the computation of the value function $v_n(\cdot)$. As an intermediate step, we first solve for the vector of decision marginal Q*s* which we denote by $\mathbf{q}(\mathbf{k}_t)$. Leveraging the expression for

$\mathbb{k}$ that we previously obtained, we can express $\mathbf{q}(\mathbf{k}_t)$ 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 7. *The PME vector of decision marginal Qs is equal to*:

$$\mathbf{q}(\mathbf{k}_t) \;=\; (1+r)\mathbf{p}^k + \mathbf{\Lambda}(\mathbb{k}-\mathbf{k}_t) \tag{61}$$

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

$$\mathbf{\Lambda} \;\stackrel{\text{def}}{=}\; \theta(\mathbf{I}-\mathbf{\Phi})+\mathbf{\Omega} \tag{62}$$

this implies that the corresponding steady-state value is

$$\mathbf{q}(\mathbb{k}) \;=\; (1+r)\mathbf{p}^k \tag{63}$$

Proof. See Appendix A. □

Because q_n is obtained by differentiating v_n, we must interpret it as the marginal value of a unit of capital for firm n holding competitors' investment plans fixed. Along the equilibrium path the matrix $\mathbf{\Lambda}$ defines a time-invariant linear mapping between the decision marginal Q of each manager and the deviation of the state vector $\mathbf{k}_t$ from its steady-state.

Armed with the vector of marginal Qs, we can now obtain the value function vector by integrating each entry $q_n(\cdot)$ with respect to k_{nt}.

Proposition 8. *The vector of PME value functions is*:

$$\mathbf{v}\left(\mathbf{k}_t;\mathbf{k}_{-n}^{(t)}\right) \;=\; \mathbf{v}(\mathbb{k})+(\mathbf{k}_t-\mathbb{k})\circ\left[(1+r)\mathbf{p}^k-\left(\mathbf{\Lambda}-\frac{1}{2}\mathbf{I}\circ\mathbf{\Lambda}\right)(\mathbf{k}_t-\mathbb{k})\right]+\boldsymbol{\kappa}_t\left(\mathbf{k}_{-n}^{(t)}\right) \tag{64}$$

where $\boldsymbol{\kappa}_t\left(\mathbf{k}_{-n}^{(t)}\right)$ *is a term that managers take as exogenous, and that in equilibrium equates the value function* $\mathbf{v}(\mathbf{k}_t)$ *to the sum of the enterprise values of the internalized business units,* $\mathbf{\Gamma}\mathbf{v}^e(\mathbf{k}_t)$*:*

$$\boldsymbol{\kappa}_t\left(\mathbf{k}_{-n}^{(t)}\right) \;=\; \mathbf{\Gamma}\mathbf{v}^e(\mathbf{k}_t)-\mathbf{v}(\mathbb{k})-(\mathbf{k}_t-\mathbb{k})\circ\left[(1+r)\mathbf{p}^k-\left(\mathbf{\Lambda}-\frac{1}{2}\mathbf{I}\circ\mathbf{\Lambda}\right)(\mathbf{k}_t-\mathbb{k})\right] \tag{65}$$

Proof. See Appendix A. □

The integration produces a term, which we denote by $\boldsymbol{\kappa}_t$, which the manager takes as given, and which depends on the forecasted path of capital stocks of competitors $\mathbf{k}_{-n}^{(t)}$. While this term does not affect the manager's decision, its equilibrium value is such that the perceived firm value (the value function v_n) is equal to the accounting value of the internalized business units, $\sum_{m\in\mathcal{M}_n} v_m^e$, in equilibrium (for a single-unit firm, simply v_n^e).

2.9 Product Market Spillovers and the Equilibrium Q Network

There is a substantial literature on product market spillovers and peer effects in corporate investment. Within this literature, we find a variety of different measures of product market spillovers. One such measure, used in the seminal paper of Bloom, Schankerman, and Van Reenen (2013, henceforth BSV), is the marginal effect of an additional unit of firm m capital on the investment of firm n.

In our framework (see eq. 21), this is simply the (n,m) entry of the matrix $\mathbf{\Phi}$:

$$\frac{\partial i_{nt}}{\partial k_{mt}} = \varphi_{nm} \quad \text{for } m \neq n \tag{66}$$

In plain language, the off-diagonal terms of the transition matrix $\mathbf{\Phi}$ capture product market investment spillovers in our model.

An alternative measure of product market spillovers is the derivative of v_n^e with respect to k_{mt} — that is, the marginal effect of a unit of capital added by business unit m on the enterprise value of its competitor n. We can collect all these derivatives (one for each n,m pair) into a matrix, which is simply the Jacobian of $\mathbf{v}^e(\mathbf{k}_t)$. We label it $\mathbf{Q}^e(\mathbf{k}_t)$:

$$\mathbf{Q}^e(\mathbf{k}_t) \overset{\text{def}}{=} \frac{\partial \mathbf{v}^e(\mathbf{k}_t)}{\partial \mathbf{k}_t} \equiv \begin{bmatrix} q_{11}^e(\mathbf{k}_t) & q_{12}^e(\mathbf{k}_t) & \dots & q_{1N}^e(\mathbf{k}_t) \\ q_{21}^e(\mathbf{k}_t) & q_{22}^e(\mathbf{k}_t) & \dots & q_{2N}^e(\mathbf{k}_t) \\ \vdots & \vdots & \ddots & \vdots \\ q_{N1}^e(\mathbf{k}_t) & q_{N2}^e(\mathbf{k}_t) & \dots & q_{NN}^e(\mathbf{k}_t) \end{bmatrix} \tag{67}$$

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 other firm in the product market network. If we think of $\mathbf{Q}^e(\mathbf{k}_t)$ as an adjacency matrix, it describes a network of firm value spillovers that operate through product market linkages. There exists an analytical formula for $\mathbf{Q}^e(\mathbf{k}_t)$, but it is relatively involved. However, as we anticipated in Proposition 6, at the steady state ($\mathbf{k}_t = \mathbb{k}$) this expression simplifies dramatically, and becomes the previously-introduced $\mathbb{Q}$, which in the single-product Cournot case takes the simple form:

$$\mathbb{Q} \overset{\text{def}}{=} \mathbf{Q}^e(\mathbb{k}) = (1+r)\mathbf{P}^k - \mathbb{K}\mathbf{Z}\mathbf{\Sigma}\mathbf{Z}\left(\mathbf{I} - \frac{1}{1+r}\mathbf{\Phi}\right)^{-1}. \tag{68}$$

Notice that, even at the steady state, the diagonal entries of $\mathbb{Q}$ differ from the vector of decision marginal Qs, which is simply equal to $(1+r)\mathbf{p}^k$. This "wedge" reflects the distinction between $\mathbf{v}^e(\cdot)$ and $\mathbf{v}(\cdot)$: it arises in the PME as a consequence of the fact that managers take as given the investment plans of their rivals, and do not internalize the feedback effects generated by their own investment. The latter are reflected in the diagonal entries of $\mathbf{Q}^e(\mathbf{k}_t)$ but not in $\mathbf{q}(\mathbf{k}_t)$.

2.10 Neoclassical Investment Equations and the Average-Marginal Q Gap

A well-known and well-tested prediction of the neoclassical investment model is that a firm's marginal Q reflects its growth options, and is therefore correlated with its investment rates (Hayashi, 1982). q_n is often called a "sufficient statistic" for firm n's investment. Echoing this literature, we can re-express our optimal investment rule in equation (21) as a vector of stacked neoclassical investment equations.

Proposition 9. *The vector of business units' investment policies in the PME obeys the following neoclassical investment equation:*

$$\boldsymbol{i}_t(\mathbf{k}_t) = \mathbf{\Delta}\mathbf{k}_t + \frac{1}{\theta}\mathbf{\Phi}\left[\frac{1}{1+r}\mathbf{q}(\mathbf{k}_t) - \mathbf{p}^k\right] \tag{69}$$

Proof. See Appendix A. □

One interesting implication of our version of the neoclassical investment equation (69) is that, in our setting, a business unit's own contemporaneous marginal Q is insufficient to predict its investment rate. While there is still an elegant linear mapping between the vector of investment $\boldsymbol{i}_t(\mathbf{k}_t)$ and the vector of marginal $\mathbf{q}(\mathbf{k}_t)$, the presence of product market spillovers implies that we must know the full vector $\mathbf{q}(\mathbf{k}_t)$ in order to predict any entry of $\boldsymbol{i}_t(\mathbf{k}_t)$.

A key contribution of the canonical study of Hayashi (1982) was to distinguish between a firm's empirically unobservable marginal Q and the associated asset market to book ratio (or average Q). We denote the latter with $\bar{q}_n$:

$$\bar{q}_n(\mathbf{k}_t) \stackrel{\text{def}}{=} \frac{v_n^e(\mathbf{k}_t)}{k_{nt}} \tag{70}$$

The wedge between average and marginal Q has been thoroughly studied in the neoclassical literature, with the primary goal of taking neoclassical investment equations to the data (see, e.g., Hayashi, 1982, Erickson and Whited, 2000, Hennessy, 2004). The wedge reflects inframarginal rents that arise from either diseconomies of scale or (as in the case of our model) monopoly power. It is easy to verify that, in the steady state of the single-product Cournot case, the wedge between average and decision marginal Q simplifies to the present value of monopoly rents per unit of capital:

$$\bar{q}_n(\mathbb{k}) - q_n(\mathbb{k}) = \frac{\rho_n}{k_{n\infty}} \tag{71}$$

2.11 Equilibrium Markups

An important contribution of our model is to investigate the relationship between markups and investment, two objects that in our model are endogenous and co-determined as part of the equilibrium.

Having solved for the PME investment policy, we can now derive an explicit expression for the endogenous markup charged by each business at time t. Here, we naturally focus on the single-product Cournot case ($\mathbf{D} = \mathbf{\Gamma} = \mathbf{I}$). We define $\mu_n(\mathbf{k}_t)$, the markup charged by business unit n at state $\mathbf{k}_t$, as the ratio of the output price to the unit labor cost:

$$\mu_n(\mathbf{k}_t) \stackrel{\text{def}}{=} \frac{p_n(\mathbf{k}_t)}{c_n} \tag{72}$$

Using the solution for the steady-state vector $\mathbb{k}$ derived in Proposition 2, we derive an explicit expression for the equilibrium vector of markups.

Proposition 10. *The equilibrium vector of markups associated with the PME in Proposition 2, for any state* $\mathbf{k}_t$*, is equal to*

$$\boldsymbol{\mu}(\mathbf{k}_t) = \boldsymbol{\mu}(\mathbb{k}) + \mathbf{C}^{-1}(\mathbf{I}+\mathbf{\Sigma})\mathbf{Z}(\mathbb{k}-\mathbf{k}_t) \tag{73}$$

where the steady-state markup vector is equal to:

$$\boldsymbol{\mu}(\mathbb{k}) = \mathbf{C}^{-1}\left\{\mathbf{b} + (\mathbf{I}+\mathbf{\Sigma})\,\mathbf{Z}\mathbf{\Omega}^{-1}\left[(r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^k - \mathbf{Z}(\mathbf{b}-\mathbf{c})\right]\right\} \tag{74}$$

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

$$\mathbf{C} \overset{\text{def}}{=} \begin{bmatrix} c_1 & 0 & \dots & 0 \\ 0 & c_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & c_N \end{bmatrix} \tag{75}$$

Proof. See Appendix A. □

The proposition above provides two important insights. First, equation (73) shows that some of the cross-sectional variation in markups is transitory: holding rivals' capital stocks fixed, a business unit that is operating below steady-state capacity will optimally reduce its markup as its own stock of capital increases. More generally the transitory component is a positively weighted combination of the entire vector of capital gaps, so rivals' adjustments along the transition can offset or reverse this own-capital effect.

Second, equation (74) shows that the markup charged over variable costs will embed, over the long run, a compensation for capital costs. The matrix $\mathbf{C}^{-1}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\boldsymbol{\Omega}^{-1}$ in equation (74) captures the extent to which business units pass the user cost of capital through to their customers via the markup.

2.12 Investment Sensitivity and Marginal Propensity to Invest

A growing literature in macroeconomics and finance studies how businesses heterogeneously adjust investment in response to changes in discount rates or shocks to fundamentals. We solve analytically for two statistics that summarize these responses, and use them to characterize how product markets shape heterogeneity in investment responses to changing financing conditions or growth opportunities.

The first statistic is the Investment Sensitivity (IS), which measures the response of investment to persistent changes in the discount rates. It plays a central role in the transmission of monetary policy (Ottonello and Winberry, 2020; Best, Born, and Menkhoff, 2026).

Proposition 11. *Let the Investment Sensitivity (IS) of business unit n be defined as the change in the investment of business unit n following a one percentage point reduction in the discount rate. In the PME, for any state* $\mathbf{k}_t$*, the vector of IS then equals:*

$$-1\%\frac{\mathrm{d}\,\boldsymbol{i}(\mathbf{k}_t)}{\mathrm{d}\,r} \;=\; 1\%\left[(\mathbf{I}-\boldsymbol{\Phi})\,\boldsymbol{\Omega}^{-1}\mathbf{p}^k + \frac{\mathrm{d}\,\boldsymbol{\Phi}}{\mathrm{d}\,r}\,(\mathbb{k}-\mathbf{k}_t)\right] \tag{76}$$

where the derivative of matrix $\boldsymbol{\Phi}$ *with respect to r is provided in the Appendix. At the steady state it is equal to*:

$$-1\%\frac{\mathrm{d}\,\boldsymbol{i}(\mathbb{k})}{\mathrm{d}\,r} \;=\; 1\%\,(\mathbf{I}-\boldsymbol{\Phi})\,\boldsymbol{\Omega}^{-1}\mathbf{p}^k \tag{77}$$

Proof. See Appendix A. □

Equation (77) shows that competitive pressure in product markets, captured by matrix $\boldsymbol{\Omega}$, modulates firms' individual response to changes in the discount rates. The matrix $\boldsymbol{\Omega}^{-1}$ maps the change in r into changes in firms' desired steady-state capital stocks, through product market spillovers. From

the point of view of a single firm n, a lower discount rate is a favorable cost shock. However, all firms receive this favorable shock, therefore the direction and degree of response of firm n's target capital stock depend on how the reduction in r affects their *relative* competitiveness. The term $(\mathbf{I}-\mathbf{\Phi})$ captures how much of this target adjustment is implemented on impact.

The second relevant statistic is firms' Marginal Propensity to Invest (MPI), which measures how a change in operating cash flows translates into additional capital expenditures. The MPI has received growing attention in the recent macroeconomics literature (e.g., Martin-Baillon, 2021; Winberry, Auclert, Rognlie, and Straub, 2025).

In our setting, the usual definition of the MPI — the change in investment following a one-dollar increase in operating cash flows — requires refinement. A one-dollar increase in operating cash flows is not a uniquely defined experiment: different perturbations of the model primitives generate the same change in cash flows with different implications for investment; absent financing frictions, transitory cash flows have no effect on investment. In addition, in our oligopolistic setting, shocks to a firm spill over to the cash flows of its competitors through product markets.

We therefore consider a specific perturbation of the primitives $(\mathbf{b},\mathbf{c})$ which permanently increases, through $k_{n\infty}$, the steady-state value of business unit n — the present value of its current and future operating cash flows — by one dollar, while leaving the steady-state values of all other units unchanged. The resulting statistic — which we label the marginal propensity to invest out of growth opportunities, or MPI^{G} — summarizes how readily a business unit converts improvements in its growth opportunities into capital accumulation.

Proposition 12. *In the single-product Cournot case, let the Marginal Propensity to Invest out of Growth Opportunities* (MPI^{G}) *of business unit n be the change in its investment expenditure following a specific perturbation in the primitives* $(\mathbf{b},\mathbf{c})$ *which increases the steady-state value of business unit n by one dollar, yet leaves its competitors' steady-state values unaffected. In the PME, for any state* $\mathbf{k}_t$*, the* MPI^{G} *of business unit n then equals:*

$$\text{MPI}^{\text{G}}_n \overset{\text{def}}{=} \frac{\partial p^k_n i_{nt}}{\partial k_{n\infty}} \Bigg/ \frac{\partial v^e_n(\mathbb{k})}{\partial k_{n\infty}} = \frac{1-\varphi_{nn}}{(1+r)+2\left[\bar{q}_n(\mathbb{k})-q_n(\mathbb{k})\right]/p^k_n} \tag{78}$$

it follows that $\text{MPI}^{\text{G}}_n \in \left(0, \frac{1}{1+r}\right)$ *for any business unit n.*

Proof. See Appendix A. □

Equation (78) admits a natural economic interpretation. The numerator, $(1-\varphi_{nn})$, is the speed at which business unit n closes the gap between its current and desired capital stocks: units with highly persistent capital (φ_{nn} close to 1) display a low MPI^{G}.

The denominator is equal to the steady-state marginal Q per dollar of installed capital, $(1+r)$, plus twice the wedge between average and marginal Q — the present value of monopoly rents per dollar of capital. The value of high-rent business units rises steeply with capital: a one dollar value gain requires little expansion and elicits little investment. Competitive units must instead expand substantially to generate the same dollar, and invest more on impact; for price takers the wedge vanishes and the MPI^{G} attains its maximum, $\left(1-\varphi_{nn}\right)/(1+r)$.

The MPI^{G} is thus increasing in the speed of capital adjustment and decreasing in monopoly rents — the two central forces of our model: the implication is that firms that have more market power and that face larger adjustment frictions will be slower to react to growth opportunities.

3 Data and Calibration

In this section, we discuss how we take our model to the data: how we assemble the data, calibrate the model, recover the firm-level primitives and test the model fit. The dataset that we use spans the near-universe of US publicly traded corporations from 1995 to 2021.

3.1 Data

Compustat/CRSP. 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 firm revenues in our model ($p_n y_n$). We also use Costs of Goods Sold, which are mapped to labor costs ($c_n y_n$). Dollar totals are deflated using GDP deflators. Finally, we also use Capex, PPE, R&D Expenditures and Selling, General and Administrative Costs (SG&A), which we use to construct a panel dataset of total capital stocks, as explained in the next subsection.

Under the assumption that COGS equal unit labor costs of production times output, we later compute markups empirically as revenues divided by COGS, in line with previous work by De Loecker, Eeckhout, and Unger (2020). We also compute, in order to validate the model, the market value of assets as total assets, minus the book value of equity, plus the market value of equity. The latter is obtained by sourcing stock price data from CRSP, and multiplying by the number of shares outstanding. We then compute Average Q as the ratio of this market value of assets to the constructed total capital stock k_{nt}.

Hoberg and Phillips (2016) Product Similarity. To measure the matrix of product cosine similarities $\mathbf{A}'\mathbf{A}$, we follow Pellegrino (2025) and rely on the text-based product cosine similarity data of Hoberg and Phillips (2016, henceforth HP). 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

TABLE 1: SUMMARY STATISTICS (1995-2021)

Description	**Notation**	**Mean**	**Std. Dev.**	**p.10**	**p.90**
Revenues (2021 US$mln)	$p_{nt}y_{nt}$	4,451	19,261	14	7,948
Costs of Goods Sold (2021 US$mln)	$c_{nt}y_{nt}$	2,959	14,449	5	4,988
Market Value of Assets (2021 US$mln)	v_{nt}	18,237	121,820	32	21,674
Total Capital Stock (2021 US$mln)	k_{nt}	4,742	21,248	22	7,735
Price of Capital	p_n^k	1.190	0.088	1.070	1.278
Depreciation Rate	δ_n	0.154	0.048	0.076	0.198
Investment Rate	i_{nt}/k_{nt}	0.273	0.303	0.085	0.473
Average Q (log)	$\log \bar{q}_{nt}$	0.737	1.468	-0.837	2.691
Capital Productivity (log)	$\log z_n$	-5.124	1.525	-6.681	-3.171

Note: following Peters and Taylor (2017), the investment rate is computed excluding firm-year observations with less than $5 million of capital.

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), compares 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, breakfast cereals, and personal computers, he finds that model-implied elasticities track the patterns in these 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 have had their variables reconciled from accounting equations and linear interpolation. This results in a sample size of 5,707 in 1995 and 3,157 in 2021.

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. These deflators are themselves finally deflated using the GDP deflator, so that all dollar amounts are expressed in consistent units.

Summary Statistics and Additional Details. Table 1 reports the summary statistics of our working sample before merging with the HP data. The distributions of firms' 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.

We provide additional details on data sources and construction in Appendix C.

3.2 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 adjustment costs (θ), productivity ($\mathbf{z}$), and the steady-state capital distribution ($\mathbb{k}$). The latter require careful treatment due to the high degree of dimensionality of our parameter space ($\mathbf{z}$ has thousands of entries) and the even larger number of equality and inequality restrictions imposed by the model, which imply our model is 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 microeconometric 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 3.1. The discount rate r is year-specific and set equal to the estimate by Gormsen and Huber (2025); its 1995 value uses the earliest available hurdle-rate spread over Treasuries.

Calibration of the adjustment cost parameter θ. The curvature parameter θ governs the convexity of adjustment costs. We anchor its level 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 first compute a year-specific value of θ_t by equating aggregate adjustment costs in our model to 3.1% of aggregate operating profits. Because these annual estimates are sensitive to transitory fluctuations in investment, we smooth them by fitting a linear time trend to $\log\theta_t$ over

1995–2021. Our quantitative implementation uses the resulting exponential trend, with θ_t remaining common across firms within each annual cross section.

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, 5,707 in 1995 and 3,157 in 2021). For each firm in the working sample, we notice that the productivity parameter z_n appears squared in the first-order conditions (34), (35) and (36). Rearranging these conditions, we can express:

$$z_n^2 = \frac{\pi_{nt}}{k_{nt}^2} - (r_t + \delta_{nt})\frac{p_{nt}^k}{k_{nt}} - \theta\left[(1+r_t)\frac{k_{nt}-k_{nt-1}}{k_{nt}} - \frac{k_{nt+1}-k_{nt}}{k_{nt}}\right] \tag{79}$$

In principle, equation (79) 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 (79). As a result, its right-hand side can take negative values in the data. These negative values are inadmissible since z_n^2 is 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 recover from equation (79). 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 (79).[17] 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 C.

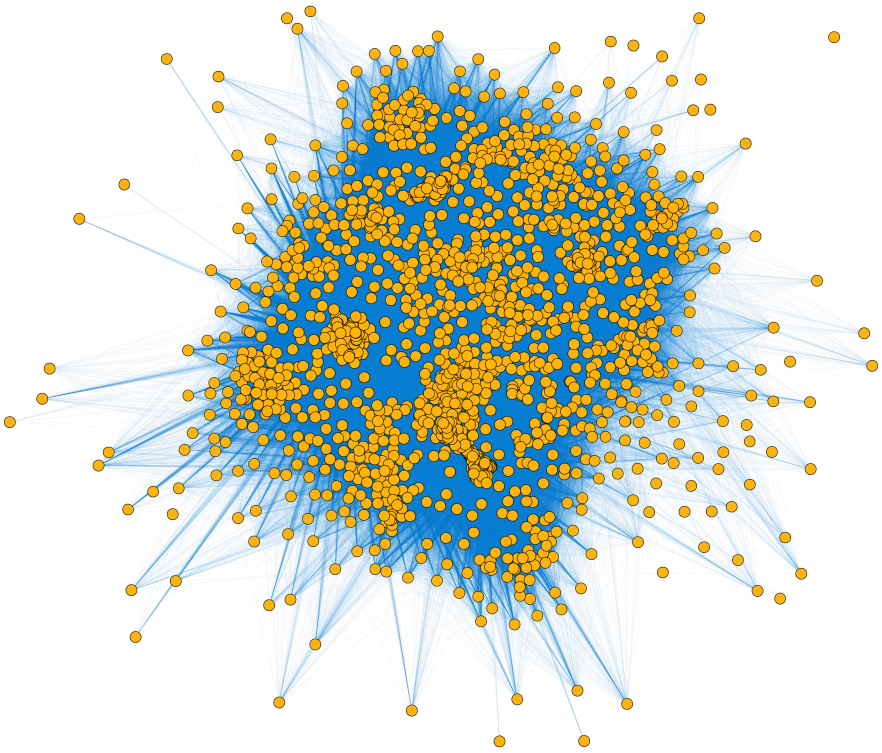
Enforcement of positive steady-state capital. While $\mathbb{k}$ is an equilibrium object rather than a primitive, we must ensure it satisfies $\mathbb{k} > 0$: this condition is both intuitive (we expect firms that we observe in the data to remain active in the long term) as well as a requirement for 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 capital stocks are always measured imperfectly, leads the direct estimate of $\mathbb{k}$ from the transition equation (20) to be negative for a small subset of firms (<16.15%). To enforce positivity while preserving the cross-sectional distribution, we fit a Gompertz distribution[18] 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 C.

Recovery of demand and cost parameters. Given $\mathbf{z}$, θ, and $\mathbb{k}$, we recover the remaining parameters directly from the data and equilibrium conditions. Unit labor costs are set to $c_n = \text{COGS}_{nt}/y_{nt}$. The demand intercept vector $\mathbf{b}$ is recovered by inverting the steady-state equation (43).

[17] In estimating the parameters of the Pareto distribution with the regression, we implement the correction suggested by Gabaix and Ibragimov (2011).

[18] We use a Gompertz distribution because it matches the empirical properties of the distribution of the ratio $k_{n\infty}/k_{nt+1}$, 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 INVESTMENT SPILLOVER NETWORK

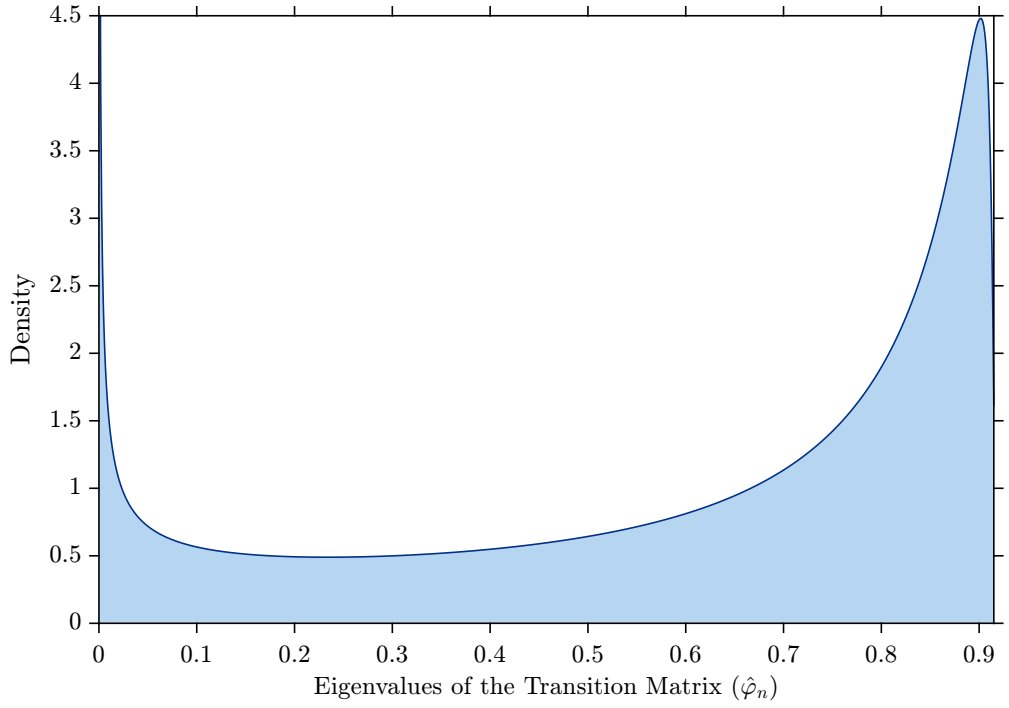


3.3 Estimated Transition Matrix

Figure 3 displays the matrix of investment spillovers (equation 66) or, equivalently, the off-diagonal entries of the transition matrix $\mathbf{\Phi}$, as a weighted network. Each node represents a firm in Compustat as of 2021; each edge represents the magnitude of investment spillovers (see equation 66 between a firm pair. Specifically, we plot edges corresponding to the top 1.2 million 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. The latter is equivalent to low reliance on capital: it reduces the incidence of capital adjustment frictions and amplifies responsiveness to competitive conditions. Peripheral firms, scattered 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 (high capital dependence) generates relatively large adjustment costs that dampen their investment response to competitive pressures.

The network displays a moderate degree of clustering: this suggests that investment spillovers organize firms into loosely defined groups of firms that affect each other's investment choices. This cluster structure proves less pronounced than the one observed in the underlying Hoberg-Phillips

FIGURE 4: TRANSITION MATRIX: DISTRIBUTION OF EIGENVALUES

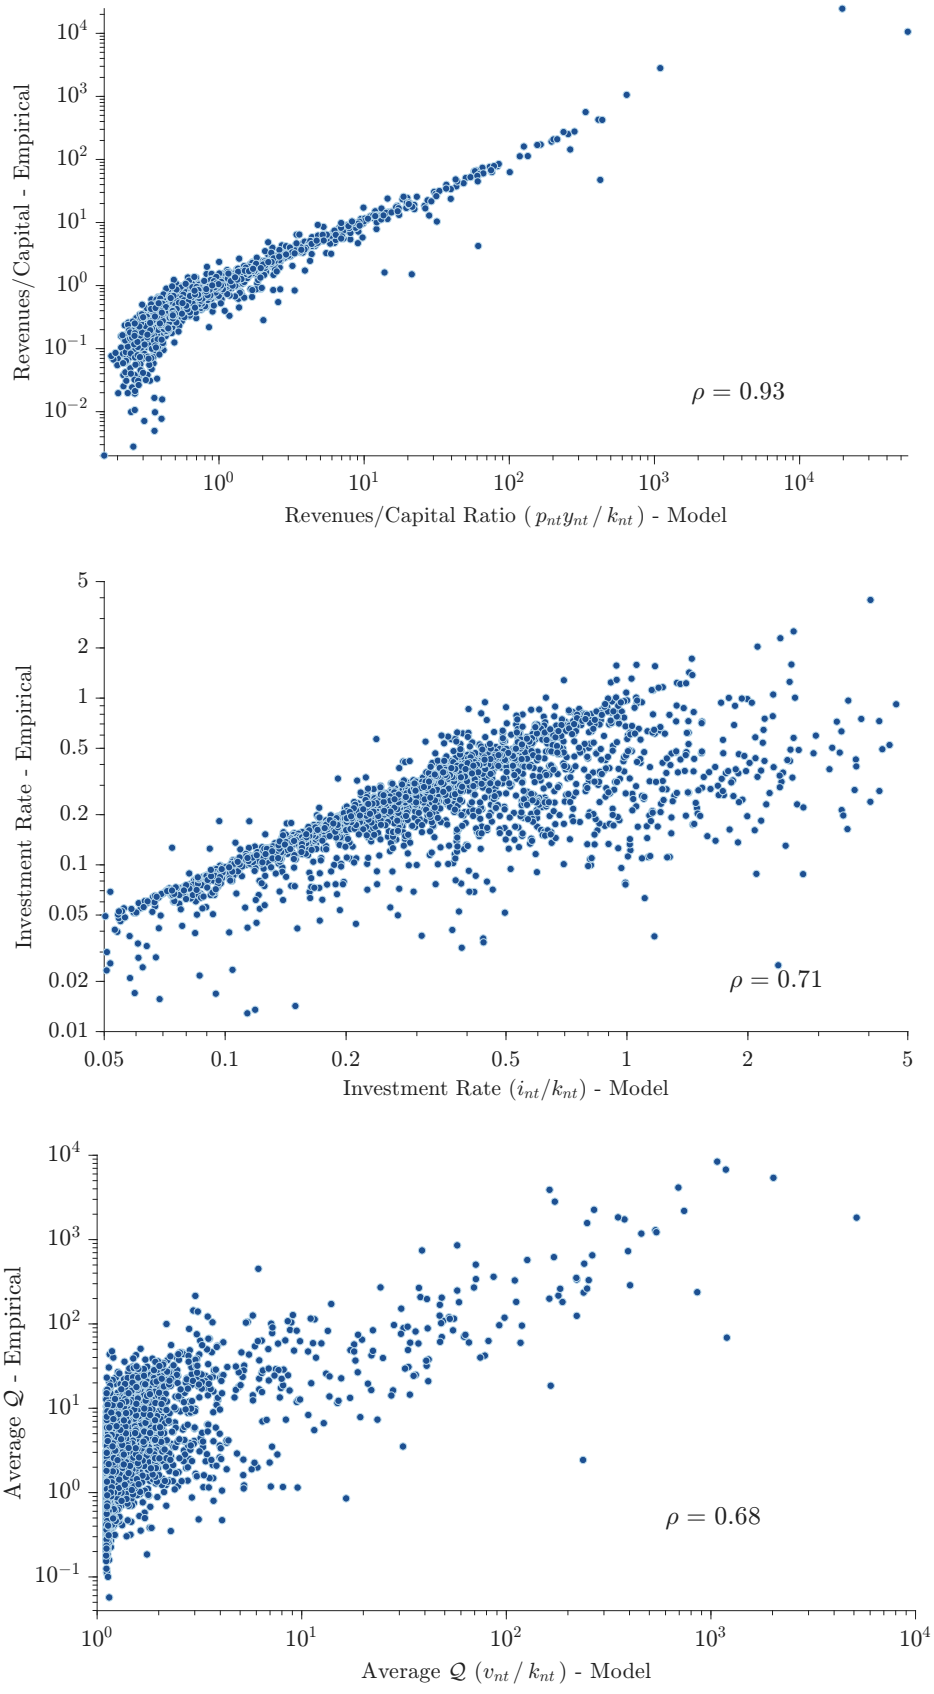


product similarity network, as documented in Pellegrino (2025). This attenuation reflects the fact that $\mathbf{\Phi}$ incorporates not only product market similarity—captured by the substitution matrix $\mathbf{\Sigma}$—but also firm-level heterogeneity in capital productivities z_n. Firms with low z_n exhibit weak investment responses even to closely positioned competitors, whereas firms with high z_n respond vigorously to product market conditions. In addition, $\mathbf{\Phi}$ also reflects indirect linkages between firms. This additional heterogeneity obscures the clean industry boundaries visible in pure product similarity networks.

Figure 4 presents the empirical distribution of eigenvalues $\hat{\varphi}_n$ of the estimated transition matrix $\mathbf{\Phi}$ for 2021. These eigenvalues govern the persistence of the capital stocks: 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 distribution: many eigenvalues cluster near zero (indicating rapid convergence along those modes) while others concentrate near the maximum, around 0.9 (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 combination of heterogeneity and state-dependence distinguishes our network model from simpler frameworks where convergence rates are uniform across all initial conditions, and underlines the importance of product markets in determining the economy's as well as the individual firms' response to shocks.

FIGURE 5: MODEL FIT



3.4 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 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 ($z_n^2 < 0$) or steady-state capital ($k_{n\infty} < 0$). 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. By comparing the model-generated versions of these variables against the corresponding empirical ones and gauging the resulting fit, we can get a sense of how binding these over-identifying restrictions are in practice.

Figure 5 displays three key validation exercises, plotting model-predicted values against their empirical counterparts for all firms of our sample in 2021. 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 =0.93 between model and data. This is reassuring: revenues are one of the most accurately measured variables in accounting data, and it is important for our model to match this variable as closely as possible.

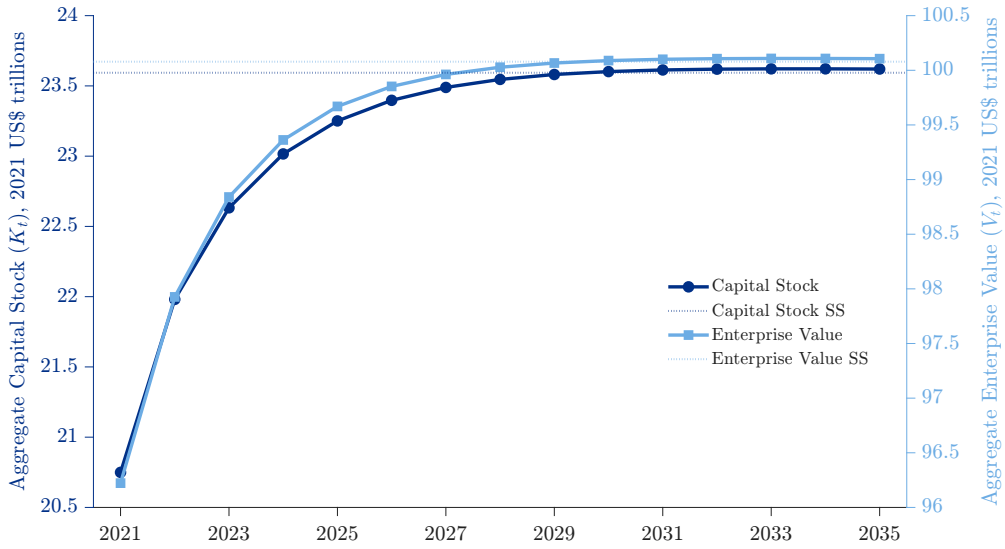
The middle panel displays the investment rate (i_{nt}/k_{nt}), with a correlation of ρ =0.71. The goodness of fit captured in the middle panel corresponds to an R^2 of 0.5, which compares favorably to that of estimated investment equations in the existing neoclassical literature. For instance PT, whose paper incorporates intangible capital and corrects for measurement error, reports R-squares ranging from 0.233 to 0.374.

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 only valuation moment we use is a single aggregate one — the scalar rescaling of the productivity vector that matches total model-implied firm value to total observed market value — which pins down the level of valuations but places no restriction on their cross-sectional allocation. The correlation reported below is in any case invariant to that scalar. Nevertheless, the model achieves a correlation of ρ =0.68 with the empirical average Q. This provides evidence that the model's structural predictions extend beyond the specific moments used in calibration, and therefore that the economic forces that underlie our model are important to explain firm heterogeneity in valuations.

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.

