Dynamic Oligopoly and Innovation

Hugo Hopenhayn
UCLA

Koki Okumura*
UCLA

June 12, 2026

Abstract

We develop and estimate a tractable many-firm endogenous growth model in which firms make forward-looking R&D decisions while interacting through product-market rivalry and technology-spillover networks. The model's linear-quadratic structure allows us to solve a dynamic oligopoly with hundreds of firms and empirically measured firm-to-firm links. Using publicly traded U.S. firms with patents, we find that competitive R&D is both underprovided in aggregate and misallocated across firms. At the observed 2017 knowledge-capital distribution, a constrained planner that reallocates R&D while holding the competitive product-market allocation fixed raises R&D to about 242% of the competitive level and consumption-equivalent welfare by 4.4%. A welfare-maximizing uniform R&D subsidy of 34% nearly replicates the constrained planner's increase in the expected economic growth rate but raises consumption-equivalent welfare by only 0.5%. The gap reflects heterogeneous firm-level social-private R&D wedges, highlighting the limits of uniform innovation subsidies.

*We thank David Baqaee, Ariel Burstein, Bruno Pellegrino, and seminar participants at UCLA and conference participants at the BSE Summer Forum for helpful comments.

1 Introduction

Product-market competition and research and development (R&D) allocations are jointly determined. Firms' R&D decisions shape future productivity, product quality, and competitive positions, while current product-market rivalry feeds back into profits and the incentives to invest in R&D. The network-mediated part of these interactions operates through two pairwise channels: business stealing, which depresses competitors' profits, and technology spillovers, which raise other firms' productivity. In addition, firms do not generally capture all non-producer surplus generated by their innovations. Because these forces are heterogeneous across firms and firm pairs, the gap between social and private returns to R&D is firm specific.

The two networks need not coincide: firms that compete closely in product markets may be distant in technology space, and firms that generate large spillovers may not be close product-market rivals. Yet common innovation-policy instruments, such as broad R&D subsidies and tax credits, change the price of research by the same proportion for every firm. The central question is how innovation policy should be designed when firms differ not only in how strongly they compete with one another, but also in how strongly they benefit from one another's research, and how much a uniform subsidy loses by ignoring this heterogeneity.

Existing endogenous growth models typically abstract from this firm-pair heterogeneity. Some models assume monopolistic competition without strategic interaction, while others allow strategic interaction among only a few firms. The difficulty is computational. In oligopoly, forward-looking R&D choices cause the dimensionality of the dynamic problem to increase rapidly with the number of strategically interacting firms. Existing quantitative frameworks therefore have had difficulty combining forward-looking strategic interaction among many firms with empirically measured networks of rivalry and spillovers. As a result, we lack firm-level measures of the social-private R&D wedges that policy should target.

This paper develops and estimates an endogenous growth model in which hundreds of firms make forward-looking R&D decisions while interacting through two empirically measured networks: product-market rivalry and technology spillovers. The contribution is twofold. Methodologically, the model's linear-quadratic structure makes a many-firm dynamic oligopoly tractable: equilibrium is characterized by a coupled system of algebraic Riccati equations rather than a high-dimensional dynamic program, so we can solve the game for the full 2017 cross section of 757 U.S. publicly listed firms with patents while preserving a continuous state space and the complete firm-to-firm networks. Substantively, the estimated model delivers firm-level social-private R&D wedges and uses them to show that a uniform R&D subsidy can replicate the constrained planner's effect on the expected economic growth rate while capturing only a small fraction of its welfare gain.

The model extends Pellegrino (2025)'s static product-market framework into a dynamic setting with knowledge accumulation, R&D investment, and technology spillovers. Conditional on the current knowledge-capital vector, a static product-market game under generalized hedonic linear demand maps firms' knowledge stocks into quantities and profits; its key property is that gross profits are quadratic in knowledge capital. Knowledge capital accumulates linearly through own R&D and through spillovers from technologically proximate firms. The dynamic R&D game is therefore linear-quadratic: value functions are quadratic in the state, Markov strategies are linear in it, and computation scales polynomially rather than exponentially in the number of firms. In the deterministic version, a technology transition matrix summarizes spillovers, obsolescence, and endogenous R&D, delivering balanced-growth implications through an eigenvalue problem.

We estimate the framework on U.S. publicly listed firms with patents over 1989–2019. Following Pellegrino (2025), we measure product-market rivalry using the product-similarity data of Hoberg and Phillips (2016); following Bloom et al. (2013) and Lucking et al. (2019), we measure technological proximity from patent-classification overlap using the DISCERN patent-firm link (Arora et al., 2024). Given these networks, the model's static equilibrium conditions recover firm-level knowledge capital from Compustat accounting data. The spillover intensity is estimated from the panel law of motion of recovered knowledge capital, with a tax-based IV robustness check confirming that the estimated spillover relationship remains positive. The R&D efficiency and depreciation parameters are jointly calibrated to aggregate R&D intensity and the observed-state expected economic growth rate. The model fits untargeted moments well: the correlation between model-implied and observed log sales is 0.97, the corresponding correlation for log R&D expenditure is 0.74, and model-implied knowledge growth evaluated at the 2010 state predicts realized 2010–2017 knowledge growth with a correlation of 0.579.

The estimated wedges are large and heterogeneous. At the observed 2017 knowledge-capital distribution, the median ratio of social to private marginal returns to R&D is 1.40, and social returns exceed private returns for 86.6% of firms with positive marginal returns, consistent with aggregate underprovision of R&D. The dispersion matters as much as the level: 36.2% of firms have ratios above 1.5, while the bottom tail includes firms for which the observed-state local diagnostic points toward taxing rather than subsidizing marginal R&D. An exact decomposition shows that each firm's wedge is primarily the net of two opposing forces: a positive non-producer-surplus component, reflecting surplus the firm creates but does not capture, and a negative rival-profit component operating through business stealing.

These wedges aggregate into sizable distortions. To separate dynamic R&D distortions from static product-market distortions, we compare the competitive equilibrium with constrained monopoly and constrained planner benchmarks that hold the competitive product-market allocation fixed and change only the dynamic R&D rule. Evaluated at the observed 2017 knowledge-capital distribution, the constrained planner raises total R&D to about 242% of the competitive level, increasing the expected economic growth rate by 0.29 percentage points and raising consumption-equivalent welfare by 4.36%. The constrained monopolist, who internalizes losses imposed on rival producers but not consumer surplus, instead cuts R&D roughly in half and lowers household welfare. Allowing production and R&D to be reallocated jointly amplifies both findings: the interior planner benchmark raises consumption-equivalent welfare by 14.54%, while full monopoly lowers it by 4.33%. Product-market misallocation and dynamic innovation incentives therefore reinforce each other.

The implementable policy benchmark is a uniform R&D subsidy financed by lump-sum taxes. On the subsidy grid, household welfare peaks at a rate of 34%. At that rate, the subsidy reproduces nearly all of the constrained planner's growth effect, increasing the expected economic growth rate by 0.292 percentage points versus the planner's 0.294, but it delivers only about one-eighth of the planner's welfare gain: 0.53% versus 4.36% in consumption-equivalent terms. The reason is firm-level allocation. A uniform subsidy lowers the common price of R&D but barely moves the distribution of R&D across firms toward the planner's: the distance between competitive and constrained-planner R&D shares is essentially unchanged under the subsidy, indeed slightly larger, and the reallocation the subsidy induces is negatively correlated with the reallocation the planner prescribes. In this economy, the expected economic growth rate is therefore a poor sufficient statistic for the welfare evaluation of innovation policy. The result identifies a limit of uniform instruments and motivates policies that condition on firms' positions in the rivalry and spillover networks.

This paper contributes to three strands of the literature. First, it contributes to work on innovation incentives, appropriability, business stealing, technology spillovers, and innovation networks (Arrow, 1962; Spence, 1984; d'Aspremont and Jacquemin, 1988; Aghion et al., 2005; Acemoglu and Akcigit, 2012; Peters, 2020; Akcigit and Ates, 2021, 2023; De Ridder, 2024). Closest to our mechanism, Bloom et al. (2013) show empirically that product-market proximity and technological proximity have opposite effects on innovation incentives, Cavenaile et al. (2026) develop an oligopolistic general-equilibrium Schumpeterian model with strategic interaction among a small number of firms, and Liu and Ma (2024) study how innovation networks shape the allocation of R&D across connected firms. Relative to this literature, we embed product-market rivalry and technology spillovers in a many-firm dynamic oligopoly disciplined by firm-level and firm-pair-level data, and we use the estimated model to quantify heterogeneous social-private R&D wedges and the limits of uniform innovation subsidies.

Second, the paper relates to the dynamic oligopoly literature in empirical industrial organization (Ericson and Pakes, 1995; Besanko and Doraszelski, 2004; Doraszelski and Pakes, 2007; Weintraub et al., 2008). This literature develops Markov perfect industry dynamics and methods for handling large state spaces. We differ by exploiting a linear-quadratic structure that makes a full-scale many-firm continuous-state R&D game tractable with empirically measured product-market and technology networks.

Third, the paper contributes to research that incorporates oligopolistic competition into macroeconomic models (Neary, 2003; Atkeson and Burstein, 2008; Azar and Vives, 2021; Pellegrino, 2025; Ederer and Pellegrino, 2025; Bustamante and Pellegrino, 2026). Our product-market environment builds on Pellegrino (2025) and Ederer and Pellegrino (2025). Closely related contemporaneous work by Bustamante and Pellegrino (2026) develops a Network-Q model of corporate investment in oligopolistic product markets using the same GHL demand system. Their focus is neoclassical capital investment, firm valuation, markups, and merger counterfactuals; we instead add a technology-spillover network and forward-looking R&D investment to study how product-market rivalry and knowledge spillovers jointly shape growth, welfare, and innovation policy.

The remainder of the paper proceeds as follows. Section 2 constructs an endogenous growth model incorporating product-market rivalry and technology-spillover networks. Section 3 describes the data and estimation strategy. Section 4 reports the wedge diagnostics, counterfactual allocations, and the uniform-subsidy evaluation. Section 5 concludes.

2 Model

We develop an endogenous growth model in which granular firms strategically interact in both product and technology spaces to examine the implications for firms' dynamic R&D investment, economic growth, and household welfare.

2.1 Environment

Our environment incorporates two networks: product-market rivalry and technology spillovers. There are n oligopolistic firms, indexed by $i \in \{1, 2, \ldots, n\}$; they choose output and dynamic R&D investment while interacting strategically through these networks. We interpret each firm as supplying one composite differentiated product, so firm and product indices coincide throughout. The set of firms, product-characteristic vectors, and technological-proximity matrix are fixed over time. Innovation changes firms' knowledge stocks, labor productivity, and product quality, but not their locations in product or technology space. Time is continuous and infinite, $t \in [0, \infty)$.

2.1.1 Demand

We use the Generalized Hedonic Linear (GHL) demand system of Pellegrino (2025) to model product-market rivalry among many oligopolistic firms. In reduced form, the final-good aggregator is

$$Y_t = \boldsymbol{q}_t^T \boldsymbol{b}_t - \frac{1}{2}\boldsymbol{q}_t^T \boldsymbol{\Sigma} \boldsymbol{q}_t \tag{1}$$

where $\boldsymbol{q}_t \in \mathbb{R}^n$ is the vector of product quantities and $\boldsymbol{b}_t \in \mathbb{R}^n$ is the vector of product-level qualities. The product-rivalry matrix is

$$\boldsymbol{\Sigma} \equiv \alpha \boldsymbol{S} + (1 - \alpha)\boldsymbol{I},$$

where $\boldsymbol{S}$ is the product-similarity matrix implied by the GHL characteristic structure, $\boldsymbol{I}$ is the identity matrix, and $\alpha \in [0, 1]$ maps product similarity into substitutability. We normalize $S_{ii} = 1$, so $\sigma_{ii} = 1$ for every firm, and for $i \neq j$, $\sigma_{ij} = \alpha S_{ij}$. From Equation (1), a larger σ_{ij} implies a larger quadratic penalty when q_i and q_j are both high, so σ_{ij} measures the strength of product-market rivalry between firms i and j. The GHL characteristic-level construction underlying $\boldsymbol{S}$ and the reduced-form aggregator in Equation (1) is collected in Section A.1.

2.1.2 Technology

Each firm has a linear production technology in labor:

$$q_{i,t} = a_{i,t} l_{i,t} \tag{2}$$

where $a_{i,t}$ is the labor productivity of firm i at time t and $l_{i,t}$ is the production labor employed by firm i at time t.

Firms possess their own knowledge capital and use it to improve labor productivity and product quality. As shown in Equation (1), $b_{i,t}$ is the value (in final-good units) generated by the first unit of product i, so we interpret $b_{i,t}$ as product quality. Firms allocate knowledge capital between labor productivity $a_{i,t}$ and product quality $b_{i,t}$:

$$\zeta a_{i,t} + b_{i,t} = z_{i,t} \tag{3}$$

where $\zeta > 0$ governs the trade-off between improving labor productivity and product quality[1].

We model the law of motion for knowledge capital, which rises through technology spillovers from technologically similar firms and through the firm's own R&D investment. Spillovers operate through source firms' knowledge stocks rather than directly through their contemporaneous R&D expenditures; a firm's R&D affects other firms only by raising its future knowledge stock. Let $\tilde{\mathbf{\Omega}} \in \mathbb{R}^{n \times n}$ denote recipient-normalized technology-spillover exposure:

$$\tilde{\mathbf{\Omega}} = \begin{bmatrix} \tilde{\omega}_{11} & \tilde{\omega}_{12} & \cdots & \tilde{\omega}_{1n} \\ \tilde{\omega}_{21} & \tilde{\omega}_{22} & \cdots & \tilde{\omega}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{\omega}_{n1} & \tilde{\omega}_{n2} & \cdots & \tilde{\omega}_{nn} \end{bmatrix}$$

We use a row-receiver convention: row i describes the sources of spillovers received by firm i. Thus, $\tilde{\omega}_{ij} \geq 0$ is the weight placed on source firm j in firm i's spillover exposure. We set the diagonal to zero, $\tilde{\omega}_{ii} = 0$, and rows with positive off-diagonal exposure are normalized so that $\sum_{j \neq i} \tilde{\omega}_{ij} = 1$. The effective spillover matrix is $\mathbf{\Omega} = \beta \tilde{\mathbf{\Omega}}$, where $\beta > 0$ scales these relative exposure weights into spillover intensities.

Let $\boldsymbol{x}_t \in \mathbb{R}^n$ denote the vector of R&D efforts:

$$\boldsymbol{x}_t \equiv \begin{bmatrix} x_{1,t} & x_{2,t} & \dots & x_{n,t} \end{bmatrix}^T,$$

and let $\boldsymbol{W}_t \in \mathbb{R}^n$ denote independent standard Wiener processes across firms:

$$\boldsymbol{W}_t \equiv \begin{bmatrix} W_{1,t} & W_{2,t} & \dots & W_{n,t} \end{bmatrix}^T$$

The control $x_{i,t}$ is an R&D effort that raises knowledge capital linearly, while the resource

[1] Assuming optimal allocation of knowledge capital between labor-productivity and quality improvements is necessary for a balanced-growth-path equilibrium. Similarly, Grossman et al. (2017) and Jones and Liu (2024) obtain a balanced growth path by assuming two productivity-enhancing technologies (capital and human capital; capital productivity and automation, respectively) and allocating resources between them.

cost of generating this effort is quadratic, as specified in the resource constraint below. Knowledge capital follows a multi-dimensional geometric Brownian motion:

$$d\boldsymbol{z}_t = \left((\boldsymbol{\Omega} - \delta \boldsymbol{I})\, \boldsymbol{z}_t + \mu \boldsymbol{x}_t\right) dt + \gamma\, \mathrm{diag}\left(\boldsymbol{z}_t\right) d\boldsymbol{W}_t \tag{4}$$

where μ, γ, and δ are positive scalars. The parameter μ governs the efficiency of R&D, γ the magnitude of shocks, and δ the depreciation rate of knowledge capital.

2.1.3 Resources and Household Preferences

We close the model with resource constraints and representative household preferences. The representative household supplies labor, owns firms, and consumes the final good. Production labor is supplied inelastically:

$$L = \sum_i l_{i,t}. \tag{5}$$

The final good is used for consumption and R&D investment. Following the endogenous growth literature (e.g. Acemoglu et al. 2016), we assume that the innovation elasticity for each firm is 0.5. Equivalently, the resource cost of firm-level R&D effort is quadratic:

$$C_t + \sum_i x_{i,t}^2 = Y_t$$

where C_t is consumption of the final good at time t. Finally, we assume that the representative household is risk-neutral and discounts the future at rate $\rho > 0$:

$$\mathbb{E}_0 \left[\int_0^\infty \exp\left(-\rho t\right) C_t \, dt \right]$$

2.2 Competitive Equilibrium

We characterize a linear Markov perfect equilibrium. The key state variable is the vector of knowledge capital $\boldsymbol{z}$. Given $\boldsymbol{z}$, the static product-market game determines quantities, technology allocations, prices, and the real wage. Firms then choose Markov R&D policies that affect future knowledge capital through Equation (4).

Formally, an equilibrium consists of Markov R&D policies $\{x_i(\boldsymbol{z})\}_{i=1}^n$, firm value functions $\{V^i(\boldsymbol{z})\}_{i=1}^n$, static allocations $\boldsymbol{q}(\boldsymbol{z})$, $\boldsymbol{a}(\boldsymbol{z})$, $\boldsymbol{b}(\boldsymbol{z})$, the real wage $w(\boldsymbol{z})/P$, and relative prices $\boldsymbol{p}(\boldsymbol{z})/P$. These objects satisfy final-good producer optimality, firms' static Cournot optimality, firms' dynamic R&D optimality, the knowledge law of motion in Equation (4), and labor-market clearing.

This structure lets us separate the problem into static and dynamic components. Conditional on the current knowledge-capital vector $\boldsymbol{z}_t$, the production and technology-allocation choices determine current quantities, prices, and gross profits but do not directly affect future knowledge capital. R&D affects current payoffs through its flow cost and affects future payoffs only through the law of motion for $\boldsymbol{z}_t$. Thus, we first solve the static product-market game as a mapping from knowledge capital into profits, and then use this mapping as the payoff flow in the dynamic R&D game.

2.2.1 Final-Good Producer's Problem

The final good is produced competitively. In each period, the final-good producer chooses the quantity of each product that maximizes profits, given prices.

$$\max_{\boldsymbol{q}_t} P_t Y_t - \boldsymbol{q}_t^T \boldsymbol{p}_t$$

where $\boldsymbol{p}_t$ is the vector of product prices at time t and P_t is the final-good price at time t. The final-good producer's profit maximization problem yields the following linear inverse demand function:

$$\frac{\boldsymbol{p}_t}{P_t} = \boldsymbol{b}_t - \boldsymbol{\Sigma} \boldsymbol{q}_t \tag{6}$$

2.2.2 Static Firm Problem

We assume that firms maximize real profits, so allocations are independent of the choice of numeraire (Azar and Vives, 2021). We characterize the interior static Cournot-Nash equilibrium in which each firm takes the real wage and rivals' quantities as given when choosing output and the allocation of knowledge capital between labor productivity and product quality. The static block maps the knowledge-capital vector into quantities:

$$\boldsymbol{q}_t = \boldsymbol{N} \boldsymbol{z}_t \tag{7}$$

where

$$\boldsymbol{N} \equiv \left(2\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma} + \boldsymbol{I}\right)^{-1},$$

and $\boldsymbol{J}$ is an $n \times n$ matrix whose entries are all one. The first term inside the inverse, $2\zeta \boldsymbol{J}/L$, captures the labor-cost channel, while $\boldsymbol{\Sigma}$ captures product-market rivalry. The additional $\boldsymbol{I}$

is the Cournot own-output strategic term: $\boldsymbol{\Sigma q}$ enters inverse demand, and the derivative of firm i's revenue with respect to q_i adds another $\sigma_{ii} q_i = q_i$ under the normalization $\sigma_{ii} = 1$. The derivation is in Section A.2.

Let $\pi_{i,t}$ denote firm i's real gross profit before subtracting R&D costs. In the interior Cournot equilibrium, the normalization $\sigma_{ii} = 1$ implies that static gross profit equals $q_{i,t}^2$. Hence the static payoff entering the dynamic R&D game is quadratic in knowledge capital:

$$\pi_{i,t} = q_{i,t}^2 = \boldsymbol{z}_t^T \boldsymbol{N}_i^T \boldsymbol{N}_i \boldsymbol{z}_t \equiv \boldsymbol{z}_t^T \boldsymbol{Q}_i \boldsymbol{z}_t$$

where $\boldsymbol{N}_i$ denotes the i-th row of $\boldsymbol{N}$, and $\boldsymbol{Q}_i \equiv \boldsymbol{N}_i^T \boldsymbol{N}_i$. This quadratic payoff representation is the static source of the model's linear-quadratic dynamic structure.

2.2.3 Dynamic R&D Game

Given a static production strategy profile, we consider a dynamic R&D game. In the dynamic game, given other players' strategies $\left\{x_{j,t}\right\}_{j \neq i, t \geq 0}$, firm i chooses R&D effort $\{x_{i,t}\}_{t \geq 0}$ to maximize the expected discounted present value of profits:

$$\max_{\left\{x_{i,t}\right\}_{t \geq 0}} \quad V^i\left(z_0\right) \equiv \mathbb{E}_0\left[\int_0^{\infty} \exp\left(-\rho t\right)\left\{\pi_{i,t} - x_{i,t}^2\right\} dt\right]$$

while the state evolves according to Equation (4). We use the fact that the interest rate is pinned down by the household's Euler equation, $r_t = \rho$. In firm i's HJB equation, x_i is the choice variable and the other components of $\boldsymbol{x}$ are evaluated at rivals' Markov strategies. The dynamic problem gives the following HJB equation:

$$\begin{aligned}\rho V^i\left(\boldsymbol{z}\right) = \max_{x_i} \Bigg\{ & \boldsymbol{z}^T \boldsymbol{Q}_i \boldsymbol{z} - x_i^2 + V_{\boldsymbol{z}}^i\left(\boldsymbol{z}\right)\left[\left(\boldsymbol{\Omega} - \delta \boldsymbol{I}\right) \boldsymbol{z} + \mu \boldsymbol{x}\right] \\ & + \frac{\gamma^2}{2} \operatorname{Tr}\left[\operatorname{diag}\left(\boldsymbol{z}\right) V_{\boldsymbol{zz}}^i\left(\boldsymbol{z}\right) \operatorname{diag}\left(\boldsymbol{z}\right)\right] \Bigg\}. \end{aligned} \tag{8}$$

The first term on the right-hand side of the HJB equation represents gross profit. The second term represents R&D costs. The first two terms therefore represent the firm's net instantaneous profit. The third term reflects the impact of changes in knowledge capital due to technology spillovers and R&D on firm i's value. The fourth term represents the expected effect of all firms' idiosyncratic shocks on firm i's value.

This dynamic game is a linear-quadratic differential game, which significantly simplifies the analysis. The simplification follows from a standard guess-and-verify argument. Guess that the value can be expressed as a quadratic form of knowledge capital:

$$V^i(\boldsymbol{z}) = \boldsymbol{z}^T \boldsymbol{X}^i \boldsymbol{z}. \tag{9}$$

Because $\boldsymbol{z}^T \boldsymbol{X}^i \boldsymbol{z}$ depends only on the symmetric part of $\boldsymbol{X}^i$, we take $\boldsymbol{X}^i$ to be symmetric without loss of generality. The derivative algebra is collected in Section A.4.

In a linear-quadratic differential game, firms' strategies are expressed as linear functions of the state variables. Let $\boldsymbol{X}_i^i$ denote the i-th column of $\boldsymbol{X}^i$. The first-order condition with respect to x_i gives the linear Markov strategy

$$x_i = \left(\mu \boldsymbol{X}_i^i\right)^T \boldsymbol{z}. \tag{10}$$

This rule is the unconstrained linear-quadratic feedback policy, which we use as a tractable interior benchmark.

Given firms' linear strategies, the law of motion for knowledge capital $\boldsymbol{z}_t$ can be rewritten as a geometric Brownian motion with constant coefficients. Define $\widetilde{\boldsymbol{X}}$ as the matrix obtained by stacking $\boldsymbol{X}_i^i$:

$$\widetilde{\boldsymbol{X}} \equiv \begin{bmatrix} \boldsymbol{X}_1^1 & \cdots & \boldsymbol{X}_n^n \end{bmatrix}^T$$

Then, the law of motion is rewritten as

$$d\boldsymbol{z}_t = \boldsymbol{\Phi} \boldsymbol{z}_t dt + \gamma \operatorname{diag}(\boldsymbol{z}_t) d\boldsymbol{W}_t$$

where the drift matrix $\boldsymbol{\Phi}$ is defined by

$$\boldsymbol{\Phi} \equiv \begin{bmatrix} \phi_{11} & \phi_{12} & \cdots & \phi_{1n} \\ \phi_{21} & \phi_{22} & \cdots & \phi_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{n1} & \phi_{n2} & \cdots & \phi_{nn} \end{bmatrix} \equiv \boldsymbol{\Omega} - \delta \boldsymbol{I} + \mu^2 \widetilde{\boldsymbol{X}}.$$

We refer to the matrix $\boldsymbol{\Phi}$ as the *technology transition matrix*. It consists of the exogenous technology-spillover matrix $\boldsymbol{\Omega}$, depreciation $-\delta \boldsymbol{I}$, and the endogenous R&D matrix $\mu^2 \widetilde{\boldsymbol{X}}$.

For any square matrix $\boldsymbol{A}$, let $\mathcal{D}(\boldsymbol{A})$ denote the diagonal matrix formed from the diagonal entries of $\boldsymbol{A}$. With independent multiplicative shocks, the diffusion term in Equation (8) contributes $\gamma^2 \mathcal{D}(\boldsymbol{X}^i)$ to the quadratic coefficient. The HJB equation can be transformed into an algebraic Riccati equation, making the many-firm dynamic oligopoly problem computationally feasible. Substituting the quadratic value function and the linear policy

Equation (10) into Equation (8), and collecting quadratic terms, we obtain stacked algebraic Riccati equations

$$0 = \boldsymbol{Q}_i - \mu^2 \boldsymbol{X}_i^i \left(\boldsymbol{X}_i^i\right)^T + \left(\boldsymbol{\Phi} - \frac{\rho}{2}\boldsymbol{I}\right)^T \boldsymbol{X}^i + \boldsymbol{X}^i \left(\boldsymbol{\Phi} - \frac{\rho}{2}\boldsymbol{I}\right) + \gamma^2 \mathcal{D}\left(\boldsymbol{X}^i\right) \tag{11}$$

for $i = 1, 2, \ldots, n$. Note that the terms $\mu^2 \boldsymbol{X}_i^i \left(\boldsymbol{X}_i^i\right)^T$ and $\boldsymbol{\Phi}$ depend on the $\boldsymbol{X}^j$ of other firms $j \neq i$. Therefore, to solve for $\boldsymbol{X}^i$, we need to simultaneously solve the algebraic Riccati equations for all firms.

We now define stability of the solution.

Definition 1. The solution of the stacked algebraic Riccati Equation (11) is *stable* if it delivers finite expected discounted payoffs and satisfies the transversality condition

$$\lim_{T\to\infty} \mathbb{E}_0 \left[\exp(-\rho T) V^i\left(\boldsymbol{z}_T\right)\right] = 0$$

for all firms i and admissible initial states. The quantitative baseline sets $\gamma = 0$. In that deterministic case, stability is equivalent to all eigenvalues of $\boldsymbol{\Phi} - \frac{\rho}{2}\boldsymbol{I}$ having negative real parts. With multiplicative shocks, the same definition requires the discounted second moments induced by the closed-loop process to be finite.

This stability requirement corresponds to the standard growth-theory condition that the discount rate be sufficiently large for infinite-horizon utility to be finite.

In the following, we assume the existence of such a solution.

Assumption 1. *There exists a stabilizing solution* $\left(\boldsymbol{X}^1, \boldsymbol{X}^2, \ldots, \boldsymbol{X}^n\right)$ *of the stacked algebraic Riccati Equation* (11)*.*

We next use the closed-loop technology transition matrix $\boldsymbol{\Phi}$ to characterize the economy's deterministic balanced growth path.

2.2.4 Deterministic Balanced Growth Path

Set $\gamma = 0$ and consider the closed-loop technology system $\dot{\boldsymbol{z}}_t = \boldsymbol{\Phi}\boldsymbol{z}_t$. If $\boldsymbol{\Phi}$ has a simple real dominant eigenvalue $g > 0$ with a strictly positive right eigenvector $\bar{\boldsymbol{z}}$, then $\boldsymbol{z}_t = e^{gt} c \bar{\boldsymbol{z}}$ is a balanced growth path. Thus, the long-run growth rate is the dominant eigenvalue of the technology transition matrix, and the balanced-growth-path cross-section is the associated eigenvector.

Because the static mapping and R&D policies are linear in $\boldsymbol{z}_t$, quantities and R&D grow at rate g along this path. Output, consumption, and profits grow at rate $2g$ because they are quadratic in knowledge capital. The formal spectral condition, proof, growth-rate table, and partial-equilibrium diagram are collected in Section A.6. The quantitative analysis below is evaluated at the observed 2017 knowledge-capital distribution, so this balanced-growth-path characterization is a long-run implication of the solved dynamic system rather than the object used for the main counterfactual comparisons.

2.2.5 Household Welfare

Once the equilibrium is solved, representative-household welfare can be computed for any initial knowledge-capital distribution with minimal additional cost. Let $V(\boldsymbol{z}_0)$ denote household welfare in the competitive equilibrium:

$$V(\boldsymbol{z}_0) = \mathbb{E}_0 \left[\int_0^\infty \exp(-\rho t) \left\{ Y_t - \sum_i x_{i,t}^2 \right\} dt \middle| \boldsymbol{z}_0 \right].$$

Total output is

$$Y_t = \boldsymbol{z}_t^T \boldsymbol{Q} \boldsymbol{z}_t$$

where

$$\boldsymbol{Q} = \boldsymbol{N}^T \left(\frac{\zeta}{L} \boldsymbol{J} + \frac{1}{2} \boldsymbol{\Sigma} + \boldsymbol{I} \right) \boldsymbol{N}$$

(see Section A.3 for the derivation). After firms' equilibrium policies are fixed, the remaining valuation equation contains no controls. Because the closed-loop process and the payoff are linear-quadratic, household welfare has the form $V(\boldsymbol{z}) = \boldsymbol{z}^T \boldsymbol{X} \boldsymbol{z}$. Substituting this quadratic form into the valuation equation gives the Lyapunov equation

$$\begin{aligned} 0 = {} & \boldsymbol{Q} - \mu^2 \widetilde{\boldsymbol{X}}^T \widetilde{\boldsymbol{X}} + \boldsymbol{X} \left(\boldsymbol{\Phi} - \frac{\rho}{2} \boldsymbol{I} \right) + \left(\boldsymbol{\Phi} - \frac{\rho}{2} \boldsymbol{I} \right)^T \boldsymbol{X} \\ & + \gamma^2 \mathcal{D}(\boldsymbol{X}) \end{aligned} \tag{12}$$

Therefore, $\boldsymbol{X}$ is obtained as the solution of Equation (12), and welfare in the competitive equilibrium is $V(\boldsymbol{z}_0) = \boldsymbol{z}_0^T \boldsymbol{X} \boldsymbol{z}_0$. The valuation equation and producer-value analogue are derived in Section A.5.

2.2.6 Producer Value

We also characterize aggregate producer value in the competitive equilibrium. This object provides a measure of producers' dynamic value that can be compared with household

Table 1: Counterfactual Allocations in the Model

Scenario	Static mapping	Dynamic R&D controller	Flow payoff before R&D cost
CC	$\boldsymbol{q} = \boldsymbol{N}\boldsymbol{z}$	Firm i maximizes own value	$\boldsymbol{z}^T\boldsymbol{Q}_i\boldsymbol{z}$
CS	$\boldsymbol{q} = \boldsymbol{N}\boldsymbol{z}$	Constrained planner maximizes household welfare	$\boldsymbol{z}^T\boldsymbol{Q}\boldsymbol{z}$
CM	$\boldsymbol{q} = \boldsymbol{N}\boldsymbol{z}$	Monopolist maximizes aggregate producer value	$\boldsymbol{z}^T\boldsymbol{N}^T\boldsymbol{N}\boldsymbol{z}$
SS	$\boldsymbol{q} = \boldsymbol{N}^*\boldsymbol{z}$	Social planner maximizes household welfare	$\frac{1}{2}\boldsymbol{z}^T\boldsymbol{N}^*\boldsymbol{z}$
MM	$\boldsymbol{q} = \boldsymbol{N}^M\boldsymbol{z}$	Monopolist maximizes aggregate producer value	$\boldsymbol{z}^T\boldsymbol{P}^M\boldsymbol{z}$

Notes: Each dynamic objective subtracts the resource cost $\boldsymbol{x}^T\boldsymbol{x}$ from the flow payoff shown in the last column. The constrained scenarios CS and CM keep the competitive static mapping fixed and change only the dynamic R&D controller.

welfare and with the counterfactual allocations below. Aggregate real profit is

$$\sum_i \pi_{i,t} = \boldsymbol{q}_t^T\boldsymbol{q}_t = \boldsymbol{z}_t^T\boldsymbol{N}^T\boldsymbol{N}\boldsymbol{z}_t.$$

Relative to the household-welfare calculation in Section 2.2.5, aggregate producer value is obtained by solving a modified version of Equation (12) in which $\boldsymbol{Q}$ is replaced with $\boldsymbol{N}^T\boldsymbol{N}$. Denoting the solution by $\boldsymbol{X}^P$, aggregate producer value is $V^P(\boldsymbol{z}_0) = \boldsymbol{z}_0^T\boldsymbol{X}^P\boldsymbol{z}_0$.

2.3 Planner, Monopoly, and Constrained Benchmarks

We examine counterfactual allocations and compare them with the competitive equilibrium. The purpose of these allocations is to decompose distortions, not to claim that each allocation is directly implementable as a policy. Each counterfactual varies two margins: the static product-market allocation and the dynamic R&D rule. The scenario labels are two-letter codes: the first letter denotes the static product-market allocation, and the second denotes the dynamic R&D rule. The label C denotes competition, S the social planner, and M monopoly. Table 1 summarizes the allocations defined in this section.

We first analyze the allocation chosen by a benevolent social planner who aims to maximize household welfare. We then consider the monopoly case, in which a monopolist determines production and R&D to maximize aggregate producer value. Finally, we define constrained counterfactuals that combine the competitive static product-market allocation with either planner or monopoly control of R&D.

2.3.1 Social Planner

To analyze the socially optimal allocation of R&D and its implications for the expected economic growth rate and household welfare, we formulate the social planner's problem. This full social-planner allocation corresponds to scenario SS in Table 1. The social planner aims to maximize the representative household's lifetime utility, subject to technological constraints. As in the competitive equilibrium, the social planner's problem can be separated into a static problem and a dynamic problem.

Static Problem First, we characterize the static allocation in the social planner's problem. The social planner maximizes final-good output given the vector of knowledge capital $\boldsymbol{z}_t$ in each period:

$$\max_{\boldsymbol{a}_t, \boldsymbol{b}_t, \boldsymbol{l}_t, \boldsymbol{q}_t} Y_t = \boldsymbol{q}_t^T \boldsymbol{b}_t - \frac{1}{2} \boldsymbol{q}_t^T \boldsymbol{\Sigma} \boldsymbol{q}_t$$

subject to

$$\begin{aligned} q_{i,t} &= a_{i,t} l_{i,t} \quad i = \{1, 2, \ldots, n\} \\ z_{i,t} &= \zeta a_{i,t} + b_{i,t} \quad i = \{1, 2, \ldots, n\} \\ L &= \sum_i l_{i,t} \end{aligned} \tag{13}$$

The first constraint represents the production technology of each firm, the second governs the allocation of each firm's knowledge capital between labor productivity and product quality, and the third ensures the labor market clears. Solving the planner's static problem yields quantities as linear functions of knowledge capital:

$$\boldsymbol{q}_t^* = \boldsymbol{N}^* \boldsymbol{z}_t$$

and a quadratic representation of output in knowledge capital

$$Y_t^* = \frac{1}{2} \boldsymbol{z}_t^T \boldsymbol{N}^* \boldsymbol{z}_t$$

where

$$\boldsymbol{N}^* = \left(2 \frac{\zeta}{L} \boldsymbol{J} + \boldsymbol{\Sigma} \right)^{-1}$$

(see Section B.1).

Dynamic Problem Next, we characterize the dynamic allocation in the social planner's problem. Given the static planner output coefficient, the planner chooses R&D to maximize

$$V^{SS}(\boldsymbol{z}_0) \equiv \max_{\{\boldsymbol{x}_t\}_{t\geq 0}} \mathbb{E}_0\left[\int_0^\infty \exp(-\rho t)\left\{\frac{1}{2}\boldsymbol{z}_t^T \boldsymbol{N}^* \boldsymbol{z}_t - \boldsymbol{x}_t^T \boldsymbol{x}_t\right\} dt\right],$$

subject to the knowledge accumulation law in Equation (4). This is again a linear-quadratic control problem, so the optimal R&D rule is linear in knowledge capital:

$$\boldsymbol{x} = \mu \boldsymbol{X}^{SS} \boldsymbol{z} \tag{14}$$

and the induced law of motion can be written as

$$d\boldsymbol{z}_t = \boldsymbol{\Phi}^{SS} \boldsymbol{z}_t dt + \gamma\, \mathrm{diag}(\boldsymbol{z}_t)\, d\boldsymbol{W}_t$$

where

$$\boldsymbol{\Phi}^{SS} \equiv \boldsymbol{\Omega} + \mu^2 \boldsymbol{X}^{SS} - \delta \boldsymbol{I}$$

Substituting the quadratic planner value and the linear R&D rule into the planner's HJB yields a single algebraic Riccati equation:

$$\begin{aligned} 0 = \frac{1}{2}\boldsymbol{N}^* - \mu^2\left(\boldsymbol{X}^{SS}\right)^2 + \boldsymbol{X}^{SS}\left(\boldsymbol{\Phi}^{SS} - \frac{\rho}{2}\boldsymbol{I}\right) + \left(\boldsymbol{\Phi}^{SS} - \frac{\rho}{2}\boldsymbol{I}\right)^T \boldsymbol{X}^{SS} \\ + \gamma^2 \mathcal{D}\left(\boldsymbol{X}^{SS}\right) \end{aligned} \tag{15}$$

Therefore, $\boldsymbol{X}^{SS}$ is obtained as the solution of Equation (15), and household welfare under the optimal allocation is $V^{SS}(\boldsymbol{z}_0) = \boldsymbol{z}_0^T \boldsymbol{X}^{SS} \boldsymbol{z}_0$. The HJB and derivative steps are in Section B.2. Unlike the competitive equilibrium, which requires one Riccati equation per firm, the social-planner problem requires only a single Riccati equation.