FIGURE 6: TRANSITION DYNAMICS OF THE AGGREGATE CAPITAL STOCK



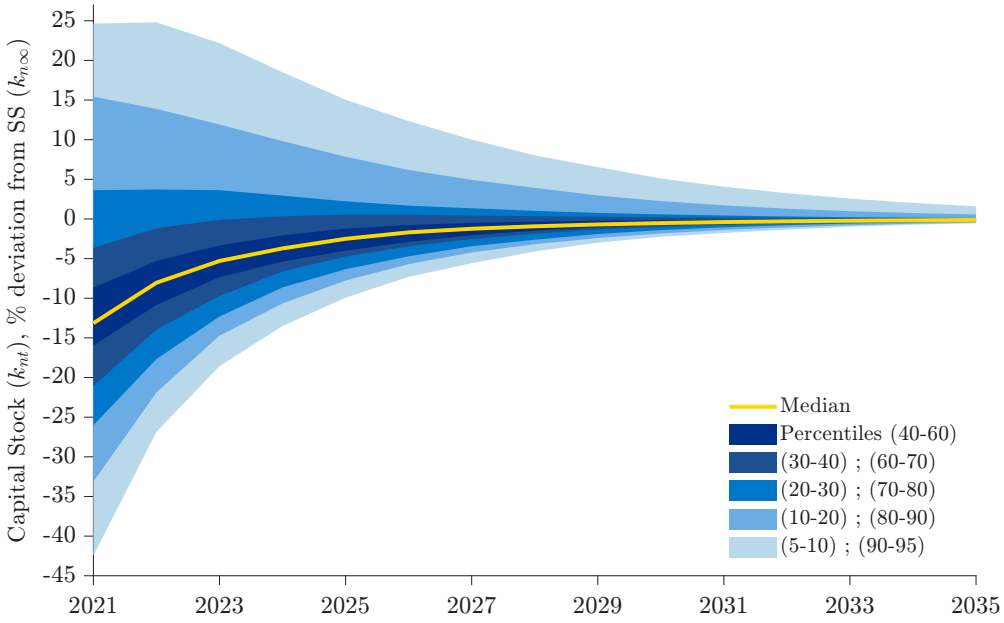
4.1 Analysis of the Factual Equilibrium

We start by taking the model to the data in 2021and examining the properties of the calibrated factual equilibrium (the PME under our baseline Cournot conduct assumption). 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 studying the properties of the PME path predicted by the model, with an emphasis on the transition dynamics. Figure 6 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 $20.75 trillion in 2021, approaching its steady-state level of $23.59 trillion by 2035, representing aggregate capital growth of roughly $2.84 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 figure also reports aggregate enterprise value against the right-hand axis. Once placed on their separate scales, the two series display nearly identical shapes: both rise rapidly at first and flatten as the economy approaches its steady state. This close comovement is a model implication. With the economic primitives held fixed, enterprise value is evaluated along the same equilibrium capital path and therefore inherits the transition dynamics governed by $\mathbf{\Phi}$; close to the steady state, the value function is locally approximately linear in deviations of capital from its long-run level. The visual overlap should therefore be interpreted as similarity in the two variables' normalized convergence profiles, rather than equality in their levels, percentage growth rates, or dollar-for-dollar changes.

The concave shape of the aggregate capital path—with faster growth in earlier years that gradually decelerates—is characteristic of VAR models. 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 steady-state level of capital stock, the investment rate naturally slows, even though adjustment continues.

FIGURE 7: FIRM-LEVEL TRANSITION DYNAMICS



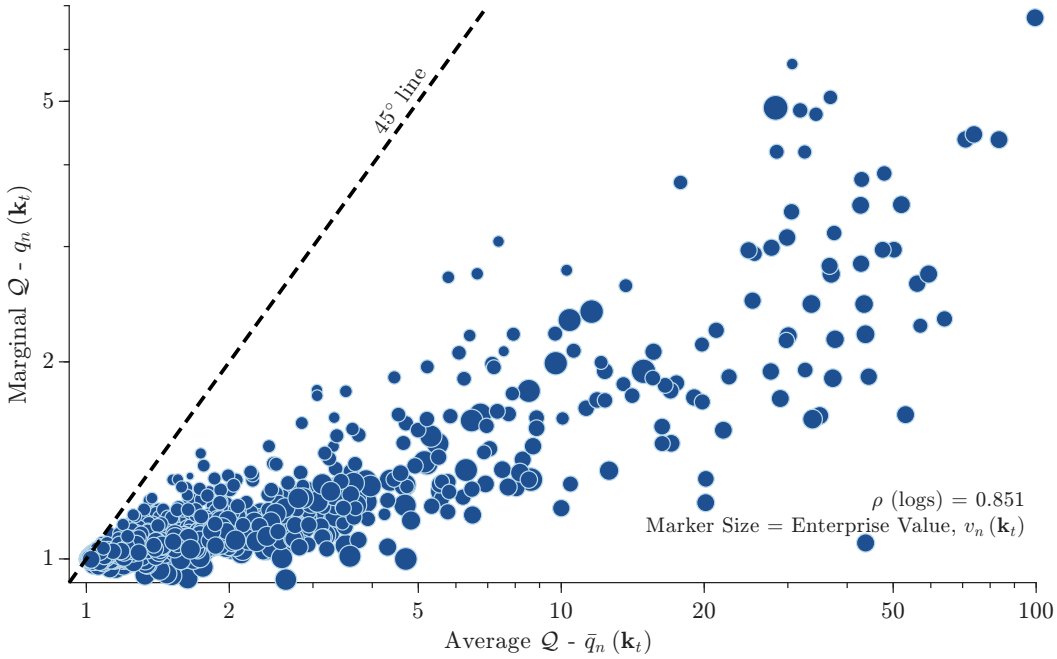
Over the transition, the aggregate capital stock increases on average by approximately 0.99% of its 2021 level per year—a pace that reflects both the fundamental forces driving capital accumulation (capital productivity, demand conditions, and competitive interactions) and the frictions that slow adjustment (convex adjustment costs parameterized by θ).

Figure 7 reveals the firm-level heterogeneity masked by the aggregate trajectory shown in Figure 6. The fan chart displays the cross-sectional distribution of capital deviations from steady state, measured as percentage gaps. Observations are weighted by 2021 capital stocks. In 2021, firms exhibit substantial dispersion, with the 5th percentile approximately 42.45% below steady state and the 95th percentile roughly 24.65% above. The distribution is notably asymmetric: more of the capital stock sits below its steady state than above. The capital-weighted median gap is approximately 13.15% below its long-run equilibrium.

The percentile bands systematically narrow and drift toward zero: over the forecast horizon, the model predicts gradual convergence of each firm towards its steady-state level of capital. 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 (or divest). By 2030, the range between the 40th and 60th percentile bands has tightened considerably, spanning only about ±0.76% around steady state.

The heterogeneous adjustment speeds visible in Figure 7 directly reflect the eigenvalue distribution analyzed in Figure 4. 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.

FIGURE 8: AVERAGE Q AND MARGINAL Q



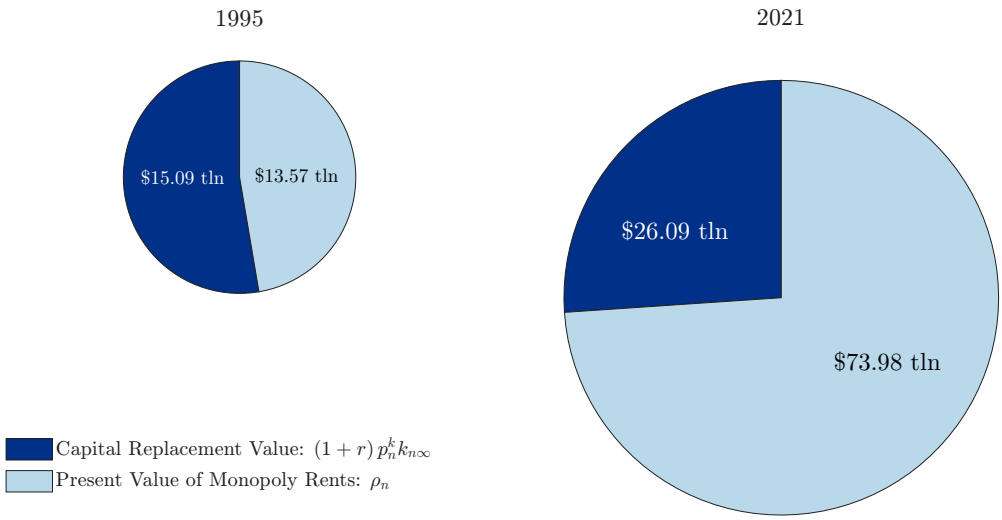
Average and Marginal Q. Figure 8 plots the model-implied firm-level estimates of average Q as of 2021 against the corresponding estimates for marginal Q. In spite of the relatively high correlation of 0.85, the fit between both quantities masks substantial heterogeneity in the cross-section. The wedge between average and marginal Q is such that most observations lie strictly below the 45-degree line, suggesting that monopoly rents contribute significantly to enterprise values for most firms. Further consistent with the workings of the model, firms associated with relatively larger dots in the plot (and, thus, with higher enterprise value) are typically further away from the 45-degree line: a higher wedge between average and marginal Q is naturally associated with higher enterprise value, consistent 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 Eeckhout, 2025; Cho, Grotteria, Kremens, and Kung, 2023). Because our model features endogenous markups and enterprise values, it can help inform this debate. In particular, Equation (56) provides a useful decomposition of firm n's steady-state market value in terms of two components: the replacement value of capital ($(1+r)p_n^k k_{n\infty}$) and the present value of monopoly rents (ρ_n). We implement this decomposition in Figure 9, which depicts the combined steady-state market value of all firms in our sample for 1995 and 2021 as two pie charts.

The first thing that we observe from the figure is that the absolute scale of firm value in steady state grew dramatically, from \$28.66 trillion in 1995 to \$100.08 trillion in 2021. This increase reflects *both* capital accumulation as well as the increase in monopoly rents associated with rising market concentration.

Next, we consider the shares of the two components. In 1995, the replacement value of capital

FIGURE 9: DECOMPOSITION OF STEADY STATE AGGREGATE FIRM VALUE



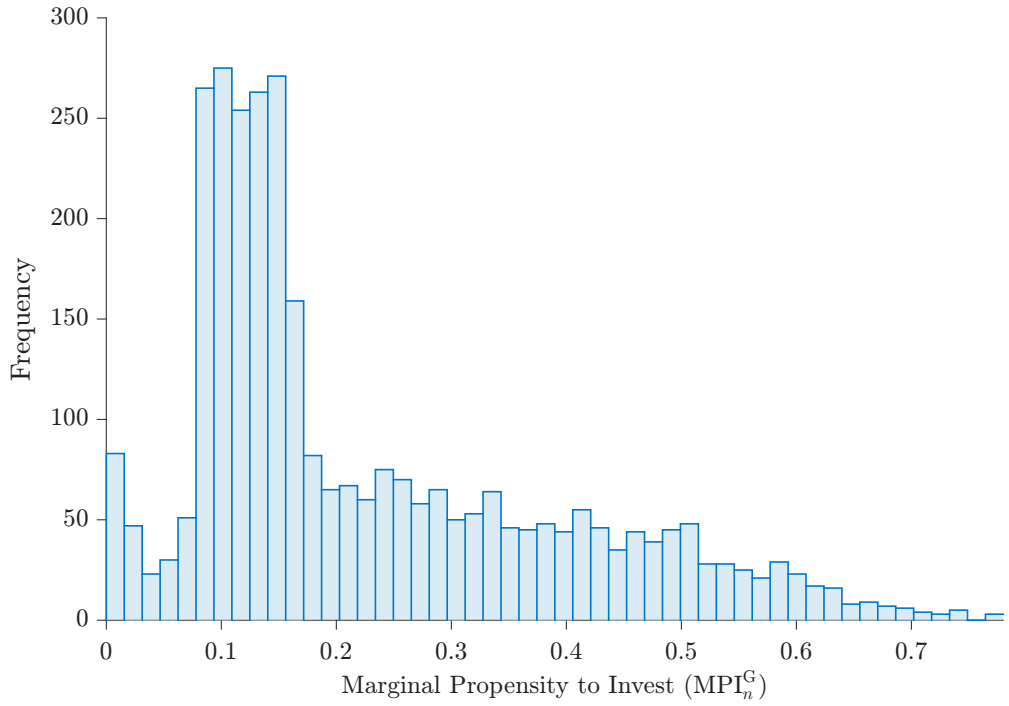
accounted for 52.67% of total enterprise value in steady state, whereas monopoly rents comprised 47.33%. By 2021, this composition changed drastically: the replacement value of capital decreased to just 26.07% of firm value in steady state, whereas monopoly rents have risen to 73.93%.

In sum, two interesting facts emerge from this analysis: 1) monopoly rents account for a significant fraction of firms' market values; 2) the relative share of this rent component has increased significantly over time. This in turn suggests that rising monopoly rents can help explain the rapidly rising market valuations of US firms.

Our findings from Figure 9 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 and that expected markups correlate with intangible investments. The latter are included in our definition of total capital. While our model does not distinguish between rents accruing to intangible capital and those accruing to tangible capital, intangible capital accounts for an increasing share of the aggregate capital stock over our sample (Eisfeldt and Papanikolaou, 2014; Crouzet and Eberly, 2021). Therefore both our findings are consistent with rising markups providing, at least in part, compensation for accumulated past intangible investments, as evinced by Crouzet and Eberly (2023).

Marginal Propensity to Invest. Figure 10 reports the cross-sectional distribution of the marginal propensity to invest out of growth opportunities, computed from equation (78) for year 2021. For the typical business unit, investment responsiveness is modest: most of the mass lies near 0.10, meaning that a permanent shock that raises steady-state firm value by one dollar increases current investment expenditure by roughly ten cents. The distribution is strongly right-skewed: a long tail of business units — those that face intense competitive pressure and, as a consequence, a low incidence of adjustment costs — converts value gains into capital expenditure at several times the typical rate, approaching the price-taker benchmark of $\left(1-\varphi_{nn}\right)/(1+r)$. At the opposite extreme, a small cluster of business units exhibits an MPI^{G} close to zero: for these firms the wedge between average and marginal Q is so large that additional value accrues almost entirely as rents on existing capacity, generating little incentive to expand. This dispersion illustrates a central implication of our framework: product market power weakens the pass-through from value creation to capital

FIGURE 10: MARGINAL PROPENSITY TO INVEST OUT OF GROWTH OPPORTUNITIES



formation, so that the business units that gain the most value are not necessarily those that invest the most.

4.2 Counterfactual Analysis

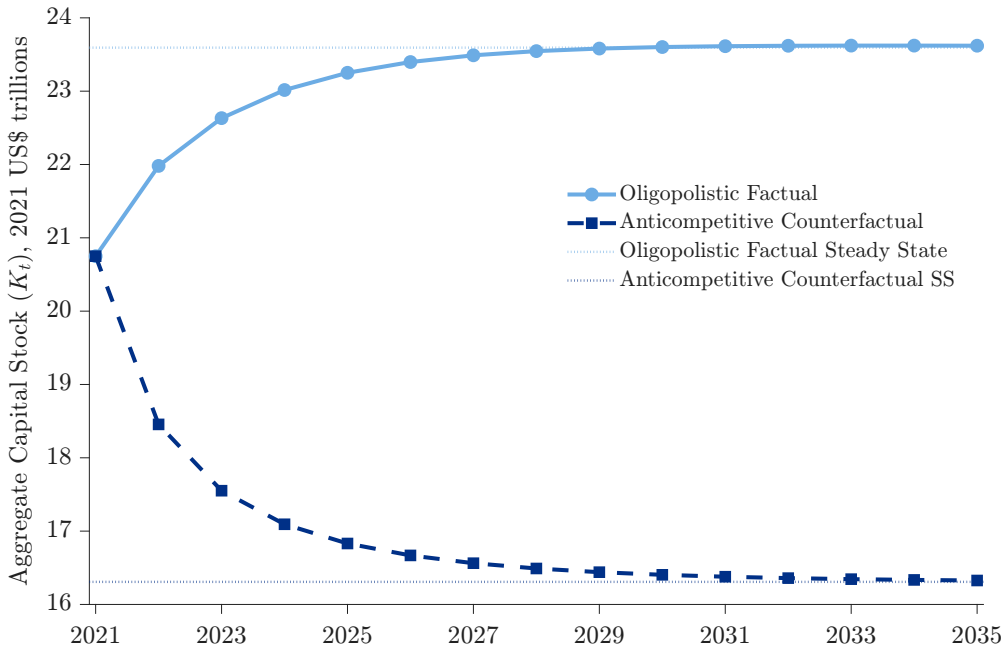
In this subsection, we exploit the counterfactual-building capabilities of our model to provide additional insights on the role of product market competition, discount rate shocks, and M&A activities. Before we do so, we provide a technical clarification on the practical implementation of the counterfactual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

Implementing Counterfactuals with Non-negativity Constraints. One minor challenge that must be dealt with when performing counterfactual exercises (mergers, changes in r or unit labor costs, etc.) is that the interior steady state implied by Proposition 2 does not preclude capital stocks from becoming negative for some business units. To perform counterfactuals, we develop a method to enforce a non-negativity constraint on $\mathbb{k}$. To build intuition for our method, we notice that the steady state presented in equation (43) corresponds to the Nash equilibrium of the corresponding *static* game obtained by shutting down adjustment costs (i.e., setting $\theta = 0$) — specifically, for the single-product Cournot case, $\mathbb{k}$ solves:

$$k_{n\infty} \equiv \underset{k_n}{\operatorname{argmax}}\, z_n k_n (b_n - c_n) - (r+\delta_n) p_n^k k_n - z_n^2 k_n^2 - \sum_{m=1}^{N} z_n z_m \sigma_{nm} k_n k_m \tag{80}$$

As discussed in Pellegrino (2025), this static game is a type of *Potential* game, which can be repre-

FIGURE 11: ANTICOMPETITIVE COUNTERFACTUAL: AGGREGATE CAPITAL STOCK



sented by a scalar function $\mathscr{P}(\mathbf{k})$ or *Potential*, defined as follows:

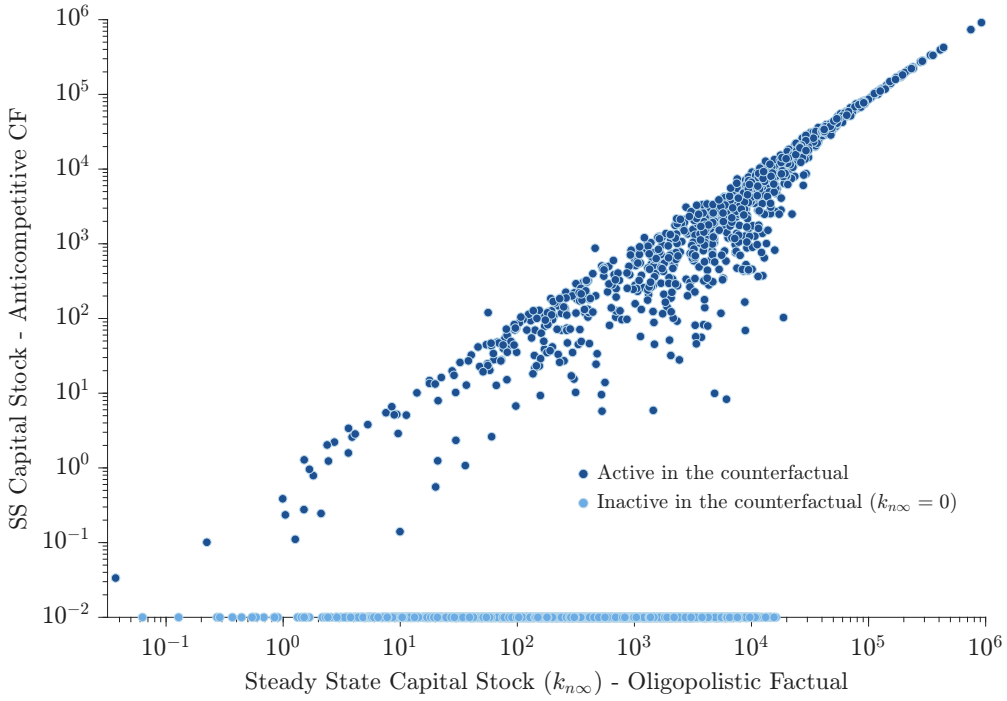
$$\mathscr{P}(\mathbf{k}) \stackrel{\text{def}}{=} \mathbf{k}'\Big[\mathbf{Z}(\mathbf{b}-\mathbf{c})-(r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^k\Big]-\frac{1}{2}\mathbf{k}'\mathbf{\Omega}\mathbf{k} \tag{81}$$

Regardless of whether a non-negativity constraint applies to $\mathbf{k}$, the solution of the static game (80) is also the maximizer of the potential function (81), 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 few seconds on an ordinary laptop even when the number of firms is in the thousands. This makes it straightforward to compute the counterfactual $\mathbb{k}$, 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 a 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 (80). 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 (44). The advantage of this approach is internal consistency: conditional on the active set $\mathscr{A}$, the interior solution in equation (43) 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

FIGURE 12: ANTICOMPETITIVE COUNTERFACTUAL: CROSS-SECTION



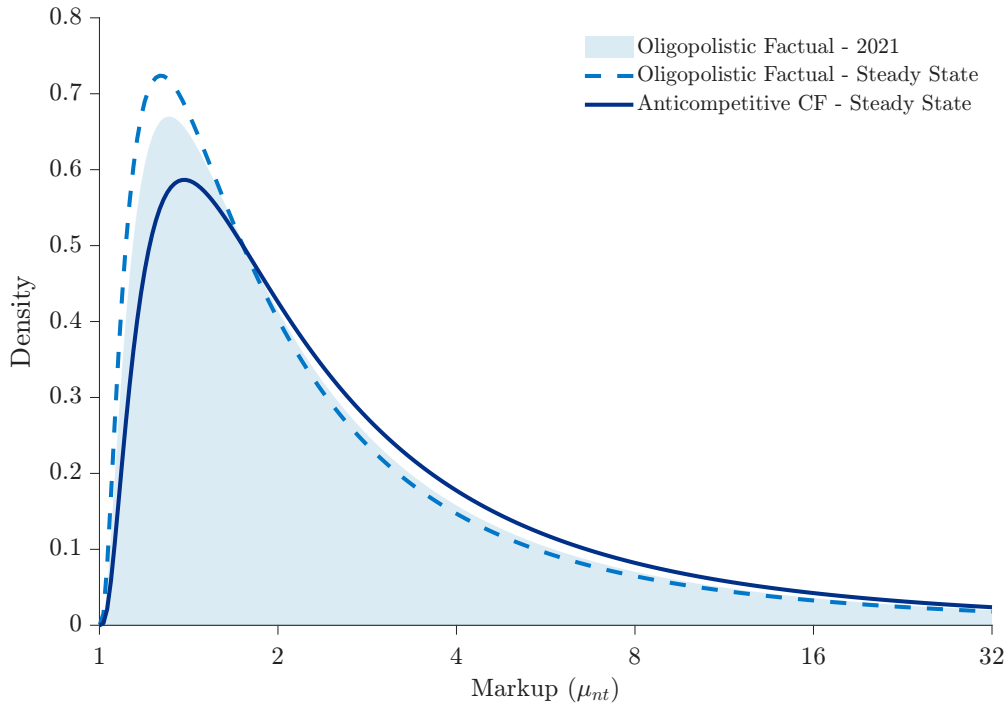
dynamics.

Full Set (interior-$\mathbf{\Phi}$) implementation. An alternative solution is to keep all firms in the state vector, use the constrained $\mathbb{k}$ from the static quadratic programming (allowing some components to be zero), and compute $\mathbf{\Phi}$ from the interior formula (44) 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 PME of the counterfactual because the law of motion consistent with equations (20) and (44) need 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 choice does not affect any of our steady-state results, which follow from the constrained problem (80) alone; it matters only for the shape of the counterfactual transition paths.

Counterfactual 1: Competition and Investment. A central question in finance and macroeconomics is the extent to which product market competition affects capital accumulation. While the literature has long recognized that product markets affect 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.

FIGURE 13: DISTRIBUTION OF MARKUPS

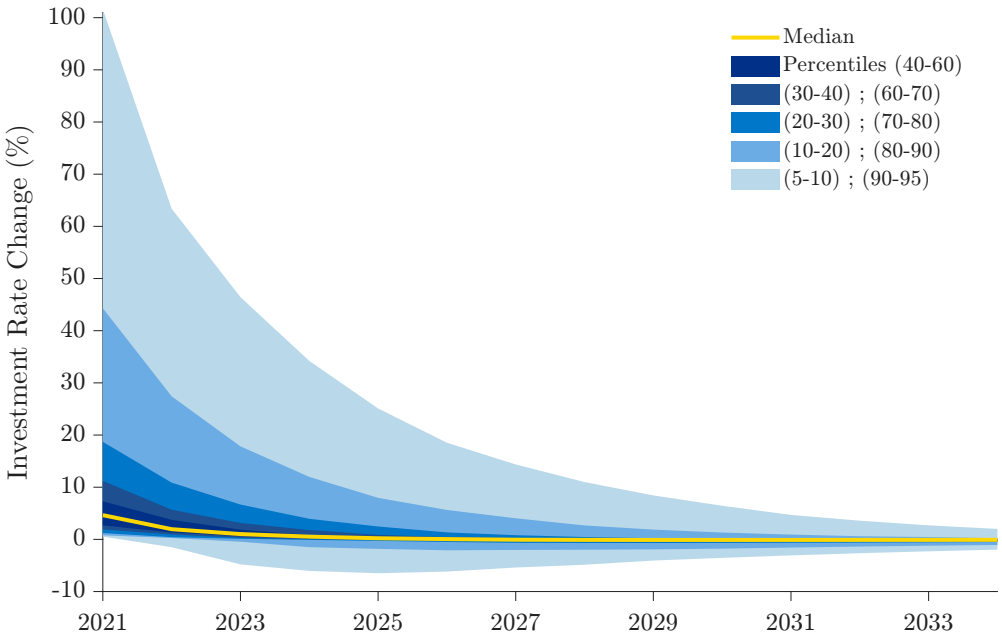


We quantify the impact of product market competition on corporate investment by comparing the factual equilibrium under Cournot conduct to an anticompetitive counterfactual in which all firms fully internalize product market spillovers. Under Cournot conduct, firms internalize their downward-sloping residual demand curves but do not account for strategic spillovers on rival firms. In our anticompetitive counterfactual all business units operate as a single cartel: we set $\mathbf{\Gamma} = \mathbf{1}\mathbf{1}'$ such that the competitive pressure 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 30.88%.

Figure 11 displays the aggregate implications. Commencing from the observed 2021 capital stock of \$20.75 trillion, the factual economy under Cournot conduct converges monotonically to its steady state of \$23.59 trillion over approximately 10 years. The anticompetitive counterfactual, in contrast, exhibits a pronounced contraction, with aggregate capital declining by 30.88% (relative to the factual steady state) to \$16.31 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 convergence completed within 5 years.

The cross-sectional heterogeneity in firm-level responses, illustrated in Figure 12, reveals important additional details. 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

FIGURE 14: RESPONSE OF INVESTMENT RATES TO DISCOUNT RATE SHOCK



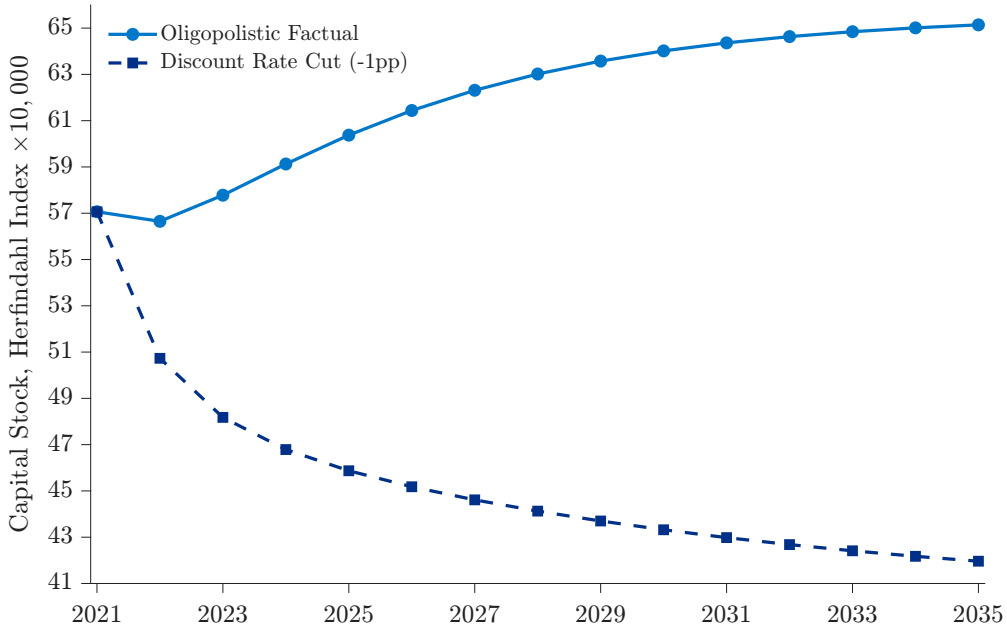
the magnitude of adjustment varies considerably across the size distribution. Notably, a non-trivial fraction of firms—depicted in light blue—exit entirely in the counterfactual steady state, with their capital stocks falling to zero.[19] These exits occur predominantly among smaller firms with high product similarity to larger competitors, for whom the cartel finds it optimal to consolidate capacity into fewer, larger producers. This pattern indicates that anticompetitive coordination not only suppresses aggregate investment but also fundamentally reshapes market structure toward greater concentration.

Figure 13 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 2.02. The factual steady state (dashed line) displays modest compression towards a median markup of 1.94 as the typical firm increases capital and supply and the economy converges to its long-run equilibrium. The anticompetitive counterfactual (solid line), however, generates a pronounced shift of the entire distribution. Markups increase for the majority of firms, with the median rising to approximately 2.18—a 12.60% increase—and the right tail extending comfortably into double digits. This reflects coordinated output restriction. 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.

These results indicate that imperfect competition in product markets has quantitatively substantial implications for aggregate and firm-level 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 (Pellegrino, 2025), imperfect competition distorts

[19] For visibility on the logarithmic scale, exiting firms are displayed at a nominal floor in Figure 12. These firms are retained in the counterfactual markup distribution of Figure 13 and are assigned their model-implied markups, which remain well-defined at zero output because the equilibrium price is well-defined.

FIGURE 15: EVOLUTION OF CONCENTRATION



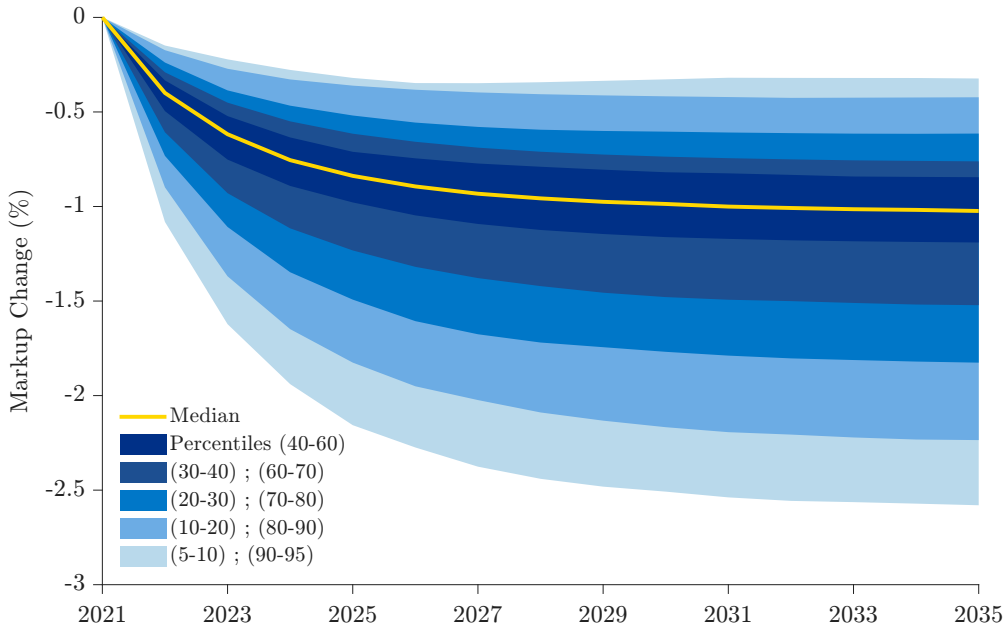
the inter-temporal allocation of resources by attenuating incentives for capital formation.

Counterfactual 2: Discount Rates and Market Structure. Our second counterfactual examines the market structure implications of changing the discount rate. Specifically, we study the case of a one percentage point permanent, unanticipated decrease in the aggregate discount rate ($r \to r - 0.01$), holding all other primitives constant, including the competitive conduct regime. This experiment helps understand how a change in aggregate financial conditions interacts with an oligopolistic market structure to affect the equilibrium path, abstracting from changes in competitive behavior.

We start by studying the impact of the shock on the investment rates. Figure 14 shows the effect of the discount-rate cut on firms' investment rates. The plot displays the % change in firm-level investment rates in the counterfactual scenario against the corresponding investment rate in the Cournot case, again grouping firms into percentiles by the magnitude of their response. It reveals significant semi-elasticities and striking cross-firm heterogeneity in the first few years following this shock: the median firm adjusts its investment rate by about 4.65%; the majority of firms' investment response on impact is between 1.21% and 18.70%. However, there is a thick right tail: as we move up in the higher percentiles of the distribution, we encounter firms that respond much more, with the 90th percentile adjusting the investment rate by 44.31% and the 95th nearly doubling it. Yet, these effects are temporary: every firm's investment rate eventually reverts towards the depreciation rate as the model converges to its new steady state. Interestingly, due to the complex product market interactions that are present in our model, some firms end up reducing their investment along the equilibrium path. Through these results, our model contributes to an emerging literature that investigates firms' heterogeneous responses to discount rate fluctuations (Ottonello and Winberry, 2020; Winberry et al., 2025; Best et al., 2026).

These asymmetric investment responses have implications for concentration, which we measure as the Herfindahl index of capital stock. Figure 15 displays the evolution of concentration, measured

FIGURE 16: RESPONSE OF MARKUPS TO DISCOUNT RATE SHOCK

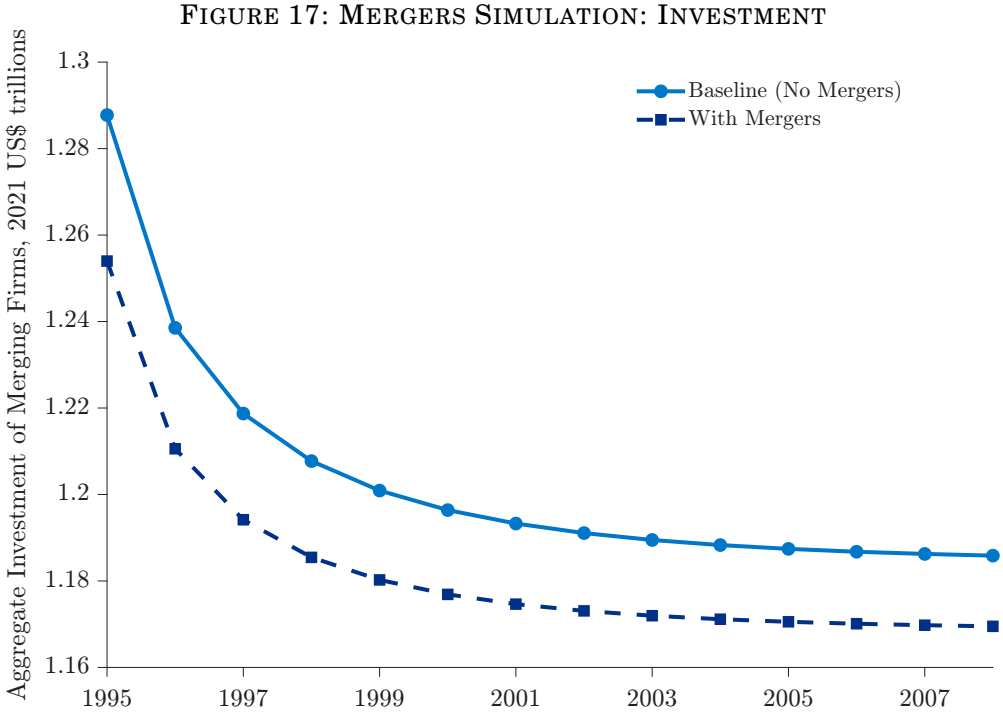


under both scenarios. The factual Cournot scenario exhibits an initial decline in the Herfindahl index from 57.06 in 2021 to approximately 56.65 by 2022, followed by a gradual increase toward 65.14 at the 2035 endpoint. Given the initial configuration of capital stocks, product market spillovers generate interesting transition dynamics, including a non-monotonic evolution of concentration. The discount rate shock counterfactual, by contrast, generates a substantially different trajectory. Concentration falls immediately and monotonically, reaching 41.96 at the 2035 endpoint, significantly lower than the factual outcome. This pronounced decrease indicates that lower 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 the aggregate discount rate shock. The discount rate shock increases the present value of future profits, strengthening investment incentives for all firms. However, this effect is not uniform: smaller firms, which charge thinner markups and operate closer to their exit thresholds, raise investment more sharply than larger incumbents. The resulting reallocation of market share toward smaller firms manifests as decreased concentration. It is again important to underscore that this channel operates holding constant the competitive conduct regime—the decrease in concentration strictly relates to the differential impact of financial conditions across the firm size distribution, rather than changes in strategic behavior.

Figure 16 illustrates the markup consequences of this reallocation effect. The fan chart plots the % change in firm-level markups in the counterfactual with respect to the corresponding firm-level markup in the Cournot factual. The median firm experiences a modest markup decrease, falling gradually to -0.93% by 2027. However, this aggregate pattern masks substantial heterogeneity. The lower percentiles of the distribution exhibit considerably greater markup decreases, with firms in the 5th-10th percentile experiencing decreases exceeding 2.02%. By contrast, firms at the 90th-95th

FIGURE 17: MERGERS SIMULATION: INVESTMENT



percentile barely adjust markups.

This dispersion reflects the oligopolistic nature of competition as well as the properties of GHL demand, which generate vast asymmetry across firms in the cost of capital pass-through. The widening fan over time indicates that these markup effects amplify along the transition path, as the cumulative impact of differential investment decisions compounds.

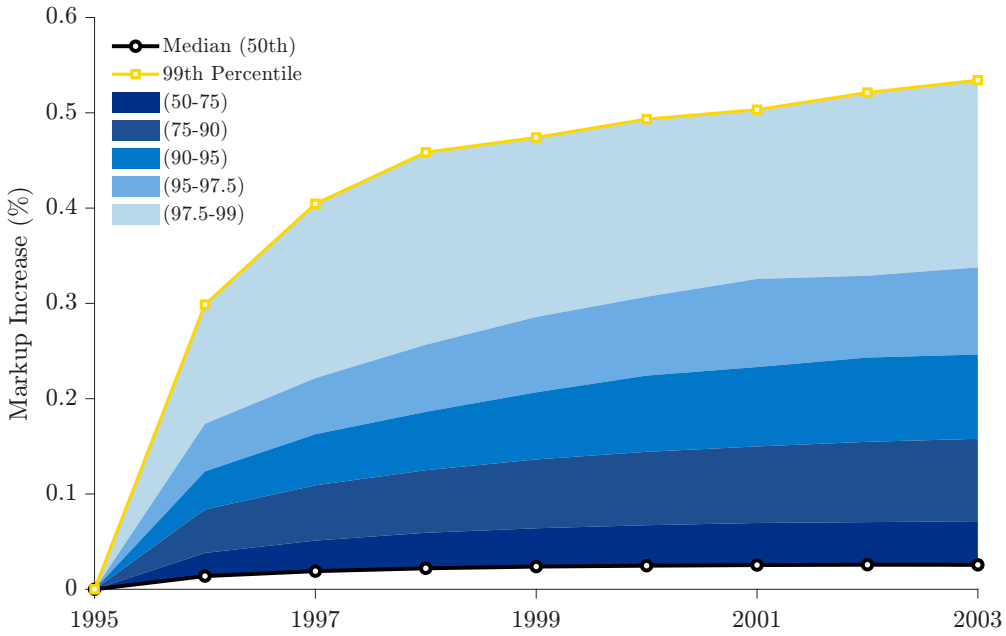
These findings demonstrate that discount rate changes have asymmetric effects across the firm distribution in oligopolistic markets, with implications extending beyond aggregate investment to market structure and pricing behavior. This in turn suggests that policies that affect the discount rate may inadvertently affect product market competition, with persistently lowered discount rates contributing to decreased market concentration and smaller markups. This mechanism is absent from standard models with perfect competition or monopolistic competition and CES demand.

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 drivers of rising market power. Grullon et al. (2019) further 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 between 1995 and 2021. This is the question that our third counterfactual exercise seeks to answer. 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.[20] This counterfactual isolates the effect of common ownership on investment decisions by comparing an economy where merged entities internalize product market

[20]We retain completed acquisitions between firms in our 1995 sample, excluding minority-stake purchases, buybacks,

FIGURE 18: MERGERS SIMULATION: MARKUPS



spillovers (setting $\mathbf{\Gamma}$ according to observed merger activity) against a baseline where all firms remain independent. The analysis encompasses 4,357 mergers involving 3,240 firms—representing 56.77% 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 transactions between competing firms. In practice, some transactions may represent vertical integration (Frésard, Hoberg, and Phillips, 2020) 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 17 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 \$1.29 trillion in 1995 to approximately \$1.19 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 the aggregate investment of merging firms is lower, reaching \$1.17 trillion by 2003.

The associated investment suppression observed in Figure 17 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

self-tenders, management buyouts, and subsidiary deals. When the SDC event history reports the post-transaction ownership stake, we require it to be at least 50%; we also retain completed acquisition events for which no final ownership percentage is reported. We collapse repeat transactions to unique acquirer–target pairs. All qualifying mergers are imposed simultaneously and permanently in 1995: the corresponding entries of $\mathbf{\Gamma}$ are set to one, acquired targets continue to operate as separate business units, and ownership links are not extended transitively across chains of deals.

entities through increased output and lower equilibrium prices. Interestingly, however, while the observed magnitude of the decline in aggregate investment is roughly $17.52 billion and therefore highly significant in absolute values, it represents a modest decline in relative terms, as it captures a 1.47% reduction in the aggregate investment of merging firms, with respect to the no-merger baseline. Since merging firms account for a majority of our sample, 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; an economy-wide figure would additionally reflect the equilibrium investment response of non-merging rivals, which our model transmits through the product market network.

Analyzing merger activity more granularly at the firm level, we next observe striking heterogeneity in the impact of mergers on markups. Figure 18 illustrates this dispersion through a percentile fan chart focused on the upper half of the distribution, plotting the % change in markups in the case with mergers with respect to the corresponding markup 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 0.16% in 2003, while those in the 95th-99th percentile range experience increases exceeding 0.25% in 2003. Most remarkably, mergers in the 99th percentile range generate a markup increase of approximately 0.53% in 2003, indicating significant anticompetitive effects 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 significant effects on markups.

5 Extensions and Sensitivity Analysis

In this section, we consider various extensions and generalizations of the Network Q model presented in Section 2, alongside a series of sensitivity analyses, where we re-run our quantitative exercises as we vary key model parameters. We begin by showing that our model can easily be generalized to accommodate arbitrarily rich investment spillovers between firms.

5.1 Arbitrarily Rich Spillovers: the General-Form Network Q Model

In the model presented in Section 2, strategic interactions between firms are governed by the matrix of product substitution $\mathbf{\Sigma}$. However, it is reasonable to ask whether the model is amenable to extensions that can accommodate additional layers of strategic interaction among firms, such as input-output linkages (Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi, 2012), labor market power (Azar and Marinescu, 2024), etc...

In this subsection, we show how the PME solution to the dynamic game presented in Section 2 actually solves an entire class of dynamic network oligopoly games, which can incorporate arbitrary richness in strategic interactions. The model solution remains virtually identical, as long as operating profits remain quadratic in the vector of capital stocks.

The General-Form Network Q Model. Consider now a more general version of the Network Q model presented in Section 2, where firms' operating profits are given by the quadratic function:

$$\boldsymbol{\pi}(\mathbf{k}_t) \;=\; \mathbf{K}_t\mathbf{Z}\boldsymbol{\beta} - \mathbf{K}_t\mathbf{Z}\mathbf{G}\mathbf{Z}\mathbf{k}_t \tag{82}$$

where the vector $\boldsymbol{\beta}$ and the matrix $\mathbf{G}$ are defined generically[21]. The matrix $\mathbf{G}$ captures arbitrary investment spillovers. The production block of the model is left unchanged. The key requirement for tractability is that we preserve the linear quadratic functional form of the operating profits function, so that the first-order conditions remain linear and the equilibrium policies are explicitly defined.

Then, identical derivations (which we omit for brevity) lead to the following equations for the PME path. The steady state is given by the following variant of equation (43).

$$\mathbb{k} \;=\; \boldsymbol{\Omega}^{-1}\left[\mathbf{Z}\boldsymbol{\beta} - (r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^{k}\right] \tag{83}$$

while the transition matrix is unchanged from equation (44). The key difference is that the competitive pressure matrix $\boldsymbol{\Omega}$ takes the more general form

$$\boldsymbol{\Omega} \;=\; \mathbf{Z}\left[\mathbf{G}\circ\left(\mathbf{1}\mathbf{1}' - \mathbf{I} + \mathbf{D} + \boldsymbol{\Gamma}\right)\right]\mathbf{Z} \tag{84}$$

it still captures how an increase in capital by any other peer firms affects each firm's incentive to invest in the next period.

It is easy to see that the model presented in Section 2 provides a product market microfoundation such that $\boldsymbol{\beta} = (\mathbf{b}-\mathbf{c})$ and $\mathbf{G} = (\mathbf{I}+\boldsymbol{\Sigma})$. In this general formulation, the Network Q model can alternatively be extended and applied to several other types of networks. We next provide a couple of concrete examples.

Input-Output Linkages. We next show how this general form game can be used to accommodate an alternative microfoundation where firms are also linked through supply chain relationships (as opposed to just final consumer demand) and where producers exercise market power in input as well as output markets.

To implement this alternative microfoundation we follow Pellegrino (2025) and modify the firms' Leontief production function to incorporate not only labor but other inputs bought from peer firms. Specifically, let o_{nm} be the quantity of input m required per unit of gross output of firm n and let $\mathbf{y}_t^c$ denote the vector of firm output that is destined to final consumption. The consumer utility in equation (5) is then re-written in terms of the vector of final output $\mathbf{y}_t^c$ instead of $\mathbf{y}_t$ so that final consumer demand is given by

$$\mathbf{p}_t \;=\; \mathbf{b} - (\mathbf{I}+\boldsymbol{\Sigma})\mathbf{y}_t^c \tag{85}$$

Collecting all o_{nm} inside the $N\times N$ matrix $\mathbf{O}$, we can write the following linear relationship between final output and total output:

$$\mathbf{y}_t^c \;=\; (\mathbf{I}-\mathbf{O})'\mathbf{y}_t \tag{86}$$

The operating profits of firm n are given by:

$$\pi_n(\mathbf{k}_t) \;=\; \left[p_n(\mathbf{k}_t) - c_n - \sum_{m=1}^{N} o_{nm}\, p_m(\mathbf{k}_t)\right] y_n(\mathbf{k}_t) \tag{87}$$

By plugging the inverse demand function (85) inside the operating profit function (87) we obtain the following expressions for $\boldsymbol{\beta}$ and $\mathbf{G}$:

$$\boldsymbol{\beta} \;=\; (\mathbf{I}-\mathbf{O})\mathbf{b} - \mathbf{c} \qquad\qquad \mathbf{G} \;=\; (\mathbf{I}-\mathbf{O})\,(\mathbf{I}+\boldsymbol{\Sigma})\,(\mathbf{I}-\mathbf{O})' \tag{88}$$

[21] We still require the resulting competitive pressure matrix $\boldsymbol{\Omega}$ to be symmetric and positive definite.

the network $\mathbf{G}$ now combines both product market as well as input-output linkages. This extended version of the model predicts that the input-output network affects the long-term level distribution of firms' capital in steady state through vector $\boldsymbol{\beta}$, as well as the transition dynamics of aggregate investment through the matrix $\boldsymbol{\Omega}$, which in turn depends on $\mathbf{G}$.

More Complex Ownership Structures. Another relevant extension involves relaxing the binary nature of the ownership matrix $\boldsymbol{\Gamma}$. In the baseline model, the entry γ_{mn} of this matrix takes 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 through matrix $\boldsymbol{\Omega}$. 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. In that case, $\boldsymbol{\Gamma}$ can be not only non-binary but also non-block-diagonal. This extension naturally accommodates the analysis of cross-shareholdings, joint ventures, and strategic alliances where firms partially internalize spillovers without full common ownership.[22]

Exploiting the case where the elements γ_{mn} range between 0 and 1, the model also lends itself to studying common ownership structures, where institutional investors hold significant stakes in multiple competing firms. Following Ederer and Pellegrino (2026), the degree of common ownership can be captured through appropriate calibration of the $\boldsymbol{\Gamma}$ matrix, allowing the analysis of how institutional ownership concentration affects investment incentives and market outcomes. By inspection of matrix $\boldsymbol{\Omega}$, we infer that common ownership, cross-shareholdings, joint ventures and strategic alliances will all affect the PME path.

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 of firms' stocks of capital.

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. While heterogeneous θ_n are redundant in our setting (as we explained in Section 2, z_n already generates massive firm heterogeneity in adjustment costs), this extension might be useful for alternative empirical applications.

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

$$\psi_n(k_{nt+1};k_{nt}) \stackrel{\text{def}}{=} p_n^k i_{nt} + \frac{\theta_n}{2}(k_{nt+1}-k_{nt})^2 \tag{89}$$

The main cost of this added heterogeneity is that we lose the elegant closed-form solution for the transition matrix $\boldsymbol{\Phi}$ presented in equation (44). As mentioned earlier, that expression is obtained as the solution of a Non-symmetric Algebraic Riccati Equation (NARE), which we derive in Appendix

[22] If ownership stakes are not reciprocal ($\gamma_{mn} \neq \gamma_{nm}$), $\boldsymbol{\Omega}$ becomes asymmetric: the interior solution applies as long as $\boldsymbol{\Omega}$ is diagonalizable with real, positive eigenvalues, but the potential game formulation that we use to compute constrained steady states requires symmetry.

A. In the general case with heterogeneous discount rates and adjustment costs, the NARE is:

$$\mathbf{\Phi}^2 - \left[\mathbf{\Theta}^{-1}\mathbf{\Omega} + (2\mathbf{I} + \mathbf{R})\right]\mathbf{\Phi} + (\mathbf{I} + \mathbf{R}) = \mathbf{0} \tag{90}$$

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.

5.3 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 competition in prices and nominal rigidities. 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 price adjustment costs à la Rotemberg (1982). We start by re-writing the vector of profits in terms of prices

$$\boldsymbol{\pi}_t(\mathbf{p}_t) \;=\; (\mathbf{p}_t - \mathbf{c}) \circ \left[(\mathbf{I} + \mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{p}_t)\right] \tag{91}$$

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

$$\psi(p_{nt+1}, p_{nt}) = \frac{\theta}{2}(p_{nt+1} - p_{nt})^2 \tag{92}$$

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

$$\frac{1}{1+r}\left[(\mathbf{I} + \mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{p}_{t+1}) - \mathbf{\Upsilon}\mathbf{p}_{t+1} + \mathbf{\Upsilon}\mathbf{c} + \theta(\mathbf{p}_{t+2} - \mathbf{p}_{t+1})\right] \;=\; \theta(\mathbf{p}_{t+1} - \mathbf{p}_t) \tag{93}$$

where $\mathbf{\Upsilon}$ is a diagonal matrix that shares its diagonal entries with the inverse of $(\mathbf{I} + \mathbf{\Sigma})$.

As for the Cournot case, we posit a stable linear PME. 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):

$$\mathbf{p}_{t+1} - \mathbb{p} \;=\; \mathbf{\Phi}(\mathbf{p}_t - \mathbb{p}) \tag{94}$$

By setting $\mathbf{p}_{t+2} = \mathbf{p}_{t+1} = \mathbf{p}_t = \mathbb{p}$ we find the PME steady-state price vector.

$$\mathbb{p} - \mathbf{c} \;=\; [\mathbf{I} + (\mathbf{I} + \mathbf{\Sigma})\,\mathbf{\Upsilon}]^{-1}(\mathbf{b} - \mathbf{c}) \tag{95}$$

As in Section 2, the equilibrium transition matrix is obtained as the solution of a Non-symmetric Algebraic Riccati Equation (NARE):

$$\mathbf{\Phi} \;=\; \frac{1}{2}\mathbf{\mho} - \sqrt[+]{\frac{1}{4}\mathbf{\mho}^2 - (1+r)\mathbf{I}} \qquad \text{where} \qquad \mathbf{\mho} \stackrel{\text{def}}{=} \frac{1}{\theta}\left(\mathbf{1}\mathbf{1}' + \mathbf{I}\right) \circ (\mathbf{I} + \mathbf{\Sigma})^{-1} + (2+r)\mathbf{I} \tag{96}$$

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.

5.4 Sensitivity Analysis

It is natural to ask how the quantitative results in Section 4 vary under alternative calibration strategies. Specifically, 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. Hence, our core quantitative findings and counterfactual exercises should be sensitive to the calibration of these parameters.

We address this concern systematically in Appendix D, where we reproduce all the figures from Sections 3 and 4 under three alternative specifications. In Appendix D.1, we replace our baseline discount rate with the 10-year US Treasury rate plus 5 percentage points. In Appendix D.2, 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 D.3, we multiply the baseline estimate of θ by 1.5. Across all three specifications, the qualitative patterns documented in our baseline calibration remain robust.

As an additional exercise and for the sake of completeness, we also winsorize the input variables of our calibration at 1% on each tail of the distribution and present the associated results in Appendix D.4. This is common practice in the neoclassical literature, which explicitly removes outliers with the goal of explaining the behavior of the average firm in the working sample. Note however that while qualitative patterns are also preserved in this case, we refrain from this practice in our baseline calibration because winsorization mechanically and arbitrarily caps the values reported by the largest and most important firms in our working sample. Our model is estimated at the firm level, and unlike previous studies we require unrestricted cross-sectional variation in revenues, total capital and firm value to derive accurate inferences on key objects such as the degree of concentration or the cross-sectional distribution of markups.

6 Conclusions

In this paper we have introduced 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 a hedonic demand system into a neoclassical investment setting, we endogenize markups and characterize in matrix form the network of investment spillovers. After obtaining an explicit solution for firms' investment policies, we then apply the model to US firm-level data: our product market network features thousands of firms and dozens of millions of firm pairs.

The quantitative analyses in the paper illustrate the multiple practical applications of the model, as a tool for both researchers and policymakers alike. Specifically, we make inferences on the aggregate

effects on investment of specific corporate events—such as cartels or mergers—and run exercises that are relevant for policymakers—such as studying the impact of a decrease 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. One key insight of our product market network approach to Q theory is to show how firms' strategic behavior (at the micro level) ultimately affects important macro objects such as aggregate capital formation or the distribution of markups.

Our Network Q model can be extended and applied to many types of networks. These alternative networks include, but are not limited to, input-output linkages, common ownership, and strategic interactions in production technologies and input markets. The many advantages and potential extensions of our model suggest a wide array of future applications in multiple fields of economics, including macroeconomics, corporate finance, and international economics.

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* 84 (5): 1369–1384.

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.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The fall of the labor share and the rise of superstar firms." *The Quarterly Journal of Economics* 135 (2): 645–709.

Azar, José, and Ioana Marinescu. 2024. "Monopsony power in the labor market." In *Handbook of Labor Economics*, Volume 5. 761–827, Elsevier.

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 Misallocation 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.

Best, Lea, Benjamin Born, and Manuel Menkhoff. 2026. "The impact of interest: Firms' investment sensitivity to interest rates." technical report, ECONtribute Discussion Paper.

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

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

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

Carvalho, Vasco, and Xavier Gabaix. 2013. "The Great Diversification and its Undoing." *American Economic Review* 103 (5): 1697–1727.

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." *Working Paper, 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.

Covarrubias, Matias. 2022. "Dynamic oligopoly and monetary policy: A deep reinforcement learning approach." *Manuscript, Banco de España*.

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

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

Dai, Min, Zhaoli Jiang, and Neng Wang. 2024. "Option Exercise Games and the q Theory of Investment." Working Paper, November 2024.

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.

Dou, Winston Wei, Yan Ji, and Wei Wu. 2021. "Competition, profitability, and discount rates." *Journal of Financial Economics* 140 (2): 582–620.

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

Ederer, Florian, and Bruno Pellegrino. 2026. "A Tale of Two Networks: Common Ownership and Product Market Rivalry." *Review of Economic Studies* 93 (3): 1746–1788.

Eeckhout, Jan. 2025. "The Value and Profits of Firms." *Journal of the European Economic Association* 23 (1): 52–118.

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

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.

Engwerda, Jacob C. 2005. *LQ Dynamic Optimization and Differential Games*. Chichester: John Wiley & Sons.

Engwerda, Jacob C. 2016. "Properties of feedback Nash equilibria in scalar LQ differential games." *Automatica* 69: 364–374.

Erickson, Timothy, and Toni M. Whited. 2000. "Measurement Error and the Relationship between Investment and q." *Journal of Political Economy* 108 (5): 1027–1057.

Erickson, Timothy, and Toni M. Whited. 2002. "Two-Step GMM Estimation of the Errors-in-Variables Model Using High-Order Moments." *Econometric Theory* 18 (3): 776–799.

Flynn, Joel P., Christina Patterson, and John Sturm. 2022. "Fiscal policy in a networked economy." technical report, National Bureau of Economic Research.

Frésard, Laurent, Gerard Hoberg, and Gordon M Phillips. 2020. "Innovation activities and integration through vertical acquisitions." *The Review of Financial Studies* 33 (7): 2937–2976.

Fudenberg, Drew, and David K. Levine. 1993. "Self-Confirming Equilibrium." *Econometrica* 61 (3): 523–545.

Gabaix, Xavier. 2009. "Power laws in economics and finance." *Annual Review of Economics* 1 (1): 255–294.

Gabaix, Xavier, and Rustam Ibragimov. 2011. "Rank-1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents." *Journal of Business & Economic Statistics* 29 (1): 24–39.

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*.

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." *Working Paper, Bocconi University*.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely. 2019. "Are US Industries Becoming More Concentrated?" *Review of Finance* 23 (4): 697–743.

Hayashi, Fumio. 1982. "Tobin's marginal q and average q: A neoclassical interpretation." *Econometrica* 50 (1): 213–224.

Hennessy, Christopher. 2004. "Tobin's Q, Debt Overhang, and Investment." *Journal of Finance* 59 (4): 1717–1742.

Higham, Nicholas J. 2008. *Functions of Matrices: Theory and Computation*. Philadelphia, PA: SIAM.

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

Hoberg, Gerard, and Gordon Phillips. 2010a. "Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis." *The Review of Financial Studies* 23 (10): 3773–3811.

Hoberg, Gerard, and Gordon Phillips. 2010b. "Real and Financial Industry Booms and Busts." *The Journal of Finance* 65 (1): 45–86.

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 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." *Working Paper*.

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.

Martin-Baillon, Alaïs. 2021. "Firms' Marginal Propensities to Invest." *Manuscript, September*.

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

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

Moll, Benjamin. 2025. "The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics." *The Economic Journal*.

Monti, Andrea, Benita Nortmann, Thulasi Mylvaganam, and Mario Sassano. 2024. "Feedback and Open-Loop Nash Equilibria for LQ Infinite-Horizon Discrete-Time Dynamic Games." *SIAM Journal on Control and Optimization* 62 (3): 1417–1436.

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.

Nortmann, Benita, Mario Sassano, and Thulasi Mylvaganam. 2025. "Feedback Nash Equilibria for Scalar N-Player Linear Quadratic Dynamic Games." *Automatica* 174: 112133.

Okumura, Koki. 2025. "Ownership Structure and Economic Growth." *Working Paper*.

Ottonello, Pablo, and Thomas Winberry. 2020. "Financial heterogeneity and the investment channel of monetary policy." *Econometrica* 88 (6): 2473–2502.

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 (2): 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* 28 (1): 69–88.

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

Rotemberg, Julio J. 1982. "Monopolistic Price Adjustment and Aggregate Output." *The Review of Economic Studies* 49 (4): 517–531.

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

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

Syverson, Chad. 2025. "Markups and markdowns." *Annual Review of Economics* 17 (1): 57–76.

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): 2815–2849.

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.

Winberry, Thomas, Adrien Auclert, Matthew Rognlie, and Ludwig Straub. 2025. "New Keynesian Economics with Household and Firm Heterogeneity." technical report, National Bureau of Economic Research.

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. We begin by explicitly solving the feedback effects internalized by managers in the MPE. Throughout this step we assume single-unit managers ($\mathbf{\Gamma} = \mathbf{I}$, so that $\mathscr{M}_n = \{n\}$ for every n), which is what allows the payoff gradient with respect to rival capital to be written as $-\boldsymbol{\omega}_{n,-n} k_{nT}$ below. We then move on to formally deriving the steady-state capital and transition matrix for both MPE and PME, using a common analytical approach.

As a first step, we solve for the feedback effect of the MPE investment strategy of a specific business unit n, using a forward iteration of the linear investment policy in equation (20):

$$\frac{\partial v_n(\mathbf{k}_{t+2})}{\partial \mathbf{k}_{-n,t+2}} \cdot \frac{\partial \mathbf{k}_{-n,t+2}}{\partial k_{n,t+1}} = -\sum_{T=t+2}^{\infty} (1+r)^{-(T-t-2)} \xi_{nT} k_{nT} \tag{97}$$

where the scalar ξ_{nT} aggregates over peer reactions at time T, to infer how the product market network jointly reacts to the investment of manager n in such period. Hence, we define vectors and matrices that implicitly remove business unit n from the computation of aggregate peer reactions, namely:

$$\xi_{nT} \overset{\text{def}}{=} \boldsymbol{\omega}_{n,-n} \mathbf{\Phi}_{-n,-n}^{T-t-2} \boldsymbol{\varphi}_{-n,n} \tag{98}$$

where $\boldsymbol{\omega}_{n,-n}$ is the row vector that is obtained by deleting column n and all rows except n from the competitive pressure matrix $\mathbf{\Omega}$; $\mathbf{\Phi}_{-n,-n}$ is the sub-matrix of the transition matrix $\mathbf{\Phi}$ that excludes the n^{th} row and the n^{th} column; and $\boldsymbol{\varphi}_{-n,n}$ is the column vector that is obtained by deleting row n and all columns except n from the transition matrix $\mathbf{\Phi}$.

We next define the vector $\boldsymbol{\xi}_T$ and matrix Ξ_T as follows:

$$\boldsymbol{\xi}_T \overset{\text{def}}{=} \begin{bmatrix} \xi_{1T} & \xi_{2T} & \dots & \xi_{NT} \end{bmatrix} \qquad \Xi_T \overset{\text{def}}{=} \text{diag}(\boldsymbol{\xi}_T) \tag{99}$$

The diagonal of matrix Ξ_T thus collects the feedback effects associated with the investment decision of each business unit n at time T. Having defined matrix Ξ_T, we move next to express in a general

form the system of stacked marginal Q as:

$$\mathbf{q}(\mathbf{k}_{t+1}) \;=\; \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \boldsymbol{\Omega}\mathbf{k}_{t+1} + (\mathbf{I}-\boldsymbol{\Delta})\mathbf{p}^{k} + \theta(\mathbf{k}_{t+2}-\mathbf{k}_{t+1}) - \mathbb{1}_{\text{MP}} \sum_{T=t+2}^{\infty} \frac{\boldsymbol{\Xi}_T \mathbf{k}_T}{(1+r)^{(T-t-1)}} \tag{100}$$

where feedback effects are only relevant to managers in the case of the MPE, for which the dummy variable $\mathbb{1}_{\text{MP}}$ is equal to one. The dummy variable $\mathbb{1}_{\text{MP}}$ equals instead zero (and therefore shuts down feedback effects) in the PME.

Evaluating the first-order condition at the steady state, we then obtain

$$\frac{1}{1+r}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - \boldsymbol{\Omega}\mathbb{k} + (\mathbf{I}-\boldsymbol{\Delta})\mathbf{p}^{k} - \mathbb{1}_{\text{MP}}\mathbf{S}_1(\boldsymbol{\Phi},\boldsymbol{\Omega})\mathbb{k}\right] \;=\; \mathbf{p}^{k} \tag{101}$$

where $\mathbf{S}_1(\boldsymbol{\Phi},\boldsymbol{\Omega})$ is a diagonal matrix that is the infinite discounted sum of $\boldsymbol{\Xi}_T$, and depends on the transition and competitive pressure matrices $(\boldsymbol{\Phi},\boldsymbol{\Omega})$:

$$\mathbf{S}_1(\boldsymbol{\Phi},\boldsymbol{\Omega}) \;=\; \sum_{T=t+2}^{\infty} \frac{\boldsymbol{\Xi}_T}{(1+r)^{(T-t-1)}} \tag{102}$$

We can then solve for the steady-state capital vector $\mathbb{k}$ as shown in Proposition 1:

$$\mathbb{k} \;=\; [\boldsymbol{\Omega} + \mathbb{1}_{\text{MP}}\mathbf{S}_1(\boldsymbol{\Phi},\boldsymbol{\Omega})]^{-1}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^{k}\right] \tag{103}$$

We proceed to determine the equilibrium transition matrix $\boldsymbol{\Phi}$. As a preliminary step, we define the infinite sum matrix:

$$\mathbf{S}_2(\boldsymbol{\Phi},\boldsymbol{\Omega}) \;=\; \sum_{T=t+2}^{\infty} \frac{\boldsymbol{\Xi}_T \boldsymbol{\Phi}^{T-t-2}}{(1+r)^{T-t-1}} \tag{104}$$

Evaluating the first-order condition in terms of differences with respect to the steady-state, we next obtain:

$$-\boldsymbol{\Omega}(\mathbf{k}_{t+1}-\mathbb{k}) + \theta(\mathbf{k}_{t+2}-\mathbf{k}_{t+1}) - \mathbb{1}_{\text{MP}}\cdot\mathbf{S}_2(\boldsymbol{\Phi},\boldsymbol{\Omega})\cdot\boldsymbol{\Phi}(\mathbf{k}_{t+1}-\mathbb{k}) = \theta(1+r)(\mathbf{k}_{t+1}-\mathbf{k}_t) \tag{105}$$

Using our guess on the linear investment rule, we have that

$$(\mathbf{k}_{t+2}-\mathbf{k}_{t+1}) \;=\; (\boldsymbol{\Phi}-\mathbf{I})(\mathbf{k}_{t+1}-\mathbb{k}) \tag{106}$$

Using the previous result, we re-write the first-order condition of investment as:

$$-\boldsymbol{\Omega}(\mathbf{k}_{t+1}-\mathbb{k}) + \theta(\boldsymbol{\Phi}-\mathbf{I})(\mathbf{k}_{t+1}-\mathbb{k}) - \mathbb{1}_{\text{MP}}\cdot\mathbf{S}_2(\boldsymbol{\Phi},\boldsymbol{\Omega})\cdot\boldsymbol{\Phi}(\mathbf{k}_{t+1}-\mathbb{k}) = \theta(1+r)\left(\mathbf{I}-\boldsymbol{\Phi}^{-1}\right)(\mathbf{k}_{t+1}-\mathbb{k}) \tag{107}$$

Re-ordering terms, we conclude that:

$$-\frac{1}{\theta}\boldsymbol{\Omega} + (\boldsymbol{\Phi}-\mathbf{I}) - \frac{\mathbb{1}_{\text{MP}}}{\theta}\cdot\mathbf{S}_2(\boldsymbol{\Phi},\boldsymbol{\Omega})\cdot\boldsymbol{\Phi} = (1+r)\left(\mathbf{I}-\boldsymbol{\Phi}^{-1}\right) \tag{108}$$

Multiplying by $\boldsymbol{\Phi}$ and rearranging, we obtain the following matrix equation, as shown in Proposition

1, which implicitly pins down the equilibrium transition matrix $\mathbf{\Phi}$:

$$\left[\mathbf{I} - \frac{\mathbb{1}_{\text{MP}}}{\theta} \cdot \mathbf{S}_2(\mathbf{\Phi}, \mathbf{\Omega})\right] \mathbf{\Phi}^2 - \left[\frac{1}{\theta}\mathbf{\Omega} + (2+r)\mathbf{I}\right] \mathbf{\Phi} + (1+r)\mathbf{I} = \mathbf{0} \tag{109}$$

Having solved for both the steady-state capital vector $\mathbb{k}$ and the transition matrix $\mathbf{\Phi}$ as shown in Proposition 1, we conclude the proof by verifying the existence of the associated equilibrium.

The verification consists of two steps. As a first step in the analysis, we prove the value function of each business unit n is indeed maximized when evaluated at the equilibrium investment strategies. Given the state variable $\mathbf{k}_t$ the optimization problem of each business unit n involves a scalar maximization over k_{nt+1}. Hence, to verify optimality, we require strict concavity of the value function of business unit n with respect to k_{nt+1}.

Given the stacked equations in (100), we derive the marginal value of capital of each business unit n with respect to k_{nt+1} to obtain:

$$\frac{\partial q_n(\mathbf{k}_{t+1})}{\partial k_{nt+1}} = -\omega_{nn} + \theta(\varphi_{nn} - 1) - \mathbb{1}_{\text{MP}}[\mathbf{S}_2(\mathbf{\Phi}, \mathbf{\Omega})\mathbf{\Phi}]_{nn} \tag{110}$$

Simultaneously, using the expression for matrix $\mathbf{\Phi}$ as defined in equation (108), and re-ordering terms, we obtain:

$$-\theta\mathbf{I} + \frac{1}{1+r}[-\mathbf{\Omega} + \theta(\mathbf{\Phi} - \mathbf{I}) - \mathbb{1}_{\text{MP}}\mathbf{S}_2(\mathbf{\Phi}, \mathbf{\Omega})\mathbf{\Phi}] = -\theta\mathbf{\Phi}^{-1} \tag{111}$$

Taking the nn^{th} element from the matrix above, we infer:

$$-\theta + \frac{1}{(1+r)}\left[-\omega_{nn} + \theta(\varphi_{nn} - 1) - \mathbb{1}_{\text{MP}}[\mathbf{S}_2(\mathbf{\Phi}, \mathbf{\Omega})\mathbf{\Phi}]_{nn}\right] = -\theta\left[\mathbf{\Phi}^{-1}\right]_{nn} \tag{112}$$

It follows that the value function of each business unit is concave in the direction of k_{nt+1} for either the MPE or PME of the model as long as:

$$-\theta + \frac{1}{(1+r)}\frac{\partial q_n(\mathbf{k}_{t+1})}{\partial k_{nt+1}} = -\theta\left[\mathbf{\Phi}^{-1}\right]_{nn} < 0 \tag{113}$$

where the condition $\left[\mathbf{\Phi}^{-1}\right]_{nn} > 0$ holds: for $\mathbb{1}_{\text{MP}} = 1$ it is among the conditions of the proposition, and for $\mathbb{1}_{\text{MP}} = 0$ it is implied by the explicit stable solution derived below, which is symmetric with all eigenvalues in $(0, 1)$. Heuristically, the condition amounts to requiring that the serial correlation or persistence of a business unit's own capital be positive in equilibrium. Conversely, at any equilibrium of either type the second-order condition requires $-\theta\left[\mathbf{\Phi}^{-1}\right]_{nn} \leq 0$; imposing the strict inequality excludes only the degenerate case $\left[\mathbf{\Phi}^{-1}\right]_{nn} = 0$, in which the Bellman objective is flat in k_{nt+1} and the investment policy would be indeterminate.

As a second step to ensure existence, we require that the transition matrix $\mathbf{\Phi}$ is stable. By Definition 2, all eigenvalues of the transition matrix $\mathbf{\Phi}$ lie in the interval $(0, 1)$; in particular, the spectral radius of $\mathbf{\Phi}$ is lower than one. We confirm the latter implication with an explicit solution for the eigenvalues of matrix $\mathbf{\Phi}$ for the PME of the model, in the proof of Proposition 2 below.

In the case of the PME, the requirement that $\mathbf{\Phi}$ has a spectral radius lower than one ensures that the transversality condition of the model holds, and therefore the value functions of all business units are bounded. In the alternative case of the MPE, the requirement by definition that $\mathbf{\Phi}_{-n,-n}$ has

spectral radius lower than $1+r$, for any business unit $n = 1,2,...,N$, ensures that the present value of feedback effects is finite. Because payoffs are quadratic and policies linear, the value functions are quadratic in capital in this case as well; stability of $\mathbf{\Phi}$ then implies that capital converges geometrically, so the value functions remain bounded and the transversality condition again holds. In either case, the boundedness of the value functions and strict concavity of the Bellman objective in k_{nt+1} at every state imply that the equilibrium policies cannot be improved further by any one-stage deviation. Because the manager's problem restricts attention to plans that satisfy the transversality condition (18), the discounted continuation value vanishes along every admissible deviation, not merely along the equilibrium path; the one-shot-deviation principle for discounted problems then delivers optimality against arbitrary deviations. □

Proof of Proposition 2 (continued). We derive most of the proof using matrices $\mathbf{\Theta}$ and $\mathbf{R}$ in general form, to accommodate the cases of heterogeneous adjustment costs and firm-specific discount rates discussed in an extension to the baseline model in Section 5. To obtain the explicit solution in Proposition 2, we further impose below $\mathbf{\Theta} = \theta\mathbf{I}$ and $\mathbf{R} = r\mathbf{I}$.

As a first step, we replace $\mathbf{k}_{t+2}$ using the investment rule (20) and re-express the system of stacked first-order conditions (34), (35) and (36) and imposing $\mathbb{1}_{\text{MP}} = 0$:

$$\begin{aligned}\mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_{t+1} + (\mathbf{I}-\mathbf{\Delta})\mathbf{p}^k + \mathbf{\Theta}[\mathbb{k} + \mathbf{\Phi}(\mathbf{k}_{t+1}-\mathbb{k}) - \mathbf{k}_{t+1}] = \\ = (\mathbf{I}+\mathbf{R})\left[\mathbf{p}^k + \mathbf{\Theta}(\mathbf{k}_{t+1}-\mathbf{k}_t)\right]\end{aligned} \tag{114}$$

and distribute:

$$\begin{aligned}\mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_{t+1} + (\mathbf{I}-\mathbf{\Delta})\mathbf{p}^k + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbb{k} - \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbf{k}_{t+1} \quad = \\ = \quad (\mathbf{I}+\mathbf{R})\mathbf{p}^k + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_{t+1} - (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_t\end{aligned} \tag{115}$$

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

$$\begin{aligned}(\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_{t+1} + \mathbf{\Omega}\mathbf{k}_{t+1} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbf{k}_{t+1} \quad = \\ = \quad \mathbf{Z}(\mathbf{b}-\mathbf{c}) - (\mathbf{R}+\mathbf{\Delta})\mathbf{p}^k + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbb{k} + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_t\end{aligned} \tag{116}$$

next, we rewrite the right hand side of the equation in terms of $\mathbf{\Omega}\mathbb{k}$:

$$\begin{aligned}(\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_{t+1} + \mathbf{\Omega}\mathbf{k}_{t+1} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbf{k}_{t+1} \quad = \\ = \quad \mathbf{\Omega}\mathbb{k} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi})\mathbb{k} + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}\mathbf{k}_t\end{aligned} \tag{117}$$

Reordering terms appropriately we obtain:

$$[\mathbf{\Omega} + \mathbf{\Theta}(\mathbf{I}-\mathbf{\Phi}) + (\mathbf{I}+\mathbf{R})\mathbf{\Theta}](\mathbf{k}_{t+1}-\mathbb{k}) = (\mathbf{I}+\mathbf{R})\mathbf{\Theta}(\mathbf{k}_t-\mathbb{k}) \tag{118}$$

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

$$(\mathbf{k}_{t+1}-\mathbb{k}) = \left[\mathbf{\Theta}^{-1}\mathbf{\Omega} + (2\mathbf{I}+\mathbf{R})\mathbf{I} - \mathbf{\Phi}\right]^{-1}(\mathbf{I}+\mathbf{R})(\mathbf{k}_t-\mathbb{k}) \tag{119}$$

Then, comparing with (20), we see that $\mathbf{\Phi}$ must satisfy:

$$\left[\mathbf{\Theta}^{-1}\mathbf{\Omega} + (2\mathbf{I} + \mathbf{R})\mathbf{I} - \mathbf{\Phi}\right]^{-1}(\mathbf{I} + \mathbf{R}) = \mathbf{\Phi} \tag{120}$$

rearranging we obtain:

$$\mathbf{\Phi}^2 - \left[\mathbf{\Theta}^{-1}\mathbf{\Omega} + (2\mathbf{I} + \mathbf{R})\right]\mathbf{\Phi} + (\mathbf{I} + \mathbf{R}) = \mathbf{0} \tag{121}$$

Equation (121) 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 (121) admits 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).[23] 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:

$$\mathbf{X} = -\frac{1}{2}\mathbf{M}_1 \pm \sqrt{\frac{1}{4}\mathbf{M}_1^2 - \mathbf{M}_2} \tag{122}$$

This condition is satisfied in our case if $\mathbf{\Theta} = \theta\mathbf{I}$ and $\mathbf{R} = r\mathbf{I}$. Applying this solution to (121) yields:

$$\mathbf{\Phi} \;=\; \frac{1}{2}\left[\frac{1}{\theta}\mathbf{\Omega} + (2+r)\mathbf{I}\right] \pm \sqrt{\frac{1}{4}\left[\frac{1}{\theta}\mathbf{\Omega} + (2+r)\mathbf{I}\right]^2 - (1+r)\mathbf{I}} \tag{123}$$

As discussed by Higham (2008), the square root of a matrix with positive eigenvalues admits $2^{\mathscr{E}}$ primary square roots, where $\mathscr{E}$ is the number of unique eigenvalues of the radicand, plus infinitely many non-primary square roots when an eigenvalue is repeated. 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 solution obtained by taking the unique principal square root (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

$$\mathbf{\Omega} \;=\; \mathbf{E}\,\mathrm{diag}(\hat{\boldsymbol{\omega}})\,\mathbf{E}^{-1} \tag{124}$$

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:

$$\mathbf{\Phi} \;=\; \mathbf{E}\,\mathrm{diag}\left(\hat{\boldsymbol{\varphi}}\right)\mathbf{E}^{-1} \tag{125}$$

and their respective n^{th} eigenvalues are related by the following function, the scalar equivalent of equation (123)

$$\hat{\varphi}_n \;=\; \frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right) \pm \sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right)^2 - (1+r)} \tag{126}$$

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

This is a known result in matrix theory that applies when $\mathbf{\Omega}$ is symmetric (it is by construction) 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 (123).