2.3.2 Monopoly

As a second counterfactual scenario, we consider the monopoly case, in which a single monopolist controls both product-market production and R&D. This full monopoly allocation corresponds to scenario MM. The monopolist determines all production and R&D investments to maximize aggregate producer value. In the static product-market block, the monopolist chooses $\boldsymbol{a}_t$, $\boldsymbol{b}_t$, and $\boldsymbol{q}_t$ to maximize aggregate real profit

$$\Pi_t \equiv \frac{\boldsymbol{q}_t^T \boldsymbol{p}_t}{P_t} - \frac{w_t}{P_t} \sum_i l_{i,t}.$$

Solving the monopolist's static problem yields

$$\boldsymbol{q}_t^M = \boldsymbol{N}^M \boldsymbol{z}_t, \qquad \boldsymbol{N}^M \equiv \frac{1}{2} \left(\frac{\zeta}{L} \boldsymbol{J} + \boldsymbol{\Sigma} \right)^{-1},$$

and aggregate firm profit can be written as

$$\Pi_t^M = \boldsymbol{z}_t^T \boldsymbol{P}^M \boldsymbol{z}_t, \qquad \boldsymbol{P}^M \equiv (\boldsymbol{N}^M)^T \boldsymbol{\Sigma} \boldsymbol{N}^M.$$

The superscript P denotes producer value, distinguishing the monopolist's objective from household welfare under the induced policy. The monopolist's dynamic producer-value problem is therefore

$$V^{P,MM}(\boldsymbol{z}_0) \equiv \max_{\{\boldsymbol{x}_t\}_{t \geq 0}} \mathbb{E}_0 \left[\int_0^\infty \exp(-\rho t) \left\{ \boldsymbol{z}_t^T \boldsymbol{P}^M \boldsymbol{z}_t - \boldsymbol{x}_t^T \boldsymbol{x}_t \right\} dt \right],$$

subject to the knowledge accumulation law. This problem is solved by replacing $\boldsymbol{N}^*/2$ with $\boldsymbol{P}^M$ in Equation (15). Let $\boldsymbol{X}^{P,MM}$ denote the solution. The monopoly R&D policy is then $\boldsymbol{x}_t = \mu \boldsymbol{X}^{P,MM} \boldsymbol{z}_t$. Because the monopolist maximizes producer value rather than household welfare, the household welfare reported for MM is computed separately from a Lyapunov equation using the monopoly output matrix $\boldsymbol{Q}^M$ and the technology transition induced by $\boldsymbol{X}^{P,MM}$. For the derivation of $\boldsymbol{N}^M$, $\boldsymbol{P}^M$, $\boldsymbol{Q}^M$, and the welfare equation, see Section B.3.

2.3.3 Separating Product-Market and R&D Control

We also consider counterfactual scenarios in which control of production differs from control of R&D. These constrained counterfactuals keep the static competitive mapping $\boldsymbol{q}_t = \boldsymbol{N} \boldsymbol{z}_t$ fixed and change only the dynamic R&D controller. They are useful because they separate innovation incentives from contemporaneous product-market reallocation.

First, consider scenario CS, in which firms competitively determine production quantities, and given this static competitive equilibrium, a benevolent constrained planner then selects the optimal allocation of R&D. Starting from an initial knowledge-capital distribution z_0, the constrained planner maximizes the value

$$V^{CS}(z_0) \equiv \max_{\{x_t\}_{t\geq 0}} \mathbb{E}_0 \left[\int_0^\infty \exp(-\rho t) \left\{ z_t^T Q z_t - x_t^T x_t \right\} dt \right],$$

subject to

$$dz_t = \left((\Omega - \delta I) z_t + \mu x_t \right) dt + \gamma \operatorname{diag}(z_t) dW_t.$$

Unlike in Section 2.3.1, the constrained planner's gross instantaneous return is $z_t^T Q z_t$ rather than $\frac{1}{2} z_t^T N^* z_t$. Therefore, we obtain the constrained planner's allocation by solving a modified version of Equation (15), in which $N^*/2$ is replaced with Q. Denoting the solution of the modified Riccati equation by X^{CS}, household welfare under the constrained-planner allocation is given by $V^{CS}(z_0) = z_0^T X^{CS} z_0$.

The constrained monopolist in scenario CM is defined analogously, except that the dynamic controller maximizes aggregate producer value rather than household welfare. Given the competitive static product-market allocation, aggregate real profit is $z_t^T N^T N z_t$. Hence the constrained monopolist chooses R&D to maximize

$$V^{P,CM}(z_0) \equiv \max_{\{x_t\}_{t\geq 0}} \mathbb{E}_0 \left[\int_0^\infty \exp(-\rho t) \left\{ z_t^T N^T N z_t - x_t^T x_t \right\} dt \right],$$

subject to the same knowledge accumulation law. This allocation is obtained by replacing $N^*/2$ with $N^T N$ in Equation (15). Let $X^{P,CM}$ denote the resulting producer-value coefficient, so the induced R&D rule is $x_t = \mu X^{P,CM} z_t$. Household welfare under this constrained monopoly R&D policy is then computed from the same Lyapunov equation used in Section 2.2.5, with the competitive output matrix Q and the technology transition induced by $X^{P,CM}$. Denoting the resulting welfare coefficient by X^{CM}, household welfare is $V^{CM}(z_0) = z_0^T X^{CM} z_0$.

In the numerical analysis, we also compare these planner and monopoly allocation benchmarks with a uniform R&D subsidy. The subsidy changes firms' private R&D costs while leaving the static competitive product-market allocation fixed.

3 Data and Estimation

We map the model's product-market and technology networks, knowledge capital, and dynamic parameters to observed data. We use a 1989–2019 panel to construct annual networks, recover firm-level knowledge capital, and estimate technology spillovers. The quantitative baseline and counterfactuals are solved at the observed 2017 cross section of

Table 2: Sample Selection Criteria

Requirement	Role in the empirical implementation
Compustat coverage	Provides accounting variables used to measure profits, sales, costs, and R&D.
Hoberg-Phillips product-similarity data	Provides the empirical product-similarity matrix $\boldsymbol{S}$ used to construct $\boldsymbol{\Sigma}$.
DISCERN-linked patent records	Links public firms to patents for technology portfolios.
Positive gross profits	Ensures that $q_i = \sqrt{\pi_i}$ is well defined under $\pi_i = q_i^2$.
At least ten patents in a five-year window	Ensures nontrivial linked patent portfolios for measuring technology shares.

757 firms.

3.1 Sample Construction and Empirical Mapping

For each year, we keep firms that satisfy the criteria in Table 2. The resulting annual samples range from 418 firms in 1989 to 802 firms in 1998, and the 2017 quantitative baseline contains 757 firms. Table 6 summarizes the estimated parameters and their sources.

Table 3 summarizes the empirical procedure: selecting the sample, constructing the two networks, recovering firm-level quantities and knowledge capital from Compustat variables, and estimating or calibrating the remaining dynamic parameters.

3.2 Product-Market Proximity and Substitutability

Following Pellegrino (2025), we measure the product-similarity matrix $\boldsymbol{S}$ using the dataset developed by Hoberg and Phillips (2016). This measure is based on the cosine similarity of the words found in the Business Description sections of 10-K filings. Hoberg and Phillips (2016) construct a vocabulary comprising 61,146 words that firms use to describe the features and characteristics of their products. For each firm i, they generate a vector of word frequencies, with each element representing the number of times a specific word appears in the firm's product description. They show that cosine similarity effectively identifies industry groupings and predicts competitive relationships between firms, outperforming other industry classification methods. This empirical matrix is the model's $\boldsymbol{S}$, the counterpart of $\boldsymbol{\Psi}^T\boldsymbol{\Psi}$ in the GHL microfoundation in Section A.1. The numerical analysis then constructs the product-rivalry matrix $\boldsymbol{\Sigma}(\alpha) = \alpha\boldsymbol{S} + (1-\alpha)\boldsymbol{I}$ before solving the static and dynamic blocks, so for $i \neq j$, $\sigma_{ij} = \alpha S_{ij}$, while the model normalization imposes $S_{ii} = \sigma_{ii} = 1$.

Table 3: Mapping Data to Model Objects

Model object	Data source or empirical object	Mapping to model
$\boldsymbol{S}, \boldsymbol{\Sigma}(\alpha)$	Hoberg-Phillips product similarity	Hoberg cosine similarities measure $\boldsymbol{S}$; the solver uses $\boldsymbol{\Sigma} = \alpha \boldsymbol{S} + (1-\alpha)\boldsymbol{I}$, with diagonal entries normalized to one.
$\tilde{\boldsymbol{\Omega}}$	DISCERN patents and CPC groups	Jaffe overlaps are computed from five-year patent shares; the diagonal is set to zero and rows are normalized by recipient.
$\boldsymbol{q}$	Gross profits, revt − cogs	The static first-order condition implies $\pi_i = q_i^2$, so quantities are recovered as $q_i = \sqrt{\pi_i}$.
ζ/L	Aggregate COGS and recovered quantities	The aggregate production-cost condition identifies the ratio ζ/L.
$\boldsymbol{z}$	Recovered $\boldsymbol{q}$, $\boldsymbol{\Sigma}$, and ζ/L	The static equilibrium condition recovers $\boldsymbol{z}_t = (2(\zeta/L)\boldsymbol{J} + \boldsymbol{\Sigma} + \boldsymbol{I})\boldsymbol{q}_t$.
β	Panel growth of recovered knowledge capital and R&D controls	Spillover intensity is estimated from the knowledge-capital law of motion; observed R&D enters only as a regression control and instrument source.
μ and δ	R&D expenditure and expected economic growth rate moments	Own-R&D productivity and depreciation are jointly calibrated in the 2017 baseline; R&D effort is endogenous in the model.

Notes: Nominal accounting variables are deflated using the GDP deflator. The effective spillover matrix used in the law of motion is $\boldsymbol{\Omega} = \beta\tilde{\boldsymbol{\Omega}}$. R&D effort is endogenous in the model; observed R&D expenditures are used for the law-of-motion regression and calibration moments. The recovered 2017 knowledge-capital vector is used as the initial state in the numerical analysis.

We take α, which governs the degree of horizontal differentiation across goods, from Pellegrino (2025). Pellegrino (2025) estimates α by matching the inverse cross-price elasticity of demand reported by Nevo (2001), as this parameter links product cosine similarity to the magnitude of cross-price elasticities. Pellegrino (2025) shows that the price elasticities implied by the demand system are well aligned with those estimated in the industrial organization literature. Specifically, Pellegrino (2025) compares them to estimated price elasticities of demand in the markets for automobiles (Berry et al., 1995), ready-to-eat cereal (Nevo, 2001), and computers (Goeree, 2008). Table 2 in Pellegrino (2025) shows that the own-price elasticities for each firm and the cross-price elasticities between firms implied by the estimated framework are similar to values estimated in the literature.

3.3 Technological Proximity

For technological proximity, we use patent data and follow Bloom et al. (2013) and Lucking et al. (2019) to construct measures of technological proximity between firms. Specifically, we use the DISCERN 2 dataset (Arora et al., 2024), which links data on U.S. publicly listed firms from the Compustat database to their patents. Following Bloom et al. (2013), we calculate technological proximity between firms using the Jaffe measure based on Cooperative Patent Classification groups. Patent shares are computed using patents granted from $t-2$ through $t+2$ around each sample year t. We use this centered grant-year window because patent grants are dated with delay relative to the underlying inventive activity: in the DISCERN-linked patents with filing years observed, the median lag between filing year and grant year is three years, and the mean lag is 3.09 years. The window therefore smooths sparse annual patent classifications and partially offsets grant-date delay; it is used to measure firms' persistent technological locations rather than contemporaneous R&D shocks.

Let T_i denote the vector of firm i's patent shares across technology classes. The raw Jaffe overlap (Jaffe, 1986) between firms i and j is

$$\hat{\omega}_{ij} = \frac{T_i T_j'}{\left(T_i T_i'\right)^{1/2} \left(T_j T_j'\right)^{1/2}}.$$

We measure T_i using Cooperative Patent Classification groups and construct firm-pair overlaps from the normalized technology vectors.

As in Liu and Ma (2024), we convert these raw overlaps into recipient-normalized spillover weights. We set the diagonal entries to zero and normalize each recipient firm's row:

$$\tilde{\omega}_{ij} = \frac{\hat{\omega}_{ij}}{\sum_{k \neq i} \hat{\omega}_{ik}}, \qquad j \neq i.$$

Rows with no off-diagonal overlap are left at zero. This normalization makes β capture the common intensity of spillovers received by a firm, while $\tilde{\omega}_{ij}$ determines how that exposure is allocated across source firms. Although the raw Jaffe overlap is symmetric, recipient normalization generally makes $\tilde{\mathbf{\Omega}}$ asymmetric. This is the empirical counterpart of the model's row-receiver convention: $\tilde{\omega}_{ij}$ is the share of firm i's spillover exposure accounted for by source firm j, and the effective spillover matrix used in the law of motion is $\mathbf{\Omega} = \beta \tilde{\mathbf{\Omega}}$.

3.4 Knowledge Capital Distribution

Next, we recover firm-level quantities, the ratio ζ/L, and knowledge capital z_t from Compustat accounting variables. The recovery proceeds in three steps.

First, we measure firm-level gross profit $\pi_{i,t}$ as total revenue (revt) minus cost of goods sold (cogs), with nominal variables deflated using the GDP deflator. In the model, the static Cournot first-order condition implies $\pi_{i,t} = q_{i,t}^2$. This identity is the reason we require positive gross profits in the sample: it lets us recover quantities as $q_{i,t} = \sqrt{\pi_{i,t}}$.

Second, we identify ζ/L from aggregate production costs. We interpret Compustat cost of goods sold as the empirical counterpart of variable production costs in the static block. Under this interpretation, the model implies

$$\frac{\zeta}{L}\boldsymbol{q}_t^T \boldsymbol{J} \boldsymbol{q}_t = \text{total production cost at time } t.$$

We calculate total production cost as the sum of cogs for all firms in the sample. Since only the ratio ζ/L enters the static equilibrium, ζ and L cannot be separately identified.

Third, we recover the knowledge-capital vector from the static equilibrium condition. Given the recovered quantities and product-rivalry matrix, Equation (7) implies

$$\boldsymbol{z}_t = \left(2\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma} + \boldsymbol{I}\right)\boldsymbol{q}_t.$$

Because $\boldsymbol{\Sigma}$ has unit diagonal in the model, the diagonal of this recovery equation contains two identity terms: one from the unit diagonal of $\boldsymbol{\Sigma}$ and one from the Cournot own-output term in the static first-order condition. Equivalently, the empirical implementation adds $2\boldsymbol{I}$ to the stored off-diagonal product-proximity matrix. The recovered $\boldsymbol{z}_t$ is then used to estimate the law of motion and, for 2017, as the initial distribution in the numerical analysis.

3.5 Knowledge-Capital Law of Motion

In the empirical implementation, $x_{i,t}$ denotes R&D effort, measured as the square root of Compustat R&D expenditure (xrd), so that $x_{i,t}^2$ corresponds to observed R&D expenditure, consistent with the model's resource constraint. We estimate the spillover intensity β from panel variation in recovered knowledge capital. Dividing the drift of the model's knowledge law of motion by $z_{i,t}$ shows that spillovers affect proportional knowledge growth through $\sum_{j\neq i}\tilde{\omega}_{ij,t}z_{j,t}/z_{i,t}$, while own R&D enters through $x_{i,t}/z_{i,t}$ and depreciation enters as a common term. The depreciation component is absorbed by fixed effects and residual common drift in the panel specification, giving the annual-difference analogue:

$$\log z_{i,t+1} - \log z_{i,t} = \beta \sum_{j\neq i} \tilde{\omega}_{ij,t}\frac{z_{j,t}}{z_{i,t}} + \text{Controls}_{i,t} + \epsilon_{i,t} \tag{16}$$

where $z_{i,t}$ denotes firm i's knowledge capital at time t, and $\tilde{\omega}_{ij,t}$ measures the recipient-normalized technological proximity between firms i and j at time t. The controls include year and firm fixed effects, log knowledge capital $(\log z_{i,t})$, and R&D-effort intensity

Table 4: Estimates of the Law of Motion for Knowledge Capital

Dependent variable	(1) OLS baseline $\Delta \log z_{i,t}$	(2) OLS own R&D $\Delta \log z_{i,t}$	(3) IV robustness $\Delta \log z_{i,t}$
Spillover exposure	0.026**	0.024**	0.073*
	(0.010)	(0.010)	(0.041)
$\log z_{i,t}$	−0.132***	−0.162***	−0.088**
	(0.013)	(0.013)	(0.038)
$\frac{x_{i,t}}{z_{i,t}}$		0.514***	
		(0.063)	
Year fixed effects	✓	✓	✓
Firm fixed effects	✓	✓	✓
IV			✓
IV first-stage F-statistic			279.765
No. of observations	14,442	14,442	14,442
R^2	0.713	0.716	–
Within R^2	0.072	0.080	–

Notes: The dependent variable is $\Delta \log z_{i,t}$. Spillover exposure is $\sum_{j \neq i} \tilde{\omega}_{ij,t} z_{j,t}/z_{i,t}$, with recipient-normalized technological proximity $\tilde{\omega}_{ij,t}$. All specifications include firm and year fixed effects; column (2) controls for R&D-effort intensity, and column (3) is an IV robustness check that instruments spillover exposure with the tax-driven source-firm exposure described in the text. Standard errors are clustered by year-by-4-digit-NAICS cells. Regression variables are trimmed at the 0.5% tails within year, and singleton fixed-effect observations are dropped. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

($x_{i,t}/z_{i,t}$). Own R&D is included as a conditioning variable, not as the source of structural identification for μ; we calibrate μ and δ below using aggregate moments in the 2017 baseline. Because the spillover regressor divides source-firm knowledge exposure by recipient knowledge, measurement error or scale misspecification in recovered $z_{i,t}$ can mechanically affect the regressor; we therefore control directly for $\log z_{i,t}$ and use tax-driven source-firm exposure as an IV robustness check that shifts the numerator rather than the recipient denominator.

Table 4 reports estimates of Equation (16). In column (1), the coefficient on technological-proximity-weighted relative knowledge capital, $\sum_{j \neq i} \tilde{\omega}_{ij,t} z_{j,t}/z_{i,t}$, is positive and statistically significant, indicating that knowledge capital from technologically proximate firms contributes to a firm's own knowledge growth. Column (2) adds firms' R&D-effort intensity, $x_{i,t}/z_{i,t}$, as a control. The spillover coefficient remains positive and significant.