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

$$\frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right) + \sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right)^2 - (1+r)} \geq 1 \tag{127}$$

$$\frac{1}{2}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right) - \sqrt{\frac{1}{4}\left(\frac{\hat{\omega}_n}{\theta} + 2 + r\right)^2 - (1+r)} \leq 1 \tag{128}$$

it follows that the unique principal square root matrix, preceded by the minus sign, yields the only stable solution $\mathbf{\Phi}$. The second inequality holds strictly whenever $\hat{\omega}_n > 0$ — as guaranteed by the positive definiteness of $\mathbf{\Omega}$ — so every eigenvalue of the stable solution lies strictly below one. Because the stable solution is a matrix function of the symmetric matrix $\mathbf{\Omega}$, it is itself symmetric; its eigenvalues are strictly positive, because the two roots of each scalar quadratic multiply to $1 + r > 0$ while the larger root is at least one by the first display, and therefore lie in $(0, 1)$. It follows that $\mathbf{\Phi}^{-1}$ is symmetric positive definite, so $\left[\mathbf{\Phi}^{-1}\right]_{nn} > 0$ for every n: the stable solution satisfies the strict-concavity condition of Proposition 1. The stable solution in Equation (123) then ensures the transversality condition in Equation (18) holds in equilibrium, because firm value is quadratic in capital and capital converges geometrically to the steady state. □

Proof of Proposition 3. Under the perturbation in the proposition, equation (37) becomes

$$\mathbf{\Omega}(\eta) = \mathbf{\Omega}_0 + \eta \mathbf{\Omega}_\Sigma, \qquad \mathbf{\Omega}_0 \equiv \mathbf{Z}(\mathbf{I} + \mathbf{D})\mathbf{Z}, \qquad \mathbf{\Omega}_\Sigma \equiv \mathbf{Z}(\mathbf{\Sigma} + \mathbf{\Gamma} \circ \mathbf{\Sigma})\mathbf{Z}. \tag{129}$$

The matrix $\mathbf{\Omega}_0$ is diagonal and positive definite. The first step is to re-write the conditions for MPE from (41) and (42) as:

$$\mathscr{R}^1_{\mathbf{\Phi}}(\mathbf{\Phi}, \eta) = \mathbf{0} \qquad \text{and} \qquad \mathscr{R}^1_{\mathbb{k}}(\mathbf{\Phi}, \mathbb{k}, \eta) = \mathbf{0} \tag{130}$$

in terms of the following "residual" functions (where the ^{1} superscript refers to MPE):

$$\mathscr{R}^1_{\mathbf{\Phi}}\left(\mathbf{\Phi}, \eta\right) \overset{\text{def}}{=} \left[\mathbf{I} - \frac{1}{\theta}\mathbf{S}_2\left(\mathbf{\Phi}, \mathbf{\Omega}(\eta)\right)\right]\mathbf{\Phi}^2 - \left[\frac{1}{\theta}\mathbf{\Omega}(\eta) + (2+r)\mathbf{I}\right]\mathbf{\Phi} + (1+r)\mathbf{I} \tag{131}$$

$$\mathscr{R}^1_{\mathbb{k}}\left(\mathbf{\Phi}, \mathbb{k}, \eta\right) \overset{\text{def}}{=} \left[\mathbf{\Omega}(\eta) + \mathbf{S}_1\left(\mathbf{\Phi}, \mathbf{\Omega}(\eta)\right)\right]\mathbb{k} - \left[\mathbf{Z}(\mathbf{b} - \mathbf{c}) - (r\mathbf{I} + \mathbf{\Delta})\mathbf{p}^k\right] \tag{132}$$

We can do the same with the PME conditions (44) and (43) (the ^{0} superscript refers to PME):

$$\mathscr{R}^0_{\mathbf{\Phi}}(\mathbf{\Phi}, \eta) = \mathbf{0} \qquad \text{and} \qquad \mathscr{R}^0_{\mathbb{k}}(\mathbb{k}, \eta) = \mathbf{0} \tag{133}$$

where the corresponding residual functions are:

$$\mathscr{R}^0_{\mathbf{\Phi}}\left(\mathbf{\Phi}, \eta\right) \overset{\text{def}}{=} \mathbf{\Phi}^2 - \left[\frac{1}{\theta}\mathbf{\Omega}(\eta) + (2+r)\mathbf{I}\right]\mathbf{\Phi} + (1+r)\mathbf{I}. \tag{134}$$

$$\mathscr{R}^0_{\mathbb{k}}(\mathbb{k},\eta) \overset{\text{def}}{=} \boldsymbol{\Omega}(\eta)\mathbb{k} - \left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^k\right]. \tag{135}$$

Next, define the set-valued correspondence $\mathscr{K}^1(\eta)$, that maps values of η to all stable pairs $(\boldsymbol{\Phi},\mathbb{k})$ that solve the MPE conditions:

$$\mathscr{K}^1(\eta) \overset{\text{def}}{=} \left\{(\boldsymbol{\Phi},\mathbb{k}):\ \boldsymbol{\Phi}\text{ stable, }\mathscr{R}^1_{\boldsymbol{\Phi}}\left(\boldsymbol{\Phi},\eta\right)=\mathbf{0}\text{ and }\mathscr{R}^1_{\mathbb{k}}\left(\boldsymbol{\Phi},\mathbb{k},\eta\right)=\mathbf{0}\right\} \tag{136}$$

Next, let $\boldsymbol{\Phi}_0$ and $\mathbb{k}_0$ denote the PME equilibrium values at the disconnected economy. For $\eta=0$, the off-diagonal elements of $\boldsymbol{\Omega}(\eta)$ vanish. Consequently, the strategic-feedback matrices $\mathbf{S}_1$ and $\mathbf{S}_2$ are zero for every transition matrix $\boldsymbol{\Phi}$. The MPE conditions therefore reduce to the PME conditions. By Proposition 2, the stable solution is unique, therefore

$$\mathscr{K}^1(0) \;=\; \{(\boldsymbol{\Phi}_0,\mathbb{k}_0)\} \tag{137}$$

For manager n, a dated strategic-feedback term evaluated at arbitrary matrices $(\boldsymbol{\Phi},\boldsymbol{\Omega})$ has the form

$$\xi_{nt}(\boldsymbol{\Phi},\boldsymbol{\Omega}) \;=\; \boldsymbol{\omega}_{n,-n}\boldsymbol{\Phi}^t_{-n,-n}\boldsymbol{\varphi}_{-n,n} \tag{138}$$

where $\boldsymbol{\omega}_{n,-n}$ is the row vector of off-diagonal entries from row n of $\boldsymbol{\Omega}$, $\boldsymbol{\Phi}_{-n,-n}$ deletes row and column n from $\boldsymbol{\Phi}$, and $\boldsymbol{\varphi}_{-n,n}$ is the column of rival-policy responses to capital k_n. At $\left(\eta=0\right)$ we must have that $\boldsymbol{\omega}_{n,-n}=\mathbf{0}$ because $\boldsymbol{\Omega}_0$ is diagonal, and $\boldsymbol{\varphi}_{-n,n}=\mathbf{0}$ because $\boldsymbol{\Phi}_0$ is diagonal. It follows that, also at $\eta=0$, the derivatives of $\mathbf{S}_1$ and $\mathbf{S}_2$ with respect to both $\boldsymbol{\Phi}$ and $\boldsymbol{\Omega}$ (expressed below as Jacobians) reduce to zero:

$$\left.\frac{\partial \operatorname{vec}\mathbf{S}_j}{\partial(\operatorname{vec}\boldsymbol{\Phi})'}\right|_{(\boldsymbol{\Phi}_0,\boldsymbol{\Omega}_0)}=\mathbf{0}, \qquad \left.\frac{\partial \operatorname{vec}\mathbf{S}_j}{\partial(\operatorname{vec}\boldsymbol{\Omega})'}\right|_{(\boldsymbol{\Phi}_0,\boldsymbol{\Omega}_0)}=\mathbf{0}, \qquad j\in\{1,2\}. \tag{139}$$

At $\eta=0$, strict stability means that the spectral radius of $\boldsymbol{\Phi}_0$ is less than 1. By continuity, matrix powers decay at a uniform geometric rate throughout a sufficiently small closed neighborhood of $(\boldsymbol{\Phi}_0,\boldsymbol{\Omega}_0)$. It follows that the discounted series defining $\mathbf{S}_1$ and $\mathbf{S}_2$, as well as the series obtained by differentiating their summands once or twice with respect to the entries of $\boldsymbol{\Phi}$ and $\boldsymbol{\Omega}$, converge uniformly there. The standard term-by-term differentiation theorem therefore implies that $\mathbf{S}_1$ and $\mathbf{S}_2$, and hence $\mathscr{R}^1_{\boldsymbol{\Phi}}\left(\boldsymbol{\Phi},\eta\right)$ and $\mathscr{R}^1_{\mathbb{k}}\left(\boldsymbol{\Phi},\mathbb{k},\eta\right)$ are twice continuously differentiable there.

Next, we compute the Jacobian of the stacked MPE residuals with respect to $(\operatorname{vec}\boldsymbol{\Phi},\mathbb{k})$. With some abuse of notation, we write $\boldsymbol{\Omega}_0=\operatorname{diag}(\omega_1,\ldots,\omega_N)$ and $\boldsymbol{\Phi}_0=\operatorname{diag}\left(\varphi_1,\ldots,\varphi_N\right)$. At $\eta=0$, equation (42) simplifies to:

$$\varphi_n^2-\left(\frac{\omega_n}{\theta}+2+r\right)\varphi_n+(1+r)=0, \qquad 0<\varphi_n<1. \tag{140}$$

With these results in hand, we can now differentiate $\mathscr{R}^1_{\boldsymbol{\Phi}}\left(\boldsymbol{\Phi},\eta\right)$ with respect to $\boldsymbol{\Phi}$ at $(\boldsymbol{\Phi},\eta)=(\boldsymbol{\Phi}_0,0)$. Again, we write it as a Jacobian of the corresponding vectorized matrices:

$$\left.\frac{\partial \operatorname{vec}\mathscr{R}^1_{\boldsymbol{\Phi}}\left(\boldsymbol{\Phi},\eta\right)}{\partial(\operatorname{vec}\boldsymbol{\Phi})'}\right|_{(\boldsymbol{\Phi}_0,0)} \;=\; \mathbf{I}_N\otimes\boldsymbol{\Phi}_0+\boldsymbol{\Phi}_0'\otimes\mathbf{I}_N-\mathbf{I}_N\otimes\left[\frac{1}{\theta}\boldsymbol{\Omega}_0+(2+r)\mathbf{I}_N\right] \tag{141}$$

Because $\boldsymbol{\Phi}_0$ and $\boldsymbol{\Omega}_0$ are diagonal, the Jacobian is also diagonal. Its off-diagonal entries are zero,

while the diagonal entry corresponding to the coordinate $[\mathbf{\Phi}]_{nm}$ is

$$\left.\frac{\partial\left[\mathscr{R}^1_{\mathbf{\Phi}}\left(\mathbf{\Phi},\eta\right)\right]_{nm}}{\partial\left[\mathbf{\Phi}\right]_{nm}}\right|_{(\mathbf{\Phi}_0,0)} = \varphi_n+\varphi_m-\frac{\omega_n}{\theta}-(2+r) = \varphi_m-\frac{1+r}{\varphi_n}. \tag{142}$$

where the second equality follows from (140), which implies $\omega_n/\theta+(2+r)=\varphi_n+(1+r)/\varphi_n$. Because $\varphi_m<1$ while $(1+r)/\varphi_n>1$, every coefficient in equation (142) is nonzero. Hence the Jacobian is invertible. By construction, the transition residual does not depend on $\mathbb{k}$, so

$$\frac{\partial \operatorname{vec}\mathscr{R}^1_{\mathbf{\Phi}}\left(\mathbf{\Phi},\eta\right)}{\partial\,\mathbb{k}'}=\mathbf{0}. \tag{143}$$

At $\eta=0$, the vanishing derivative of $\mathbf{S}_1$ with respect to $\mathbf{\Phi}$ gives the remaining derivatives

$$\left.\frac{\partial\mathscr{R}^1_{\mathbb{k}}\left(\mathbf{\Phi},\mathbb{k},\eta\right)}{\partial(\operatorname{vec}\mathbf{\Phi})'}\right|_{(\mathbf{\Phi}_0,\mathbb{k}_0,0)}=\mathbf{0},\qquad \left.\frac{\partial\mathscr{R}^1_{\mathbb{k}}\left(\mathbf{\Phi},\mathbb{k},\eta\right)}{\partial\,\mathbb{k}'}\right|_{(\mathbf{\Phi}_0,\mathbb{k}_0,0)}=\mathbf{\Omega}_0. \tag{144}$$

The joint Jacobian therefore has the block-diagonal form

$$\frac{\partial}{\partial[(\operatorname{vec}\mathbf{\Phi})',\,\mathbb{k}']}\left[\begin{array}{c}\operatorname{vec}\mathscr{R}^1_{\mathbf{\Phi}}\left(\mathbf{\Phi},\eta\right)\\ \mathscr{R}^1_{\mathbb{k}}\left(\mathbf{\Phi},\mathbb{k},\eta\right)\end{array}\right]\Bigg|_{(\mathbf{\Phi}_0,\mathbb{k}_0,0)}=\left[\begin{array}{cc}\left.\frac{\partial\operatorname{vec}\mathscr{R}^1_{\mathbf{\Phi}}(\mathbf{\Phi},\eta)}{\partial(\operatorname{vec}\mathbf{\Phi})'}\right|_{(\mathbf{\Phi}_0,0)} & \mathbf{0}\\ \mathbf{0} & \mathbf{\Omega}_0\end{array}\right]. \tag{145}$$

Both diagonal blocks are invertible, therefore the joint Jacobian is invertible. Positive definiteness of $\mathbf{\Omega}_0$ and continuity of $\mathbf{\Omega}(\eta)$ ensure that the competitive pressure matrix remains positive definite on an open interval of η around zero. Because (as established above) the residual system is twice continuously differentiable in a neighborhood of $(\mathbf{\Phi}_0,\mathbb{k}_0,0)$ and the joint Jacobian in equation (145) is invertible, the implicit-function theorem provides a number $\bar{\eta}>0$, a neighborhood $\mathscr{N}$ of $(\mathbf{\Phi}_0,\mathbb{k}_0)$, and unique twice continuously differentiable functions $\mathbf{\Phi}^1(\eta)$ and $\mathbb{k}^1(\eta)$ such that

$$\mathscr{K}^1(\eta)\cap\mathscr{N}=\left\{\left(\mathbf{\Phi}^1(\eta),\mathbb{k}^1(\eta)\right)\right\},\qquad \eta\in(-\bar{\eta},\bar{\eta}). \tag{146}$$

At this point, we have established that the MPE necessary conditions admit a locally unique solution inside $\mathscr{N}$: the candidate pair $\left(\mathbf{\Phi}^1(\eta),\mathbb{k}^1(\eta)\right)$. What remains is to show that the candidate is indeed an MPE — that is, that it satisfies the additional conditions under which the necessary conditions become sufficient. The sufficiency part of Proposition 1 applies to the candidate provided three conditions hold at $\left(\mathbf{\Phi}^1(\eta),\mathbb{k}^1(\eta)\right)$: first, that $\left[\mathbf{\Phi}^1(\eta)^{-1}\right]_{nn}>0$ for every n, which is equivalent to strict concavity of each business unit's value function with respect to its own next-period capital; second, that $\mathbf{\Phi}^1(\eta)$ is stable, which delivers the transversality condition and bounded value functions; and third, that $\left[\mathbf{\Phi}^1(\eta)\right]_{-n,-n}$ has spectral radius lower than $1+r$ for every n, so that the discounted sums defining $\mathbf{S}_1$ and $\mathbf{S}_2$ converge and the present value of feedback effects is finite. At $\eta=0$ all three conditions hold strictly: $\left[\mathbf{\Phi}_0^{-1}\right]_{nn}=1/\varphi_n>1$, the spectral radius of $\mathbf{\Phi}_0$ equals $\max_n\varphi_n<1$, and the spectral radius of $[\mathbf{\Phi}_0]_{-n,-n}$ is below one, hence below $1+r$. Because $\mathbf{\Phi}^1(\eta)$ is continuous in η, each condition persists after further shrinking $\bar{\eta}$ and $\mathscr{N}$ if necessary. The second condition requires an additional argument, because stability in the sense of Definition 2 asks for eigenvalues that are real as well as smaller than one, whereas continuity alone preserves only the spectral radius. Since $\mathbf{\Sigma}=\mathbf{0}$ at $\eta=0$, both $\mathbf{\Omega}(0)$ and $\mathbf{\Phi}_0$ are diagonal, the latter with entries $\varphi(\hat{\omega}_n)$, where $\varphi(\cdot)$ denotes the stable

root map in equation (51). That map is strictly decreasing, as shown in the proof of Proposition 4, and therefore injective: the assumed distinctness of the eigenvalues of $\mathbf{\Omega}(0)$ implies that $\mathbf{\Phi}_0$ has N simple eigenvalues. A simple real eigenvalue of a real matrix cannot leave the real line under a small real perturbation, since the complex eigenvalues of a real matrix occur in conjugate pairs; hence, after further shrinking $\bar{\eta}$ if necessary, the eigenvalues of $\mathbf{\Phi}^1(\eta)$ remain real and, by continuity, remain inside $(0,1)$. Thus the candidate is an MPE, which proves both existence and local uniqueness, as well as the differentiability asserted in the first part of the proposition.

We next compare this locally unique MPE with the PME solution. Let $\left(\mathbf{\Phi}^0(\eta), \mathbb{k}^0(\eta)\right)$ denote the unique stable pair that sets $\mathscr{R}^0_{\mathbf{\Phi}}\left(\mathbf{\Phi}, \eta\right)$ and $\mathscr{R}^0_{\mathbb{k}}(\mathbb{k}, \eta)$ equal to zero. At $\eta = 0$, the MPE and PME residual systems have the same joint Jacobian with respect to $(\operatorname{vec}\mathbf{\Phi}, \mathbb{k})$, equal to the expression on the right side of equation (145). They also have the same derivative with respect to η at that point. Both statements follow because $\mathbf{S}_1$, $\mathbf{S}_2$, and their first derivatives with respect to $\mathbf{\Phi}$ and $\mathbf{\Omega}$ vanish at $(\mathbf{\Phi}_0, \mathbf{\Omega}_0)$. In particular, the implicit-function theorem applies verbatim to the PME residual system and yields, after further shrinking $\bar{\eta}$ and $\mathscr{N}$ if necessary, a unique twice continuously differentiable solution of the PME conditions inside $\mathscr{N}$. Because all transition matrices in $\mathscr{N}$ are stable, and the stable PME pair is unique by Proposition 2, this solution coincides with $\left(\mathbf{\Phi}^0(\eta), \mathbb{k}^0(\eta)\right)$, which is therefore twice continuously differentiable in η. Implicit differentiation of the MPE and PME residual systems of equations therefore yields

$$\left.\frac{\mathrm{d}\mathbf{\Phi}^1(\eta)}{\mathrm{d}\eta}\right|_{\eta=0} = \left.\frac{\mathrm{d}\mathbf{\Phi}^0(\eta)}{\mathrm{d}\eta}\right|_{\eta=0}, \qquad \left.\frac{\mathrm{d}\mathbb{k}^1(\eta)}{\mathrm{d}\eta}\right|_{\eta=0} = \left.\frac{\mathrm{d}\mathbb{k}^0(\eta)}{\mathrm{d}\eta}\right|_{\eta=0} \tag{147}$$

Because both locally defined solutions are twice continuously differentiable and coincide at zero, Taylor's theorem gives

$$\mathbf{\Phi}^1\left(\eta\right) - \mathbf{\Phi}^0\left(\eta\right) = O\left(\eta^2\right), \qquad \mathbb{k}^1\left(\eta\right) - \mathbb{k}^0\left(\eta\right) = O\left(\eta^2\right) \tag{148}$$

This proves equation (47) and completes the proof. □

Proof of Proposition 4. Consider first the single-product, symmetric Cournot oligopoly case. Matrix $\mathbf{\Omega}$ is then equal to an $N \times N$ 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:

$$\mathbf{\Omega} = z^2\sigma\mathbf{1}\mathbf{1}' + z^2(2-\sigma)\mathbf{I}. \tag{149}$$

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 eigenvalue of matrix $\mathbf{\Omega}$ associated with the unit eigenvector, for the symmetric, single-product Cournot oligopoly case, is equal to $z^2[2+(N-1)\sigma]$, as stated in equation (52). It remains to sign the comparative statics. By the Corollary to Proposition 2, the eigenvalues of $\mathbf{\Phi}$ and $\mathbf{\Omega}$ are related by the scalar map (51). Writing $a \overset{\text{def}}{=} \hat{\omega}_{\mathbf{1}}/\theta + 2 + r$, that map reads $\hat{\varphi}_{\mathbf{1}} = a/2 - \sqrt{a^2/4 - (1+r)}$, so

that

$$\frac{\mathrm{d}\hat{\varphi}_{\mathbf{1}}}{\mathrm{d}a} = \frac{1}{2} - \frac{a/4}{\sqrt{a^2/4-(1+r)}} < 0, \tag{150}$$

where the inequality follows from $\sqrt{a^2/4-(1+r)} < a/2$, which holds because $r > 0$. Since a is strictly increasing in $\hat{\omega}_{\mathbf{1}}$,$\hat{\varphi}_{\mathbf{1}}$ is strictly decreasing in $\hat{\omega}_{\mathbf{1}}$. Differentiating (52) gives

$$\frac{\partial\hat{\omega}_{\mathbf{1}}}{\partial z} = 2z[2+(N-1)\sigma] > 0, \qquad \frac{\partial\hat{\omega}_{\mathbf{1}}}{\partial\sigma} = z^2(N-1) \geq 0, \qquad \frac{\partial\hat{\omega}_{\mathbf{1}}}{\partial N} = z^2\sigma \geq 0, \tag{151}$$

and the three signs stated in the proposition follow by the chain rule. The second and third inequalities are strict provided $N > 1$ and $\sigma > 0$, respectively. □

Proof of Proposition 5. As a first step, we expand the expression for $\boldsymbol{\pi}(\mathbb{k})$ and collect $\mathbb{K}$:

$$\frac{r}{1+r}\mathbf{v}^e(\mathbb{k}) = \mathbb{K}\left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{Z}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\mathbb{k} - \boldsymbol{\Delta}\mathbf{p}^k\right] \tag{152}$$

Then rearrange equation (43), noting that in the single-product Cournot case ($\mathbf{D} = \boldsymbol{\Gamma} = \mathbf{I}$) the competitive pressure matrix reduces to $\boldsymbol{\Omega} = \mathbf{Z}(2\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}$:

$$\mathbf{Z}(2\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\mathbb{k} = \mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\boldsymbol{\Delta})\mathbf{p}^k \tag{153}$$

to obtain the term in square brackets in equation (152)

$$\mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{Z}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\mathbb{k} - \boldsymbol{\Delta}\mathbf{p}^k = r\mathbf{p}^k + \mathbf{Z}^2\mathbb{k} \tag{154}$$

And replace it inside equation (152) to obtain:

$$\frac{r}{1+r}\mathbf{v}^e(\mathbb{k}) = r\,\mathbf{P}^k\mathbb{k} + \mathbb{K}\mathbf{Z}^2\mathbb{k} \tag{155}$$

Then divide both sides by $r/(1+r)$ to obtain equation (56). Finally, since $\boldsymbol{\Gamma} = \mathbf{I}$ implies $\mathscr{M}_n = \{n\}$, equation (55) gives $v_n(\mathbb{k}) = v_n^e(\mathbb{k})$. □

Proof of Proposition 6. For the derivation, let $\tilde{\mathbf{k}}_t \equiv \mathbf{k}_t - \mathbb{k}$. Iterating the PME policy and using the law of motion for capital gives

$$\tilde{\mathbf{k}}_{t+T} = \boldsymbol{\Phi}^T\tilde{\mathbf{k}}_t, \qquad \boldsymbol{i}_t = \boldsymbol{\Delta}\mathbb{k} + (\boldsymbol{\Phi}-\mathbf{I}+\boldsymbol{\Delta})\tilde{\mathbf{k}}_t, \qquad \mathbf{k}_{t+1} - \mathbf{k}_t = (\boldsymbol{\Phi}-\mathbf{I})\tilde{\mathbf{k}}_t. \tag{156}$$

Expanding operating profits around $\mathbb{k}$ gives (for any conduct, since the profit function itself does not depend on $\mathbf{D}$ or $\boldsymbol{\Gamma}$)

$$\begin{aligned}\boldsymbol{\pi}\left(\mathbb{k}+\tilde{\mathbf{k}}_t\right) = \boldsymbol{\pi}(\mathbb{k}) &+ \{\mathrm{diag}[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{Z}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\mathbb{k}] - \mathbb{K}\mathbf{Z}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\}\tilde{\mathbf{k}}_t \\ &- \tilde{\mathbf{k}}_t \circ \left[\mathbf{Z}(\mathbf{I}+\boldsymbol{\Sigma})\mathbf{Z}\tilde{\mathbf{k}}_t\right],\end{aligned} \tag{157}$$

while investment costs are

$$\boldsymbol{\psi}(\mathbf{k}_{t+1};\mathbf{k}_t) = \mathbf{p}^k \circ \boldsymbol{\Delta}\mathbb{k} + \mathbf{P}^k(\boldsymbol{\Phi} - \mathbf{I} + \boldsymbol{\Delta})\tilde{\mathbf{k}}_t + \frac{\theta}{2}\left[(\boldsymbol{\Phi} - \mathbf{I})\tilde{\mathbf{k}}_t\right]^{\circ 2}. \tag{158}$$

The enterprise value vector is therefore

$$\mathbf{v}^e(\mathbf{k}_t) = \sum_{T=0}^{\infty} \frac{\boldsymbol{\pi}(\mathbf{k}_{t+T}) - \boldsymbol{\psi}(\mathbf{k}_{t+T+1};\mathbf{k}_{t+T})}{(1+r)^T}. \tag{159}$$

The constant terms reproduce $\mathbf{v}^e(\mathbb{k})$ in equation (54). Summing the discounted linear terms gives

$$\begin{aligned}\mathbb{Q} = \{&\operatorname{diag}[\mathbf{Z}(\mathbf{b} - \mathbf{c}) - \mathbf{Z}(\mathbf{I} + \boldsymbol{\Sigma})\mathbf{Z}\mathbb{k}] - \mathbb{K}\mathbf{Z}(\mathbf{I} + \boldsymbol{\Sigma})\mathbf{Z} \\ &- \mathbf{P}^k(\boldsymbol{\Phi} - \mathbf{I} + \boldsymbol{\Delta})\Big\}\left(\mathbf{I} - \frac{\boldsymbol{\Phi}}{1+r}\right)^{-1}.\end{aligned} \tag{160}$$

The steady-state equation (43) implies

$$\operatorname{diag}[\mathbf{Z}(\mathbf{b} - \mathbf{c}) - \mathbf{Z}(\mathbf{I} + \boldsymbol{\Sigma})\mathbf{Z}\mathbb{k}] = (r\mathbf{I} + \boldsymbol{\Delta})\mathbf{P}^k + \operatorname{diag}[\mathbf{Z}(\mathbf{D} + \boldsymbol{\Gamma} \circ \boldsymbol{\Sigma})\mathbf{Z}\mathbb{k}] \tag{161}$$

so that the matrix inside braces is

$$(1+r)\mathbf{P}^k - \mathbf{P}^k\boldsymbol{\Phi} - \mathbb{K}\mathbf{Z}\boldsymbol{\Sigma}\mathbf{Z} + \operatorname{diag}[\mathbf{Z}(\mathbf{D} + \boldsymbol{\Gamma} \circ \boldsymbol{\Sigma} - \mathbf{I})\mathbf{Z}\mathbb{k}]. \tag{162}$$

Since $(1+r)\mathbf{P}^k - \mathbf{P}^k\boldsymbol{\Phi} = (1+r)\mathbf{P}^k(\mathbf{I} - \boldsymbol{\Phi}/(1+r))$, this proves equation (59). In the single-product Cournot case, $\mathbf{D} + \boldsymbol{\Gamma} \circ \boldsymbol{\Sigma} = \mathbf{I}$ and the diagonal correction term vanishes, yielding the simpler expression reported in the main text.

It remains to sum the quadratic terms. Define

$$\mathbf{W}(\mathbf{k}_t) \equiv \sum_{T=0}^{\infty} \frac{\left(\boldsymbol{\Phi}^T\tilde{\mathbf{k}}_t\right)\left(\boldsymbol{\Phi}^T\tilde{\mathbf{k}}_t\right)'}{(1+r)^T}. \tag{163}$$

The series converges absolutely because $\hat{\varphi}_{\max}^2/(1+r) < 1$. Separating the term for $T = 0$ shows that it is the unique solution of the discrete Lyapunov–Stein equation

$$\mathbf{W}(\mathbf{k}_t) = \tilde{\mathbf{k}}_t\tilde{\mathbf{k}}_t' + \frac{1}{1+r}\boldsymbol{\Phi}\mathbf{W}(\mathbf{k}_t)\boldsymbol{\Phi}. \tag{164}$$

Recall $\mathbf{E}$ captures the eigenvectors in the spectral decomposition of matrix $\boldsymbol{\Phi}$. Premultiplying by $\mathbf{E}^*$ and postmultiplying by $\mathbf{E}$ gives, element by element,

$$\left[\mathbf{E}^*\mathbf{W}(\mathbf{k}_t)\mathbf{E}\right]_{ij} = \frac{\left(\mathbf{E}^*\tilde{\mathbf{k}}_t\right)_i\left(\mathbf{E}^*\tilde{\mathbf{k}}_t\right)_j}{1 - \hat{\varphi}_i\hat{\varphi}_j/(1+r)}. \tag{165}$$

Reassembling the eigendecomposition and substituting $\tilde{\mathbf{k}}_t = \mathbf{k}_t - \mathbb{k}$ yields equation (60). For any vector $\mathbf{y}$ and symmetric matrix $\mathbf{S}$, $\mathbf{y} \circ (\mathbf{S}\mathbf{y}) = \operatorname{diag}\left(\mathbf{y}\mathbf{y}'\mathbf{S}\right)$. Equation (163) consequently implies

$$\sum_{T=0}^{\infty} \frac{\left(\boldsymbol{\Phi}^T\tilde{\mathbf{k}}_t\right) \circ \left[\mathbf{Z}(\mathbf{I} + \boldsymbol{\Sigma})\mathbf{Z}\boldsymbol{\Phi}^T\tilde{\mathbf{k}}_t\right]}{(1+r)^T} = \operatorname{diag}[\mathbf{W}(\mathbf{k}_t)\mathbf{Z}(\mathbf{I} + \boldsymbol{\Sigma})\mathbf{Z}] \tag{166}$$

and

$$\sum_{T=0}^{\infty} \frac{\left[(\mathbf{\Phi}-\mathbf{I})\mathbf{\Phi}^T\tilde{\mathbf{k}}_t\right]^{\circ 2}}{(1+r)^T} = \operatorname{diag}[(\mathbf{\Phi}-\mathbf{I})\mathbf{W}(\mathbf{k}_t)(\mathbf{\Phi}-\mathbf{I})]. \tag{167}$$

Substitution proves equation (58). Stability of $\mathbf{\Phi}$ makes the capital and cash-flow paths bounded from every finite initial state. Discounting therefore makes the present-value sum converge. Iterating the policy-evaluation recursion shows that any two solutions satisfying the transversality condition have a difference converging to zero; hence the solution is unique. □

Proof of Proposition 7. Rearrange equation (20) as follows

$$\mathbf{k}_{t+1} - \mathbf{k}_t \;=\; (\mathbf{\Phi}-\mathbf{I})(\mathbf{k}_t - \mathbb{k}) \tag{168}$$

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

$$\mathbf{q}(\mathbf{k}_t) \;=\; \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_t + (\mathbf{I}-\mathbf{\Delta})\mathbf{p}^k + \theta(\mathbf{k}_{t+1}-\mathbf{k}_t) \tag{169}$$

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

$$\mathbf{q}(\mathbf{k}_t) \;=\; \mathbf{Z}(\mathbf{b}-\mathbf{c}) - \mathbf{\Omega}\mathbf{k}_t - \mathbf{\Delta}\mathbf{p}^k + \mathbf{p}^k + \theta(\mathbf{\Phi}-\mathbf{I})(\mathbf{k}_t - \mathbb{k}) \tag{170}$$

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

$$\mathbf{q}(\mathbf{k}_t) \;=\; \left[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^k\right] - \mathbf{\Omega}\mathbf{k}_t + (1+r)\mathbf{p}^k + \theta(\mathbf{\Phi}-\mathbf{I})(\mathbf{k}_t - \mathbb{k}) \tag{171}$$

rearrange equation (43) as follows:

$$\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\mathbf{I}+\mathbf{\Delta})\mathbf{p}^k \;=\; \mathbf{\Omega}\,\mathbb{k} \tag{172}$$

and substitute inside equation (171) to obtain:

$$\mathbf{q}(\mathbf{k}_t) \;=\; \mathbf{\Omega}(\mathbb{k} - \mathbf{k}_t) + (1+r)\mathbf{p}^k + \theta(\mathbf{I}-\mathbf{\Phi})(\mathbb{k} - \mathbf{k}_t) \tag{173}$$

finally, collect terms to obtain the following version of equation (61):

$$\mathbf{q}(\mathbf{k}_t) \;=\; (1+r)\mathbf{p}^k + [\theta(\mathbf{I}-\mathbf{\Phi}) + \mathbf{\Omega}](\mathbb{k} - \mathbf{k}_t) \tag{174}$$

□

Proof of Proposition 8. In the PME, the manager solves the conditional Bellman equation (30), taking the entire sequence of competitors' capital stocks $\mathbf{k}_{-n}^{(t)}$ as given. Then, the marginal value of capital equals:

$$\frac{\partial v_n\left(k_{nt};\mathbf{k}_{-n}^{(t)}\right)}{\partial k_{nt}} = (1+r)\,p_n^k - \lambda_{nn}(k_{nt}-k_{n\infty}) - \sum_{m\neq n}\lambda_{nm}(k_{mt}-k_{m\infty}) \tag{175}$$

Integrating this expression with respect to k_{nt} while holding $\mathbf{k}_{-n}^{(t)}$ fixed, results in the expression:

$$v_n\left(k_{nt};\mathbf{k}_{-n}^{(t)}\right) = v_n(\mathbb{k}) + (k_{nt} - k_{n\infty})\left[(1+r)\,p_n^k - \sum_{m=1}^{N} \frac{\lambda_{nm}}{2^{\mathbb{I}\{m=n\}}}(k_{mt} - k_{m\infty})\right] + \kappa_n\left(\mathbf{k}_{-n}^{(t)}\right) \tag{176}$$

where by construction the term κ_n does not depend on k_{nt}. Stacking the preceding expression across managers gives equation (64). Finally, equation (30) is the recursive representation of the manager's objective (17), conditional on the correctly forecasted sequence $\mathbf{k}_{-n}^{(t)}$. Evaluated at the PME policies, it therefore generates the same discounted sequence of cash flows as the equilibrium enterprise values of the internalized business units. Hence, using compact notation:

$$\mathbf{v}(\mathbf{k}_t) = \mathbf{\Gamma}\mathbf{v}^e(\mathbf{k}_t) \tag{177}$$

Equation (65) follows immediately from equations (64) and (177) which completes the proof. □

Proof of Proposition 9. Rewrite the first-order condition (34) as

$$\mathbf{q}(\mathbf{k}_t) \;=\; (1+r)\Big[\mathbf{p}^k + \theta(\mathbf{k}_t - \mathbf{k}_{t-1})\Big] \tag{178}$$

next, use the fact that

$$\mathbf{k}_t - \mathbf{k}_{t-1} \;=\; \mathbf{\Phi}^{-1}(\mathbf{k}_{t+1} - \mathbf{k}_t) \tag{179}$$

substitute inside (178) and divide by $(1+r)$ to obtain (69). □

Proof of Proposition 10. Consider the derivation of steady-state markups in Equation (74). Multiply both sides of the steady-state capital equation by $\mathbf{Z}$

$$\mathbf{y}(\mathbb{k}) \;=\; \mathbf{Z}\mathbf{\Omega}^{-1}\Big[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\,\mathbf{I} + \mathbf{\Delta})\mathbf{p}^k\Big] \tag{180}$$

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

$$\mathbf{b} - (\mathbf{I}+\mathbf{\Sigma})\mathbf{y}(\mathbb{k}) \;=\; \mathbf{b} - (\mathbf{I}+\mathbf{\Sigma})\mathbf{Z}\mathbf{\Omega}^{-1}\Big[\mathbf{Z}(\mathbf{b}-\mathbf{c}) - (r\,\mathbf{I} + \mathbf{\Delta})\mathbf{p}^k\Big] \tag{181}$$

finally we multiply both sides by $\mathbf{C}^{-1}$ and rearrange to obtain Equation (74):

$$\boldsymbol{\mu}(\mathbb{k}) \;=\; \mathbf{C}^{-1}\mathbf{b} + \mathbf{C}^{-1}(\mathbf{I}+\mathbf{\Sigma})\mathbf{Z}\mathbf{\Omega}^{-1}\Big[(r\,\mathbf{I} + \mathbf{\Delta})\mathbf{p}^k - \mathbf{Z}(\mathbf{b}-\mathbf{c})\Big] \tag{182}$$

The solution outside steady state in (73) follows from combining the definition of markup with the law of motion of investment in equilibrium, and rearranging terms. □

Proof of Proposition 11. Consider the equilibrium investment policy functions:

$$\boldsymbol{i}(\mathbf{k}_t) = \boldsymbol{\Delta}\mathbf{k}_t + (\boldsymbol{\Phi} - \mathbf{I})(\mathbf{k}_t - \mathbb{k}) \tag{183}$$

Considering the dependence of matrix $\boldsymbol{\Phi}$ and vector $\mathbb{k}$ on the discount rate, and holding the current state $\mathbf{k}_t$ fixed, we have that:

$$\frac{\mathrm{d}\,\boldsymbol{i}(\mathbf{k}_t)}{\mathrm{d}\,r} = (\mathbf{I} - \boldsymbol{\Phi})\frac{\mathrm{d}\,\mathbb{k}}{\mathrm{d}\,r} - \frac{\mathrm{d}\,\boldsymbol{\Phi}}{\mathrm{d}\,r}(\mathbb{k} - \mathbf{k}_t) \tag{184}$$

In the PME, the derivative of vector $\mathbb{k}$ with respect to the discount rate is:

$$\frac{\mathrm{d}\,\mathbb{k}}{\mathrm{d}\,r} = -\boldsymbol{\Omega}^{-1}\mathbf{p}^k \tag{185}$$

whereas the derivative of matrix $\boldsymbol{\Phi}$ with respect to the discount rate equals:

$$\frac{\mathrm{d}\,\boldsymbol{\Phi}}{\mathrm{d}\,r} = \mathbf{E}\frac{\mathrm{d}\,\mathrm{diag}\left(\hat{\boldsymbol{\varphi}}\right)}{\mathrm{d}r}\mathbf{E}^* = \frac{1}{2}\mathbf{I} - \left[\frac{1}{4}\left(\frac{1}{\theta}\boldsymbol{\Omega} + (2+r)\mathbf{I}\right) - \frac{1}{2}\mathbf{I}\right]\left[\frac{1}{4}\left(\frac{1}{\theta}\boldsymbol{\Omega} + (2+r)\mathbf{I}\right)^2 - (1+r)\mathbf{I}\right]^{-\frac{1}{2}} \tag{186}$$

Because the common orthogonal eigenvectors do not depend on r, the first equivalence follows; differentiating the eigenvalues in equation (51) yields the second. Substitution into the expression above gives equation (76), and setting $\mathbf{k}_t = \mathbb{k}$ while holding the state fixed gives equation (77). □

Proof of Proposition 12. Consider again equation (183). At the firm level, the derivative of the investment expenditure of business unit n with respect to its own steady-state capital stock $k_{n\infty}$ equals:

$$\frac{\mathrm{d}\left[p_n^k\, i_n(\mathbf{k}_t)\right]}{\mathrm{d}\,k_{n\infty}} = p_n^k\left(1 - \varphi_{nn}\right) \tag{187}$$

which corresponds to the numerator of the $\mathrm{MPI}^{\mathrm{G}}$ of business unit n at time t. Note that matrix $\boldsymbol{\Phi}$ and the price of capital p_n^k do not depend on the primitives $(\mathbf{b}, \mathbf{c})$, and are therefore unaffected by the perturbation, which moves $\mathbb{k}$ only. To solve the expression in the denominator, consider equation (56), by which the steady-state enterprise value of business unit n equals $v_n^e(\mathbb{k}) = (1+r)\,p_n^k k_{n\infty} + \rho_n$, with $\rho_n = \frac{1+r}{r} z_n^2 k_{n\infty}^2$ by the definition of ρ_n. Note that the steady-state value of each business unit depends only on its own steady-state capital stock; hence, the perturbation leaves competitors' steady-state values unaffected if and only if it leaves their steady-state capital stocks unaffected. Differentiating with respect to $k_{n\infty}$:

$$\frac{\mathrm{d}\,v_n^e(\mathbb{k})}{\mathrm{d}\,k_{n\infty}} = (1+r)\,p_n^k + \frac{2\rho_n}{k_{n\infty}} = (1+r)\,p_n^k + 2\left[\bar{q}_n(\mathbb{k}) - q_n(\mathbb{k})\right] \tag{188}$$

where the last equality follows from equation (71) evaluated at the steady state, by which $\bar{q}_n(\mathbb{k}) - q_n(\mathbb{k}) = \rho_n / k_{n\infty}$. This gives the result. □

B Additional Discussion of Proximate Markov Equilibrium

In this appendix we formally define *unrecoverability* — the PME condition on the transition matrix $\mathbf{\Phi}$ that makes the equilibrium not self-disproving. That is, the condition that makes managers unable to infer their rivals' investment strategies when the number of players is sufficiently large. We also discuss the intuition behind this criterion in more detail.