Column (3) provides an IV robustness check for the spillover exposure rather than the baseline estimate used in the quantitative model. The strategy exploits tax-induced variation in the effective user cost of R&D, following Wilson (2009), Bloom et al. (2013), and Lucking et al. (2019). We first regress a firm's R&D effort on the user cost of R&D, constructed using tax credits in Lucking et al. (2019), to obtain predicted effort. We then

Table 5: First Stage: Predicting R&D Effort from Tax Credits

	R&D effort $x_{i,t}$ (1)
State tax credit component of R&D user cost	-1.16^{***} (0.30)
Federal tax credit component of R&D user cost	-34.30^{***} (3.64)
Firm fixed effects	✓
Year fixed effects	✓
Partial F-statistic	49.293
No. of observations	16,196

Notes: Outcome is firm-level R&D effort, measured as the square root of Compustat R&D expenditure (xrd). Regressors are the federal and state tax-credit components of the R&D user cost, following Lucking et al. (2019); the specification includes firm and year fixed effects. The partial F-statistic tests their joint significance. Standard errors are clustered by year-by-4-digit-NAICS cells. Regression variables are trimmed at the 0.5% tails within year, and singleton fixed-effect observations are dropped. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

average the five lagged values of predicted effort, from $t-1$ through $t-5$, to construct a tax-driven source-firm innovation proxy, $z_{j,t}^{\text{TAX}}$, and form the unscaled source-firm exposure $\sum_{j\neq i}\tilde{\omega}_{ij,t}z_{j,t}^{\text{TAX}}$. Table 5 reports the corresponding first-stage regression. This excluded instrument targets variation in the source-firm knowledge exposure that enters the endogenous spillover regressor, $\sum_{j\neq i}\tilde{\omega}_{ij,t}z_{j,t}/z_{i,t}$; the recipient firm's own knowledge scale is controlled for directly through $\log z_{i,t}$. The excluded instrument is strong in the IV first stage, with a first-stage F-statistic of 279.8, and the IV estimate of the spillover coefficient remains positive and statistically significant. In the quantitative baseline, we use the column (2) estimate, $\beta = 0.024$, because it conditions on firms' own R&D-effort intensity and maps directly into the law of motion used in the model. Column (3) confirms that the estimated spillover relationship remains positive when spillover exposure is instrumented.

3.6 Remaining Dynamic Parameters

Having estimated the spillover intensity β, we set the discount rate ρ externally and use aggregate moments to discipline the remaining dynamic parameters, the R&D coefficient μ and the depreciation rate δ. The baseline sets $\rho = 10\%$.

We estimate μ and δ jointly in the 2017 competitive-equilibrium baseline by matching two aggregate moments: R&D intensity of 2.6% and average real GDP per capita growth of 1.5% over 1989–2019. For any candidate pair (μ, δ), we solve the 2017 competitive equilibrium at the observed knowledge-capital vector and compute the model-implied R&D intensity and observed-state expected economic growth rate, $d\log Y_t/dt$. This expected

Table 6: Parameters

Notation	Description	Value	Source
$\boldsymbol{S}$	Product-market proximity	–	Hoberg and Phillips (2016)
$\tilde{\boldsymbol{\Omega}}$	Technological proximity	–	DISCERN patents, CPC groups
α	Product proximity to substitutability	0.12	Pellegrino (2025)
β	Technology proximity to spillover	0.024	Estimated from the law of motion
γ	Std. dev. of idiosyncratic shocks	0.00	Deterministic economy in the baseline
ζ/L	Labor augmentation efficiency	0.0040	Compustat, cost of goods sold
ρ	Discount rate	0.10	Baseline choice
μ	R&D coefficient	0.054	Jointly match R&D share and growth
δ	Depreciation rate	0.015	Jointly match R&D share and growth

Notes: Dashes indicate matrix-valued empirical objects rather than scalar parameter estimates. The baseline sets $\gamma = 0$ for the deterministic economy. The ratio ζ/L is identified from aggregate production costs; ζ and L are not separately identified. The parameters μ and δ are jointly calibrated in the 2017 competitive baseline to match aggregate R&D intensity and the observed-state expected economic growth rate.

economic growth rate moment is not the balanced-growth-path eigenvalue characterized in the model section. We measure model-implied R&D intensity as aggregate R&D expenditure, $\sum_i x_i^2$, divided by labor payments plus gross operating profits. The objective minimizes squared relative deviations from the two targets. This procedure yields $\mu = 0.054$ and $\delta = 0.015$.

3.7 Model Fit and Validation

We evaluate the model's fit using both targeted and non-targeted moments. Figure 1 compares firm-level profits, sales, and R&D expenditures in the model and in the data. Profits are targeted by construction because we recover quantities from the profit identity implied by the static equilibrium. The more informative checks are therefore sales and R&D. The model reproduces the cross-sectional distribution of sales closely, with a correlation of 0.97 between model-implied and observed log sales. It also captures a substantial part of the dispersion in innovation activity: the correlation between model-implied and observed log R&D expenditures is 0.74.

Figure 2 compares firm-level growth rates of knowledge capital in the model and in the data. This is a stricter validation exercise because firm-level knowledge growth is not

Figure 1: Firm-Level Profits, Sales, and R&D Expenditures: Model vs. Data

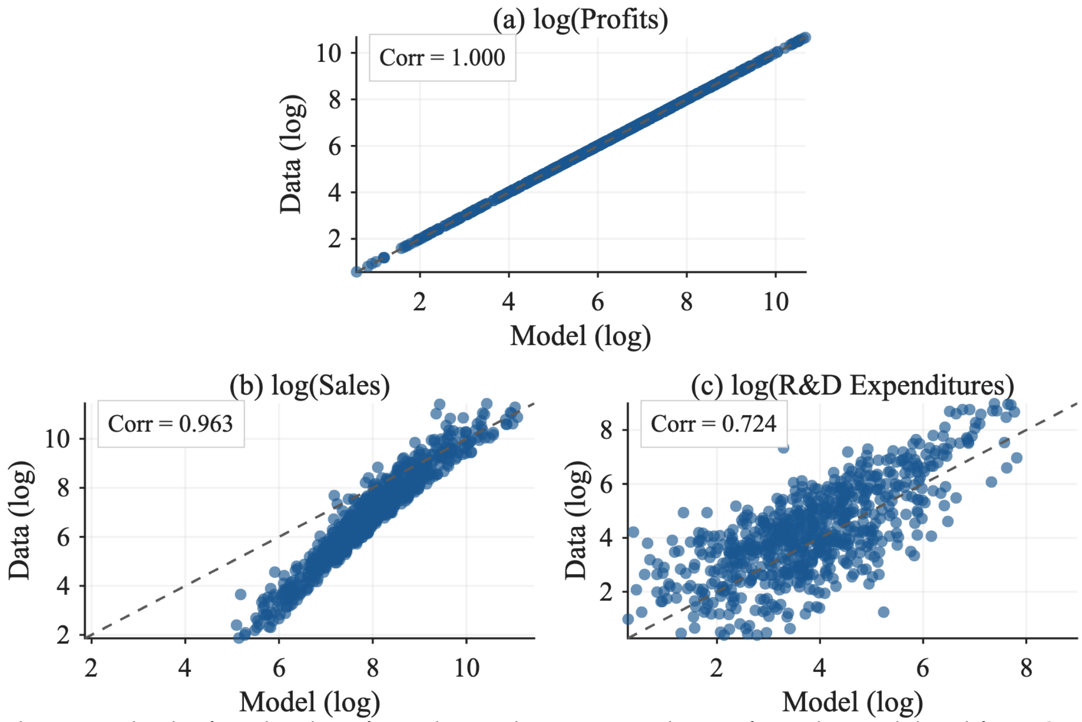

Notes: These panels plot firm-level profits, sales, and R&D expenditures from the model and from Compustat. Profits in the data are measured as "Revenue – Total" (revt) minus "Cost of Goods Sold" (cogs) in Compustat. The model is calibrated so that firm-level profits match those in the data.

matched directly in the calibration of the static block. Although firms grow at a common rate on the deterministic balanced growth path, this validation exercise is not evaluated at that asymptotic object. It computes model-implied local growth rates at the observed 2010 knowledge-capital distribution and network, so cross-sectional growth heterogeneity reflects transitional dynamics from the empirical state. The data-based growth rate is computed from changes between 2010 and 2017 using the 2010 network and initial knowledge-capital distribution. The correlation between model-implied and data-based growth rates is 0.579, indicating that the estimated technology-spillover network and R&D incentives generate meaningful cross-sectional variation in subsequent knowledge growth, even though the model is not designed to match every firm-level growth realization.

4 Numerical Analysis

This section reports the quantitative implications of the estimated model. The main result is that competitive R&D is too low in aggregate and misallocated across firms: the constrained planner more than doubles R&D and raises welfare substantially, while a

Figure 2: Knowledge-Capital Growth Rate Comparison: Model vs. Data

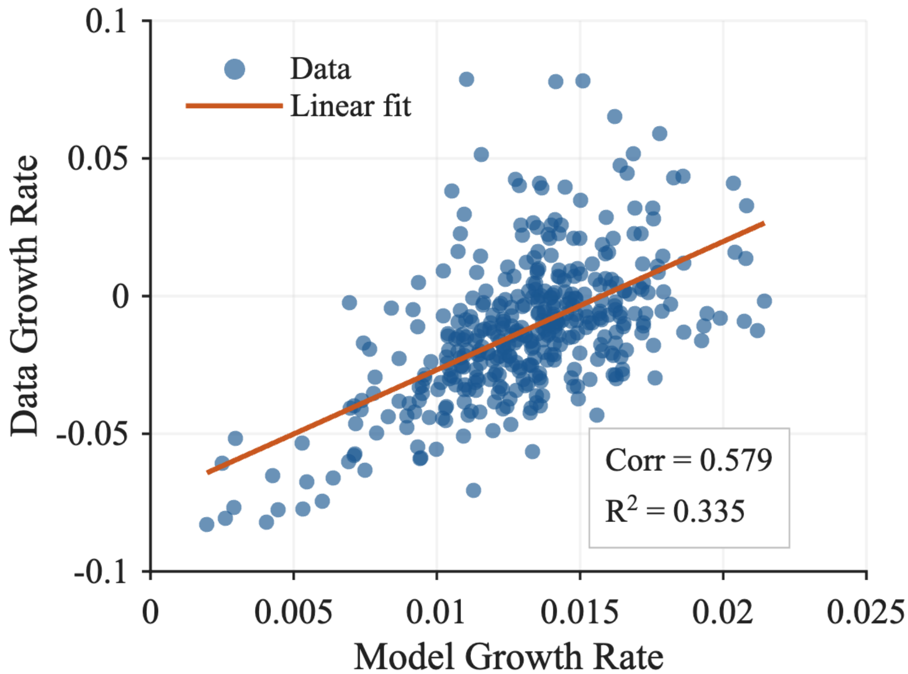

Notes: This figure compares firm-level growth rates of knowledge capital implied by the model and computed from the data. The data-based growth rate is measured between 2010 and 2017 and is constructed using the 2010 network and the initial level of knowledge capital.

welfare-maximizing uniform subsidy nearly reproduces the constrained planner's increase in the expected economic growth rate but not its welfare gain. We first summarize the computational setup. We then examine the baseline allocation by summarizing social-private R&D wedges and decomposing their components. Next, we use counterfactual allocations to distinguish the effects of dynamic R&D control from those of static product-market reallocation. Finally, we evaluate a uniform R&D subsidy and compare it with the constrained planner benchmark.

4.1 Computation and Scale

The numerical implementation follows the model's separation between static product-market mappings and dynamic R&D rules. For each static allocation, we construct the closed-form matrices that map knowledge capital into quantities and flow payoff matrices, including output and gross producer profit. Conditional on a static block, we solve the relevant dynamic R&D problem, which delivers the feedback rule and the value matrices used in producer-value and household-welfare calculations. Scenario labels have two letters: the first denotes the static product-market mapping and the second denotes the dynamic R&D controller. The letter C denotes competition, M monopoly, and S the social planner.

The counterfactuals below are evaluated at the observed 2017 knowledge-capital distribution. They are therefore comparisons of expected household welfare and the expected economic growth rate starting from the same empirical state, not comparisons of asymptotic balanced-growth-path distributions. The main scenario table should be read as an unconstrained interior benchmark. The full static reallocations use closed-form interior mappings; imposing quantity non-negativity would define a different constrained static problem and cannot be implemented by truncating quantities ex post. Section C.3 reports observed-state diagnostics for negative interior quantities and R&D efforts.

The competitive dynamic problem solves a coupled system of Riccati equations for all firms' quadratic value functions. The implementation uses blockwise dense matrix updates, which lets the 2017 baseline solve with 757 firms while preserving a continuous state space and the full empirical product-market and technology-spillover networks. The monopoly and planner dynamic problems reduce to single Riccati equations because one decision maker internalizes all R&D choices. After solving the R&D rule, we compute household welfare and aggregate producer value from continuous-time Lyapunov equations. Section C.2 and Table 13 give the update equation, numerical settings, and computational scale.

4.2 Social-Private R&D Wedges

Before turning to counterfactual allocations, we first compare private and social innovation incentives in the competitive baseline allocation. We compare private and social marginal returns to R&D at the observed 2017 knowledge-capital vector. The private marginal return is the marginal increase in firm value induced by one additional unit of firm i's R&D effort before subtracting the marginal cost of that effort. Since one additional unit of x_i raises the drift of z_i by μ, this private marginal benefit is

$$\mathrm{PMR}_i(\boldsymbol{z}) \equiv \mu \, \frac{\partial V^i(\boldsymbol{z})}{\partial z_i} = 2\mu \left(\boldsymbol{X}_i^i\right)^\top \boldsymbol{z}.$$

In vector form, $\mathrm{PMR}(\boldsymbol{z}) = 2\mu \, \widetilde{\boldsymbol{X}} \, \boldsymbol{z}$. Similarly, define the social marginal return as the increase in household welfare from one additional unit of firm i's R&D effort,

$$\mathrm{SMR}_i(\boldsymbol{z}) \equiv \mu \, \frac{\partial V(\boldsymbol{z})}{\partial z_i}, \qquad \mathrm{SMR}(\boldsymbol{z}) = 2\mu \, \boldsymbol{X} \, \boldsymbol{z}.$$

Figure 3: Social-Private R&D Wedges: Marginal Returns and Return Ratios

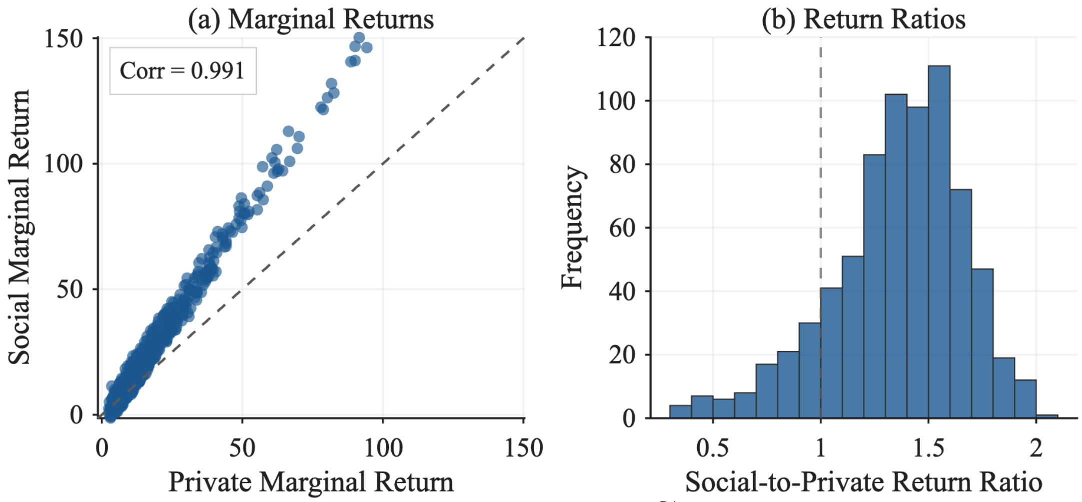

Notes: Panel (a) plots firm-level private marginal returns, $2\mu\widetilde{X}z$, against social marginal returns, $2\mu Xz$. Panel (b) plots the distribution of the social-to-private marginal-return ratio among firms with positive private and social returns; the distribution trims the top and bottom 2 percent for readability, and the dashed vertical line marks equality of social and private returns. Returns are computed using the estimated networks and firm-level knowledge capital in 2017.

The social-private R&D wedge and the social-to-private return ratio are

$$\Delta_i(z) \equiv \mathrm{SMR}_i(z) - \mathrm{PMR}_i(z), \qquad \mathcal{W}_i(z) \equiv \frac{\mathrm{SMR}_i(z)}{\mathrm{PMR}_i(z)}.$$

When both marginal returns are positive, we also report the local FOC-subsidy diagnostic

$$s_i^{\mathrm{loc}}(z) \equiv 1 - \frac{\mathrm{PMR}_i(z)}{\mathrm{SMR}_i(z)}.$$

This object is an observed-state first-order-condition diagnostic, not a solved firm-specific subsidy or targeted-policy counterfactual; we use it only to summarize the sign and magnitude of the local wedge.

Figure 3 summarizes the firm-level wedge diagnostic by plotting marginal-return levels and the social-to-private return-ratio distribution. The model-implied correlation between private and social marginal returns is high (0.991), but social returns tend to exceed private returns, especially among firms with high R&D returns.

The wedge distribution previews the uniform-subsidy result below. The fact that social returns exceed private returns for many firms is consistent with aggregate R&D underprovision, while the dispersion indicates that the underprovision is heterogeneous across firms. Table 7 establishes the size and dispersion of the wedge. It is computed on the 744 of 757 firms with positive private and social marginal returns. In this sample, the

Table 7: Social-Private R&D Wedge at the Observed 2017 State

Scenario	SMR>PMR (%)	Ratio>1.5 (%)	Ratio>2 (%)	Median ratio	Median local FOC subsidy (%)
CC	86.6	36.2	1.1	1.40	28.3

Notes: PMR and SMR are $\text{PMR}_i(z) = 2\mu(X_i^i)^T z$ and $\text{SMR}_i(z) = 2\mu e_i^T X z$. Ratios and local FOC subsidies use the sample with positive PMR and positive SMR, which contains 744 of 757 firms in CC. The local FOC subsidy is $1 - \text{PMR}_i/\text{SMR}_i$ and is an observed-state diagnostic, not a solved firm-specific policy.

Table 8: Social-Private R&D Wedge Components by Return-Ratio Decile

Decile	Median SMR/PMR	Mean wedge	NPS	RP	RRC
1	0.73	-1.92	12.15	-14.95	0.88
2	1.03	0.28	14.30	-14.90	0.87
3	1.19	2.25	16.35	-14.98	0.87
4	1.28	3.98	17.98	-14.85	0.85
5	1.35	5.62	19.63	-14.85	0.84
6	1.43	8.64	22.22	-14.37	0.79
7	1.51	12.83	26.50	-14.46	0.78
8	1.57	21.45	34.91	-14.17	0.72
9	1.64	20.75	33.56	-13.46	0.65
10	1.80	13.12	25.48	-12.96	0.60

Notes: Deciles sort firms by log(SMR/PMR) in the 744-firm positive marginal-return sample at the observed 2017 state. The table averages exact deterministic Lyapunov components defined in Section B.5. Mean wedge is SMR − PMR. NPS, RP, and RRC denote non-producer-surplus, rival-profit, and rival R&D resource-cost components. The decomposition is an accounting diagnostic, not a targeted-policy counterfactual.

median social-to-private return ratio is 1.40, and social returns exceed private returns for 86.6% of firms. The dispersion is economically meaningful: 36.2% of firms have ratios above 1.5, while the bottom tail includes firms for which the observed-state local diagnostic points toward taxing rather than subsidizing marginal R&D.

Having established the size and dispersion of the wedge, Table 8 turns to the mechanism. To interpret the table, we decompose the wedge as

$$\Delta_i(z) = \Delta_i^{\text{NPS}}(z) + \Delta_i^{\text{RP}}(z) + \Delta_i^{\text{RRC}}(z),$$

where the three terms are the non-producer-surplus (NPS), rival-profit (RP), and rival R&D resource-cost (RRC) components. This is a flow-payoff decomposition, not a decomposition by primitive network channel. The RRC component is the rival resource-cost effect, not the technology-spillover channel. Technology spillovers enter through the closed-loop state dynamics and through the solved R&D feedback rule, so they shape the discounted NPS, RP, and RRC components rather than appearing as a separate component. The Lyapunov definitions of these components are in Section B.5.

Economically, NPS measures the part of the household value created by firm i's marginal R&D that is not captured by producer value. It is therefore the appropriability component of the wedge, including consumer surplus and any other non-producer surplus generated along the induced transition path. RP measures the effect of firm i's marginal R&D on other firms' gross profits. When an innovation mainly reallocates demand or profits away from rivals, this component is negative and corresponds to business stealing. The rival R&D resource-cost component accounts for R&D resource costs borne by other firms under their competitive policies: firm i's innovation changes the future state and hence the discounted R&D costs borne by other firms. The signs in Table 8 follow the wedge convention: positive entries raise social marginal returns relative to private marginal returns.

The decomposition shows why these wedges are economically structured. Because deciles are sorted by the social-to-private return ratio rather than by the level wedge, the mean wedge level need not be monotone across rows. The NPS component is positive in every return-ratio decile and is much larger in the upper part of the distribution, increasing from 12.15 in the bottom decile to 34.91 in the eighth decile. The RP component is negative in every decile, roughly between -14.95 and -12.96, so business stealing offsets a substantial part of the positive appropriability force. The RRC component is smaller, between 0.60 and 0.88 in this benchmark. Quantitatively, this channel is much smaller than the NPS and RP components.

Appendix Table 12 connects these exact components to model-implied network exposure measures. The diagnostic shows that firms linked to more valuable spillover destinations have larger positive non-producer-surplus components, whereas firms whose R&D more strongly reallocates profits from product-market rivals have more negative rival-profit components.

4.3 Counterfactual Allocations

The counterfactual exercises vary static product-market behavior and dynamic R&D behavior separately. This distinction is useful because product-market behavior changes the current allocation of labor and output, whereas dynamic behavior changes the incentives to accumulate knowledge capital. The baseline expected economic growth rate is decomposed in Section B.7; in the deterministic benchmark, R&D policy affects the endogenous R&D component of growth, while product-market reallocations change the output matrix that maps knowledge capital into aggregate output. The benchmark is CC: competitive product

Table 9: Scenario Comparison in 2017

Scenario	Output (CC=100)	R&D expenditure (CC=100)	Expected economic growth rate (%)	Welfare (CC=100)	Producer value share (%)
Baseline					
CC	100.00	100.00	1.51	100.00	26.54
Constrained R&D control, competitive product-market allocation					
CM	100.00	51.64	1.36	98.89	27.39
CS	100.00	241.95	1.80	104.36	24.09
Full product-market and R&D control					
MM	96.33	63.22	1.40	95.67	43.09
SS	109.29	348.72	1.83	114.54	-5.55

Notes: Columns labeled CC=100 are evaluated at the observed 2017 knowledge-capital vector and normalized to the competitive equilibrium. Total R&D expenditure is computed as $\sum_i \max\{0, x_i^u(z)\}^2$, where x_i^u is the unconstrained linear R&D rule. The expected economic growth rate is $d \log Y_t/dt$ under the signed unconstrained drift, not a balanced-growth eigenvalue. Producer value share is $100 \times V^P/V$, net of R&D costs. With linear utility, the consumption-equivalent welfare change for scenario k is $100 \times (V^k/V^{CC} - 1)$, the welfare column minus 100.

markets and competitive R&D. The constrained counterfactuals CM and CS keep the static competitive allocation fixed and change only the dynamic R&D rule. The full counterfactuals MM and SS let the monopolist or planner choose both production and R&D within the interior static benchmark. Thus, comparing CC with CM or CS isolates the dynamic R&D rule, holding the competitive product-market allocation fixed. Comparing MM or SS with CC gives the benchmark effect of full monopoly or planner control relative to the competitive equilibrium. When useful, we use CM versus MM and CS versus SS as intermediate comparisons to decompose how much of those benchmark effects comes from allowing production to be reallocated in addition to R&D.

Table 9 compiles five counterfactual allocations that retain the estimated 2017 product-market rivalry, technology-spillover networks, and knowledge-capital distribution. Output is contemporaneous output at the observed 2017 state, while total R&D expenditure and household welfare are evaluated under the scenario-specific dynamic policy. Output, total R&D expenditure, and household welfare are normalized so that the competitive equilibrium (CC) equals 100, and the producer value share reports total producer value divided by aggregate welfare. The table highlights that the constrained counterfactuals isolate the dynamic R&D margin, while the full counterfactuals add the static product-market reallocation margin.

4.3.1 Constrained R&D Control

We first examine the allocations chosen by the constrained monopolist (CM) and the constrained planner (CS), holding fixed the product-market allocation from the competitive equilibrium. With the static product-market allocation kept at its competitive level and the same initial knowledge-capital distribution, production quantities are unchanged. Consequently, contemporaneous output remains $Y = 100$ across these cases (CC, CM, CS).

Both constrained controllers internalize cross-firm effects within their objective functions, but they attach different weights to the surplus created by innovation. The monopolist (CM) evaluates innovation through aggregate producer value; it internalizes losses imposed on rival producers but not consumer surplus, and cuts investment to 51.64, roughly half of the competitive equilibrium value. The constrained planner (CS) instead raises total R&D expenditure to 241.95, more than doubling the competitive equilibrium allocation, because it values innovation through household welfare, including surplus that firms do not appropriate. Because the product-market matrix $\boldsymbol{Q}$, spillover matrix $\boldsymbol{\Omega}$, obsolescence rate δ, and initial knowledge-capital vector $\boldsymbol{z}$ are held fixed across CC, CM, and CS, the difference in the expected economic growth rate comes from the endogenous R&D component. The expected economic growth rate therefore tracks these R&D choices: it falls to 1.36% under the monopolist, 0.15 percentage points below the competitive equilibrium, and rises to 1.80% under the constrained planner, 0.29 points above the competitive equilibrium.

Because R&D is already undersupplied in the competitive equilibrium allocation, the constrained monopolist's lower R&D expenditure pushes welfare down, leading to a 1.11% consumption-equivalent welfare loss relative to the competitive equilibrium outcome. The constrained planner, by contrast, lifts consumption-equivalent welfare 4.36% above competition. The producer value share moves in the opposite direction: prioritizing aggregate producer value raises the monopolist's share to 27.39%, whereas the constrained planner steers more surplus toward consumers and reduces the share to 24.09%. This constrained-planner result is the aggregate counterpart of the wedge diagnostics above: many firms have social marginal R&D returns above private marginal returns, but the planner also reallocates effort across firms according to these heterogeneous gaps.

4.3.2 Full Product-Market Reallocation

We now turn to the full counterfactual allocations in which the monopolist or the full social planner chooses both production and dynamic R&D. The main comparisons are MM and SS relative to CC, the competitive equilibrium, because these comparisons give the interior-benchmark effect of joint control over production and innovation. The constrained allocations CM and CS then serve as intermediate benchmarks for decomposing how much of the full benchmark effect comes from allowing production to be reallocated in addition to R&D. These full reallocations should be read as interior benchmarks that measure the force of static product-market reallocation, not as active-set production policies.

Giving monopoly control over production reduces output by 3.67% relative to the competitive equilibrium allocation. The monopolist chooses static quantities to maximize current aggregate profit rather than household welfare, so it does not internalize consumer surplus when reallocating production. By contrast, the interior planner benchmark reallocates production to maximize final-good output, so output rises by 9.29%. As noted above, these full static reallocations are unconstrained interior benchmarks; Table 14 shows that negative interior quantities arise only in MM and SS and account for less than 1 percent of the q^2 aggregate.

Changes in the product market feed directly into R&D. Relative to the competitive equilibrium allocation, the monopolist's joint choice of production and innovation (MM) reduces total R&D expenditure to 63.22, roughly 37% below the competitive level. This is nevertheless higher than the 51.64 level in CM because higher markups raise the private return to new ideas. In contrast, the planner's joint allocation (SS) raises total R&D expenditure to 348.72, up from 241.95 in CS and about three and a half times the competitive equilibrium allocation.[2] The increase from CS to SS reflects the fact that the interior static planner mapping expands output and lifts the social value of innovation.

These joint changes in production and R&D affect the expected economic growth rate. When the monopolist selects both margins, the economy grows at 1.40%, which is 0.11 percentage points below the competitive equilibrium allocation but 0.04 points above the CM outcome with fixed production. Allowing the planner to pick both margins raises the expected economic growth rate to 1.83%, 0.32 points above competition and 0.03 points higher than the CS scenario. The planner's joint reallocation accelerates knowledge accumulation, whereas the monopolist's restraint keeps growth below the competitive equilibrium allocation.

Letting both margins adjust magnifies the welfare stakes. When the monopolist chooses production and R&D jointly, consumption-equivalent welfare falls 4.33% below the competitive equilibrium allocation, a much steeper loss than the 1.11% decline observed when only R&D is distorted. The extra damage comes from product-market misallocation, which erodes static efficiency on top of the innovation shortfall. In contrast, the planner raises consumption-equivalent welfare by 14.54% relative to competition when allowed to

[2]This size accords with Bloom et al. (2013), who estimate that the socially optimal level of R&D is roughly three times the observed level.

reallocate labor and R&D, far exceeding the 4.36% gain secured by adjusting R&D alone. The interior static reallocation delivers an immediate boost to output and simultaneously amplifies the social return to each unit of research.

Producer value shares swing even more when both margins move. A monopolist that sets production and R&D captures 43.09% of aggregate welfare, up sharply from the 26.54% competitive share. In contrast, the planner drives the share to −5.55 because firms relinquish all static profits yet still finance positive R&D outlays, pushing nearly all surplus to consumers. This negative value should be read as an accounting decomposition under the planner allocation rather than as a market valuation: aggregate producer value net of R&D costs is negative, while consumer surplus more than offsets it.

4.4 Uniform R&D Subsidy

We model a uniform R&D subsidy as reducing each firm's private R&D cost from $x_{i,t}^2$ to $(1-s)x_{i,t}^2$, where $s \in [0,1)$ is the subsidy rate. The subsidy changes firms' dynamic incentives but leaves the static product-market allocation unchanged. Given a value-function coefficient matrix, the feedback rule becomes

$$\boldsymbol{x}_t = \frac{1}{1-s}\mu\widetilde{\boldsymbol{X}}(s)\boldsymbol{z}_t.$$

We treat subsidy payments as financed by lump-sum taxes, so the subsidy changes private incentives but not the aggregate resource cost of R&D; household welfare therefore continues to subtract the full resource cost $\boldsymbol{x}_t^T\boldsymbol{x}_t$. In the numerical exercise, we solve the competitive equilibrium over a grid of subsidy rates and choose the rate that maximizes household welfare. At the 2017 state, a uniform subsidy can improve welfare because aggregate equilibrium R&D falls short of the constrained planner's allocation. Welfare reaches its maximum at a subsidy of 34% on the grid used for the subsidy exercise (Figure 4). At that rate, the expected economic growth rate rises to 1.80%, 0.292 percentage points above competition, nearly matching the constrained planner's 0.294-point increase. Welfare gains differ sharply: the subsidy raises consumption-equivalent welfare by 0.53%, versus 4.36% under the constrained planner. Higher subsidies eventually lower welfare as additional R&D costs outweigh growth benefits.

The wedge diagnostic explains this gap. A uniform subsidy changes the common R&D price, whereas the constrained planner's firm-level allocation reflects heterogeneous wedges. Table 10 groups firms by deciles of the baseline social-to-private return ratio and compares each decile's share of total R&D costs under competition, the solved uniform subsidy, and the constrained planner.

Figure 4: Expected Economic Growth Rate and Household Welfare under R&D Subsidy

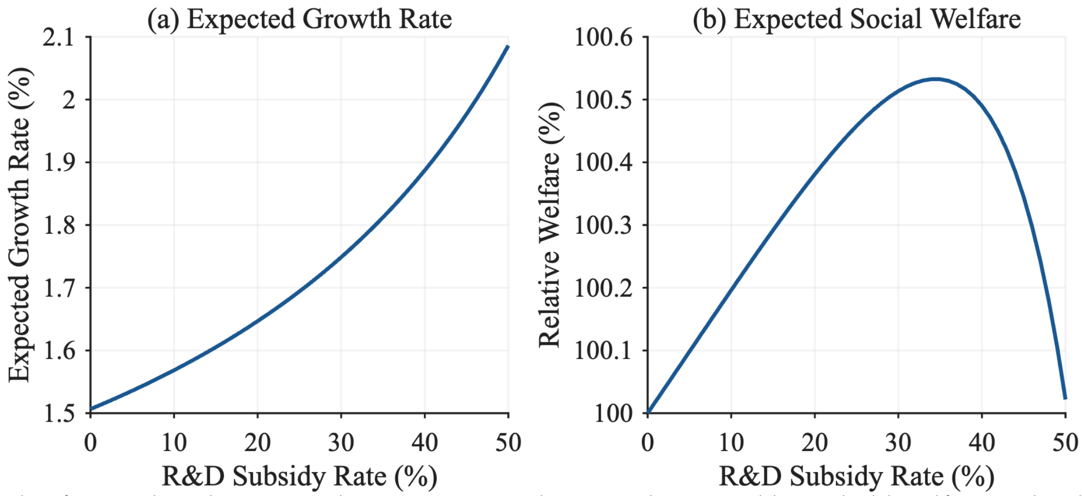

Notes: This figure plots the expected economic growth rate and expected household welfare under different levels of R&D subsidy. The networks of technology spillovers and product-market rivalry, as well as the distribution of knowledge capital, are estimated using data from 2017.

Table 10: R&D Cost Shares by Baseline Social-to-Private Return-Ratio Decile

	Baseline wedge	R&D cost share (%)			Allocation gap
Decile	Median SMR/PMR	Competitive	Uniform subsidy	Planner (CS)	Planner − uniform (pp)
1	0.73	0.6	0.6	0.6	0.06
2	1.03	1.5	1.4	1.5	0.09
3	1.19	2.3	2.2	2.2	0.03
4	1.28	3.3	3.3	3.5	0.26
5	1.35	4.5	4.4	4.6	0.15
6	1.43	7.1	7.0	7.4	0.33
7	1.51	13.9	13.9	13.7	-0.23
8	1.57	33.6	33.9	33.2	-0.64
9	1.64	27.5	27.7	27.4	-0.31
10	1.80	5.6	5.5	5.8	0.26

Notes: Deciles sort firms by the baseline social-to-private return ratio, $\text{SMR}_i/\text{PMR}_i$, among firms with positive marginal returns. R&D cost shares are $x_i^2/\sum_j x_j^2$ at the observed 2017 state. "Uniform" is the solved 34% competitive-subsidy equilibrium. "Planner" is the constrained-planner allocation (CS), holding the competitive product-market allocation fixed. The last column is planner minus uniform-subsidy R&D share in percentage points. The table is an allocation diagnostic, not a targeted-policy welfare estimate.

The uniform subsidy does not move R&D shares toward the constrained-planner allocation. If it reproduced the constrained planner's allocation, the uniform-subsidy shares would move toward the constrained-planner shares across baseline return-ratio deciles; instead, they remain close to the competitive shares and sometimes move in the opposite direction. The total-variation distance between competitive and constrained-planner R&D cost shares is 0.0268. Replacing the competitive allocation with the solved 34% uniform-subsidy allocation leaves the distance to the constrained planner at 0.0275, slightly larger rather than smaller. The reallocation vector induced by the uniform subsidy is negatively correlated with the constrained planner's reallocation vector, with correlation −0.241. Thus the uniform subsidy raises the expected economic growth rate without reallocating R&D toward the firms that receive larger R&D shares in the constrained-planner benchmark. Section B.4 reports the HJB equations underlying the subsidy calculation. This mechanism motivates evaluating instruments that can condition on heterogeneity in product-market rivalry and technology spillovers rather than relying on uniform subsidies alone.

5 Conclusion

This paper develops and estimates an endogenous growth model in which firms interact through product-market rivalry and technology-spillover networks. The main contribution is a tractable many-firm dynamic oligopoly that maps the two empirical firm-to-firm networks into the expected economic growth rate, household welfare, and firm-level social-private R&D wedge diagnostics.

At the observed 2017 knowledge-capital distribution, competition underprovides and misallocates R&D. Holding the product-market allocation fixed, the constrained planner raises R&D to about 242% of the competitive level, increases the expected economic growth rate by 0.29 percentage points, and raises consumption-equivalent welfare by 4.36%; the constrained monopolist instead cuts R&D and lowers household welfare. Allowing production and R&D to move jointly, the planner raises consumption-equivalent welfare by 14.54%, while monopoly control reduces it by 4.33%. These results show that product-market misallocation and innovation misallocation reinforce each other.

Uniform R&D subsidies partly correct the innovation distortion but deliver much smaller welfare gains. The welfare-maximizing grid subsidy is 34% and raises consumption-equivalent welfare by only 0.53%, even though it reproduces nearly all of the constrained planner's increase in the expected economic growth rate. The gap reflects heterogeneous social-private R&D wedges. In the 2017 competitive benchmark, the median social-to-private R&D return ratio is 1.40, social returns exceed private returns for 86.6% of firms with positive private and social marginal returns, and the exact decomposition shows that the wedge is mainly the net of positive non-producer-surplus gains and negative rival-profit losses. A uniform subsidy changes the average price of R&D but cannot reproduce the constrained planner's firm-level allocation. This diagnostic, not a targeted-policy counterfactual, motivates instruments that account for spillover and rivalry networks.

References

Acemoglu, Daron, and Ufuk Akcigit. 2012. "Intellectual Property Rights Policy, Competition and Innovation." *Journal of the European Economic Association* 10 (1): 1–42.

Acemoglu, Daron, Ufuk Akcigit, Douglas Hanley, and William Kerr. 2016. "Transition to Clean Technology." *Journal of Political Economy* 124 (1): 52–104.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. "Competition and Innovation: An Inverted-U Relationship." *Quarterly Journal of Economics* 120 (2): 701–728.

Akcigit, Ufuk, and Sina T. Ates. 2021. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *American Economic Journal: Macroeconomics* 13 (1): 257–298.

Akcigit, Ufuk, and Sina T. Ates. 2023. "What Happened to US Business Dynamism?" *Journal of Political Economy* 131 (8): 2059–2124.

Arora, Ashish, Sharon Belenzon, Larisa Cioaca, Lia Sheer, Hyun Moh (John) Shin, and Dror Shvadron. 2024. "DISCERN 2.0: Duke Innovation & SCientific Enterprises Research Network." `10.5281/zenodo.13619821`, Dataset, Version 2.0.1.

Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–626, Princeton, NJ: Princeton University Press.

Atkeson, Andrew, and Ariel Burstein. 2008. "Pricing-to-Market, Trade Costs, and International Relative Prices." *American Economic Review* 98 (5): 1998–2031.

Azar, Jose, and Xavier Vives. 2021. "General Equilibrium Oligopoly and Ownership Structure." *Econometrica* 89 (3): 999–1048.

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

Besanko, David, and Ulrich Doraszelski. 2004. "Capacity Dynamics and Endogenous Asymmetries in Firm Size." *RAND Journal of Economics* 35 (1): 23–49.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4): 1347–1393.

Bustamante, Maria Cecilia, and Bruno Pellegrino. 2026. "Dynamic Investment and Product Market Rivalry: The Network Q Model." February, CESifo Working Paper No. 12548.

Cavenaile, Laurent, Murat Alp Celik, and Xu Tian. 2026. "Are Markups Too High?

Competition, Strategic Innovation, and Industry Dynamics." Accepted, Review of Economic Studies.

d'Aspremont, Claude, and Alexis Jacquemin. 1988. "Cooperative and Noncooperative R&D in Duopoly with Spillovers." *American Economic Review* 78 (5): 1133–1137.

De Ridder, Maarten. 2024. "Market Power and Innovation in the Intangible Economy." *American Economic Review* 114 (1): 199–251.