Auxiliary Definitions and Identification of the Policy Matrix. We begin by constructing $\tilde{\mathbf{K}}_{t,T}$, defined as the matrix of equilibrium capital stocks, net of the corresponding PME steady-state values, from time t up to a given horizon T:

$$\tilde{\mathbf{K}}_{t,T} \stackrel{\text{def}}{=} [(\mathbf{k}_t - \mathbb{k}), (\mathbf{k}_{t+1} - \mathbb{k}), \ldots, (\mathbf{k}_{t+T} - \mathbb{k})]. \tag{189}$$

Also, fix a precision threshold $\varepsilon \in (0,1)$. For any nonzero matrix $\mathbf{A}$, let $\text{rank}_\varepsilon(\mathbf{A})$ denote the number of singular values of $\mathbf{A}/\|\mathbf{A}\|_2$ that exceed ε, and set $\text{rank}_\varepsilon(\mathbf{0}) = 0$.

In order to evaluate whether an Open-Loop equilibrium (a candidate for PME) is self-disproving, we need to ask the question of when and how knowledge of $\tilde{\mathbf{K}}_{t,T}$ (for some horizon T) allows an econometrician to recover the policy transition matrix $\mathbf{\Phi}$.

If $\tilde{\mathbf{K}}_{t,T}\tilde{\mathbf{K}}'_{t,T}$ is invertible, the policy matrix $\mathbf{\Phi}$ can be recovered using the linear regression formula

$$\mathbf{\Phi} = \mathbb{F}(\tilde{\mathbf{K}}_{t,T}) \cdot \tilde{\mathbf{K}}'_{t,T}(\tilde{\mathbf{K}}_{t,T}\tilde{\mathbf{K}}'_{t,T})^{-1} \tag{190}$$

where $\mathbb{F}$ denotes the $t \to t+1$ forward operator, applied through $t+T+1$. $\tilde{\mathbf{K}}_{t,T}\tilde{\mathbf{K}}'_{t,T}$ is invertible if $\tilde{\mathbf{K}}_{t,T}$ is full rank (it has N rows, therefore it needs to have rank N).

Convergence Horizon and Unrecoverability. Next, we define the concept of convergence horizon. The convergence horizon T_ε is the maximum numerical rank that can be attained by $\tilde{\mathbf{K}}_{t,T}$. It captures the maximum number of numerically independent directions revealed by the equilibrium capital path before the economy converges to steady state. When convergence makes additional observations numerically redundant, they cannot increase this rank. Hence, subsequent observations beyond T_ε do not help the managers recover the transition matrix $\mathbf{\Phi}$ econometrically.

Definition 5. *For a stable linear Markov strategy profile* $(\mathbb{k}, \mathbf{\Phi}; \mathbf{k}_t)$ *we define* $T_\varepsilon(\mathbb{k}, \mathbf{\Phi}; \mathbf{k}_t)$*, the convergence horizon with precision* ε*, as follows*

$$T_\varepsilon(\mathbb{k}, \mathbf{\Phi}; \mathbf{k}_t) \stackrel{\text{def}}{=} \sup_{T \in \mathbb{N}_+} \text{rank}_\varepsilon\left(\tilde{\mathbf{K}}_{t,T}\right) \tag{191}$$

This sup exists because (as we show in the later Lemma) the convergence horizon obeys a finite upper bound. Technically, we show that dynamic stability places an upper bound on the number of numerically independent state directions that can be learned from the equilibrium transition path.

It follows that, if the number of players N is sufficiently large with respect to the convergence horizon, an econometrician cannot recover the policy transition matrix $\mathbf{\Phi}$ from the equilibrium investment plans $\mathbf{k}^{(t)}$. Having shown formally the convergence horizon has an upper bound, any sufficiently large number of players (or rank of matrix $\tilde{\mathbf{K}}_{t,T}$) yields the prediction. We are now ready to provide a formal definition of unrecoverability.

Definition 6. *Given a strategy profile* $(\mathbb{k}, \mathbf{\Phi}; \mathbf{k}_t)$*, we say that the transition matrix* $\mathbf{\Phi}$ *is* ε*-unrecoverable from equilibrium path* $\mathbf{k}^{(t)}$ *if:*

$$T_\varepsilon(\mathbb{k}, \mathbf{\Phi}; \mathbf{k}_t) < N. \tag{192}$$

Intuitively, when $T_\varepsilon < N$, agents cannot recover the opponents' policy functions at all, even if they have correctly forecasted the entire future sequence of capital stocks infinitely into the future.

We conclude by proving an upper bound to the convergence horizon.

Lemma. *Let* $\mathbf{\Phi}$ *be the symmetric PME transition matrix and let* $0 < \varepsilon < 1$*. Then:*

$$T_\varepsilon(\mathbb{k}, \mathbf{\Phi}; \mathbf{k}_t) \le \left\lceil \frac{\log \varepsilon}{\log \hat{\varphi}_{\max}} \right\rceil \tag{193}$$

where $\hat{\varphi}_{\max}$ *is the largest eigenvalue of matrix* $\mathbf{\Phi}$*.*

Proof. If $\tilde{\mathbf{k}}_t = \mathbf{0}$, then $\tilde{\mathbf{K}}_{t,T} = \mathbf{0}$ and the result follows from $\operatorname{rank}_\varepsilon(\mathbf{0}) = 0$. Otherwise, normalize $\tilde{\mathbf{K}}_{t,T}$ as in the definition above so that $\|\tilde{\mathbf{K}}_{t,T}\|_2 = 1$.

Because $\mathbf{\Phi}$ is symmetric and all of its eigenvalues lie in $(0,1)$, we have:

$$\|\mathbf{\Phi}^h\|_2 = \hat{\varphi}_{\max}^h, \qquad h \ge 1. \tag{194}$$

Fix any integer $h \ge 1$. If $T < h$, the matrix $\tilde{\mathbf{K}}_{t,T}$ has only $T+1 \le h$ columns, and therefore its numerical rank cannot exceed h. Consider then the alternative case in which $T \ge h$, and approximate $\tilde{\mathbf{K}}_{t,T}$ by retaining its first h columns and replacing all subsequent columns by zero. This approximation has rank at most h, while the approximation error is:

$$\left[\mathbf{0}, \ldots, \mathbf{0}, \tilde{\mathbf{k}}_{t+h}, \ldots, \tilde{\mathbf{k}}_{t+T}\right] = \mathbf{\Phi}^h \left[\mathbf{0}, \ldots, \mathbf{0}, \tilde{\mathbf{k}}_t, \ldots, \tilde{\mathbf{k}}_{t+T-h}\right]. \tag{195}$$

Since the spectral norm of the second matrix on the right hand side cannot exceed $\|\tilde{\mathbf{K}}_{t,T}\|_2 = 1$, we have that:

$$\left\|\tilde{\mathbf{K}}_{t,T} - [\tilde{\mathbf{k}}_t, \ldots, \tilde{\mathbf{k}}_{t+h-1}, \mathbf{0}, \ldots, \mathbf{0}]\right\|_2 \le \hat{\varphi}_{\max}^h. \tag{196}$$

By the Eckart-Young-Mirsky theorem, we know that

$$\sigma_{h+1}(\tilde{\mathbf{K}}_{t,T}) \le \hat{\varphi}_{\max}^h \tag{197}$$

where $\sigma_{h+1}(\tilde{\mathbf{K}}_{t,T})$ is the $(h+1)^{th}$ singular value of $\tilde{\mathbf{K}}_{t,T}$. Next, choose h

$$h = \left\lceil \frac{\log \varepsilon}{\log \hat{\varphi}_{\max}} \right\rceil \tag{198}$$

Since $0 < \hat{\varphi}_{\max} < 1$ due to dynamic stability, this choice implies $\hat{\varphi}_{\max}^h \le \varepsilon$. Hence $\operatorname{rank}_\varepsilon(\tilde{\mathbf{K}}_{t,T}) \le h$ for every horizon T. Taking the supremum yields

$$T_\varepsilon(\mathbb{k}, \mathbf{\Phi}; \mathbf{k}_t) \le \left\lceil \frac{\log \varepsilon}{\log \hat{\varphi}_{\max}} \right\rceil. \tag{199}$$

□

C 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 (79)).

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 (79). Third, we compute each firm's within-year percentile rank and estimate the Pareto scale and shape parameters by regressing the noisy estimate of $\log z_{nt}$ on the log of 1 minus the corresponding quantile, using only firms in the top 5% of each year's distribution. Both the intercept and the slope of this regression are specified as linear functions of the year, so that the Pareto scale parameter follows an exponential time trend, while the slope on the log survival rank — the reciprocal of the Pareto shape parameter — follows a linear one: the intercept of each trend identifies the level of the corresponding parameter, and its slope the parameter's drift over the sample period.

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 letting the regression coefficients evolve along smooth time trends rather than vary freely year by year, we accommodate secular changes in productivity dispersion over our 1995–2021 sample period. Within the estimation sample, the Pareto assumption provides an excellent fit: the R^2 of the regression is typically 98 to 99%, with roughly 200 tail observations entering per year.

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, choosing $\mathcal{Z}$ to satisfy the aggregate moment condition:

$$\sum_{n=1}^{N} v_n^e(\mathbf{k}_t; \mathcal{Z}\mathbf{z}) \;=\; \sum_{n}(1+r)v_{nt}^{\text{data}} \tag{200}$$

where $v_n^e(\mathbf{k}_t; \mathcal{Z}\mathbf{z})$ denotes the model-predicted enterprise value for firm n as a function of the rescaled productivity vector, and v_{nt}^{data} is the observed market value of assets. The factor $(1+r)$ accounts for the cum-dividend adjustment in the 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 $\mathbb{k}$. The steady-state capital distribution $\mathbb{k}$ is an equilibrium object rather than a primitive parameter, but we must ensure it satisfies the positivity constraint $\mathbb{k} > 0$. From the transition dynamics in Proposition 2, we can express the steady state as:

$$\mathbb{k} = (\mathbf{I} - \mathbf{\Phi})^{-1}(\mathbf{k}_{t+1} - \mathbf{\Phi}\mathbf{k}_t) \tag{201}$$

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 90th percentile) 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.

Recovery of demand and cost parameters. Given the calibrated values of $\mathbf{z}$, θ, and $\mathbb{k}$, we recover the remaining structural parameters. Output y_{nt} is simply obtained as $z_n k_{nt}$. The unit labor cost c_n is equal to the ratio of cost 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 (43):

$$\mathbf{b} = \mathbf{c} + \mathbf{Z}^{-1}\left[\mathbf{\Omega}\mathbb{k} + (r\mathbf{I} + \mathbf{\Delta})\mathbf{p}^{k}\right] \tag{202}$$

where $\mathbf{\Omega} = \mathbf{Z}(2\mathbf{I} + \mathbf{\Sigma})\mathbf{Z}$ is the competitive pressure 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 (it features more structural equation restrictions—equalities and inequalities—than degrees of freedom) 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.

D Sensitivity Analysis

This appendix evaluates the sensitivity of our main quantitative results to alternative calibration choices. We consider four variations.

First, in Appendix D.1, 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.

Second, in Appendix D.2 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.

Third, in Appendix D.3, 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.

Last, in Appendix D.4, we winsorize six key inputs—revenues, costs, firm value, and three capital-stock measures—at 1% in each tail within each year. We also reproduce the full set of figures from the main case in this case for completeness.

D.1 Alternate Discount Rate: r = 10y US Treasury + 5%

This appendix reproduces all figures from the main paper using, as an alternative value of the discount rate r, the 10-year US Treasury rate plus 5 percentage points.

FIGURE 4A: TRANSITION MATRIX: DISTRIBUTION OF EIGENVALUES

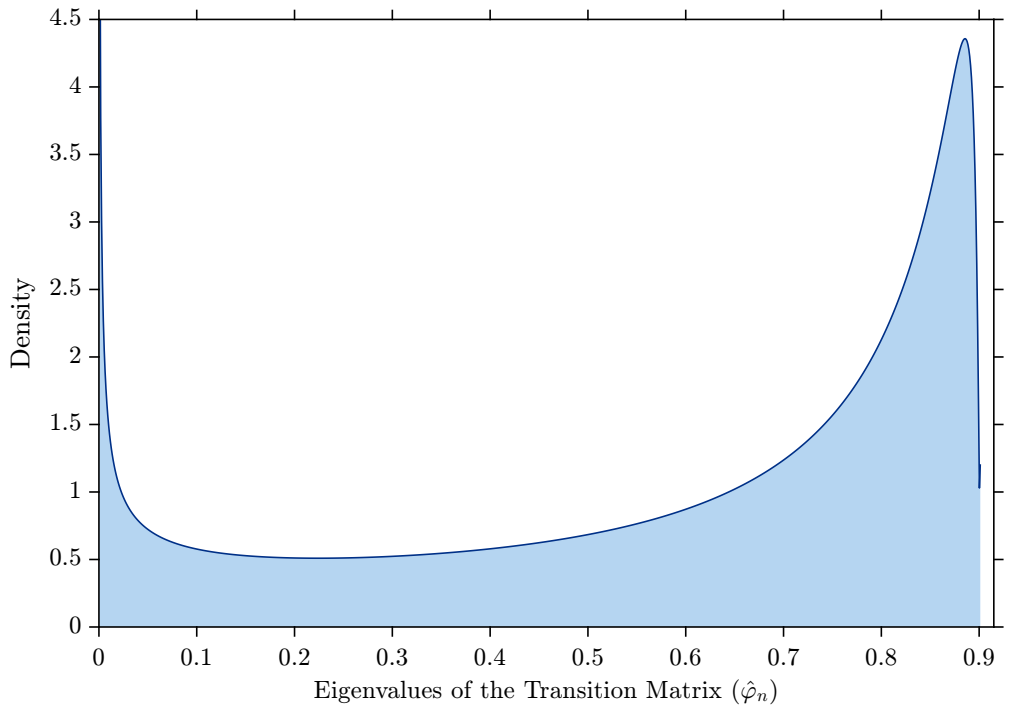


FIGURE 5A: MODEL FIT

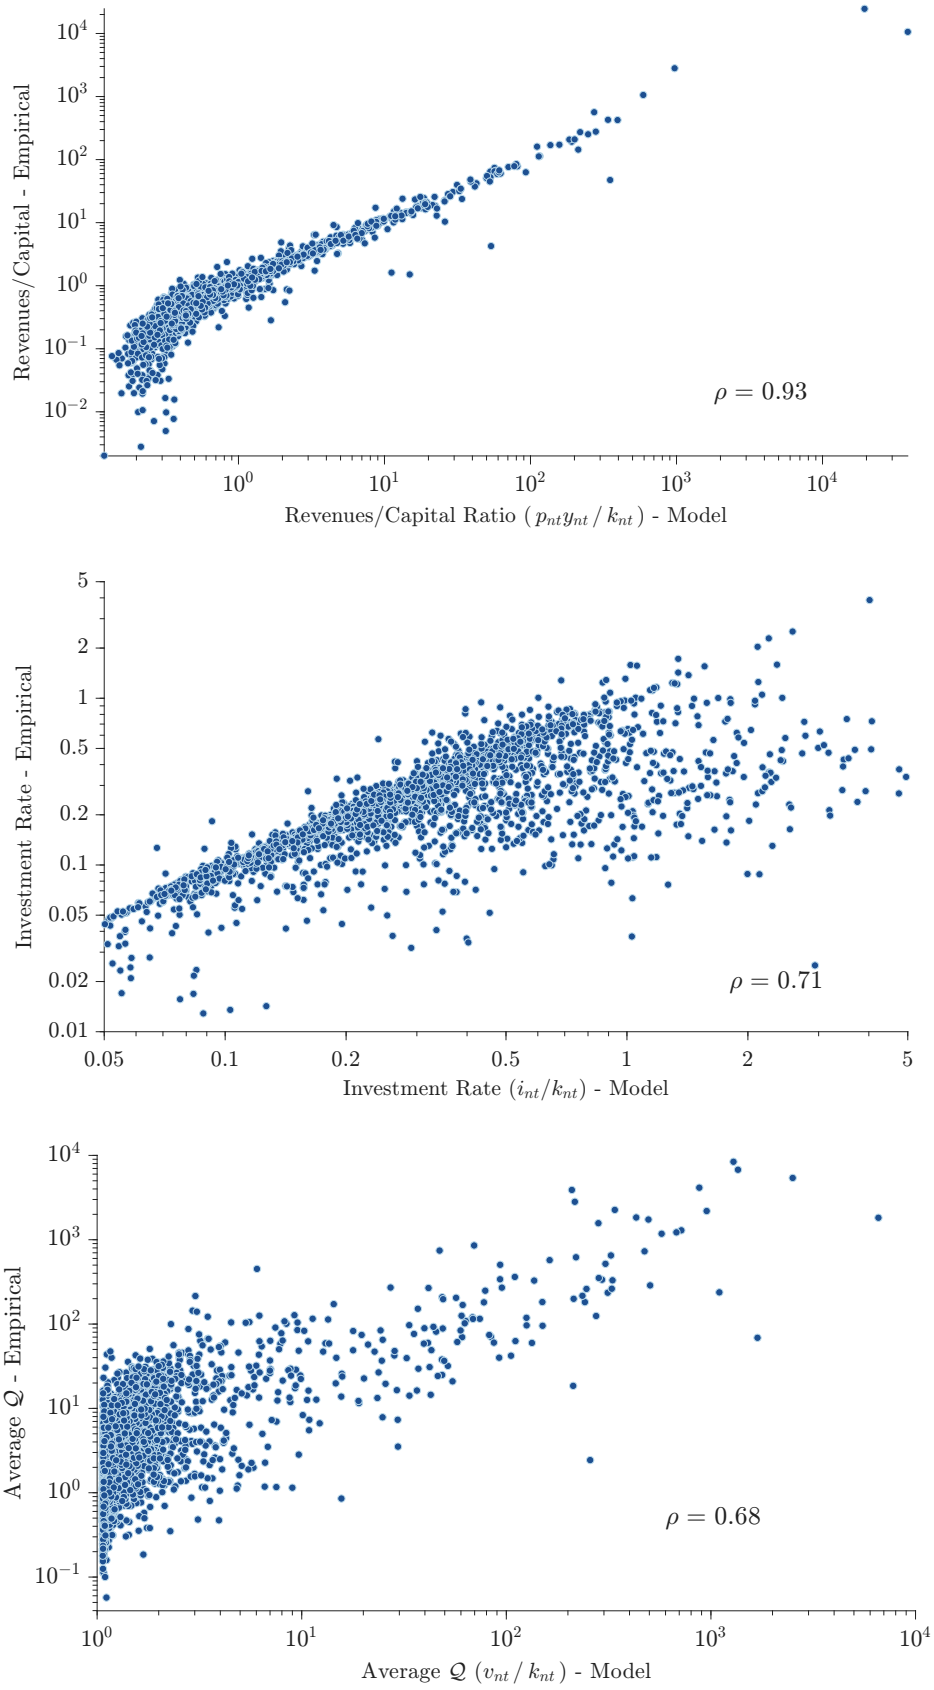


FIGURE 6A: TRANSITION DYNAMICS OF THE AGGREGATE CAPITAL STOCK

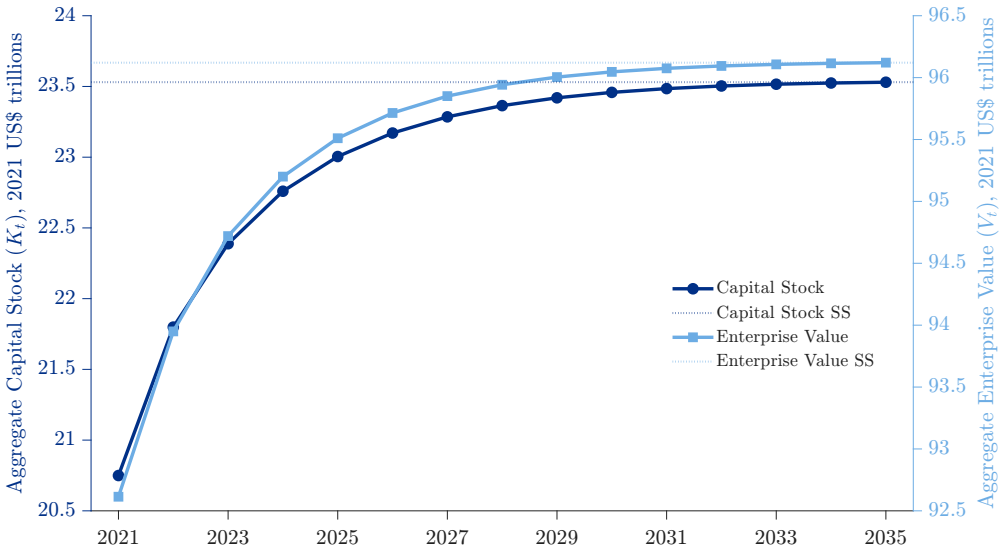


FIGURE 7A: FIRM-LEVEL TRANSITION DYNAMICS

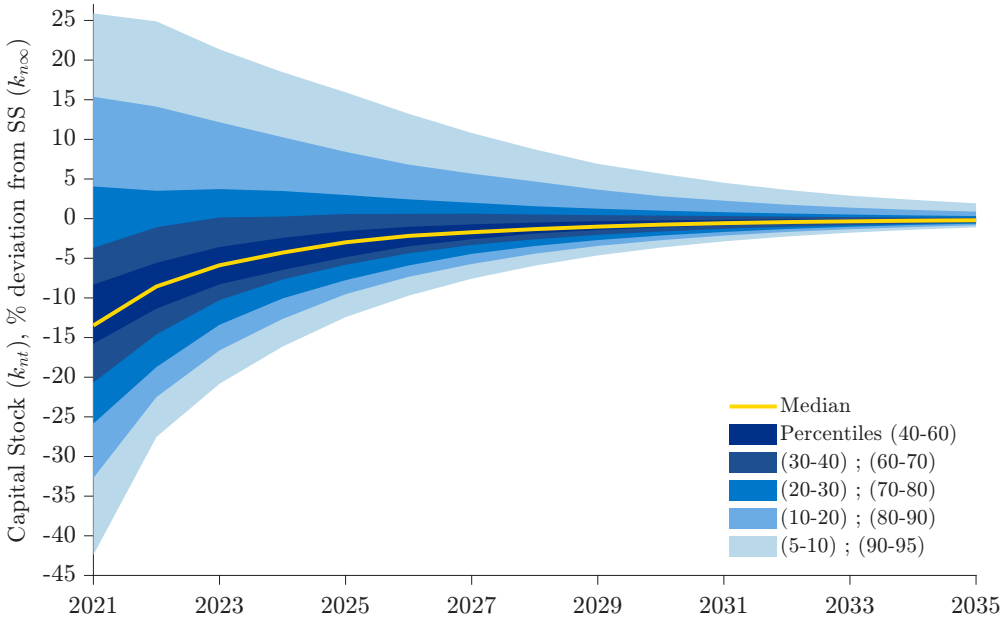


FIGURE 8A: AVERAGE Q AND MARGINAL Q

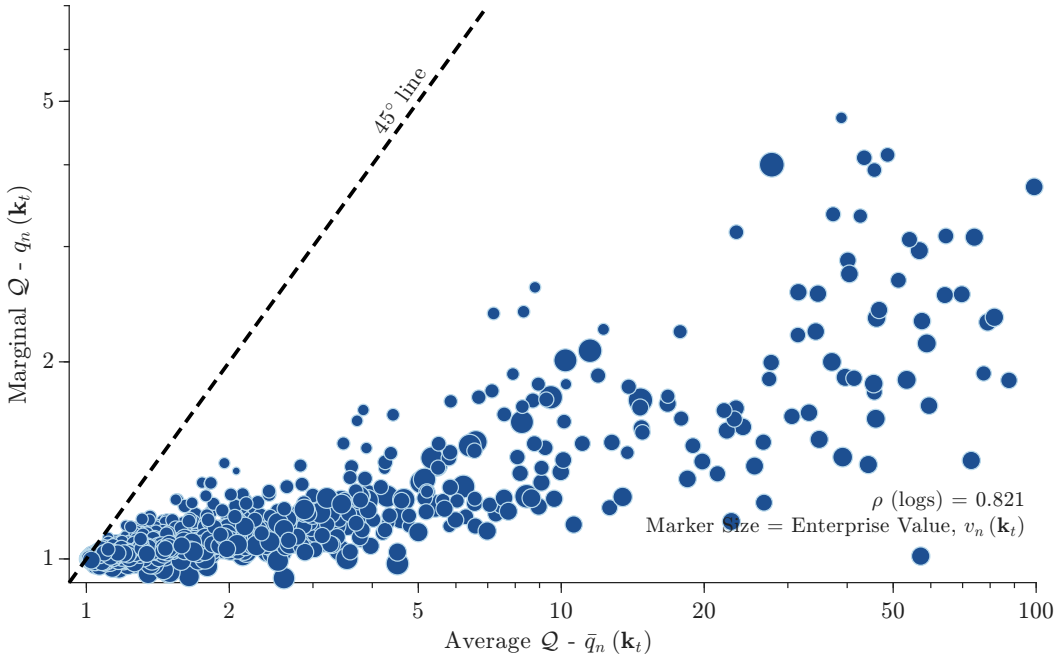


FIGURE 9A: DECOMPOSITION OF STEADY STATE AGGREGATE FIRM VALUE

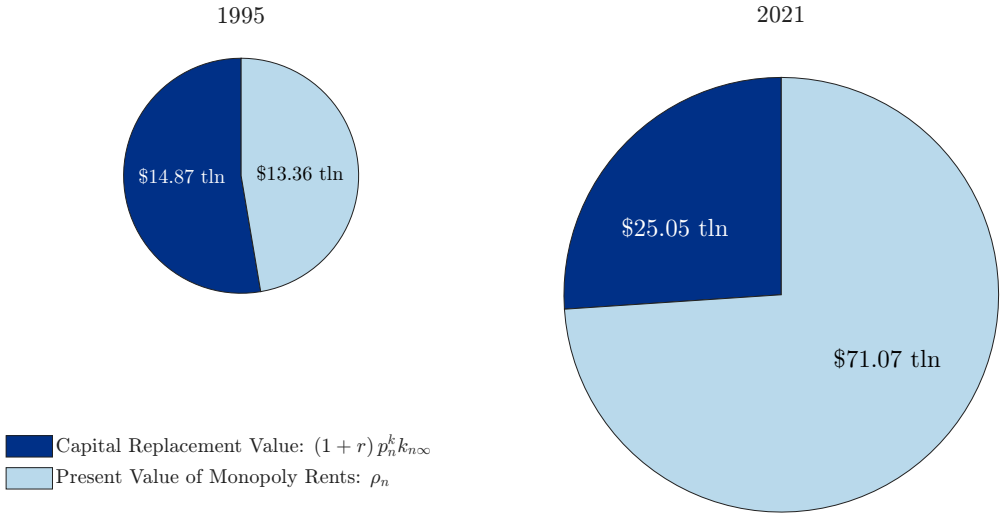


FIGURE 10A: MARGINAL PROPENSITY TO INVEST OUT OF GROWTH OPPORTUNITIES

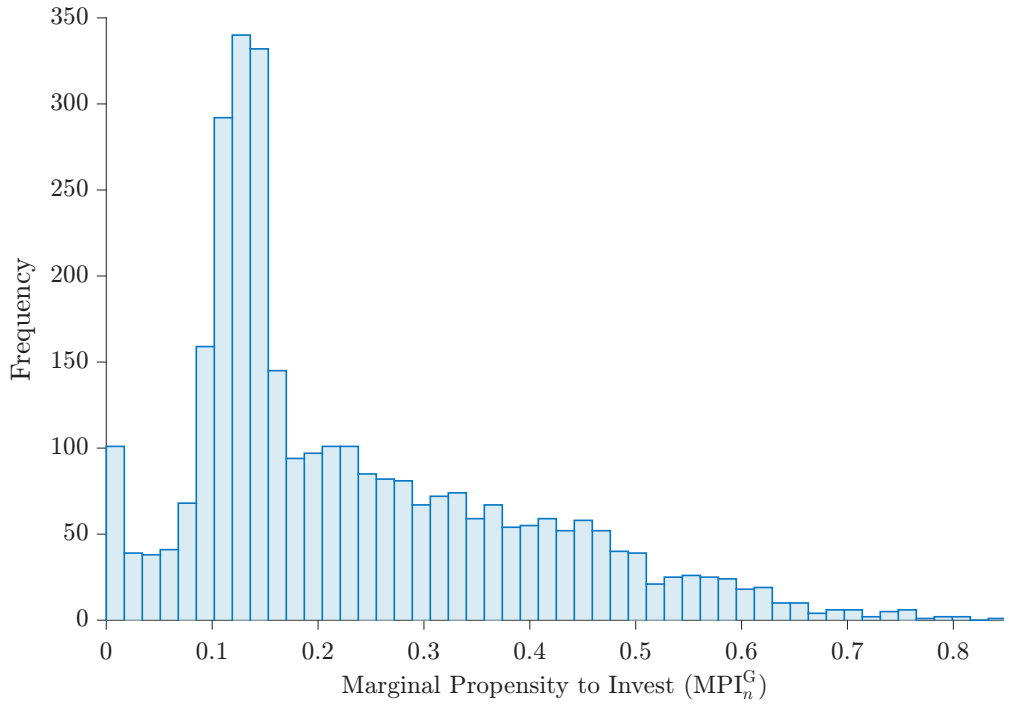


FIGURE 11A: ANTICOMPETITIVE COUNTERFACTUAL: AGGREGATE CAPITAL STOCK

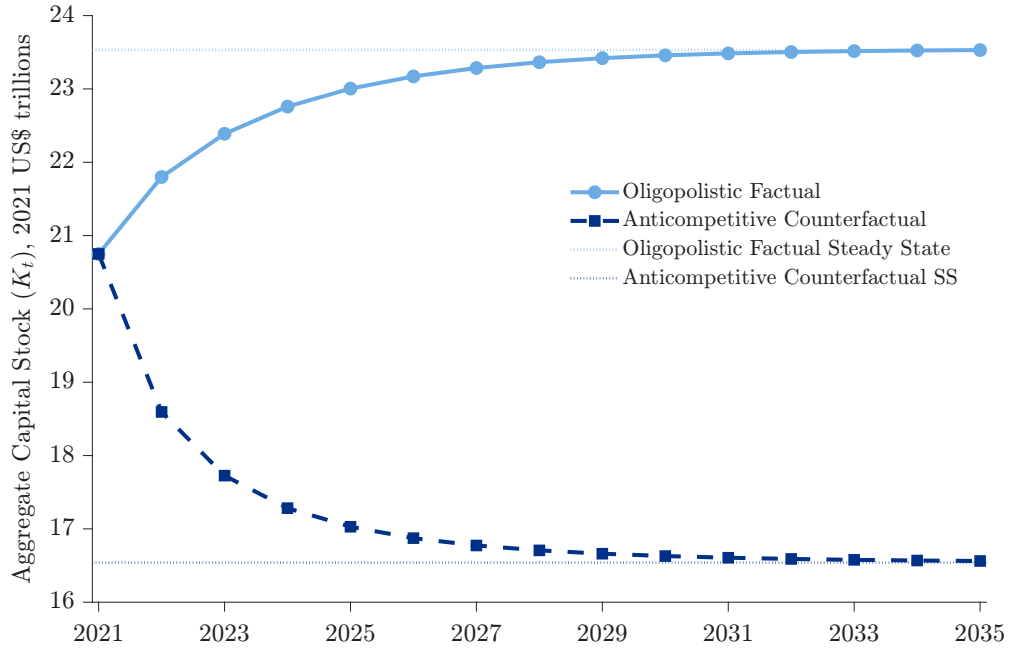


FIGURE 12A: ANTICOMPETITIVE COUNTERFACTUAL: CROSS-SECTION

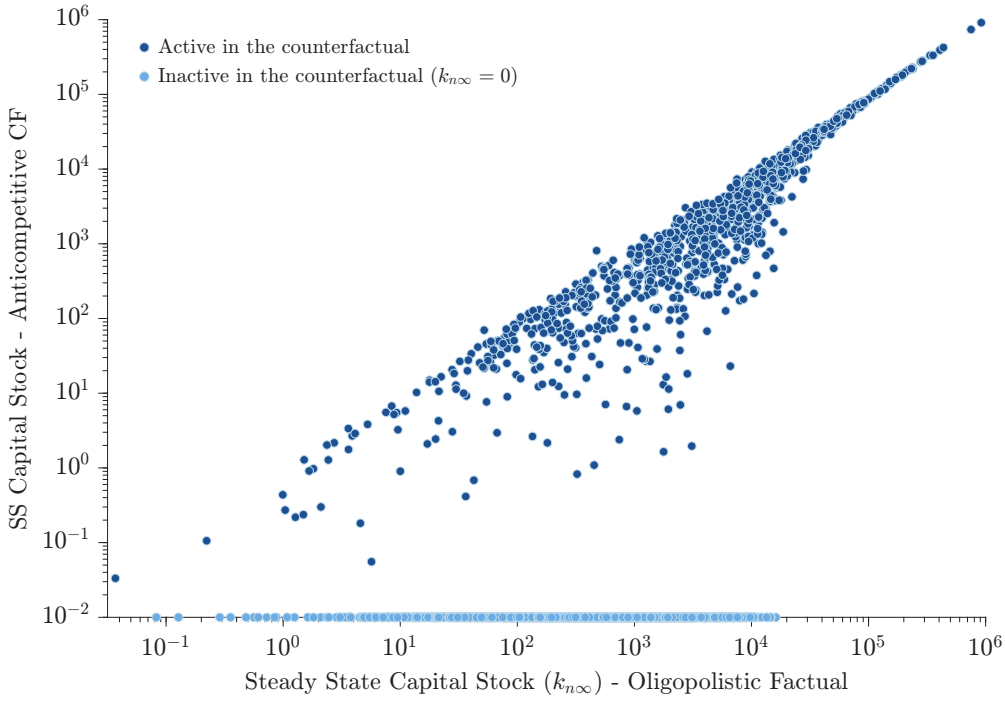


FIGURE 13A: DISTRIBUTION OF MARKUPS

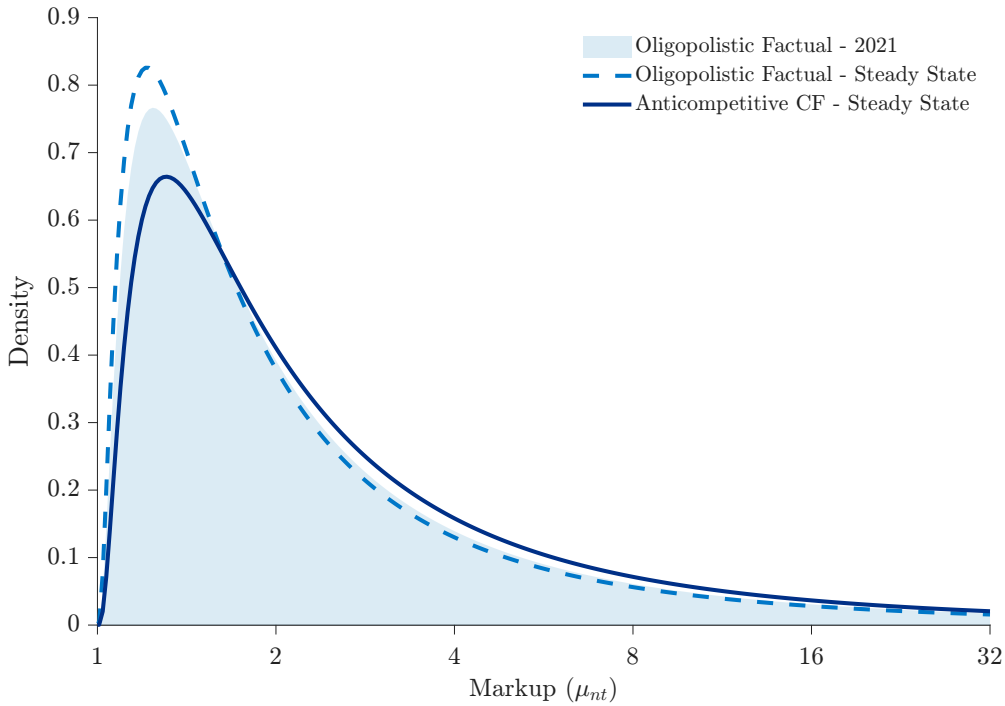


FIGURE 14A: RESPONSE OF INVESTMENT RATES TO DISCOUNT RATE SHOCK

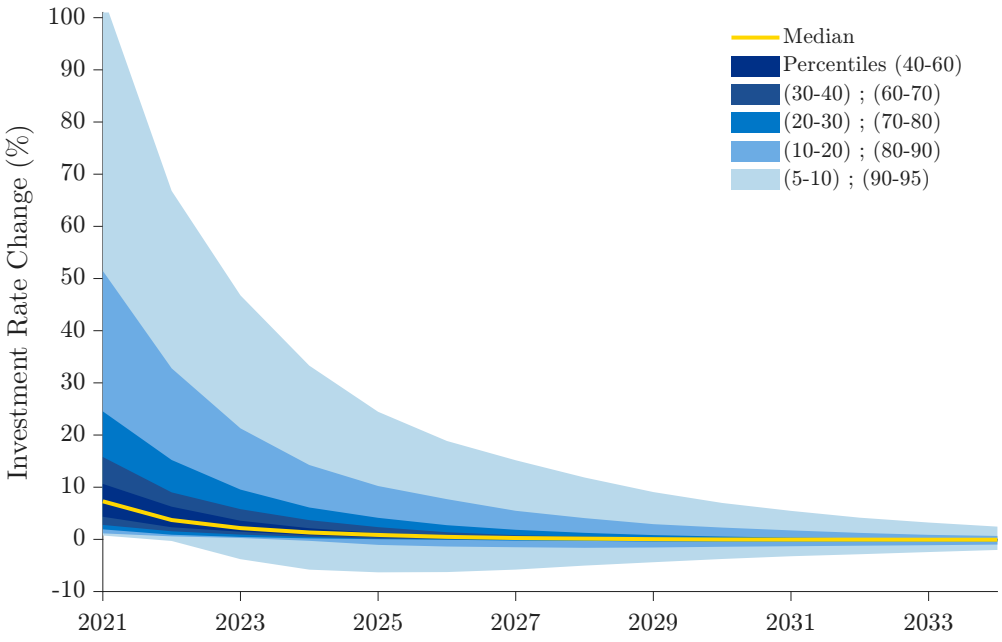


FIGURE 15A: EVOLUTION OF CONCENTRATION

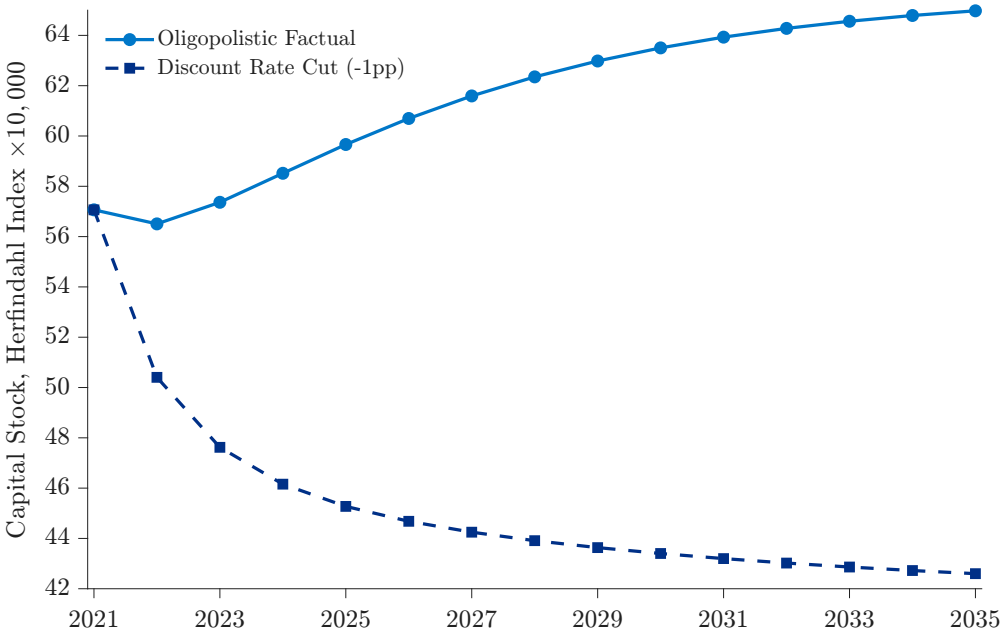


FIGURE 16A: RESPONSE OF MARKUPS TO DISCOUNT RATE SHOCK

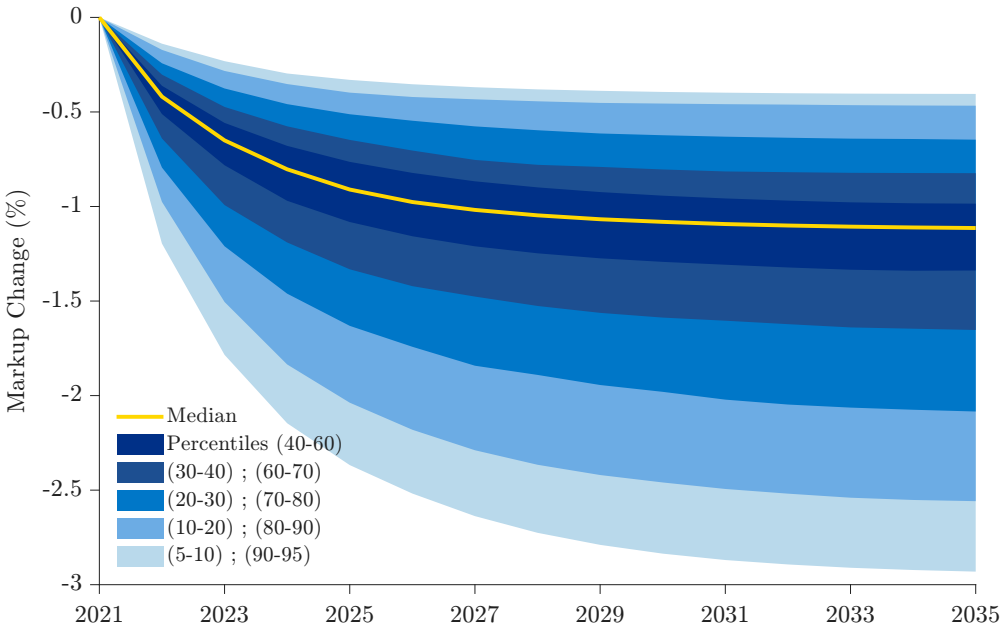


FIGURE 17A: MERGERS SIMULATION: INVESTMENT

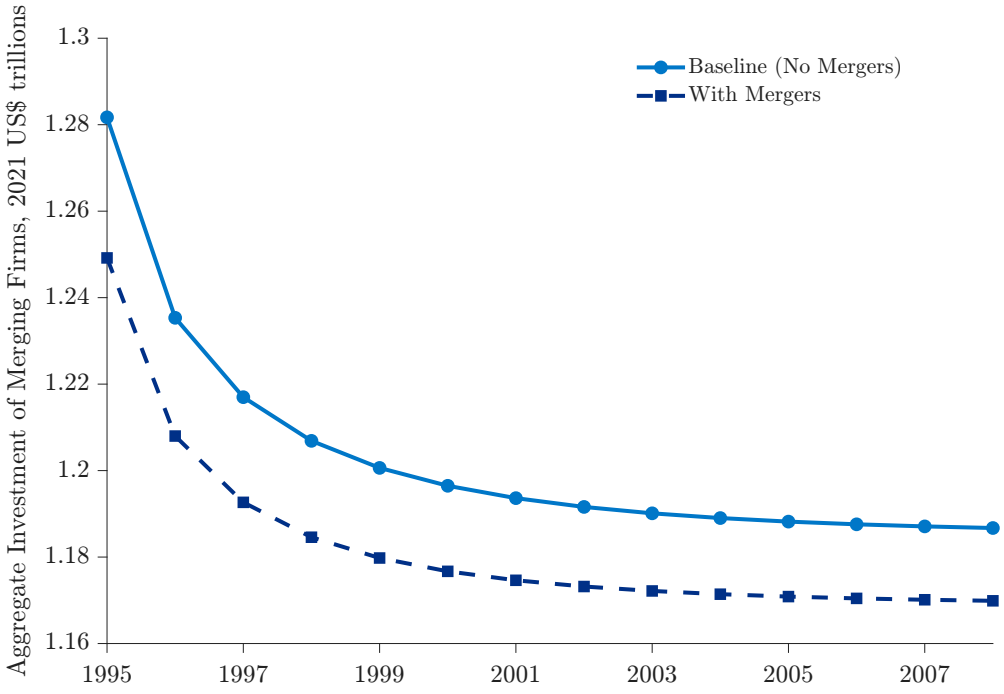
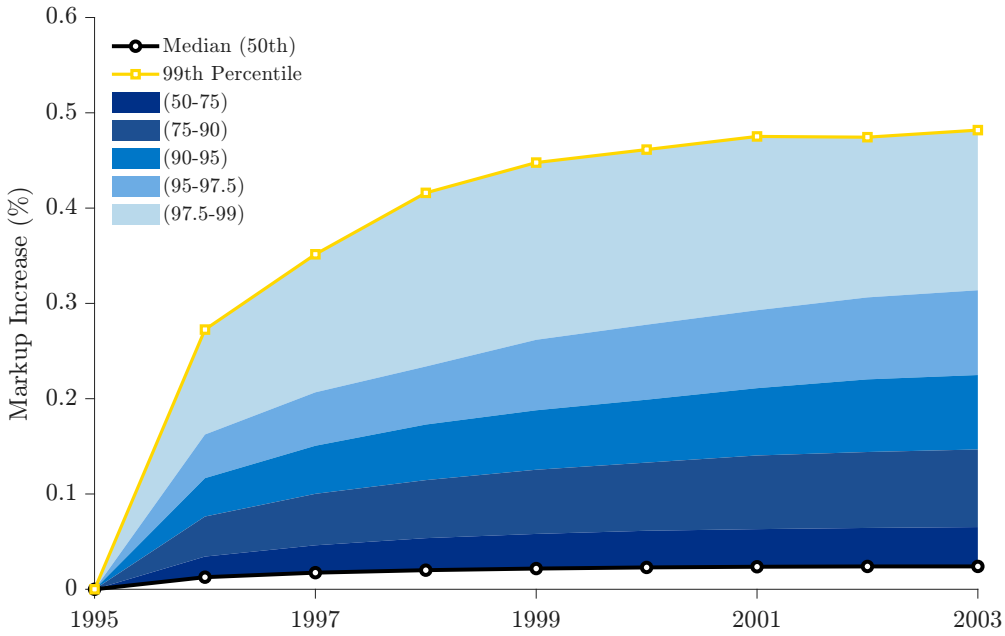


FIGURE 18A: MERGERS SIMULATION: MARKUPS



D.2 Alternate Discount Rate: r = Lee, So, and Wang (2021) ICC

This appendix reproduces all figures from the main paper using, as an alternative value of the discount rate r, analyst-based estimates of the implied cost of capital from Lee, So, and Wang (2021). As in the baseline calibration, r remains a common year-specific scalar: we take the cross-sectional median of their firm-level analyst-based estimates in each year, so the equilibrium is the one characterized in Proposition 2 rather than the heterogeneous-rate system of Section 5.2.