Doraszelski, Ulrich, and Ariel Pakes. 2007. "A Framework for Applied Dynamic Analysis in IO." In *Handbook of Industrial Organization*, edited by Mark Armstrong, and Robert Porter, Volume 3. Chap. 30, 1887–1966, Elsevier.

Ederer, Florian, and Bruno Pellegrino. 2025. "A Tale of Two Networks: Common Ownership and Product Market Rivalry." *Review of Economic Studies* rdaf057. `10.1093/restud/rdaf057`.

Ericson, Richard, and Ariel Pakes. 1995. "Markov-Perfect Industry Dynamics: A Framework for Empirical Work." *Review of Economic Studies* 62 (1): 53–82. `10.2307/2297841`.

Goeree, Michelle Sovinsky. 2008. "Limited Information and Advertising in the U.S. Personal Computer Industry." *Econometrica* 76 (5): 1017–1074.

Grossman, Gene M., Elhanan Helpman, Ezra Oberfield, and Thomas Sampson. 2017. "Balanced Growth Despite Uzawa." *American Economic Review* 107 (4): 1293–1312.

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

Jaffe, Adam B. 1986. "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value." *American Economic Review* 76 (5): 984–1001.

Jones, Benjamin F., and Xiaojie Liu. 2024. "A Framework for Economic Growth with Capital-Embodied Technical Change." *American Economic Review* 114 (5): 1448–1487.

Liu, Ernest, and Song Ma. 2024. "Innovation Networks and R&D Allocation." Revise and resubmit, American Economic Review; NBER Working Paper No. 29607.

Lucking, Brian, Nicholas Bloom, and John Van Reenen. 2019. "Have R&D Spillovers Declined in the 21^{st} Century?" *Fiscal Studies* 40 (4): 561–590.

Neary, J. Peter. 2003. "Globalization and Market Structure." *Journal of the European Economic Association* 1 (2-3): 245–271.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica* 69 (2): 307–342.

Pellegrino, Bruno. 2025. "Product Differentiation and Oligopoly: A Network Approach." *American Economic Review* 115 (4): 1170–1225. `10.1257/aer.20201425`.

Peters, Michael. 2020. "Heterogeneous Markups, Growth, and Endogenous Misallocation." *Econometrica* 88 (5): 2037–2073.

Spence, Michael. 1984. "Cost Reduction, Competition, and Industry Performance." *Econometrica* 52 (1): 101.

Weintraub, Gabriel Y., C. Lanier Benkard, and Benjamin Van Roy. 2008. "Markov Perfect Industry Dynamics with Many Firms." *Econometrica* 76 (6): 1375–1411.

Wilson, Daniel J. 2009. "Beggar Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits." *Review of Economics and Statistics* 91 (2): 431–436.

Appendix

Appendix A Model Derivations

A.1 GHL Demand Reduction

The GHL demand system begins with common characteristics and idiosyncratic product characteristics. Product i is described by an m-dimensional nonnegative vector $\boldsymbol{\psi}_i$ of common characteristics, normalized so that $\boldsymbol{\psi}_i^T\boldsymbol{\psi}_i = 1$. Stacking these vectors gives

$$\boldsymbol{\Psi} \equiv \left[\begin{array}{cccc} \boldsymbol{\psi}_1 & \boldsymbol{\psi}_2 & \cdots & \boldsymbol{\psi}_n \end{array} \right],$$

and the product-similarity matrix is

$$\boldsymbol{S} \equiv \boldsymbol{\Psi}^T\boldsymbol{\Psi}.$$

The normalization implies $S_{ii} = 1$, while for $i \neq j$, $S_{ij} = \boldsymbol{\psi}_i^T\boldsymbol{\psi}_j$ is the cosine similarity between products' common-characteristic vectors.

Let $\boldsymbol{q}_t$ be the vector of product quantities. Common and idiosyncratic characteristic quantities are

$$\boldsymbol{y}_t^C = \boldsymbol{\Psi}\boldsymbol{q}_t \tag{17}$$

and

$$\boldsymbol{y}_t^I = \boldsymbol{q}_t. \tag{18}$$

The final-good producer aggregates these characteristics according to

$$Y_t = \alpha \sum_{k=1}^{m} \left(b_{k,t}^C y_{k,t}^C - \frac{1}{2}\left(y_{k,t}^C\right)^2 \right) + (1-\alpha) \sum_{i=1}^{n} \left(b_{i,t}^I y_{i,t}^I - \frac{1}{2}\left(y_{i,t}^I\right)^2 \right), \tag{19}$$

where $\alpha \in [0,1]$ is the weight on common characteristics. Substituting Equations (17) and (18) into Equation (19) gives

$$Y_t = \boldsymbol{q}_t^T \left[\alpha\boldsymbol{\Psi}^T\boldsymbol{b}_t^C + (1-\alpha)\boldsymbol{b}_t^I\right] - \frac{1}{2}\boldsymbol{q}_t^T \left[\alpha\boldsymbol{\Psi}^T\boldsymbol{\Psi} + (1-\alpha)\boldsymbol{I}\right]\boldsymbol{q}_t.$$

Thus the reduced-form product-quality vector and rivalry matrix are

$$\boldsymbol{b}_t \equiv \alpha\boldsymbol{\Psi}^T\boldsymbol{b}_t^C + (1-\alpha)\boldsymbol{b}_t^I, \qquad \boldsymbol{\Sigma} \equiv \alpha\boldsymbol{S} + (1-\alpha)\boldsymbol{I},$$

which yields the reduced-form aggregator in Equation (1).

A.2 Competitive Static Equilibrium

Each firm's real gross profit before subtracting R&D cost is

$$\begin{aligned}\pi_{i,t} &= \frac{p_{i,t}}{P_t} q_{i,t} - \frac{w_t}{P_t} l_{i,t} \\ &= \left(b_{i,t} - \frac{1}{a_{i,t}} \frac{w_t}{P_t} - \sum_j \sigma_{ij} q_{j,t} \right) q_{i,t},\end{aligned}$$

where the second line uses the inverse demand in Equation (6) and the production function in Equation (2). Firm i chooses $q_{i,t}$, $a_{i,t}$, and $b_{i,t}$, taking rivals' quantities, the real wage, and its own knowledge capital as given. The FOC with respect to quantity is

$$0 = b_{i,t} - \frac{1}{a_{i,t}} \frac{w_t}{P_t} - 2q_{i,t} - \sum_{j \neq i} \sigma_{ij} q_{j,t}. \tag{20}$$

The FOCs with respect to $a_{i,t}$ and $b_{i,t}$ imply

$$a_{i,t} = \sqrt{\frac{w_t}{\zeta P_t}} \tag{21}$$

$$b_{i,t} = z_{i,t} - \sqrt{\frac{\zeta w_t}{P_t}}. \tag{22}$$

Inserting Equation (21) into the labor-market-clearing condition in Equation (5) gives

$$\sqrt{\frac{\zeta w_t}{P_t}} = \frac{\zeta}{L} \sum_i q_{i,t}. \tag{23}$$

Inserting the optimality condition for technology choice Equation (21) and Equation (22) into the optimality condition for quantity choice Equation (20), and using $\sigma_{ii} = 1$ to write $2q_{i,t} + \sum_{j \neq i} \sigma_{ij} q_{j,t} = q_{i,t} + \sum_j \sigma_{ij} q_{j,t}$, we obtain

$$z_{i,t} = 2\sqrt{\zeta \frac{w_t}{P_t}} + q_{i,t} + \sum_j \sigma_{ij} q_{j,t}$$

Inserting the expression for real wage Equation (23), we obtain

$$z_{i,t} = q_{i,t} + \sum_j \left(2\frac{\zeta}{L} + \sigma_{ij} \right) q_{j,t}.$$

In vector form,

$$\boldsymbol{z}_t = \left(2\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma} + \boldsymbol{I} \right) \boldsymbol{q}_t.$$

Therefore, quantities are linear in knowledge capital:

$$\boldsymbol{q}_t = \boldsymbol{N}\boldsymbol{z}_t$$

where

$$\boldsymbol{N} \equiv \left(2\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma} + \boldsymbol{I} \right)^{-1}.$$

Combining the quantity FOC with $\sigma_{ii} = 1$ gives

$$b_{i,t} - \frac{1}{a_{i,t}}\frac{w_t}{P_t} - \sum_j \sigma_{ij} q_{j,t} = q_{i,t},$$

so the equilibrium static gross profit is $\pi_{i,t} = q_{i,t}^2$.

A.3 Output in Competitive Equilibrium

Let $\mathbf{1}$ denote an $n \times 1$ vector with all elements equal to 1. Rewrite Equation (22) in vector form:

$$\boldsymbol{b}_t = \boldsymbol{z}_t - \sqrt{\zeta \frac{w_t}{P_t}}\mathbf{1}$$

Inserting Equation (23), we obtain

$$\begin{aligned}\boldsymbol{b}_t &= \boldsymbol{z}_t - \frac{\zeta}{L}\left(\sum_i q_{i,t} \right)\mathbf{1} \\ &= \boldsymbol{z}_t - \frac{\zeta}{L}\boldsymbol{J}\boldsymbol{q}_t.\end{aligned}$$

From Equation (7), we obtain

$$\boldsymbol{b}_t = \left(\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma} + \boldsymbol{I}\right)\boldsymbol{q}_t$$

Therefore, from Equation (1), the total output is given by

$$\begin{aligned} Y_t &= \boldsymbol{q}_t^T\left(\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma} + \boldsymbol{I}\right)\boldsymbol{q}_t - \frac{1}{2}\boldsymbol{q}_t^T\boldsymbol{\Sigma}\boldsymbol{q}_t \\ &= \boldsymbol{q}_t^T\left(\frac{\zeta}{L}\boldsymbol{J} + \frac{1}{2}\boldsymbol{\Sigma} + \boldsymbol{I}\right)\boldsymbol{q}_t \\ &= \boldsymbol{z}_t^T\boldsymbol{N}^T\left(\frac{\zeta}{L}\boldsymbol{J} + \frac{1}{2}\boldsymbol{\Sigma} + \boldsymbol{I}\right)\boldsymbol{N}\boldsymbol{z}_t \\ &= \boldsymbol{z}_t^T\boldsymbol{Q}\boldsymbol{z}_t \end{aligned}$$

where

$$\boldsymbol{Q} \equiv \boldsymbol{N}^T\left(\frac{\zeta}{L}\boldsymbol{J} + \frac{1}{2}\boldsymbol{\Sigma} + \boldsymbol{I}\right)\boldsymbol{N}$$

A.4 Competitive Dynamic R&D Derivation

For the quadratic value-function guess in Equation (9), with $\boldsymbol{X}^i$ symmetric, the first and second derivatives are

$$\frac{\partial V^i(\boldsymbol{z})}{\partial \boldsymbol{z}} = 2\left(\boldsymbol{X}^i\boldsymbol{z}\right)^T \tag{24}$$

$$\frac{\partial^2 V^i(\boldsymbol{z})}{\partial \boldsymbol{z}\partial \boldsymbol{z}^T} = 2\boldsymbol{X}^i. \tag{25}$$

The derivative of Equation (8) with respect to x_i is $-2x_i + \mu\left[V_z^i(\boldsymbol{z})\right]_i = 0$. Since $\left[V_z^i(\boldsymbol{z})\right]_i = 2\left(\boldsymbol{X}_i^i\right)^T\boldsymbol{z}$, the first-order condition yields the linear policy in Equation (10).

The diffusion term in Equation (8) simplifies to

$$\frac{\gamma^2}{2}\operatorname{Tr}\left[\operatorname{diag}(\boldsymbol{z})\,2\boldsymbol{X}^i\operatorname{diag}(\boldsymbol{z})\right] = \gamma^2\boldsymbol{z}^T\mathcal{D}\left(\boldsymbol{X}^i\right)\boldsymbol{z}.$$

Substituting Equations (9), (10), (24) and (25) into Equation (8) and collecting the quadratic terms in $\boldsymbol{z}$ gives the stacked Riccati system in Equation (11).

Table 11: Growth Rate of Variables in Deterministic Economy

Growth rate	Variables
0	$l_{i,t}$
g	$a_{i,t}, b_{i,t}, z_{i,t}, q_{i,t}, x_{i,t}, p_{i,t}/P_t$
$2g$	$\pi_{i,t}, C_t, Y_t, w_t/P_t$

Notes: Growth rates are continuous-time instantaneous rates evaluated along the balanced-growth-path equilibrium in the deterministic economy with no shocks.

A.5 Competitive Welfare and Producer Value

After firms' equilibrium policies are fixed, household welfare is obtained from the valuation equation

$$\rho V(z) = z^T Q z - \mu^2 z^T \widetilde{X}^T \widetilde{X} z + V_z(z)\, \Phi z + \frac{\gamma^2}{2} \mathrm{Tr}\left[\mathrm{diag}(z)\, V_{zz}(z)\, \mathrm{diag}(z)\right].$$

The first two terms are household consumption, $Y_t - \sum_i x_{i,t}^2$, and the last two terms capture expected changes in continuation value. With $V(z) = z^T X z$, substituting the derivatives from Equations (24) and (25) with X^i replaced by X gives the Lyapunov equation in Equation (12).

Aggregate producer value in the competitive equilibrium is

$$V^P(z_0) = \mathbb{E}_0 \left[\int_0^\infty \exp(-\rho t) \left(\sum_i \pi_{i,t} - \sum_i x_{i,t}^2 \right) dt \,\middle|\, z_0 \right].$$

Because aggregate real profit is $\sum_i \pi_{i,t} = z_t^T N^T N z_t$, the producer-value coefficient X^P solves the same Lyapunov equation as Equation (12), replacing Q with $N^T N$.

A.6 Deterministic Balanced Growth Path

We characterize the existence of a balanced growth path in the deterministic economy, $\gamma = 0$, using the closed-loop system $\dot{z}_t = \Phi z_t$. When $\gamma > 0$, the same transition matrix governs the drift of knowledge capital. Table 11 summarizes the implied growth rates of the main variables along a deterministic balanced growth path.

We impose the following spectral condition on Φ, which is sufficient for the existence of a balanced-growth-path equilibrium.

Assumption 2. *The deterministic closed-loop technology transition matrix* $\mathbf{\Phi}$ *has a simple real dominant eigenvalue* $g > 0$ *with an associated strictly positive right eigenvector, and all other eigenvalues have real parts strictly below* g.

This assumption imposes the spectral property needed for balanced growth: the closed-loop technology system has a unique positive long-run growth direction. It rules out cases with multiple dominant growth directions or a non-positive balanced-growth-path knowledge-capital distribution.

Definition 2. Balanced Growth Path Equilibrium: A *balanced-growth-path equilibrium* is a linear Markov perfect equilibrium such that all firms' knowledge capital grows at the same constant rate.

Proposition 1. *Consider the deterministic economy,* $\gamma = 0$. *Let Assumptions 1 and 2 hold, and let* $\bar{z} \gg 0$ *be the right eigenvector of* $\mathbf{\Phi}$ *associated with the dominant eigenvalue* g. *Then, for any scalar* $c > 0$, *there exists a balanced-growth-path equilibrium whose knowledge-capital path is*

$$z_t = e^{gt} c \bar{z}.$$

Along this path, all firms' knowledge capital grows at rate g, *and the cross-sectional distribution of knowledge capital is proportional to* $\bar{z}$.

Proof. Assumption 1 gives a stabilizing solution to the stacked Riccati equations and hence linear Markov R&D policies, together with the static equilibrium mappings characterized above. In the deterministic economy, these policies imply the closed-loop system $\dot{z}_t = \mathbf{\Phi} z_t$. Since $\mathbf{\Phi}\bar{z} = g\bar{z}$, the path $z_t = e^{gt} c \bar{z}$ satisfies this system. Because $\bar{z} \gg 0$ and $c > 0$, all firms' knowledge stocks remain positive and grow at the common rate g. The induced static allocations and R&D policies therefore constitute a linear Markov perfect equilibrium on a balanced growth path. □

Proposition 1 implies that the equilibrium growth rate of knowledge capital is the dominant eigenvalue of $\mathbf{\Phi}$, and the balanced-growth-path knowledge-capital distribution is the associated eigenvector. Since $q_t = N z_t$ and equilibrium R&D strategies are linear in z_t, quantities and R&D grow at rate g. The static optimality conditions imply that product quality, labor productivity, and real product prices also grow at rate g, while the real wage, output, consumption, and profits grow at rate $2g$ because they are pinned down by quadratic terms in knowledge capital. In the deterministic case, Assumption 1 requires $\rho > 2g$, so discounted payoffs are finite along the balanced growth path.

Figure 5 illustrates the mechanics of a balanced-growth-path equilibrium in the deterministic economy. In equilibrium, quality, quantity, real price, and real marginal cost

Figure 5: Partial Equilibrium Diagram and Its Proportional Growth

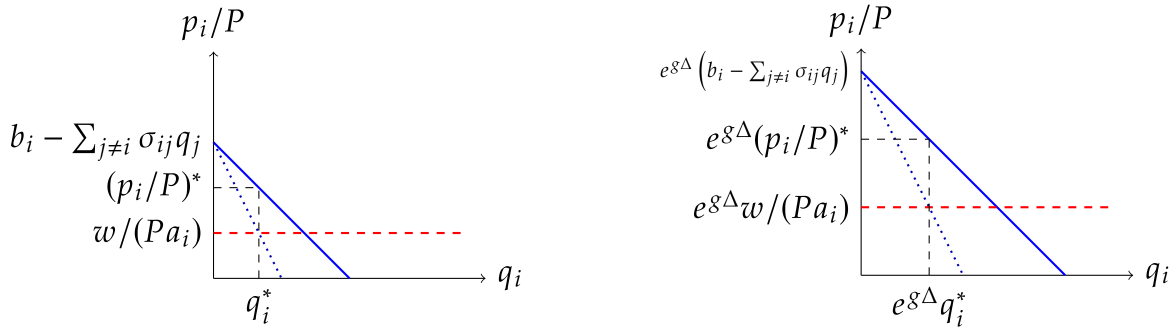

Notes: The left panel depicts the static equilibrium for a firm's product, showing the residual demand (blue), marginal revenue (dotted line), and real marginal cost (red) curves. The right panel illustrates how the real equilibrium price, quantity, real marginal cost, and the intercept of the residual demand scale proportionally with the growth factor $\exp(g\Delta)$ over an interval of length Δ.

grow at the same continuous-time rate g, so the partial-equilibrium diagram expands homothetically along both axes.