FIGURE 4B: TRANSITION MATRIX: DISTRIBUTION OF EIGENVALUES

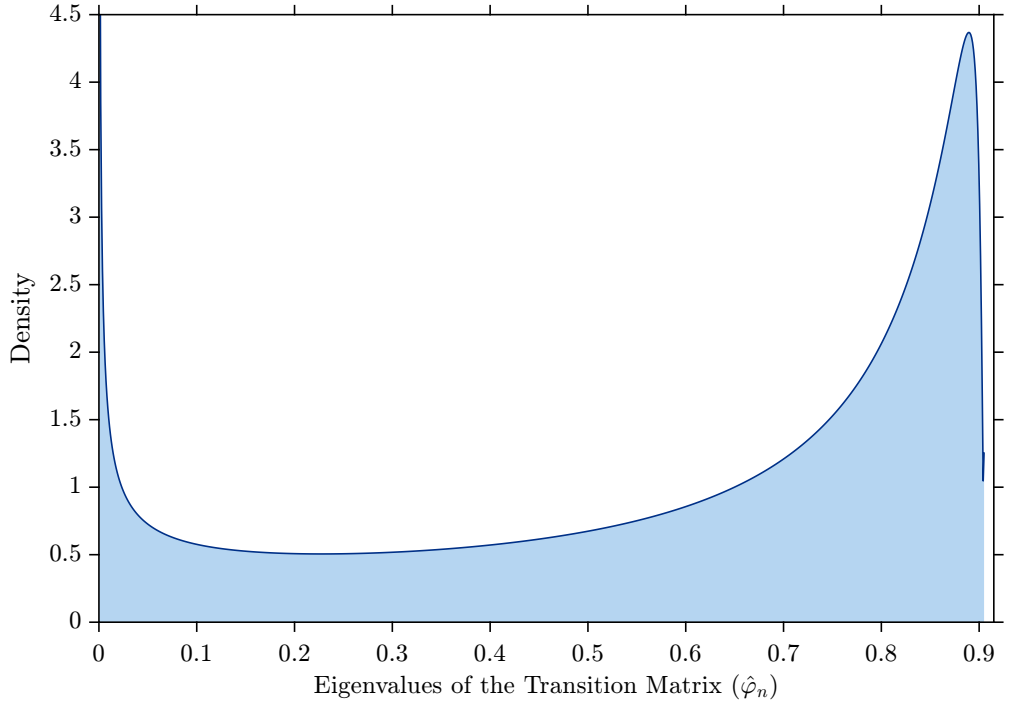


FIGURE 5B: MODEL FIT

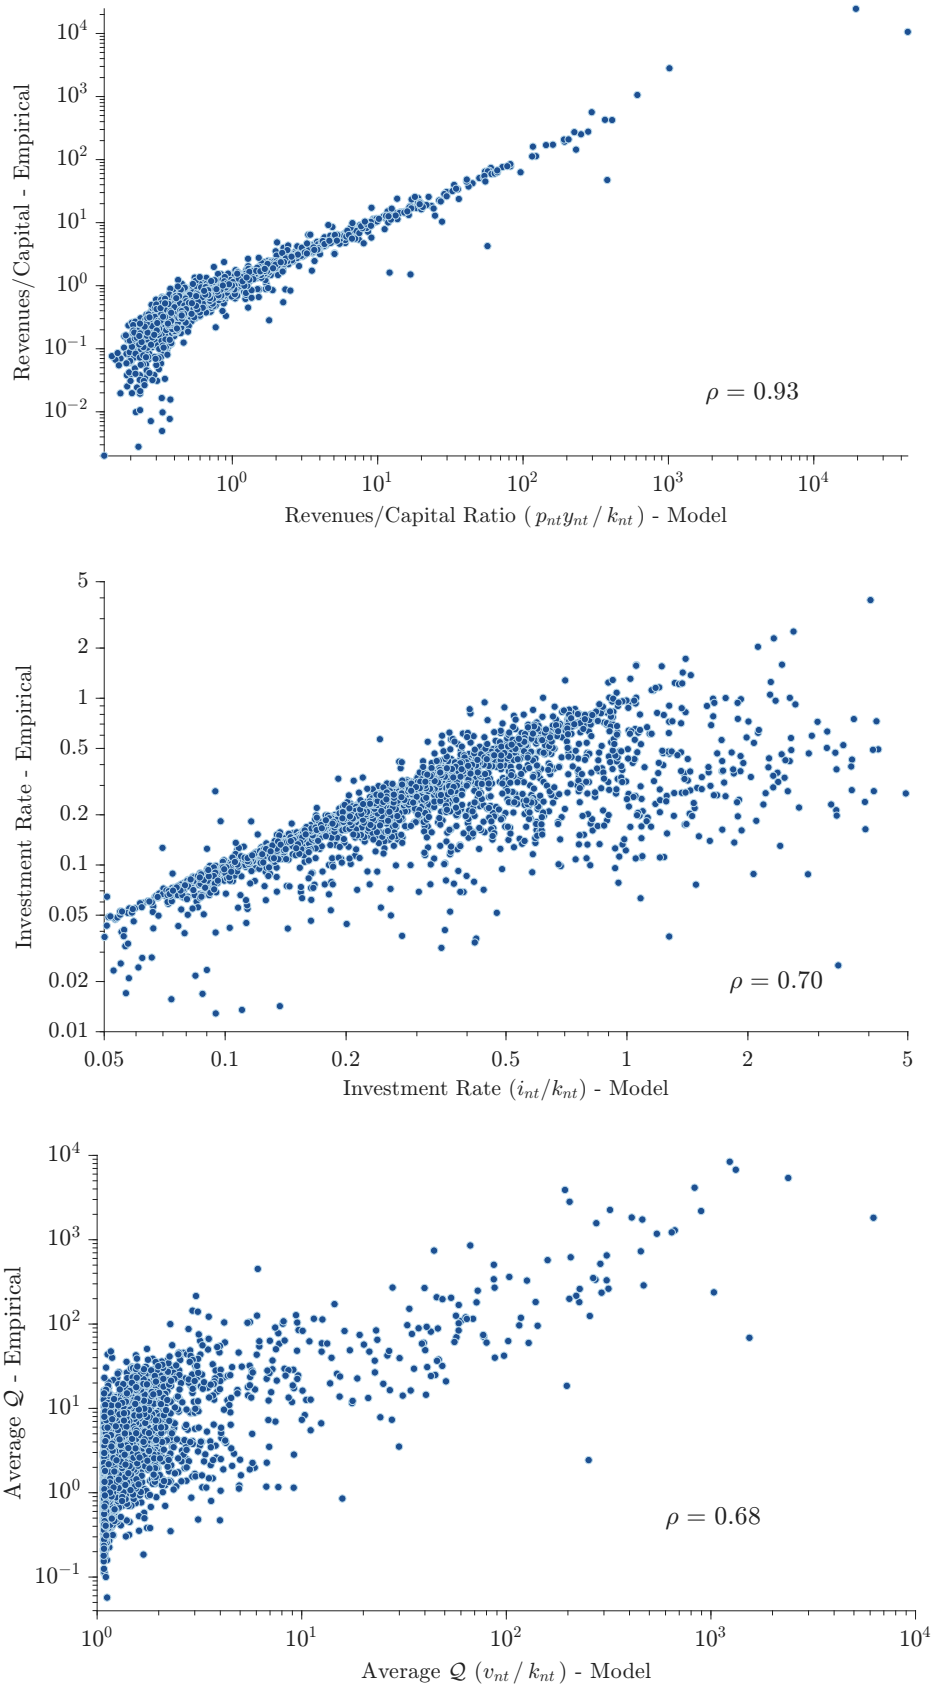


FIGURE 6B: TRANSITION DYNAMICS OF THE AGGREGATE CAPITAL STOCK

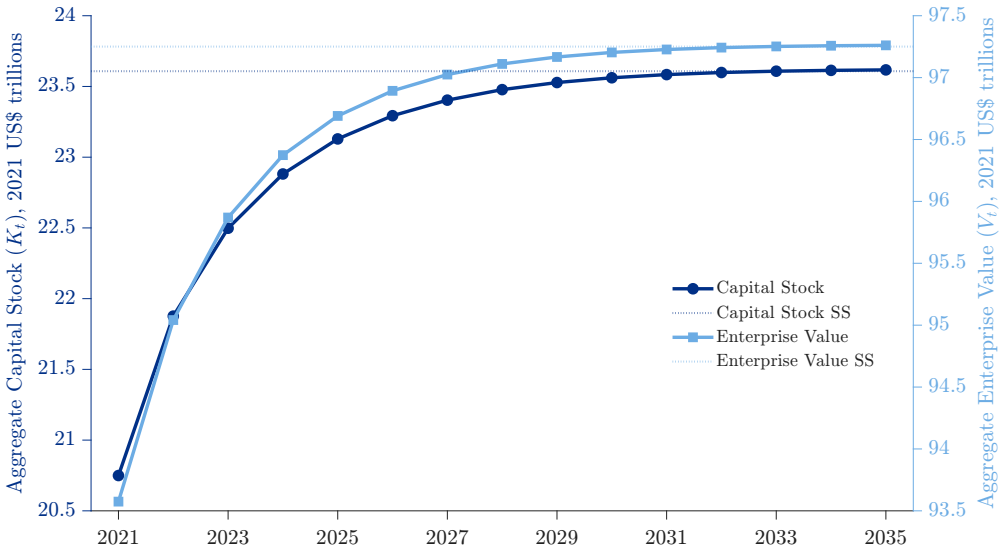


FIGURE 7B: FIRM-LEVEL TRANSITION DYNAMICS

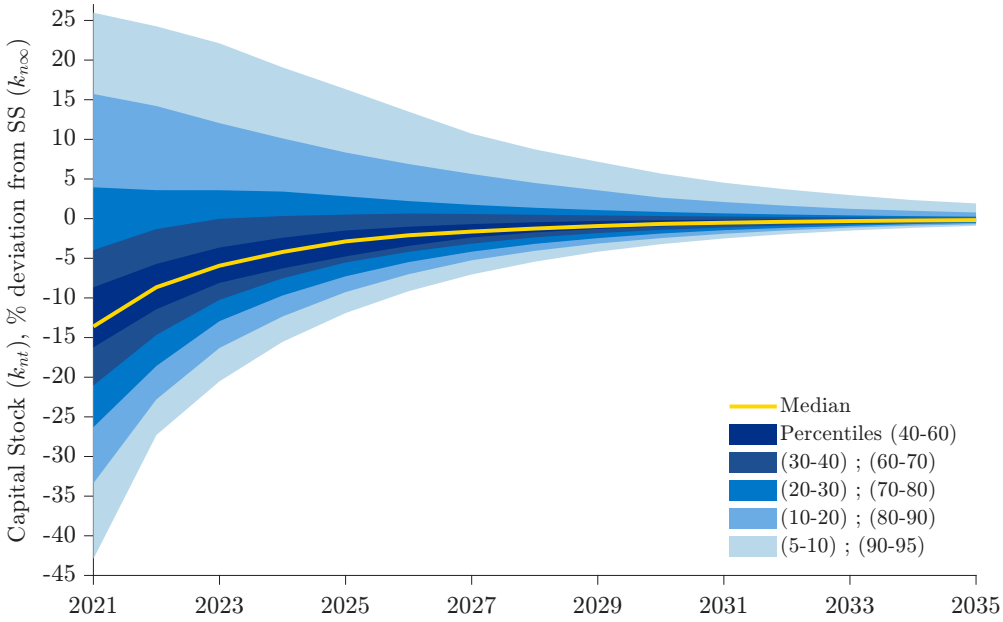


FIGURE 8B: AVERAGE Q AND MARGINAL Q

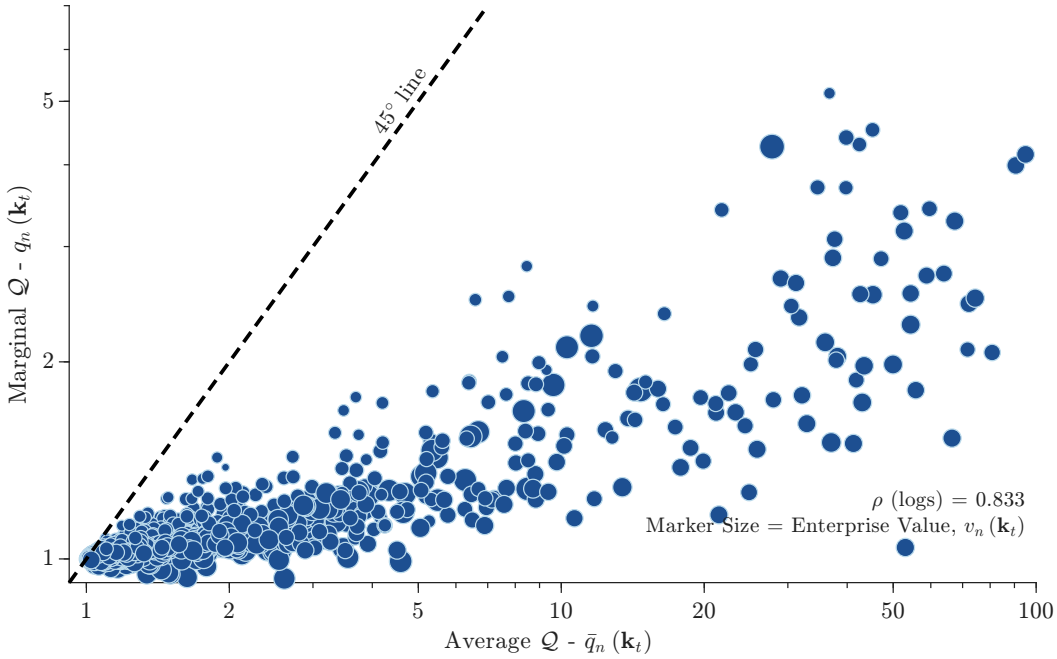


FIGURE 9B: DECOMPOSITION OF STEADY STATE AGGREGATE FIRM VALUE

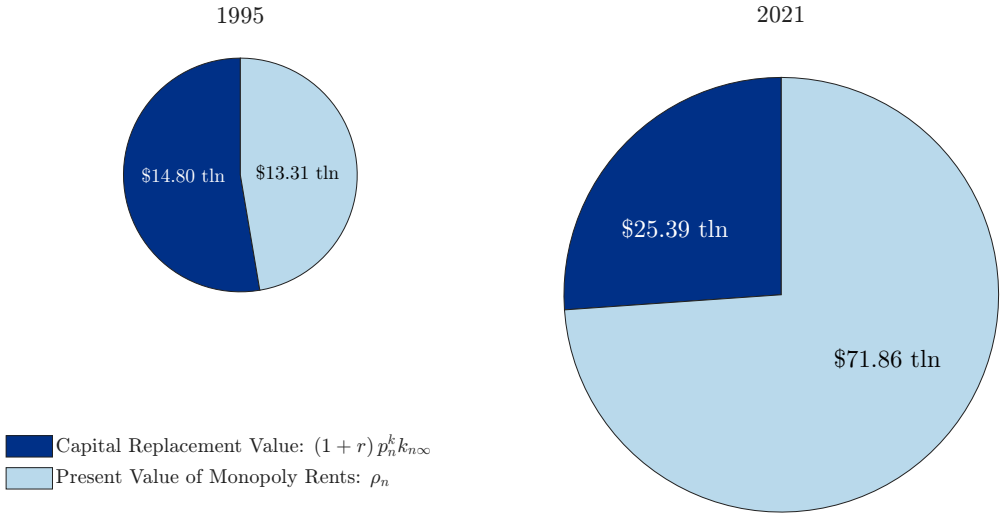


FIGURE 10B: MARGINAL PROPENSITY TO INVEST OUT OF GROWTH OPPORTUNITIES

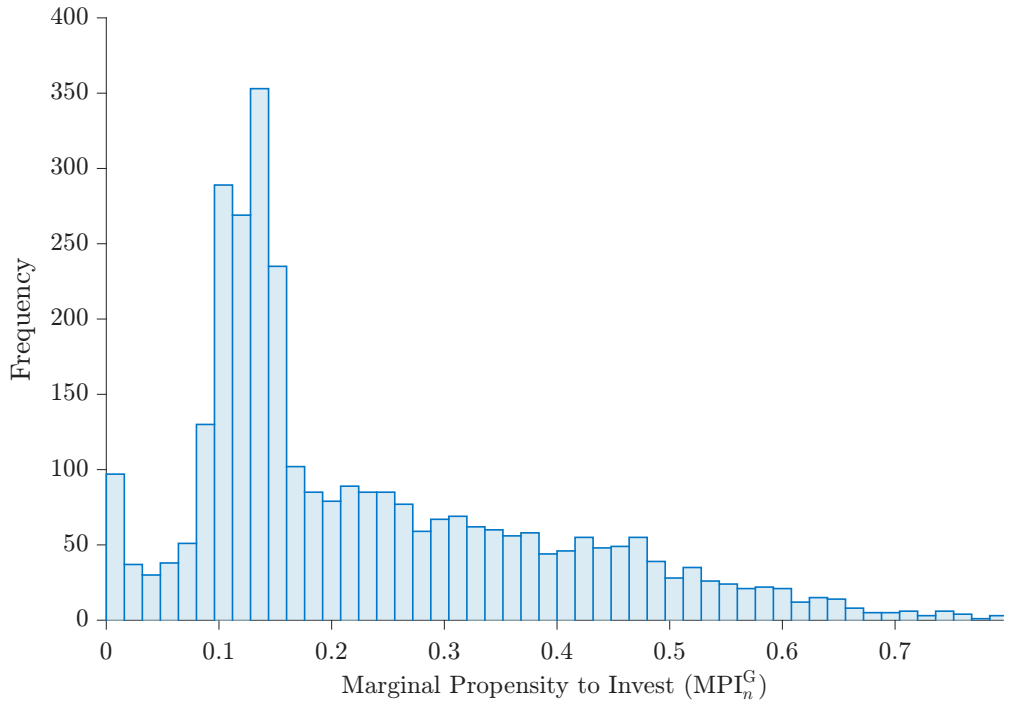


FIGURE 11B: ANTICOMPETITIVE COUNTERFACTUAL: AGGREGATE CAPITAL STOCK

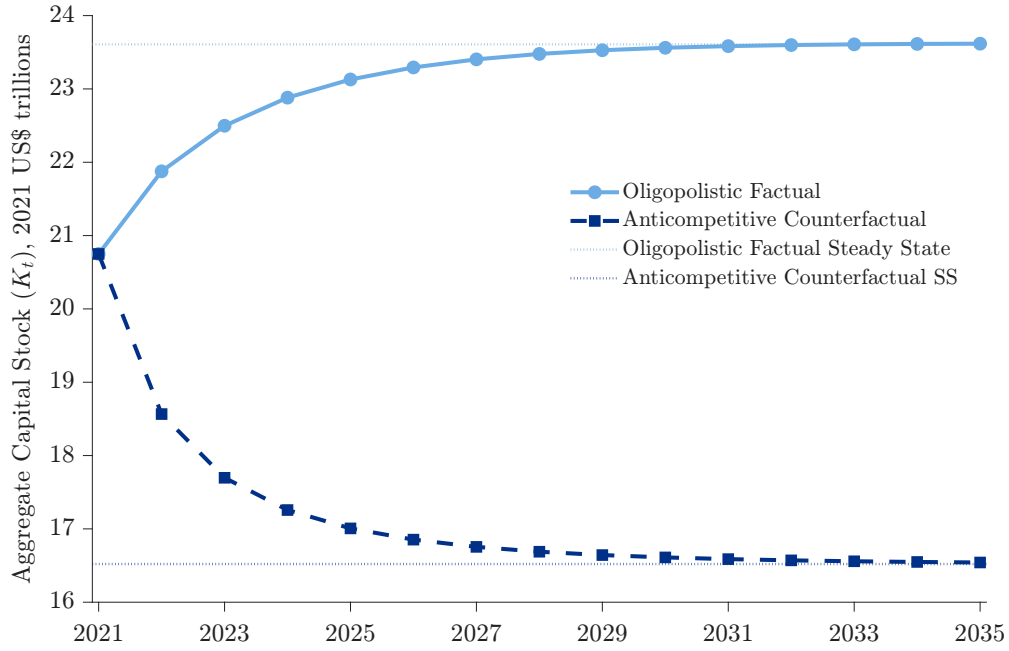


FIGURE 12B: ANTICOMPETITIVE COUNTERFACTUAL: CROSS-SECTION

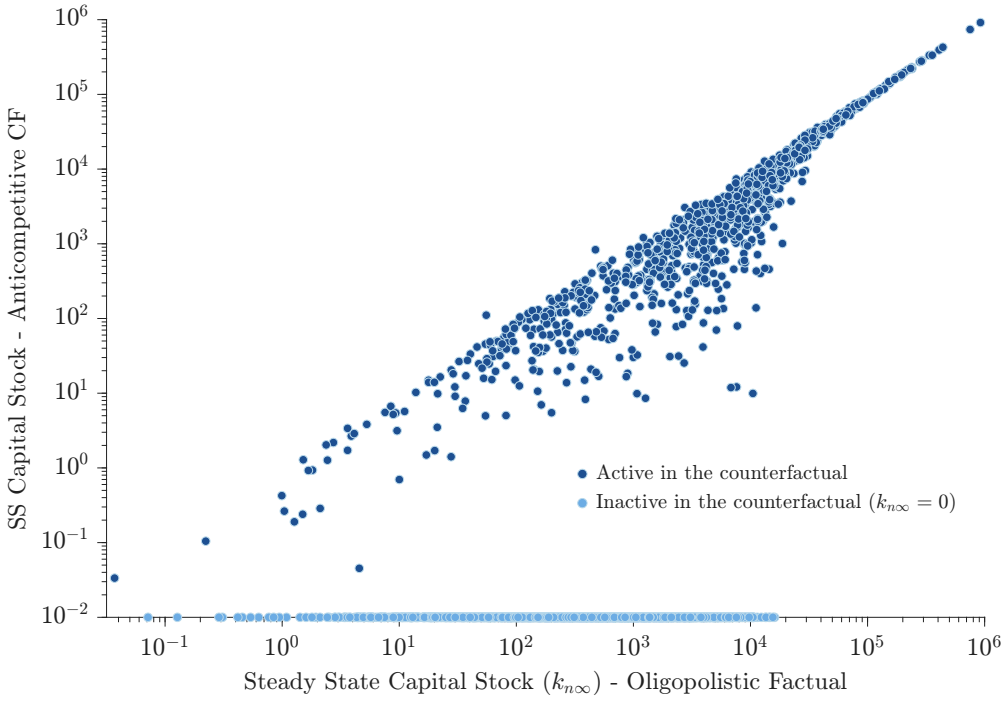


FIGURE 13B: DISTRIBUTION OF MARKUPS

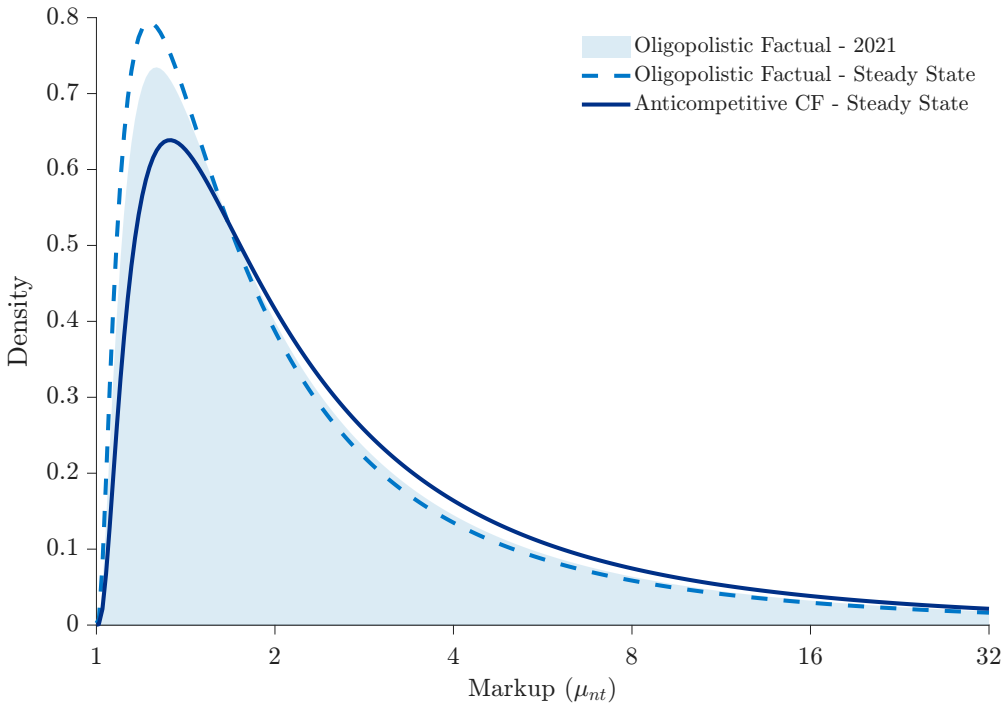


FIGURE 14B: RESPONSE OF INVESTMENT RATES TO DISCOUNT RATE SHOCK

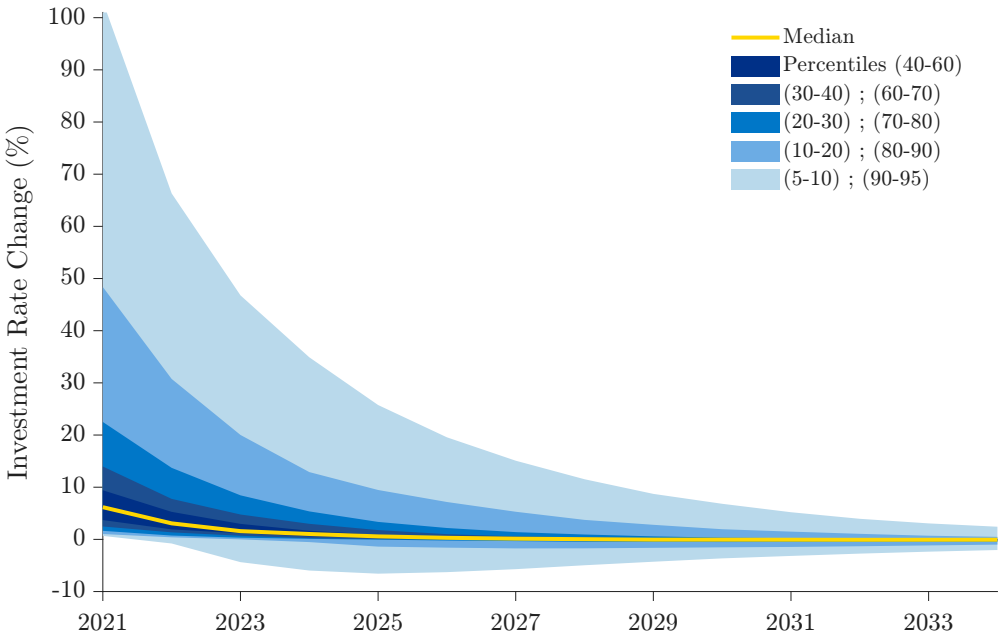


FIGURE 15B: EVOLUTION OF CONCENTRATION

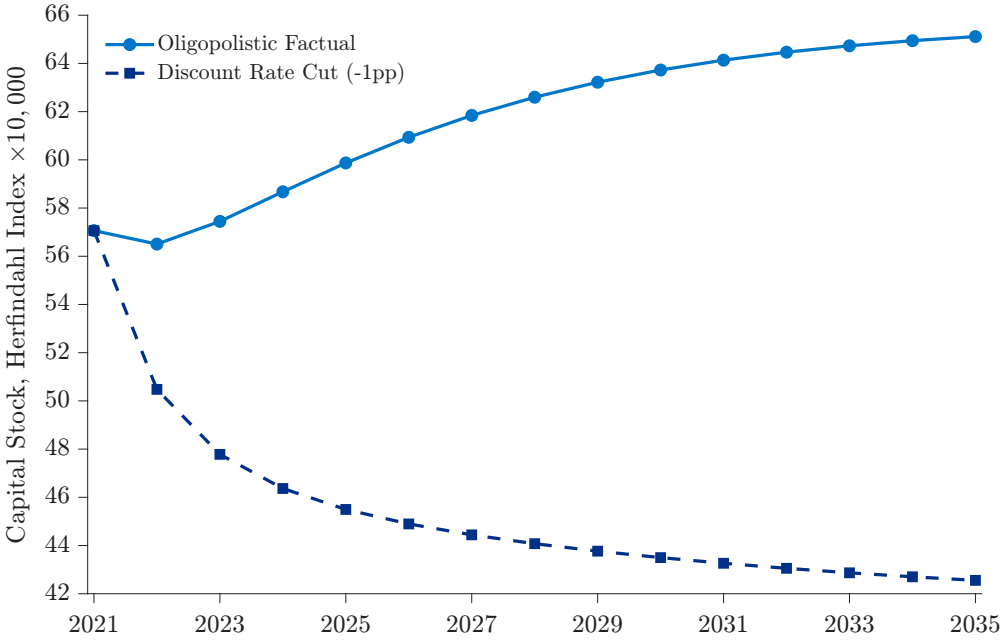


FIGURE 16B: RESPONSE OF MARKUPS TO DISCOUNT RATE SHOCK

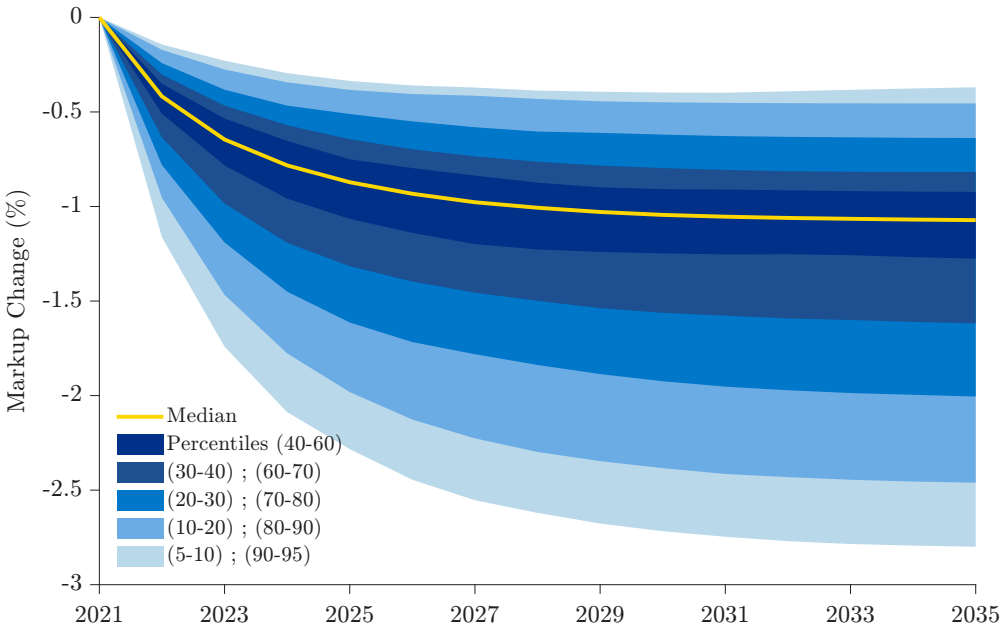


FIGURE 17B: MERGERS SIMULATION: INVESTMENT

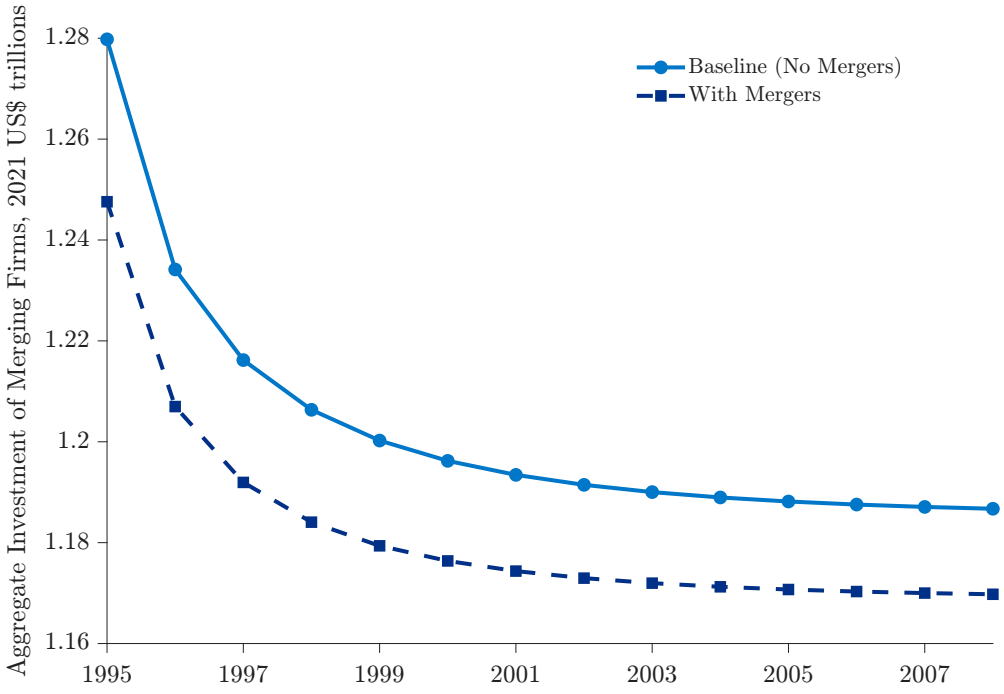
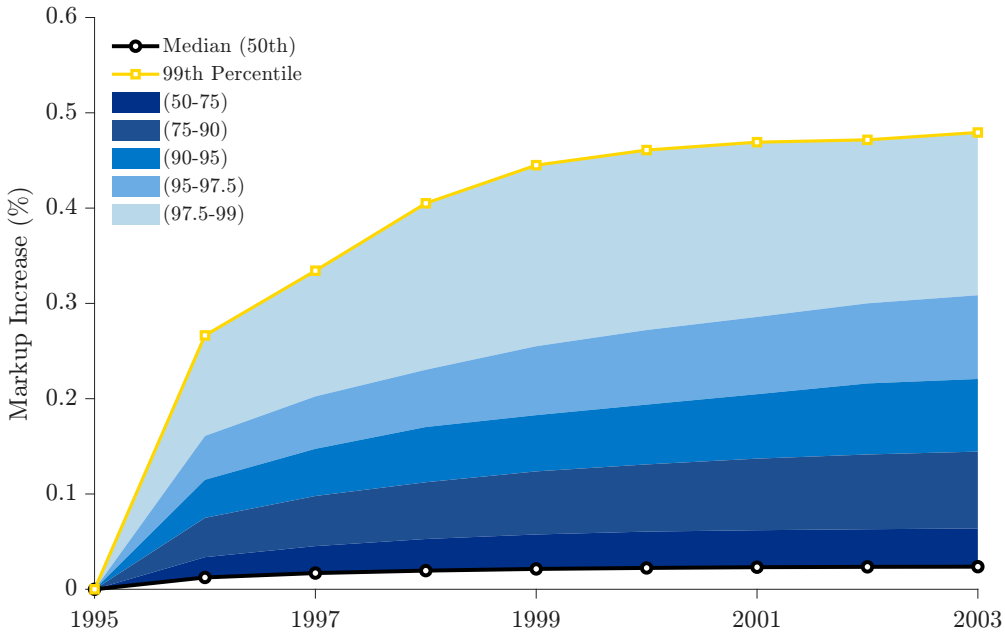


FIGURE 18B: MERGERS SIMULATION: MARKUPS



D.3 High Adjustment Cost Parameter: $\theta \times 1.5$

This appendix reproduces all figures from the main paper using, as an alternative value of the adjustment cost parameter θ, the baseline estimate multiplied by 1.5.

FIGURE 4C: TRANSITION MATRIX: DISTRIBUTION OF EIGENVALUES

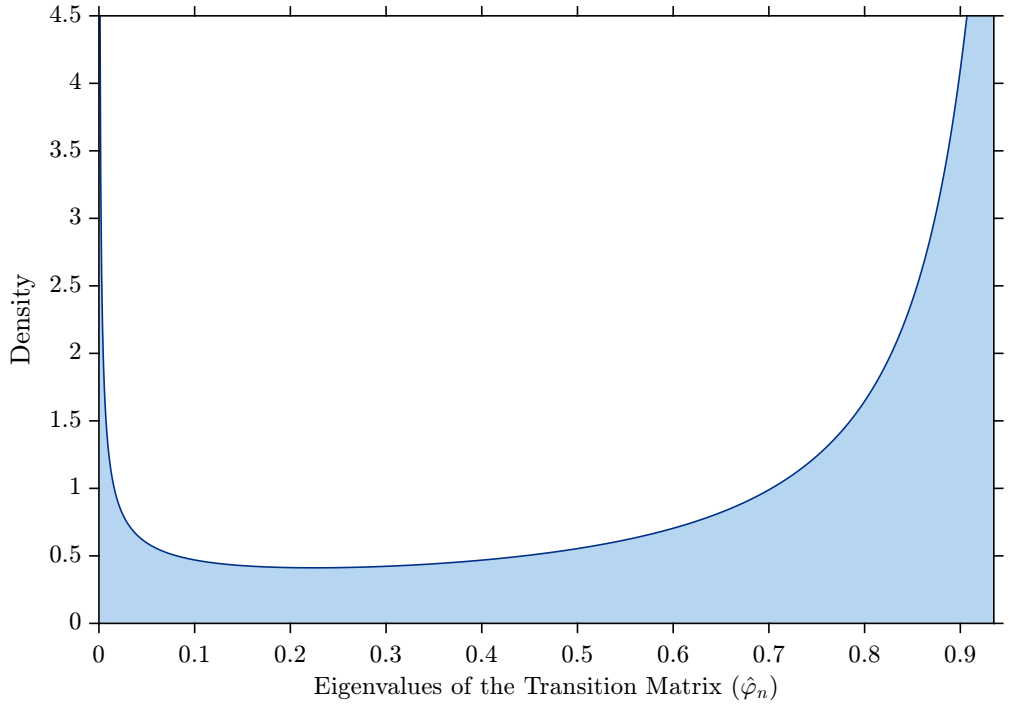


FIGURE 5C: MODEL FIT

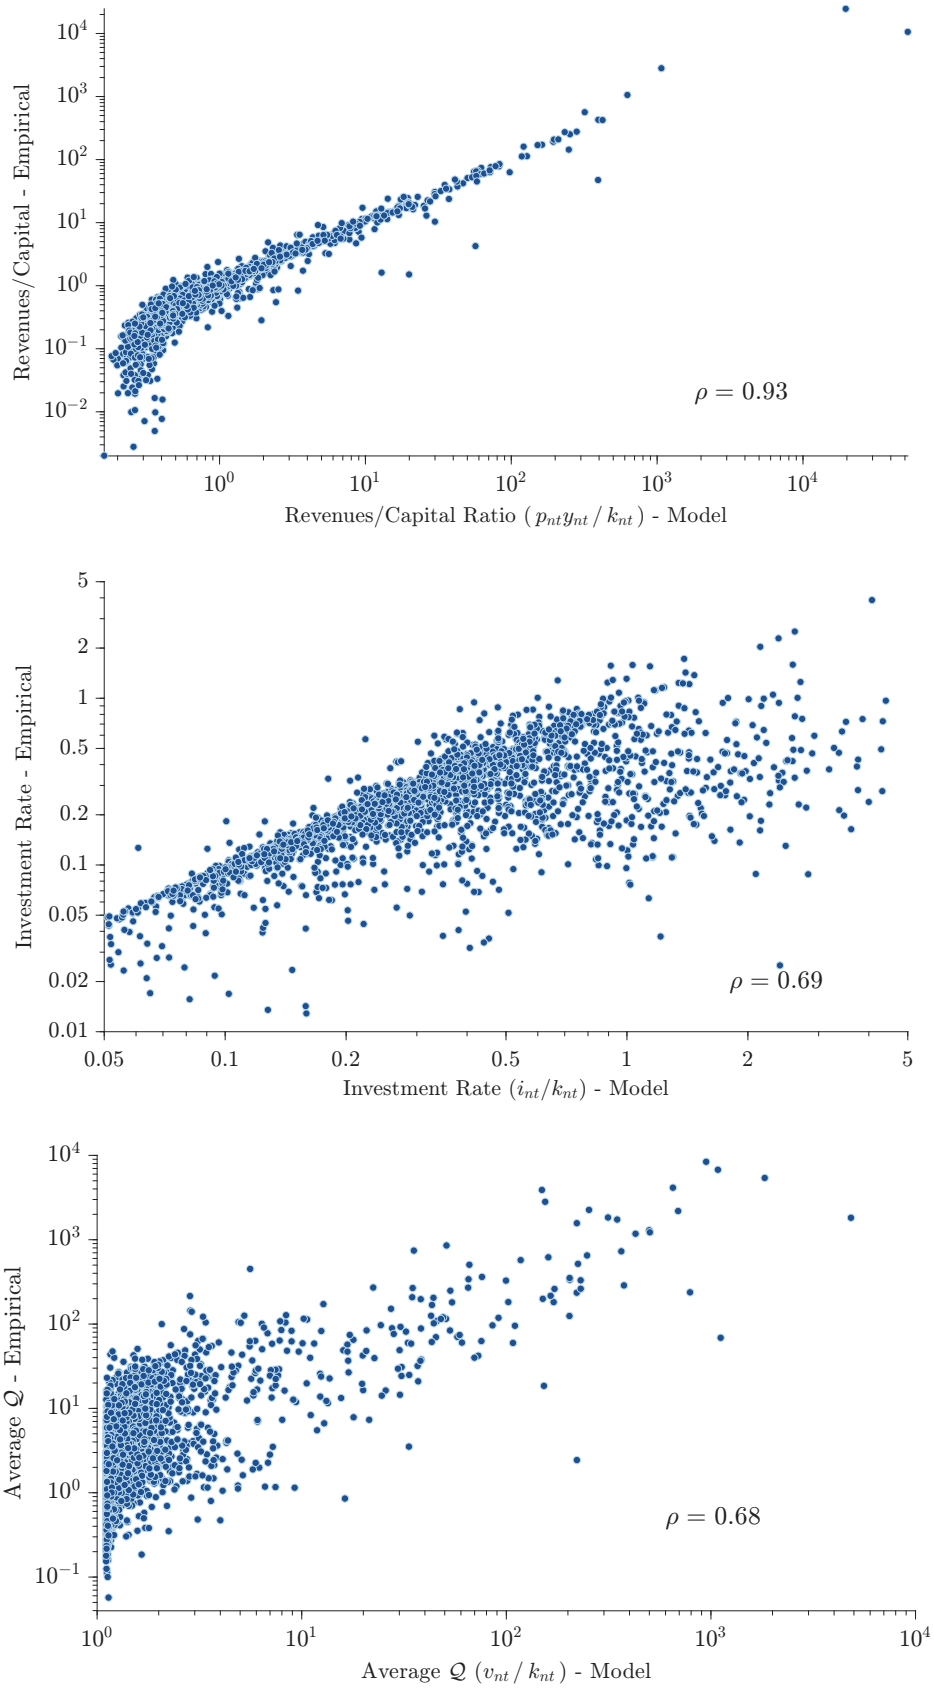


FIGURE 6C: TRANSITION DYNAMICS OF THE AGGREGATE CAPITAL STOCK

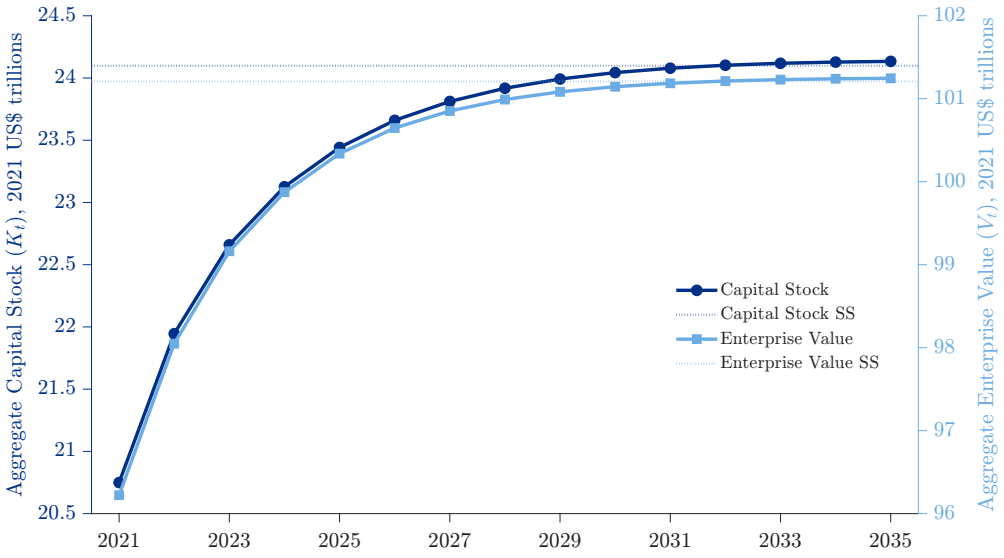


FIGURE 7C: FIRM-LEVEL TRANSITION DYNAMICS

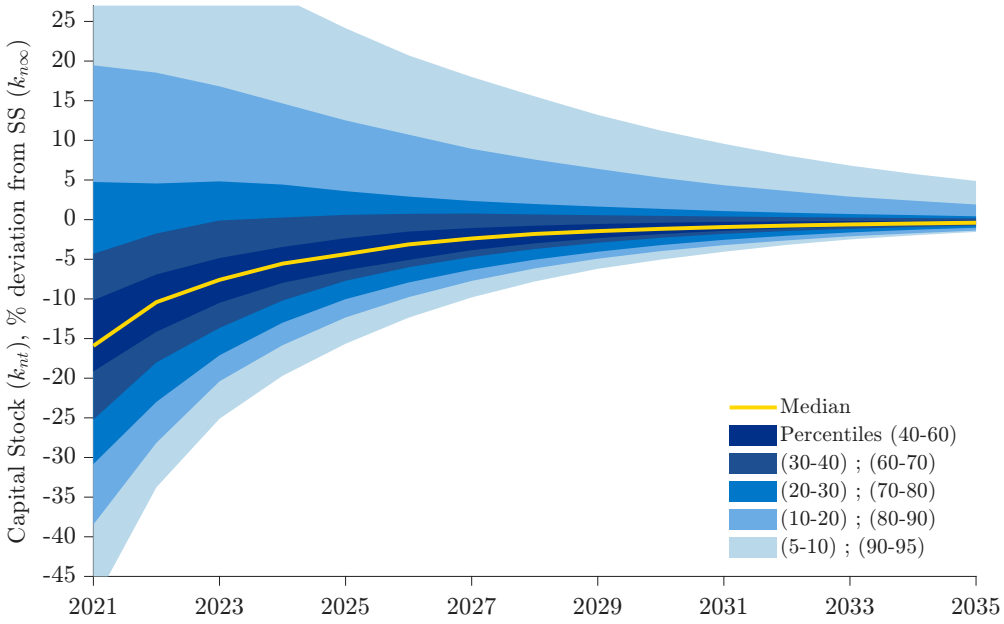


FIGURE 8C: AVERAGE Q AND MARGINAL Q

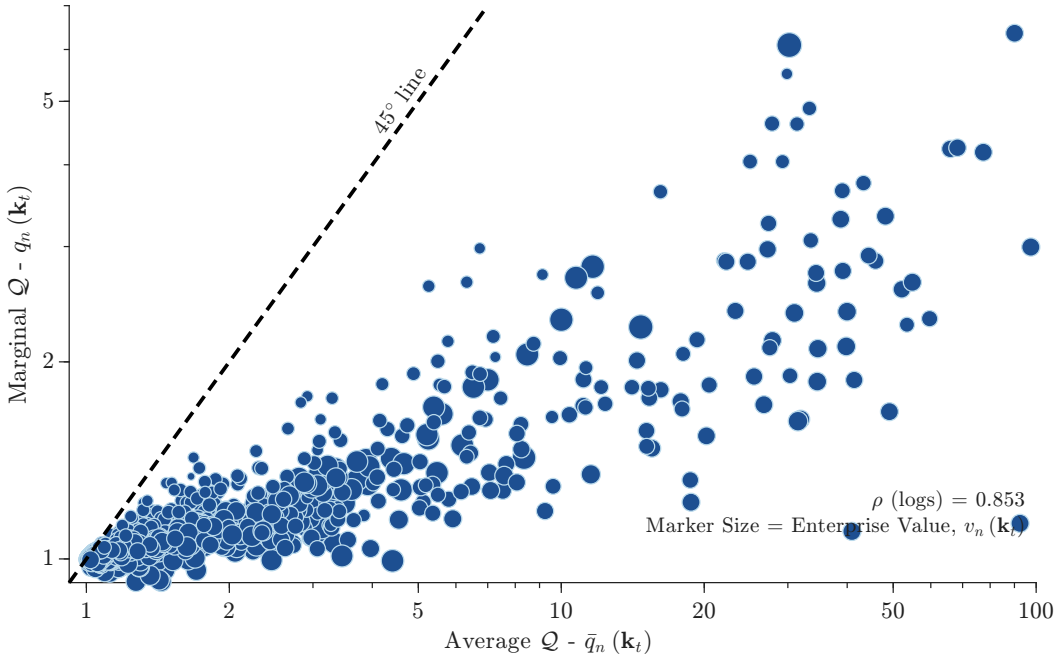


FIGURE 9C: DECOMPOSITION OF STEADY STATE AGGREGATE FIRM VALUE

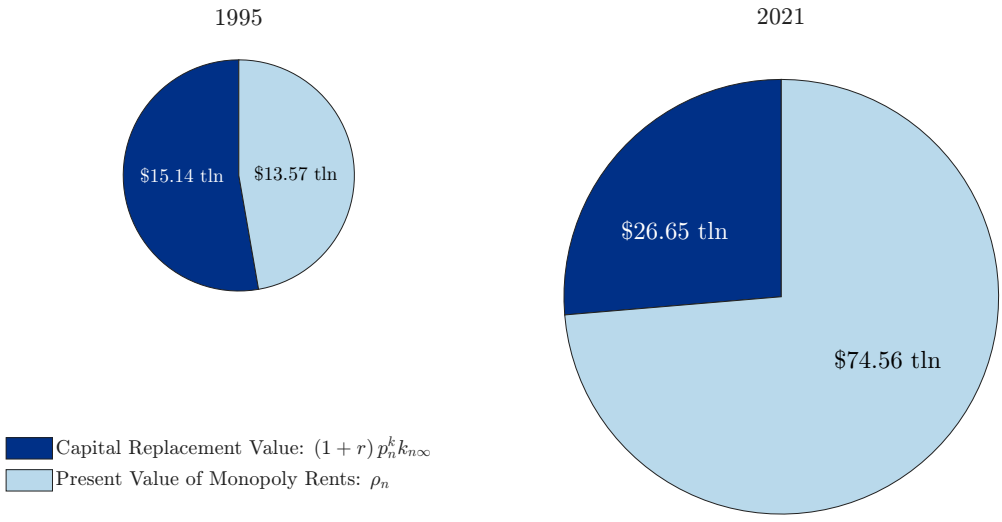


FIGURE 10C: MARGINAL PROPENSITY TO INVEST OUT OF GROWTH OPPORTUNITIES

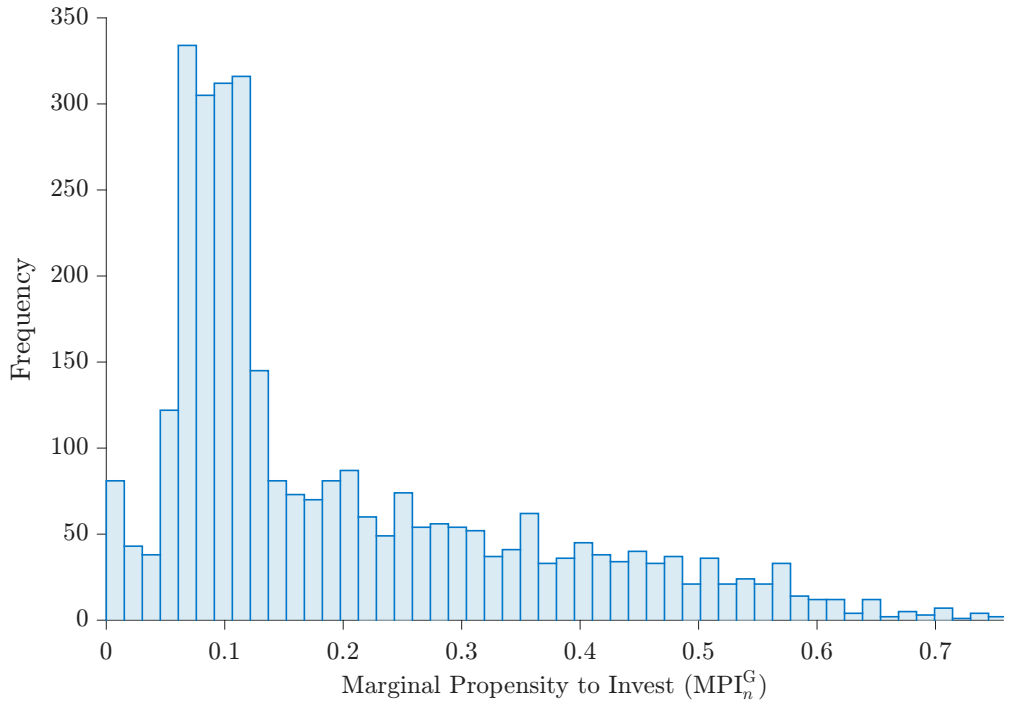


FIGURE 11C: ANTICOMPETITIVE COUNTERFACTUAL: AGGREGATE CAPITAL STOCK

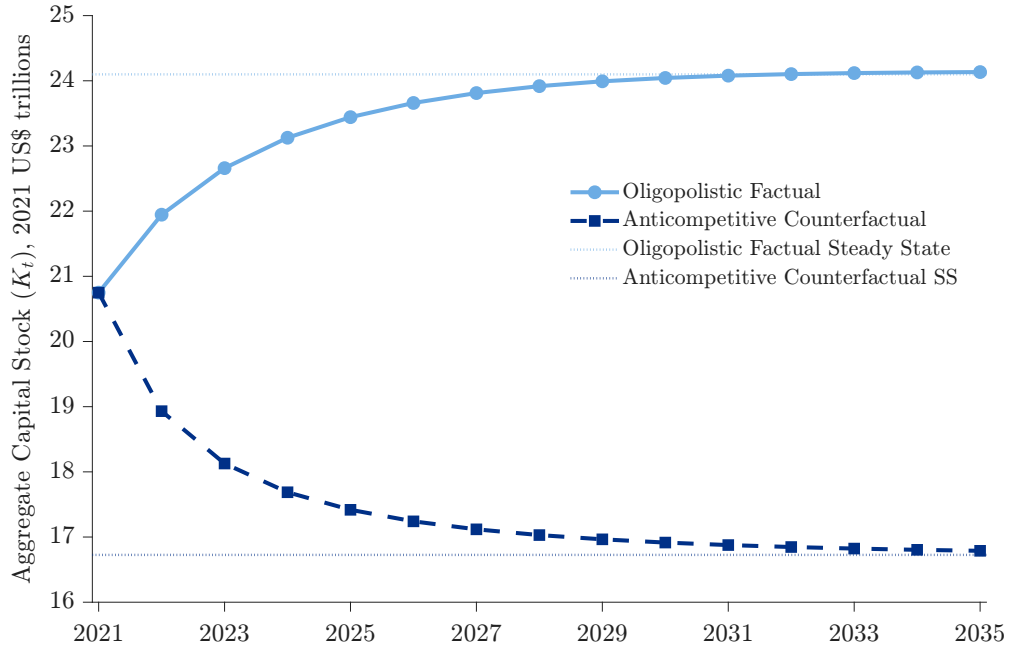


FIGURE 12C: ANTICOMPETITIVE COUNTERFACTUAL: CROSS-SECTION

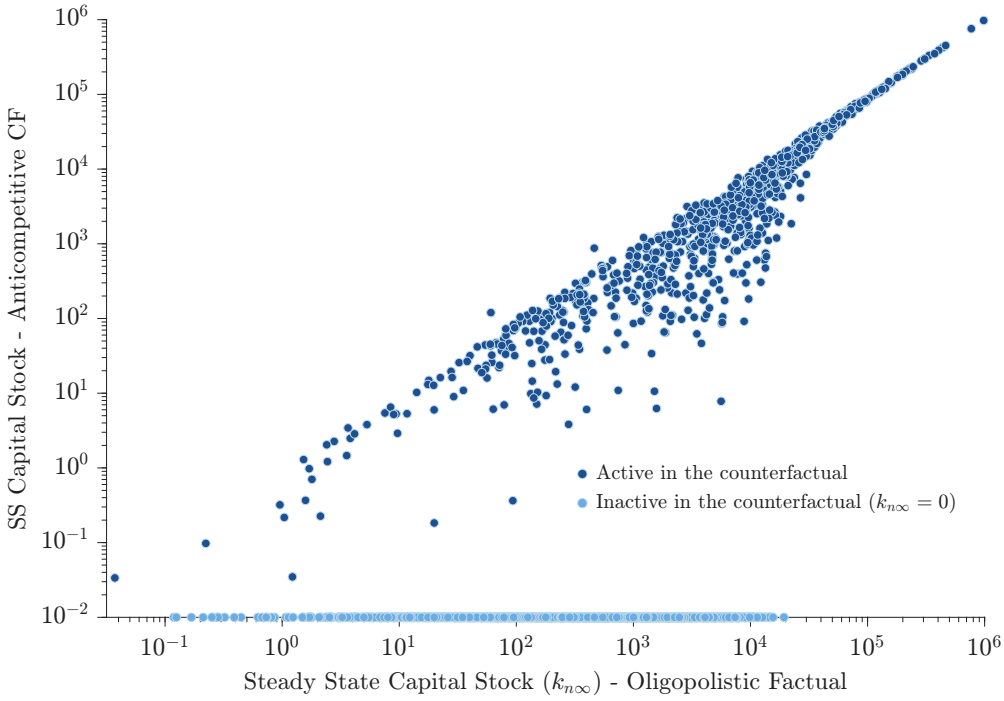


FIGURE 13C: DISTRIBUTION OF MARKUPS

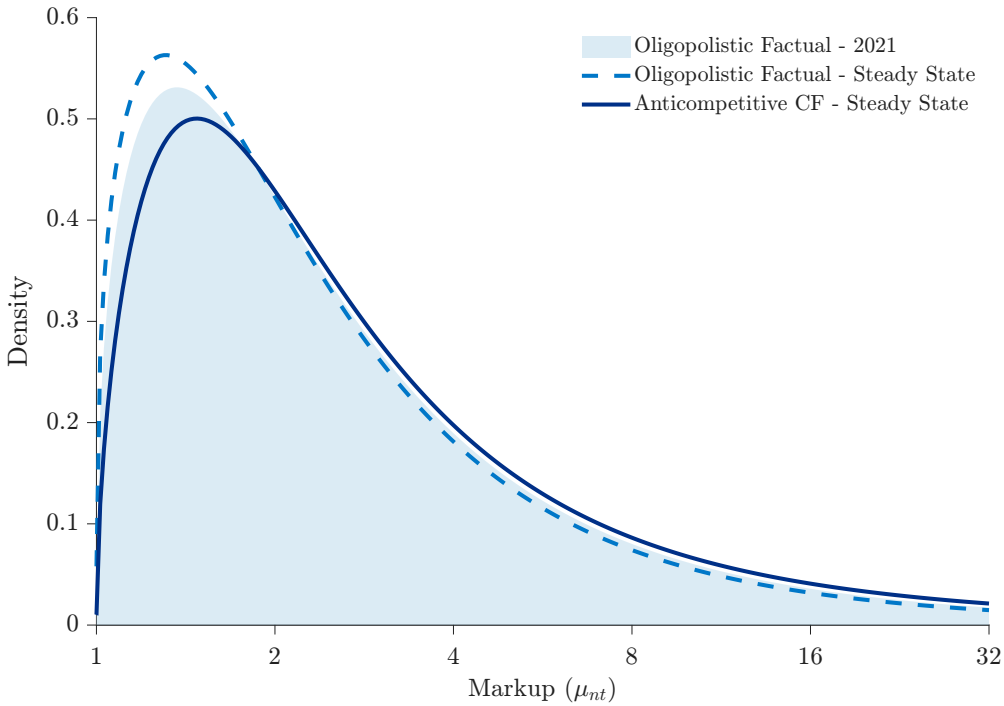


FIGURE 14C: RESPONSE OF INVESTMENT RATES TO DISCOUNT RATE SHOCK

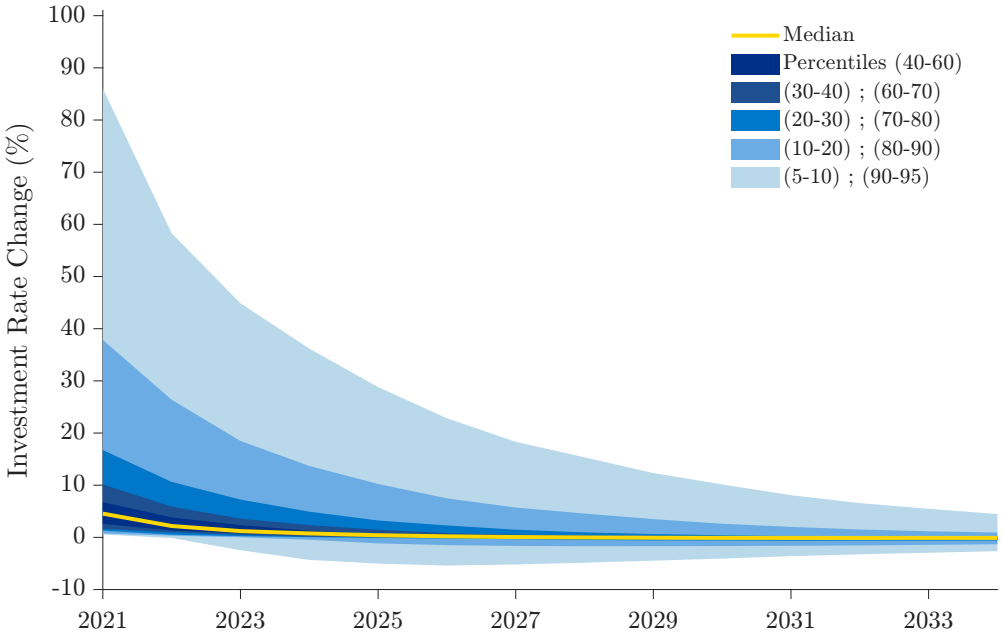


FIGURE 15C: EVOLUTION OF CONCENTRATION

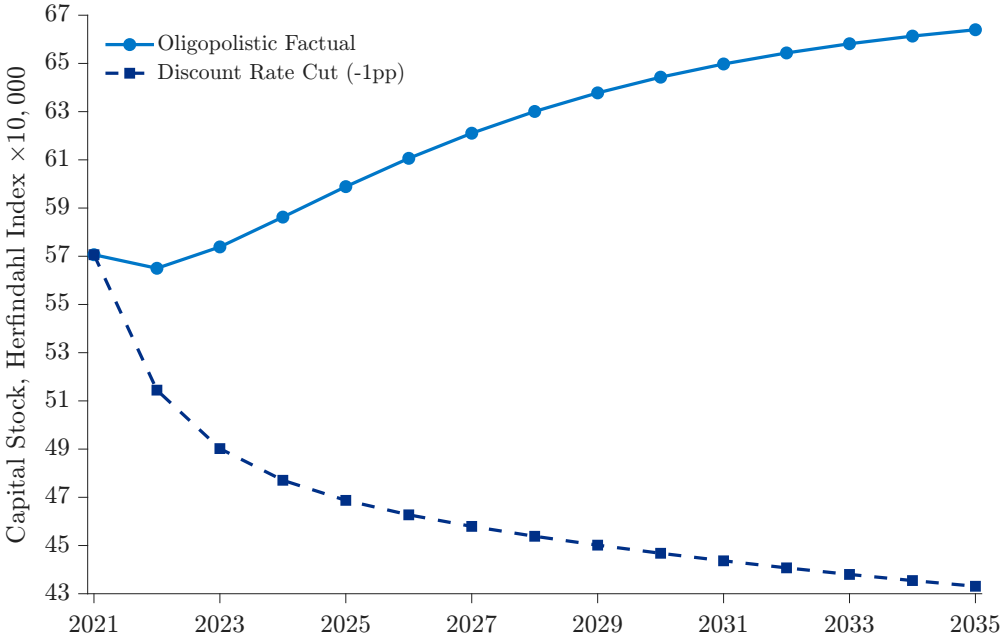


FIGURE 16C: RESPONSE OF MARKUPS TO DISCOUNT RATE SHOCK

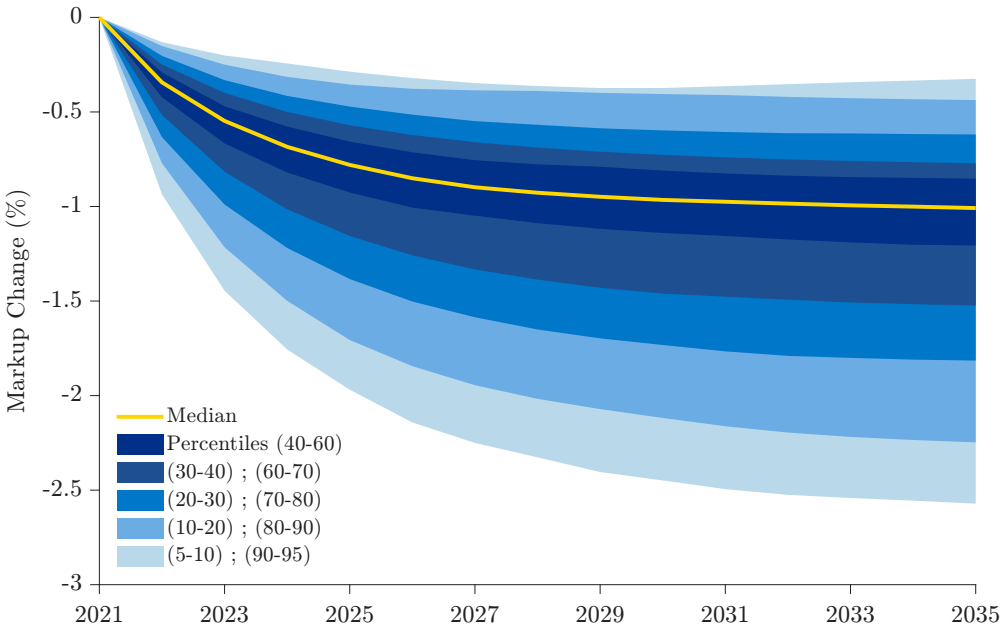


FIGURE 17C: MERGERS SIMULATION: INVESTMENT

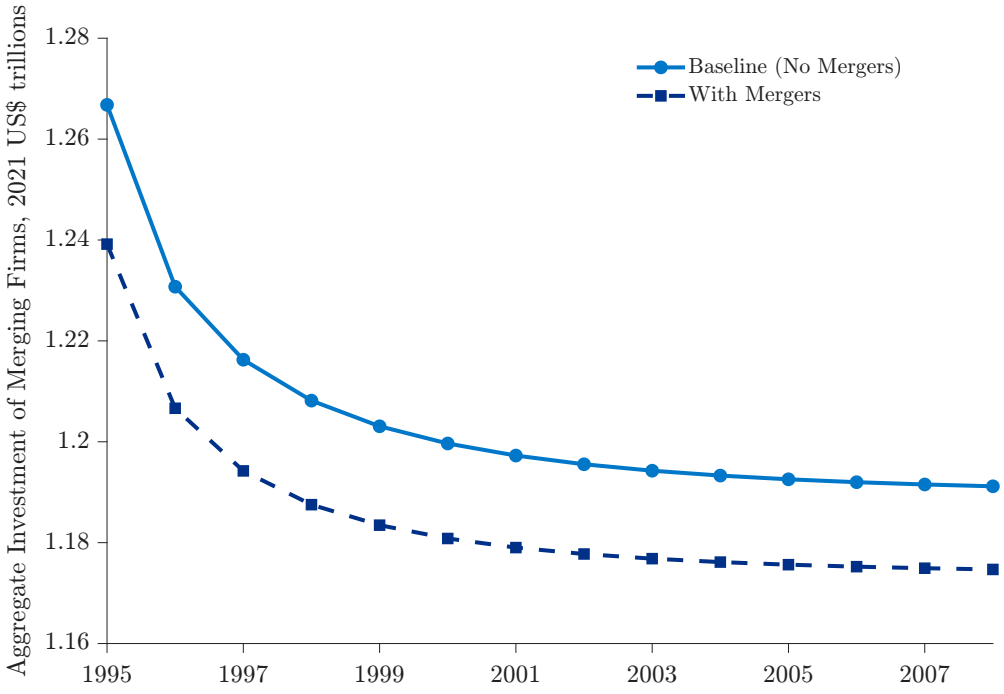
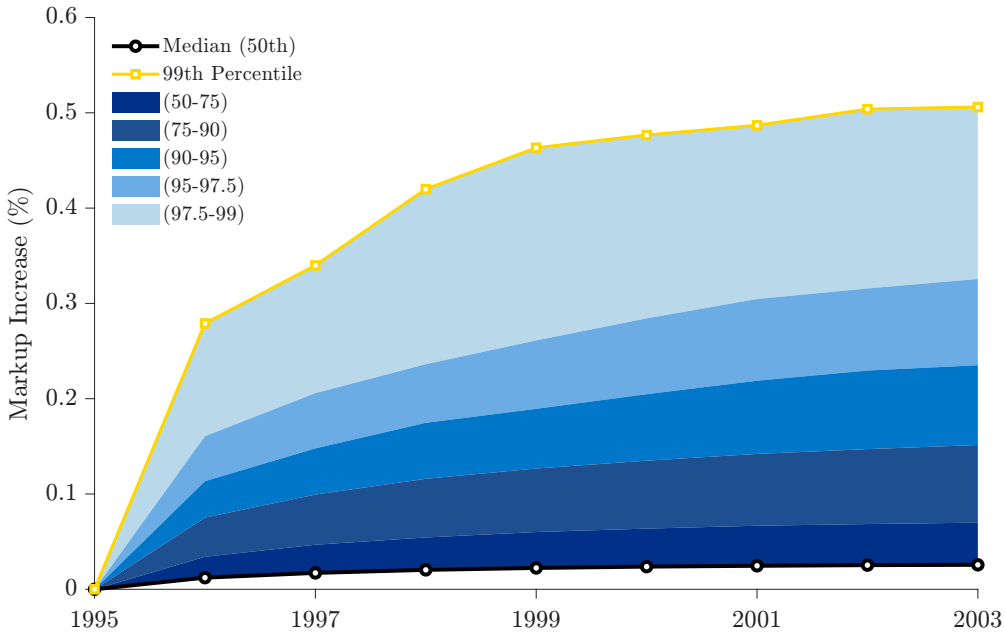


FIGURE 18C: MERGERS SIMULATION: MARKUPS



D.4 Winsorized Input Variables

This appendix reproduces all figures from the main paper after winsorizing six key inputs—revenues, costs, firm value, and three capital-stock measures—at 1% in each tail within each year.

FIGURE 4D: TRANSITION MATRIX: DISTRIBUTION OF EIGENVALUES

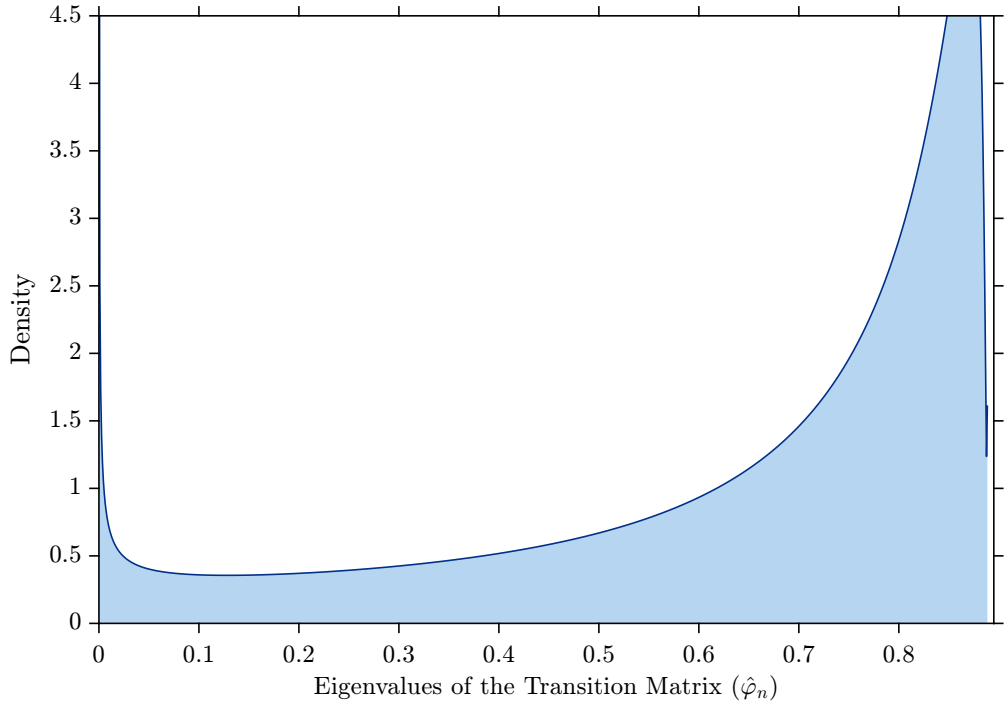


FIGURE 5D: MODEL FIT

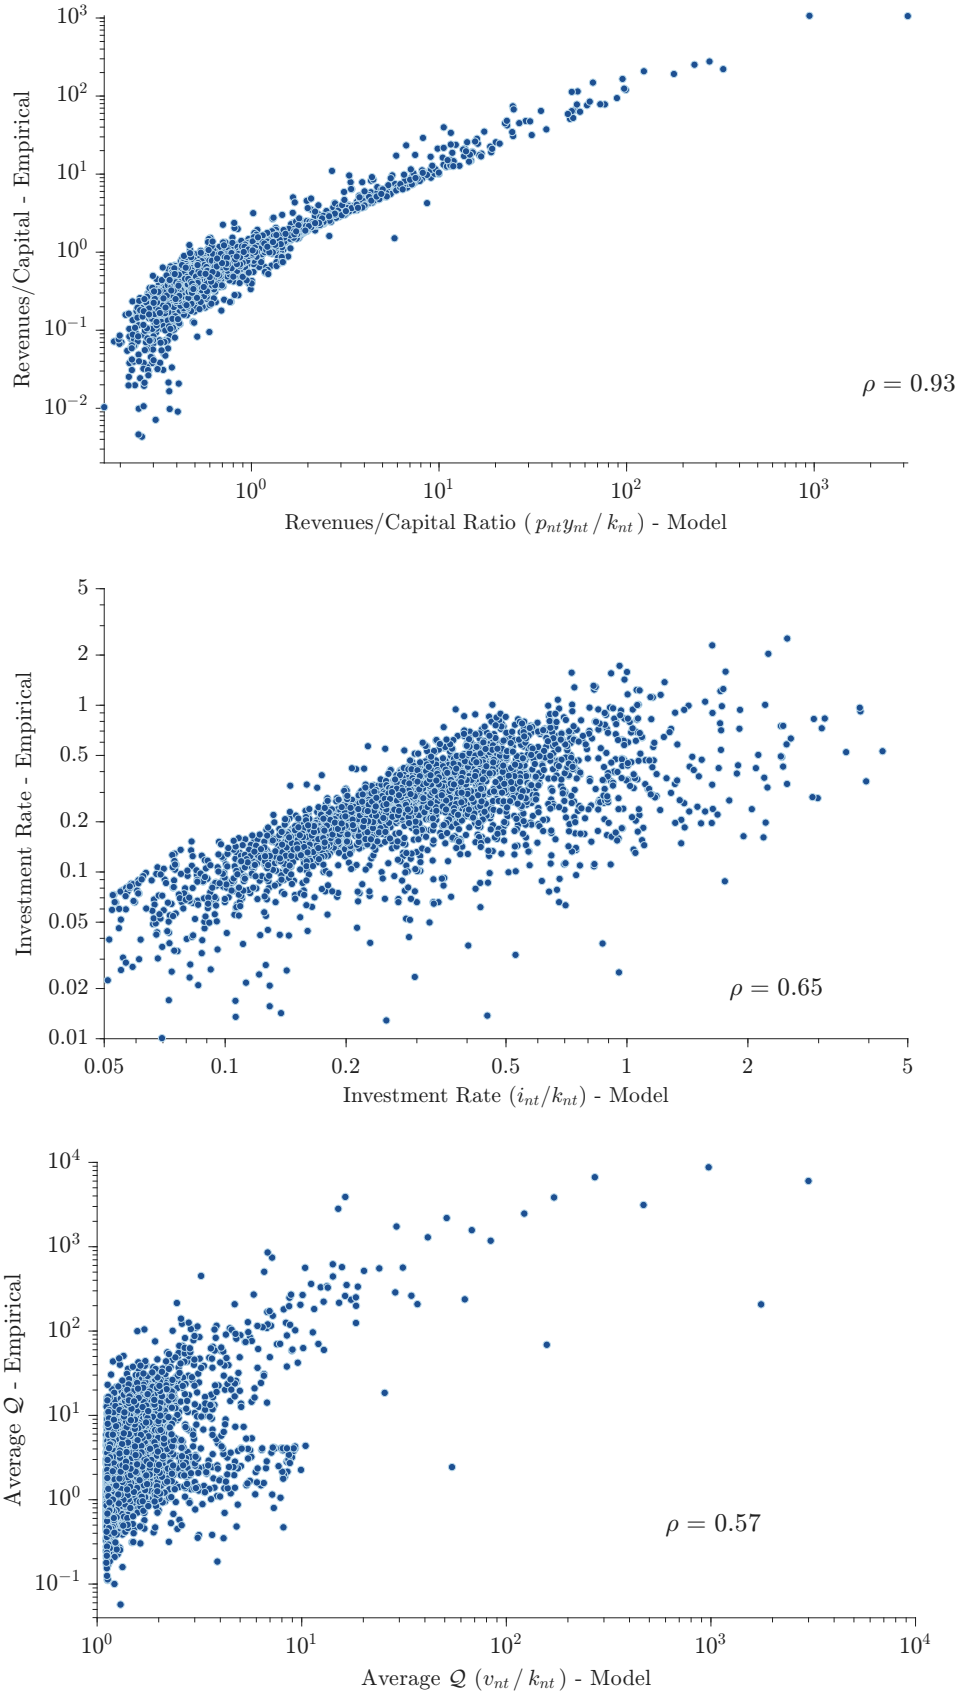


FIGURE 6D: TRANSITION DYNAMICS OF THE AGGREGATE CAPITAL STOCK

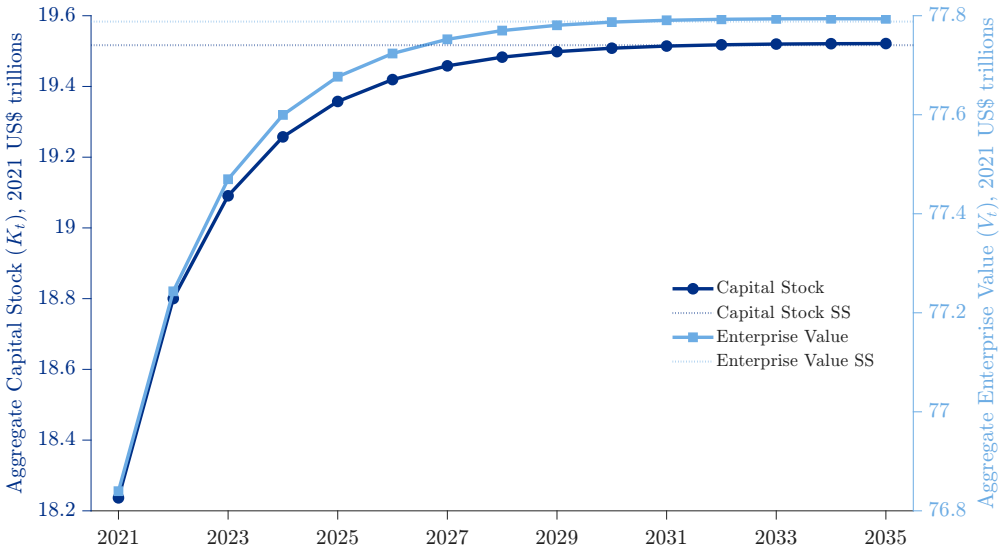


FIGURE 7D: FIRM-LEVEL TRANSITION DYNAMICS

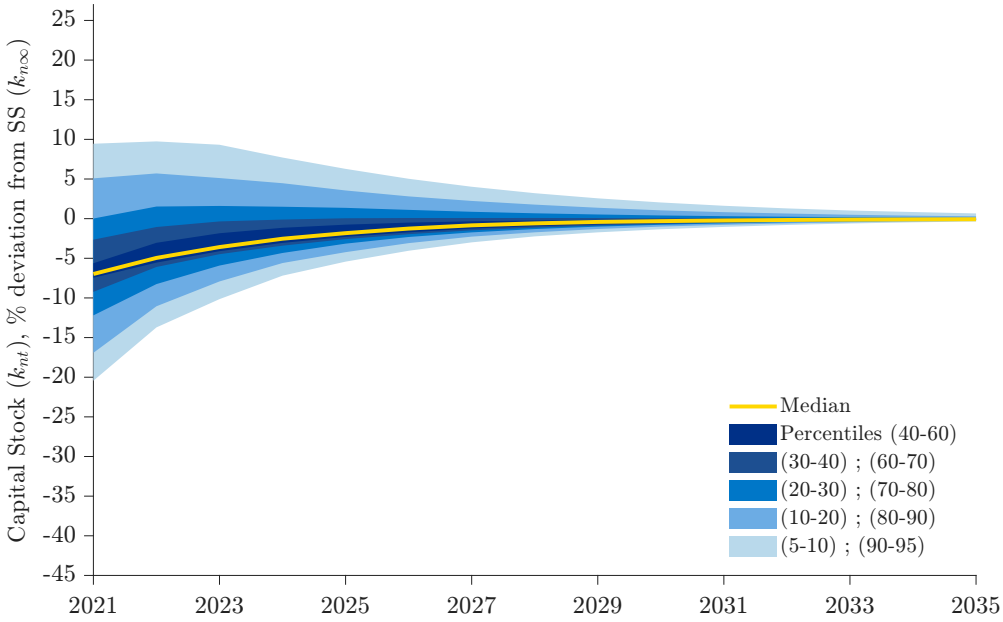


FIGURE 8D: AVERAGE Q AND MARGINAL Q

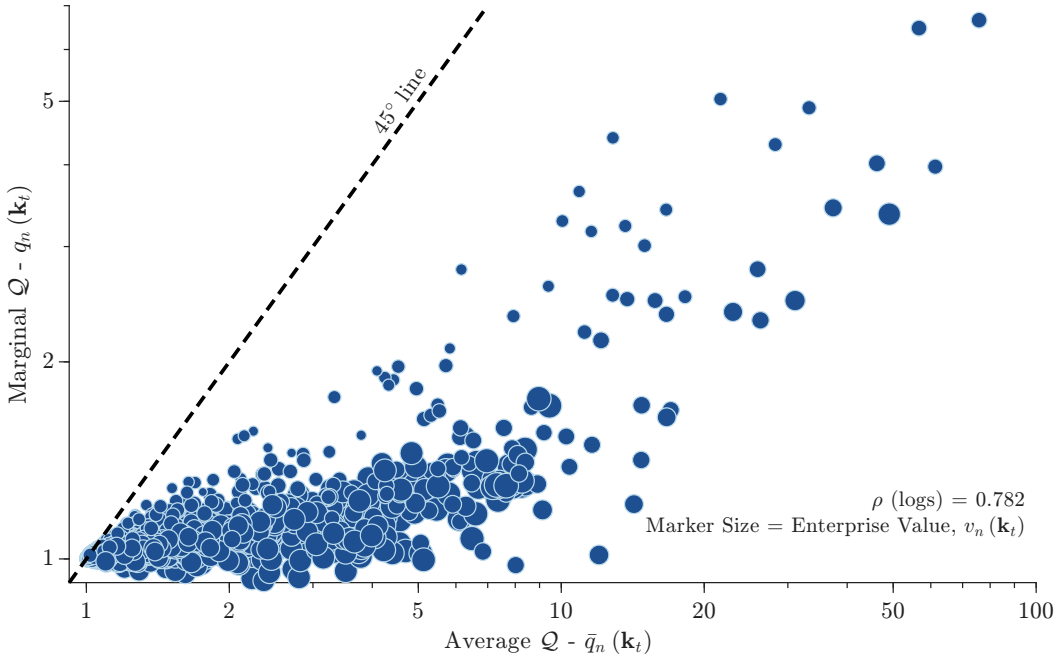


FIGURE 9D: DECOMPOSITION OF STEADY STATE AGGREGATE FIRM VALUE

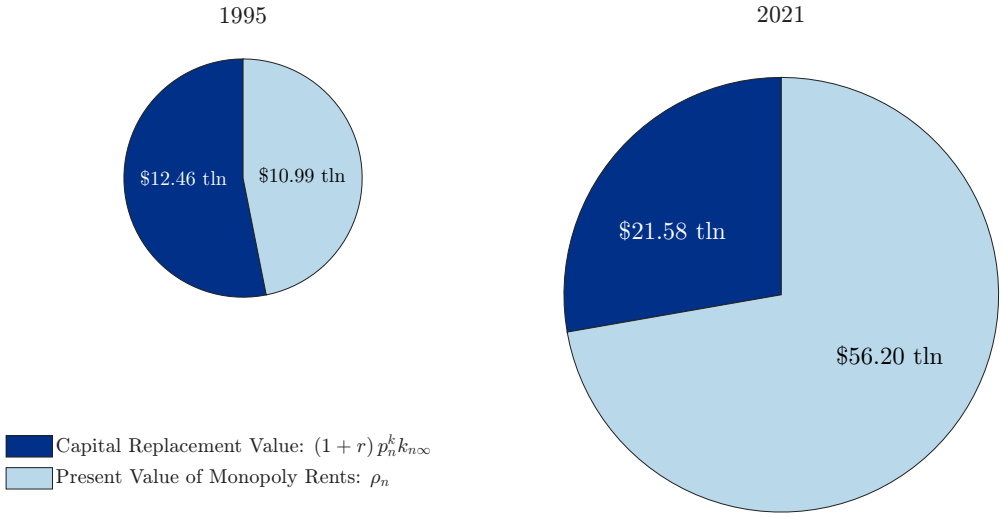


FIGURE 10D: MARGINAL PROPENSITY TO INVEST OUT OF GROWTH OPPORTUNITIES

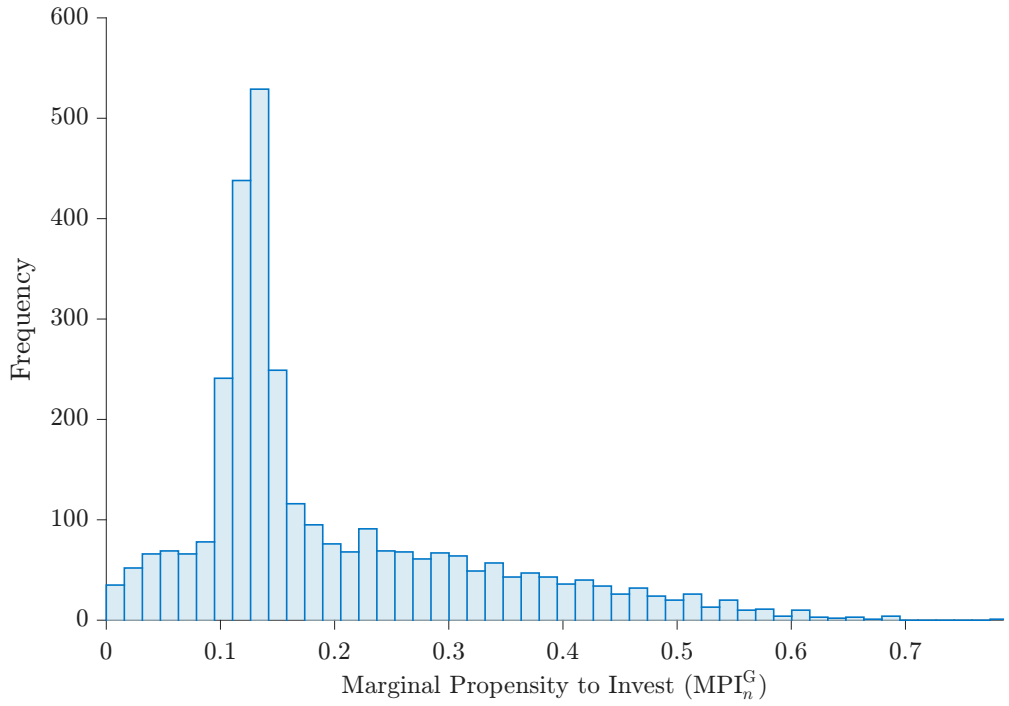


FIGURE 11D: ANTICOMPETITIVE COUNTERFACTUAL: AGGREGATE CAPITAL STOCK

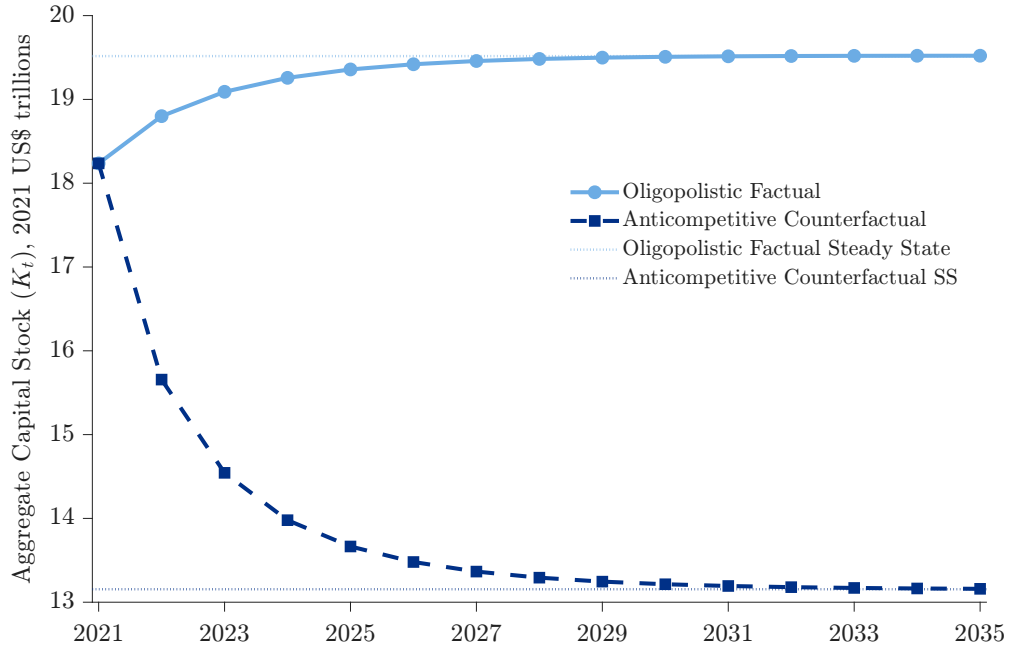


FIGURE 12D: ANTICOMPETITIVE COUNTERFACTUAL: CROSS-SECTION

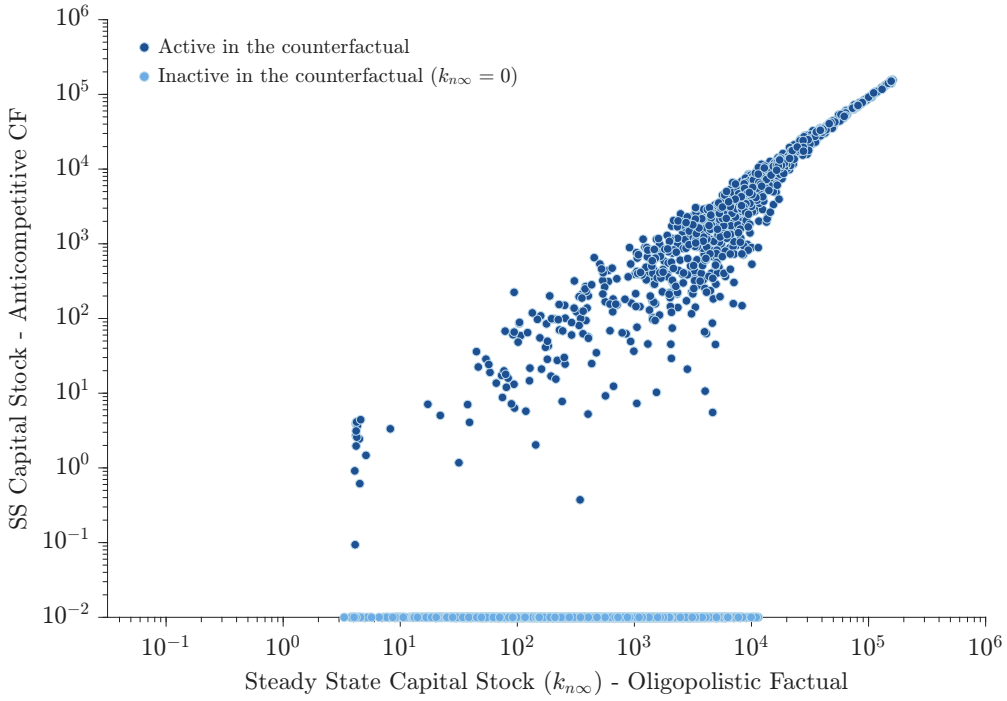


FIGURE 13D: DISTRIBUTION OF MARKUPS

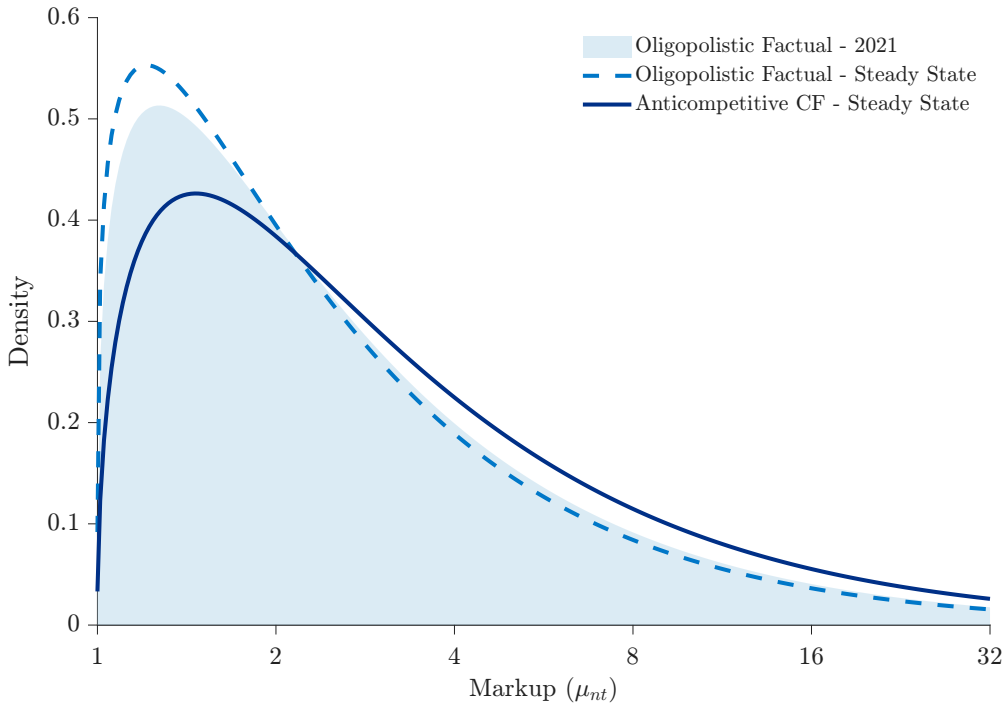


FIGURE 14D: RESPONSE OF INVESTMENT RATES TO DISCOUNT RATE SHOCK

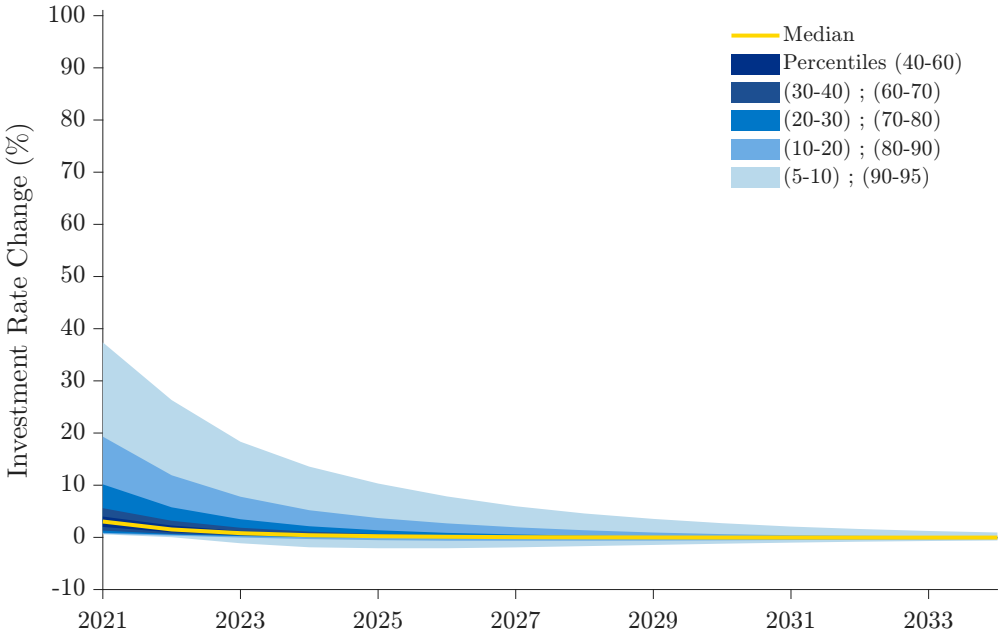


FIGURE 15D: EVOLUTION OF CONCENTRATION

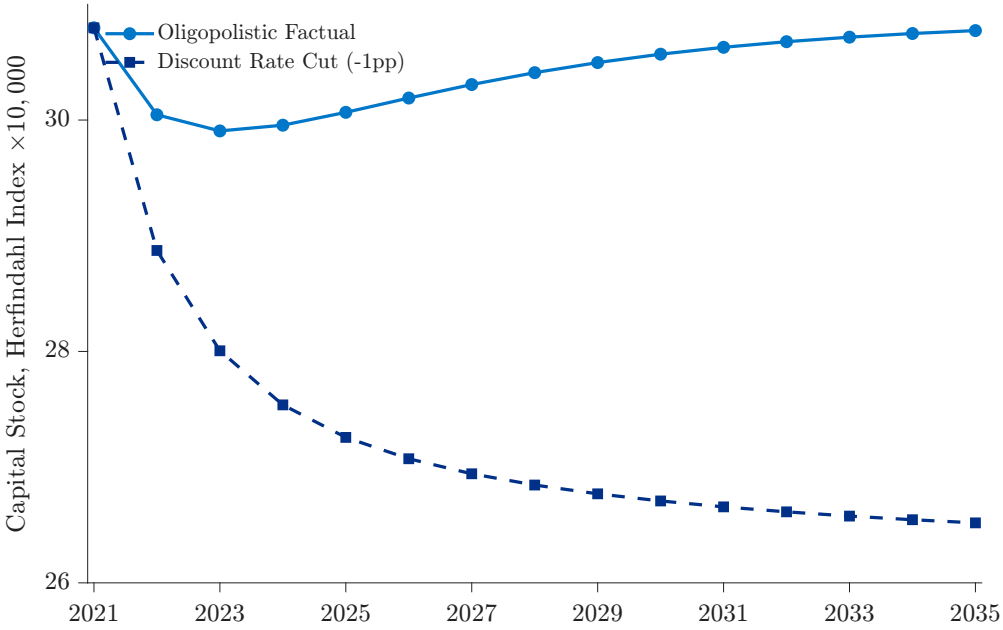


FIGURE 16D: RESPONSE OF MARKUPS TO DISCOUNT RATE SHOCK

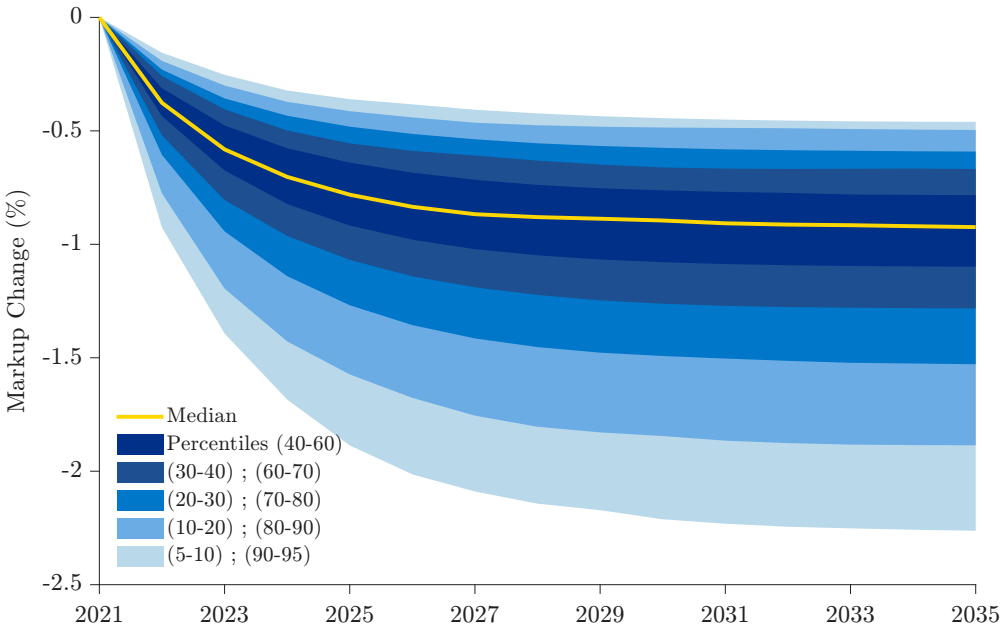


FIGURE 17D: MERGERS SIMULATION: INVESTMENT

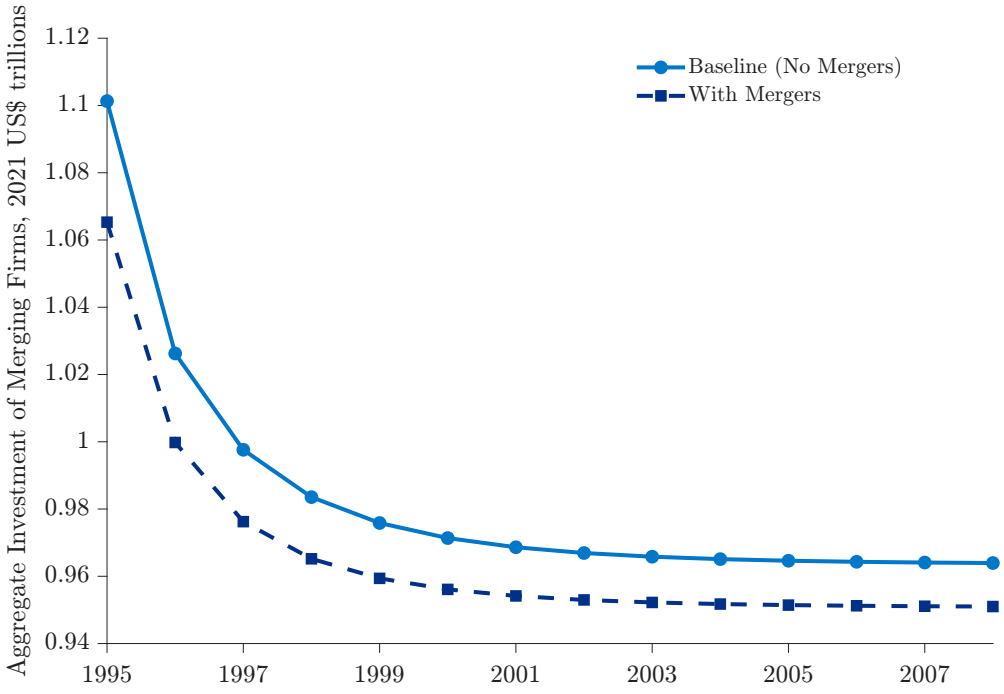


FIGURE 18D: MERGERS SIMULATION: MARKUPS