Appendix B Counterfactual and Wedge Derivations

B.1 Social Planner's Static Problem

We reduce the planner's static problem by substituting the knowledge-allocation constraint $b_{i,t} = z_{i,t} - \zeta a_{i,t}$ and the production technology $l_{i,t} = q_{i,t}/a_{i,t}$ into the planner's objective and labor-market constraint. Define the Lagrangian as

$$\mathcal{L} = \boldsymbol{q}_t^T(\boldsymbol{z}_t - \zeta\boldsymbol{a}_t) - \frac{1}{2}\boldsymbol{q}_t^T\boldsymbol{\Sigma}\boldsymbol{q}_t - \lambda_t\left\{\sum_i \frac{q_{i,t}}{a_{i,t}} - L\right\}$$

We focus on the interior solution in which the labor constraint binds. The FOC for $a_{i,t}$ implies $\lambda_t = \zeta a_{i,t}^2 > 0$, so the square root below is well defined. The FOC w.r.t. $a_{i,t}$ gives

$$a_{i,t} = \sqrt{\frac{\lambda_t}{\zeta}} \tag{26}$$

and therefore,

$$b_{i,t} = z_{i,t} - \sqrt{\zeta\lambda_t} \tag{27}$$

Inserting Equation (26) into Equation (13), we obtain

$$\sqrt{\zeta\lambda_t} = \frac{\zeta}{L}\sum_i q_{i,t}$$

Note that

$$\sqrt{\zeta\lambda_t}\mathbf{1} = \frac{\zeta}{L}\left(\sum_i q_{i,t}\right)\mathbf{1} = \frac{\zeta}{L}\boldsymbol{J}\boldsymbol{q}_t. \tag{28}$$

The FOC w.r.t. $\boldsymbol{q}_t$ gives

$$\boldsymbol{z}_t - \zeta\boldsymbol{a}_t - \boldsymbol{\Sigma}\boldsymbol{q}_t = \lambda_t \begin{bmatrix} 1/a_{1,t} \\ 1/a_{2,t} \\ \vdots \\ 1/a_{n,t} \end{bmatrix}$$

Substituting Equation (26), we obtain

$$\boldsymbol{z}_t = 2\sqrt{\zeta\lambda_t}\mathbf{1} + \boldsymbol{\Sigma}\boldsymbol{q}_t$$

Using Equation (28), we obtain

$$\boldsymbol{z}_t = \left(2\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma}\right)\boldsymbol{q}_t$$

Therefore,

$$\boldsymbol{q}_t = \boldsymbol{N}^*\boldsymbol{z}_t \tag{29}$$

where

$$\boldsymbol{N}^* = \left(2\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma}\right)^{-1}$$

Because $\boldsymbol{J}$ and $\boldsymbol{\Sigma}$ are symmetric, $\boldsymbol{N}^*$ is symmetric as well. To obtain the quadratic form of output, substitute Equation (27) into Equation (1):

$$Y_t = \boldsymbol{q}_t^T\left(\boldsymbol{z}_t - \sqrt{\zeta\lambda_t}\mathbf{1}\right) - \frac{1}{2}\boldsymbol{q}_t^T\boldsymbol{\Sigma}\boldsymbol{q}_t$$

Using Equation (28) together with Equation (29),

$$\begin{aligned} Y_t &= \boldsymbol{q}_t^T\left(\left(2\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma}\right)\boldsymbol{q}_t - \frac{\zeta}{L}\boldsymbol{J}\boldsymbol{q}_t\right) - \frac{1}{2}\boldsymbol{q}_t^T\boldsymbol{\Sigma}\boldsymbol{q}_t \\ &= \boldsymbol{q}_t^T\left(\frac{\zeta}{L}\boldsymbol{J} + \frac{1}{2}\boldsymbol{\Sigma}\right)\boldsymbol{q}_t \end{aligned}$$

$$= \frac{1}{2}\boldsymbol{q}_t^T (\boldsymbol{N}^*)^{-1} \boldsymbol{q}_t$$
$$= \frac{1}{2}\boldsymbol{z}_t^T \boldsymbol{N}^* \boldsymbol{z}_t$$

B.2 Planner Dynamic R&D Derivation

The HJB equation for the social planner's dynamic problem is

$$\rho V^{SS}(\boldsymbol{z}) = \max_{\boldsymbol{x}} \left\{ \frac{1}{2}\boldsymbol{z}^T \boldsymbol{N}^* \boldsymbol{z} - \boldsymbol{x}^T \boldsymbol{x} + V_z^{SS}(\boldsymbol{z}) \left[(\boldsymbol{\Omega} - \delta \boldsymbol{I}) \boldsymbol{z} + \mu \boldsymbol{x} \right] + \frac{\gamma^2}{2} \mathrm{Tr} \left[\mathrm{diag}(\boldsymbol{z}) V_{zz}^{SS}(\boldsymbol{z}) \, \mathrm{diag}(\boldsymbol{z}) \right] \right\}. \tag{30}$$

Using the quadratic guess

$$V^{SS}(\boldsymbol{z}) = \boldsymbol{z}^T \boldsymbol{X}^{SS} \boldsymbol{z}, \tag{31}$$

the first-order condition with respect to $\boldsymbol{x}$ is $-2\boldsymbol{x} + \mu \left(V_z^{SS}(\boldsymbol{z}) \right)^T = 0$, which gives the linear R&D rule in Equation (14). The induced transition is $\boldsymbol{\Phi}^{SS} = \boldsymbol{\Omega} + \mu^2 \boldsymbol{X}^{SS} - \delta \boldsymbol{I}$. Substituting the quadratic value function, the linear R&D rule, and this transition into Equation (30) gives the planner Riccati equation in Equation (15).

B.3 Monopoly Allocation and Welfare

This subsection derives the full monopoly allocation MM used in the counterfactual analysis. We keep the superscript M for static monopoly matrices, such as $\boldsymbol{N}^M$, $\boldsymbol{P}^M$, and $\boldsymbol{Q}^M$, and use the superscript MM for the full scenario in which the monopolist controls both production and R&D. Since one decision maker controls the full R&D vector, the monopoly R&D policy is obtained from a single Riccati equation; household welfare under that policy is then computed separately from a Lyapunov equation.

Static Mapping

The monopolist chooses $\boldsymbol{a}_t$, $\boldsymbol{b}_t$, and $\boldsymbol{q}_t$ to maximize aggregate real profit. Differentiating Π_t with respect to $q_{i,t}$, and using the symmetry of $\boldsymbol{\Sigma}$ so that $\partial(\boldsymbol{q}_t^T \boldsymbol{\Sigma} \boldsymbol{q}_t)/\partial q_{i,t} = 2\sum_j \sigma_{ij} q_{j,t}$, gives

$$0 = b_{i,t} - \frac{1}{a_{i,t}} \frac{w_t}{P_t} - 2 \sum_j \sigma_{ij} q_{j,t}. \tag{32}$$

The real profit of each product is therefore

$$\begin{aligned}\pi_{i,t} &= \left(b_{i,t} - \frac{1}{a_{i,t}}\frac{w_t}{P_t}\right) q_{i,t} - \sum_j \sigma_{ij} q_{i,t} q_{j,t} \\ &= 2\sum_j \sigma_{ij} q_{j,t} q_{i,t} - \sum_j \sigma_{ij} q_{i,t} q_{j,t} \\ &= \sum_j \sigma_{ij} q_{i,t} q_{j,t}. \end{aligned} \tag{33}$$

The FOC with respect to $a_{i,t}$ gives

$$a_{i,t} = \sqrt{\frac{1}{\zeta}\frac{w_t}{P_t}}, \tag{34}$$

and therefore,

$$b_{i,t} = z_{i,t} - \sqrt{\zeta \frac{w_t}{P_t}}. \tag{35}$$

Inserting Equation (34) into the labor-market-clearing condition Equation (5) gives

$$\sqrt{\zeta \frac{w_t}{P_t}} = \frac{\zeta}{L} \sum_i q_{i,t}. \tag{36}$$

Combining Equations (32), (35) and (36) yields

$$\boldsymbol{z}_t = 2\left(\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma}\right)\boldsymbol{q}_t, \tag{37}$$

so monopoly quantities are linear in knowledge capital:

$$\boldsymbol{q}_t^M = \boldsymbol{N}^M \boldsymbol{z}_t, \qquad \boldsymbol{N}^M \equiv \frac{1}{2}\left(\frac{\zeta}{L}\boldsymbol{J} + \boldsymbol{\Sigma}\right)^{-1}. \tag{38}$$

From Equation (33), aggregate real profit is

$$\Pi_t^M = \boldsymbol{q}_t^T \boldsymbol{\Sigma} \boldsymbol{q}_t = \boldsymbol{z}_t^T \boldsymbol{P}^M \boldsymbol{z}_t, \qquad \boldsymbol{P}^M \equiv (\boldsymbol{N}^M)^T \boldsymbol{\Sigma} \boldsymbol{N}^M.$$

Output and Producer Surplus

As in the competitive equilibrium,

$$\boldsymbol{b}_t = \boldsymbol{z}_t - \frac{\zeta}{L}\boldsymbol{J}\boldsymbol{q}_t.$$

Using Equation (37), this becomes

$$\boldsymbol{b}_t = \left(\frac{\zeta}{L}\boldsymbol{J} + 2\boldsymbol{\Sigma}\right)\boldsymbol{q}_t.$$

Therefore, from Equation (1), monopoly output is

$$\begin{aligned} Y_t^M &= \boldsymbol{q}_t^T\left(\frac{\zeta}{L}\boldsymbol{J} + 2\boldsymbol{\Sigma}\right)\boldsymbol{q}_t - \frac{1}{2}\boldsymbol{q}_t^T\boldsymbol{\Sigma}\boldsymbol{q}_t \\ &= \frac{1}{2}\boldsymbol{q}_t^T\left(2\frac{\zeta}{L}\boldsymbol{J} + 3\boldsymbol{\Sigma}\right)\boldsymbol{q}_t = \boldsymbol{z}_t^T\boldsymbol{Q}^M\boldsymbol{z}_t, \end{aligned}$$

where

$$\boldsymbol{Q}^M \equiv \frac{1}{2}(\boldsymbol{N}^M)^T\left(2\frac{\zeta}{L}\boldsymbol{J} + 3\boldsymbol{\Sigma}\right)\boldsymbol{N}^M.$$

Dynamic R&D

The monopolist chooses R&D to maximize aggregate producer value:

$$V^{P,MM}(\boldsymbol{z}_0) \equiv \max_{\{\boldsymbol{x}_t\}_{t\geq 0}} \mathbb{E}_0\left[\int_0^\infty \exp(-\rho t)\left\{\boldsymbol{z}_t^T\boldsymbol{P}^M\boldsymbol{z}_t - \boldsymbol{x}_t^T\boldsymbol{x}_t\right\}dt\right],$$

subject to the knowledge accumulation law in Equation (4). We obtain the monopoly R&D rule by solving Equation (15) with $\frac{1}{2}\boldsymbol{N}^*$ replaced by $\boldsymbol{P}^M$. Let $\boldsymbol{X}^{P,MM}$ denote the symmetric coefficient matrix in the monopolist's producer-value function. The R&D allocation is

$$\boldsymbol{x}_t = \mu\boldsymbol{X}^{P,MM}\boldsymbol{z}_t,$$

and the induced transition is

$$\boldsymbol{\Phi}^{MM} \equiv \boldsymbol{\Omega} - \delta\boldsymbol{I} + \mu^2\boldsymbol{X}^{P,MM}.$$

Household Welfare

In contrast to the social planner, the monopolist maximizes aggregate producer value rather than household utility. Hence household welfare under the monopoly policy is computed separately:

$$V^{MM}\left(\boldsymbol{z}_0\right) = \mathbb{E}_0\left[\int_0^\infty \exp\left(-\rho t\right)\left(\boldsymbol{z}_t^T \boldsymbol{Q}^M \boldsymbol{z}_t - \boldsymbol{x}_t^T \boldsymbol{x}_t\right) dt \middle| \boldsymbol{z}_0\right],$$

subject to $d\boldsymbol{z}_t = \boldsymbol{\Phi}^{MM}\boldsymbol{z}_t dt + \gamma\, \mathrm{diag}\left(\boldsymbol{z}_t\right) d\boldsymbol{W}_t$. Guessing $V^{MM}(\boldsymbol{z}) = \boldsymbol{z}^T \boldsymbol{X}^{MM} \boldsymbol{z}$ gives the Lyapunov equation

$$\begin{aligned} 0 = \boldsymbol{Q}^M - \mu^2\left(\boldsymbol{X}^{P,MM}\right)^T \boldsymbol{X}^{P,MM} + \boldsymbol{X}^{MM}\left(\boldsymbol{\Phi}^{MM} - \frac{\rho}{2}\boldsymbol{I}\right) \\ + \left(\boldsymbol{\Phi}^{MM} - \frac{\rho}{2}\boldsymbol{I}\right)^T \boldsymbol{X}^{MM} + \gamma^2 \mathcal{D}\left(\boldsymbol{X}^{MM}\right). \end{aligned} \tag{39}$$

Therefore, $\boldsymbol{X}^{MM}$ is obtained as the solution of Equation (39), and welfare in the monopoly problem is $V^{MM}(\boldsymbol{z}_0) = \boldsymbol{z}_0^T \boldsymbol{X}^{MM} \boldsymbol{z}_0$.

B.4 Uniform R&D Subsidy Equations

This appendix records the equations used in the uniform R&D subsidy exercise. Let $s \in [0, 1)$ denote a subsidy rate that reduces the private cost of R&D effort from $x_{i,t}^2$ to $(1-s)x_{i,t}^2$. Holding the static product-market allocation fixed, firm i solves

$$V^i\left(\boldsymbol{z}_0\right) = \max_{\{x_{i,t}\}_{t\geq 0}} \mathbb{E}_0\left[\int_0^\infty e^{-\rho t}\left\{\boldsymbol{z}_t^T \boldsymbol{Q}_i \boldsymbol{z}_t - (1-s)x_{i,t}^2\right\} dt\right],$$

taking other firms' R&D rules as given. With the quadratic guess $V^i(\boldsymbol{z}) = \boldsymbol{z}^T \boldsymbol{X}^i \boldsymbol{z}$, the first-order condition is

$$x_i = \frac{\mu}{1-s}\left(\boldsymbol{X}_i^i\right)^T \boldsymbol{z}.$$

Stacking firms' policies into the subsidy-specific policy matrix $\widetilde{\boldsymbol{X}}(s)$ gives

$$\boldsymbol{x}_t = \frac{\mu}{1-s}\widetilde{\boldsymbol{X}}(s)\boldsymbol{z}_t, \qquad \boldsymbol{\Phi}(s) = \boldsymbol{\Omega} - \delta \boldsymbol{I} + \frac{\mu^2}{1-s}\widetilde{\boldsymbol{X}}(s).$$

The competitive Riccati equations under the subsidy become

$$0 = \boldsymbol{Q}_i - \frac{\mu^2}{1-s}\boldsymbol{X}^i_i \left(\boldsymbol{X}^i_i\right)^T + \left(\boldsymbol{\Phi}(s) - \frac{\rho}{2}\boldsymbol{I}\right)^T \boldsymbol{X}^i + \boldsymbol{X}^i \left(\boldsymbol{\Phi}(s) - \frac{\rho}{2}\boldsymbol{I}\right) + \gamma^2 \mathcal{D}\left(\boldsymbol{X}^i\right).$$

The subsidy changes firms' private cost but not the resource cost borne by the household. Hence household welfare subtracts the full cost

$$\boldsymbol{x}_t^T \boldsymbol{x}_t = \frac{\mu^2}{(1-s)^2} \boldsymbol{z}_t^T \widetilde{\boldsymbol{X}}(s)^T \widetilde{\boldsymbol{X}}(s) \boldsymbol{z}_t.$$

Given the equilibrium policy, household welfare is obtained from the Lyapunov equation

$$0 = \boldsymbol{Q} - \frac{\mu^2}{(1-s)^2}\widetilde{\boldsymbol{X}}(s)^T \widetilde{\boldsymbol{X}}(s) + \boldsymbol{X}(s)\left(\boldsymbol{\Phi}(s) - \frac{\rho}{2}\boldsymbol{I}\right) + \left(\boldsymbol{\Phi}(s) - \frac{\rho}{2}\boldsymbol{I}\right)^T \boldsymbol{X}(s) + \gamma^2 \mathcal{D}\left(\boldsymbol{X}(s)\right).$$

The numerical exercise solves this system for $s = 0, 0.01, \ldots, 0.50$ and selects the value of s that maximizes $\boldsymbol{z}_0^T \boldsymbol{X}(s) \boldsymbol{z}_0$ at the 2017 knowledge-capital vector.

The main text reports the allocation diagnostic that compares R&D cost shares under competition, the solved 34% uniform subsidy, and the constrained planner.

B.5 Social-Private R&D Wedge Decomposition

This appendix gives the component definitions used in Table 8. Let $\boldsymbol{e}_i$ denote the ith standard basis vector. Let $\boldsymbol{P}$ denote the aggregate competitive gross-profit matrix,

$$\boldsymbol{P} \equiv \sum_{j=1}^{n} \boldsymbol{Q}_j = \boldsymbol{N}^T \boldsymbol{N}.$$

Thus aggregate competitive gross profit is $\sum_j \pi_j(\boldsymbol{z}) = \boldsymbol{z}^T \boldsymbol{P} \boldsymbol{z}$. The social-private R&D wedge can be written as

$$\Delta_i(\boldsymbol{z}) = 2\mu\, \boldsymbol{e}_i^T \left(\boldsymbol{X} - \boldsymbol{X}^i\right) \boldsymbol{z}. \tag{40}$$

Define $\boldsymbol{M}_i \equiv \boldsymbol{X} - \boldsymbol{X}^i$ and

$$\boldsymbol{C}_{-i} \equiv \widetilde{\boldsymbol{X}}^T \left(\boldsymbol{I} - \boldsymbol{e}_i \boldsymbol{e}_i^T\right) \widetilde{\boldsymbol{X}}.$$

Since the competitive feedback rule is $\boldsymbol{x}^{\mathrm{C}}(\boldsymbol{z}) = \mu\widetilde{\boldsymbol{X}}\boldsymbol{z}$, the quadratic form $\mu^2\boldsymbol{C}_{-i}$ gives the R&D resource costs of all firms other than i under the competitive policy. Subtracting firm i's Riccati equation from the household welfare Lyapunov equation gives

$$\boldsymbol{0} = \boldsymbol{F}_i + \boldsymbol{M}_i\left(\boldsymbol{\Phi} - \frac{\rho}{2}\boldsymbol{I}\right) + \left(\boldsymbol{\Phi} - \frac{\rho}{2}\boldsymbol{I}\right)^T \boldsymbol{M}_i + \gamma^2\mathcal{D}(\boldsymbol{M}_i), \tag{41}$$

where the flow wedge is

$$\begin{aligned}\boldsymbol{F}_i &\equiv \boldsymbol{Q} - \boldsymbol{Q}_i - \mu^2\boldsymbol{C}_{-i} \\ &= \underbrace{\boldsymbol{Q} - \boldsymbol{P}}_{\text{non-producer surplus}} + \underbrace{\boldsymbol{P} - \boldsymbol{Q}_i}_{\text{rival profit}} - \underbrace{\mu^2\boldsymbol{C}_{-i}}_{\text{rival R\&D resource-cost component}}.\end{aligned} \tag{42}$$

This decomposition separates the flow-payoff consequences of a marginal increase in firm i's knowledge drift under the competitive closed-loop policy. The term $\boldsymbol{Q} - \boldsymbol{P}$ is the part of household output value that is not captured by aggregate producer gross profit, so it is the non-producer-surplus, or appropriability, component. The term $\boldsymbol{P} - \boldsymbol{Q}_i$ is the gross-profit value accruing to firms other than i; when firm i's innovation reallocates demand away from rivals, this component is negative after propagation and corresponds to business stealing. Finally, $-\mu^2\boldsymbol{C}_{-i}$ subtracts the R&D resource costs borne by other firms under their competitive feedback rules. This decomposition is by flow payoff, not by primitive network channel. The RRC component is the rival resource-cost effect, not the technology-spillover channel. Technology spillovers enter through the closed-loop state dynamics and through the solved R&D feedback rule, so they shape the discounted NPS, RP, and RRC components rather than appearing as a separate component.

For each component $m \in \{\mathrm{NPS}, \mathrm{RP}, \mathrm{RRC}\}$, solve Equation (41) with the corresponding flow term

$$\boldsymbol{F}^{\mathrm{NPS}} = \boldsymbol{Q} - \boldsymbol{P}, \qquad \boldsymbol{F}_i^{\mathrm{RP}} = \boldsymbol{P} - \boldsymbol{Q}_i, \qquad \boldsymbol{F}_i^{\mathrm{RRC}} = -\mu^2\boldsymbol{C}_{-i}.$$

Denoting the resulting solutions by $\boldsymbol{M}^{\mathrm{NPS}}$, $\boldsymbol{M}_i^{\mathrm{RP}}$, and $\boldsymbol{M}_i^{\mathrm{RRC}}$, the component wedges are

$$\begin{aligned}\Delta_i^{\mathrm{NPS}}(\boldsymbol{z}) &= 2\mu\, \boldsymbol{e}_i^T \boldsymbol{M}^{\mathrm{NPS}} \boldsymbol{z}, \\ \Delta_i^{\mathrm{RP}}(\boldsymbol{z}) &= 2\mu\, \boldsymbol{e}_i^T \boldsymbol{M}_i^{\mathrm{RP}} \boldsymbol{z}, \\ \Delta_i^{\mathrm{RRC}}(\boldsymbol{z}) &= 2\mu\, \boldsymbol{e}_i^T \boldsymbol{M}_i^{\mathrm{RRC}} \boldsymbol{z},\end{aligned} \tag{43}$$

and they add up to the total wedge,

$$\Delta_i(\boldsymbol{z}) = \Delta_i^{\mathrm{NPS}}(\boldsymbol{z}) + \Delta_i^{\mathrm{RP}}(\boldsymbol{z}) + \Delta_i^{\mathrm{RRC}}(\boldsymbol{z}). \tag{44}$$

Table 12: Network Diagnostics for Social-Private R&D Wedge Components

	Total wedge	NPS	RP	RRC
Value-weighted spillover exposure	0.361	0.282	0.796	-1.003
Business-stealing exposure	-0.220	-0.121	-0.892	0.485
log z	0.767	0.787	-0.049	0.064
R^2	0.857	0.851	0.879	0.822
Observations	744	744	744	744

Notes: The table reports standardized coefficients from flow-proxy regressions for the 744-firm positive marginal-return sample at the observed 2017 state. Dependent variables are the exact deterministic Lyapunov wedge components defined in Section B.5; both dependent variables and predictors are standardized. Value-weighted spillover exposure and business-stealing exposure are defined in Equations (45) and (46). NPS is the non-producer-surplus component, RP is the rival-profit component, and RRC is the rival R&D resource-cost component. These regressions are model-internal diagnostics rather than causal estimates. The RRC component remains a secondary mechanism because its level contribution is small in Table 8.

For the component-level network diagnostic, define two firm-level flow-proxy exposures, evaluated at the observed state before standardization:

$$\text{Value-weighted spillover exposure}_i(z) \equiv 2 \sum_{j \neq i} \Omega_{ji}\, \boldsymbol{e}_j^T \boldsymbol{X} \boldsymbol{z}, \tag{45}$$

$$\text{Business-stealing exposure}_i(z) \equiv - \sum_{j \neq i} \frac{\partial \pi_j(z)}{\partial z_i} = - \sum_{j \neq i} 2 q_j N_{ji}. \tag{46}$$

Because rows of $\boldsymbol{\Omega}$ are recipients, $[\boldsymbol{\Omega}]_{ji}$ is the effective spillover exposure of recipient j to source firm i. The value weight $2\boldsymbol{e}_j^T \boldsymbol{X} \boldsymbol{z}$ is the household-welfare marginal value of recipient j's knowledge capital, so the measure emphasizes spillovers to recipients whose knowledge is more valuable in the model. Since $\boldsymbol{q} = \boldsymbol{N}\boldsymbol{z}$ and $\pi_j = q_j^2$, business-stealing exposure measures the extent to which an increase in firm i's knowledge capital reduces rivals' gross profits.

Table 12 reports a component-level network diagnostic for the same observed-state competitive-policy wedge. The table relates the exact Lyapunov components to value-weighted spillover exposure and business-stealing exposure, controlling for firm knowledge capital. The signs provide a compact interpretation of the decomposition: firms with greater exposure to valuable spillover destinations have larger positive non-producer-surplus components, whereas firms whose R&D more strongly reallocates sales from product-market rivals have more negative rival-profit components. These regressions are not causal estimates; they check whether the computed components line up with the model-implied network forces.

B.6 Expected Log Output Growth

Apply Itô's lemma to the closed-loop knowledge law

$$d\boldsymbol{z}_t = \boldsymbol{\Phi}\boldsymbol{z}_t dt + \gamma\, \mathrm{diag}\left(\boldsymbol{z}_t\right) d\boldsymbol{W}_t$$

and to output $Y_t = f(\boldsymbol{z}_t) = \boldsymbol{z}_t^T \boldsymbol{Q} \boldsymbol{z}_t$. The matrix $\boldsymbol{Q}$ is symmetric by construction, and the log-growth formula below is evaluated on states with $Y_t > 0$. The required derivatives are

$$\begin{aligned} f_z\left(\boldsymbol{z}_t\right) &= 2\boldsymbol{Q}\boldsymbol{z}_t, \\ f_{zz}\left(\boldsymbol{z}_t\right) &= 2\boldsymbol{Q}. \end{aligned}$$

Therefore,

$$\begin{aligned} dY_t &= \left\{2\boldsymbol{z}_t^T \boldsymbol{Q}\boldsymbol{\Phi}\boldsymbol{z}_t + \frac{1}{2}\mathrm{Tr}\left[2\gamma^2\, \mathrm{diag}\left(\boldsymbol{z}_t\right) \boldsymbol{Q}\, \mathrm{diag}\left(\boldsymbol{z}_t\right)\right]\right\} dt + 2\gamma \boldsymbol{z}_t^T \boldsymbol{Q}\, \mathrm{diag}\left(\boldsymbol{z}_t\right) d\boldsymbol{W}_t \\ &= \left\{\boldsymbol{z}_t^T \left(\boldsymbol{Q}\boldsymbol{\Phi} + \boldsymbol{\Phi}^T \boldsymbol{Q}\right) \boldsymbol{z}_t + \gamma^2 \sum_i z_{i,t}^2 \boldsymbol{Q}_{ii}\right\} dt + 2\gamma \boldsymbol{z}_t^T \boldsymbol{Q}\, \mathrm{diag}\left(\boldsymbol{z}_t\right) d\boldsymbol{W}_t \end{aligned}$$

Next, applying Itô's lemma to $g\left(Y_t\right) = \log Y_t$ and using the diffusion term in dY_t gives

$$\begin{aligned} d\log Y_t =& \left[\frac{\boldsymbol{z}_t^T \left(\boldsymbol{Q}\boldsymbol{\Phi} + \boldsymbol{\Phi}^T \boldsymbol{Q}\right) \boldsymbol{z}_t}{\boldsymbol{z}_t^T \boldsymbol{Q} \boldsymbol{z}_t} + \gamma^2 \left\{\frac{\sum_i z_{i,t}^2 \boldsymbol{Q}_{ii}}{\boldsymbol{z}_t^T \boldsymbol{Q}\boldsymbol{z}_t} - 2\frac{\boldsymbol{z}_t^T \boldsymbol{Q}\, \mathrm{diag}\left(\boldsymbol{z}_t^2\right) \boldsymbol{Q}\boldsymbol{z}_t}{\left(\boldsymbol{z}_t^T \boldsymbol{Q}\boldsymbol{z}_t\right)^2}\right\}\right] dt \\ &+ \frac{2\gamma}{\boldsymbol{z}_t^T \boldsymbol{Q}\boldsymbol{z}_t} \boldsymbol{z}_t^T \boldsymbol{Q}\, \mathrm{diag}\left(\boldsymbol{z}_t\right) d\boldsymbol{W}_t. \end{aligned}$$

Because the technology transition matrix can be written as

$$\boldsymbol{\Phi} = \boldsymbol{\Omega} - \delta \boldsymbol{I} + \mu^2 \widetilde{\boldsymbol{X}},$$

the first drift term in log output can be decomposed as

$$\boldsymbol{Q}\boldsymbol{\Phi} + \boldsymbol{\Phi}^T \boldsymbol{Q} = \left(\boldsymbol{Q}\boldsymbol{\Omega} + \boldsymbol{\Omega}^T \boldsymbol{Q}\right) + \mu^2 \left(\boldsymbol{Q}\widetilde{\boldsymbol{X}} + \widetilde{\boldsymbol{X}}^T \boldsymbol{Q}\right) - 2\delta \boldsymbol{Q}.$$

Thus, the contributions of technology spillovers, R&D, obsolescence, and the Itô term to the expected economic growth rate are

$$\frac{d\log Y_t|_{\mathrm{spillover}}}{dt} = \frac{\boldsymbol{z}_t^T \left(\boldsymbol{Q}\boldsymbol{\Omega} + \boldsymbol{\Omega}^T \boldsymbol{Q}\right) \boldsymbol{z}_t}{\boldsymbol{z}_t^T \boldsymbol{Q}\boldsymbol{z}_t} \tag{47}$$

$$\frac{d \log Y_t|_{\text{R\&D}}}{dt} = \mu^2 \frac{z_t^T \left(Q\widetilde{X} + \widetilde{X}^T Q \right) z_t}{z_t^T Q z_t} \tag{48}$$

$$\frac{d \log Y_t|_{\text{Obsolescence}}}{dt} = -2\delta \tag{49}$$

$$\frac{d \log Y_t|_{\text{Itô}}}{dt} = \gamma^2 \left\{ \frac{\sum_i z_{i,t}^2 Q_{ii}}{z_t^T Q z_t} - 2 \frac{z_t^T Q \,\text{diag}\left(z_t^2\right) Q z_t}{\left(z_t^T Q z_t\right)^2} \right\}.$$

B.7 Baseline Expected Economic Growth Rate Decomposition

As a diagnostic for interpreting the expected economic growth rate column in the counterfactual table, we decompose the baseline expected economic growth rate into technology spillovers, obsolescence, R&D, and the contribution of shocks. The components are computed from the primitive matrices Q, Ω, $\widetilde{X}$, and the observed 2017 knowledge-capital vector so that they add up to the expected economic growth rate under the baseline policy.

The baseline expected economic growth rate is 1.51%, close to the expected economic growth rate target used in the calibration. It is the net of a 3.87 percentage-point contribution from technology spillovers, a 0.54 point contribution from direct R&D, and a 2.90 point drag from knowledge obsolescence. In the baseline deterministic specification, the Itô term is zero because $\gamma = 0$. This decomposition guides the counterfactuals in the main text: changing R&D incentives affects the expected economic growth rate through the endogenous R&D component, while changing product-market allocations changes the output matrix Q and hence how a given knowledge-capital distribution maps into aggregate output.

Appendix C Numerical Implementation and Diagnostics

C.1 Computational Scale

For the competitive dynamic problem, we solve the coupled system of Riccati equations by iterating on the coefficient matrices of firms' quadratic value functions. The implementation stores the firm-level value matrices as a three-dimensional array, recomputes the policy matrix $\widetilde{X}$ and closed-loop transition matrix Φ at each iteration, and updates firm pages using blockwise vectorized matrix products. For the monopoly and social-planner dynamic problems, the corresponding dynamic block is a single Riccati equation because one decision maker internalizes all R&D choices. After solving the R&D rule, we compute household welfare and aggregate producer value from continuous-time Lyapunov equations under

Table 13: Computational Complexity and Scale

Model	Computational complexity	Number of firms	Productivity state space
Cavenaile et al. (2026)	$O(k^n)$	4	Six grid points per firm
Our model	$O(n^4)$	418–802	Continuous

Notes: Complexity entries describe the core state-space or dense-matrix operations that determine scalability. In our algorithm, storage of the competitive value-function coefficients scales as $O(n^3)$ and the $O(n^4)$ term refers to dense matrix operations within a Riccati update.

the induced closed-loop law of motion.

When the number of firms is n, each firm has an $n \times n$ matrix of value-function coefficients, $\boldsymbol{X}^i$. The competitive dynamic problem therefore stores one matrix page per firm, or n^3 coefficients in total. In the 2017 quantitative baseline, $n = 757$, so this amounts to about 434 million coefficients. The main arithmetic cost comes from Riccati updates such as $\boldsymbol{\Phi}^T \boldsymbol{X}^i$ and $\boldsymbol{X}^i \boldsymbol{\Phi}$: each dense matrix product costs order $O(n^3)$ for one firm page, and updating all n pages gives an $O(n^4)$ operation.

Within existing endogenous growth frameworks, solving a dynamic oligopoly problem involving all publicly traded U.S. firms with patents is computationally challenging. With a discretized productivity state for each firm, the state space grows exponentially in the number of firms, at order $O(k^n)$ for k grid points per firm. For instance, Cavenaile et al. (2026) introduce a dynamic oligopoly into a Schumpeterian growth model, similar to ours. However, due to the curse of dimensionality, they analyze oligopolies with a maximum of four firms. In contrast, our model has a computational complexity of order $O\left(n^4\right)$ and preserves a continuous state space. This tractability is what makes it possible to evaluate counterfactual allocations for the full empirical network rather than for a small calibrated industry. Table 13 summarizes this computational contrast and the empirical scale of our application.

C.2 Riccati Solver and Numerical Settings

For the competitive dynamic problem, the numerical algorithm iterates on the coefficient matrices of firms' quadratic value functions. Let k denote the iteration index. Given the current firm-level value matrices $\boldsymbol{X}_k^1, \ldots, \boldsymbol{X}_k^n$, we recover $\widetilde{\boldsymbol{X}}_k$, construct the implied closed-loop matrix $\boldsymbol{\Phi}_k$, and update the matrices using the forward-Euler step

$$\frac{\boldsymbol{X}_{k+1}^i - \boldsymbol{X}_k^i}{\Delta} = \boldsymbol{Q}_i - \mu^2 \boldsymbol{X}_{i,k}^i \left(\boldsymbol{X}_{i,k}^i\right)^T$$

$$
\begin{aligned}
&+ \left(\boldsymbol{\Phi}_k - \frac{\rho}{2}\boldsymbol{I}\right)^T \boldsymbol{X}_k^i + \boldsymbol{X}_k^i \left(\boldsymbol{\Phi}_k - \frac{\rho}{2}\boldsymbol{I}\right) \\
&+ \gamma^2 \mathcal{D}\left(\boldsymbol{X}_k^i\right).
\end{aligned}
$$

Here Δ is a pseudo-time step for the fixed-point iteration, not a model time step. The vector $\boldsymbol{X}_{i,k}^i$ is the ith column of firm i's value-function matrix at iteration k, equivalently the transpose of the ith row of $\widetilde{\boldsymbol{X}}_k$. The baseline quantitative specification sets $\gamma = 0$, but the displayed update keeps the diffusion term to match the general model. The implementation follows these steps:

1. Initialize $\boldsymbol{X}_0^i = \boldsymbol{0}$ for all firms.
2. After each update, recover $\widetilde{\boldsymbol{X}}_k$ from the firm-level value matrices and construct the closed-loop transition matrix $\boldsymbol{\Phi}_k$.
3. Update the competitive firm-value matrices in page blocks, which keeps the memory footprint manageable while using vectorized dense matrix products within each block.
4. Stop when the maximum absolute change in $\widetilde{\boldsymbol{X}}_k$ falls below the convergence tolerance. The pseudo-time step, tolerance, maximum iteration count, and page-block size are recorded in the replication code.
5. As a diagnostic separate from the convergence criterion, check stability using the largest real eigenvalue of $\boldsymbol{\Phi}_k - \rho\boldsymbol{I}/2$.
6. For the monopoly and social-planner benchmarks, solve the corresponding single-agent Riccati equation with the same convergence criterion.

C.3 Interior-Benchmark Non-Negativity Diagnostics

The main scenario comparison uses unconstrained interior static mappings and the unconstrained linear-quadratic R&D equilibrium evaluated at the observed 2017 knowledge-capital vector. This subsection is a diagnostic for that tractable benchmark; it does not solve an active-set static equilibrium, impose an R&D non-negativity constraint, or replace the counterfactual allocations reported in Table 9. The distinction matters because the static block maps knowledge capital into quantities linearly,

$$
\boldsymbol{q}^u(\boldsymbol{z}) = \boldsymbol{N}\boldsymbol{z}, \tag{50}
$$

Table 14: Negative Interior Static Quantities at the Observed 2017 State

Scenario	Neg. firms	Neg. $\lvert q\rvert$ (%)	Neg. q^2 (%)	Sales (%)	Obs. R&D (%)
CC	0/757	0.00	0.00	0.00	0.00
CM	0/757	0.00	0.00	0.00	0.00
CS	0/757	0.00	0.00	0.00	0.00
MM	69/757	1.08	0.07	0.13	0.62
SS	234/757	7.91	0.79	1.03	2.58

Notes: The table evaluates the unconstrained interior static mapping $q^u(z) = Nz$ at the observed 2017 knowledge-capital vector. Negative $|q|$ is $\sum_{i:q_i^u<0} |q_i^u| / \sum_i |q_i^u|$, and negative q^2 is $\sum_{i:q_i^u<0} (q_i^u)^2 / \sum_i (q_i^u)^2$. Sales and observed R&D are observed 2017 shares accounted for by firms with negative interior static quantities. The table is a diagnostic and does not recompute an active-set allocation with $q_i \geq 0$.

and the dynamic block implies a linear R&D policy,

$$\boldsymbol{x}^u(\boldsymbol{z}) = \mu \widetilde{\boldsymbol{X}} \boldsymbol{z}. \tag{51}$$

Because these rules are linear, some counterfactual allocations can assign negative quantities or negative effort to some firms. We report observed-state checks for the unconstrained benchmark: the incidence and economic mass of negative static quantities, the incidence and mass of negative R&D effort, and the sensitivity of the expected economic growth rate to projecting negative R&D effort to zero.

The main scenario table uses two reporting conventions. Reported R&D expenditure is computed using the positive-effort resource-cost convention, $\sum_i \max\{0, x_i^u(\boldsymbol{z})\}^2$, while the expected economic growth rate is computed from the signed linear drift implied by the solved policy at the observed state. The positive-effort convention prevents negative effort components from contributing positive squared costs to the expenditure aggregate.

Static Quantity Diagnostics

Table 14 reports negative interior static quantities at the observed 2017 knowledge-capital vector. There are no negative quantities when the static allocation is competitive (CC, CM, and CS), so the constrained R&D-control counterfactuals do not raise this issue. Negative quantities arise only in the full static reallocation benchmarks: 69 firms in MM and 234 firms in SS. The affected firms are economically small in the quadratic quantity aggregate: they account for 0.07 percent of q^2 in MM and 0.79 percent in SS. These diagnostics measure the extent of the interior-benchmark sign violation; they do not impose $q_i \geq 0$ or recompute active-set static matrices.

Table 15: Negative Unconstrained R&D Effort at the Observed 2017 State

Scenario	Neg. firms	Neg. R&D cost (%)	Sales (%)	Obs. R&D (%)
CC	0/757	0.00	0.00	0.00
CM	345/757	3.24	3.27	5.12
CS	3/757	< 0.01	< 0.01	0.01
MM	135/757	0.32	0.45	0.79
SS	36/757	0.03	0.04	0.25

Notes: Negative R&D cost share is $\sum_{i:x_i<0} x_i^2 / \sum_i x_i^2$. Sales and observed R&D are shares of observed 2017 totals accounted for by firms with negative unconstrained R&D effort.

R&D Effort Diagnostics

Table 15 reports the number and severity of negative unconstrained R&D effort at the observed 2017 knowledge-capital vector. The competitive benchmark has no negative R&D effort. Negative effort appears in some counterfactuals, especially CM, but the affected firms account for a small share of the R&D cost aggregate used in the scenario table. In CM, for example, 345 firms have negative unconstrained effort, while their negative R&D cost share is 3.24 percent, their sales share is 3.27 percent, and their observed-R&D share is 5.12 percent. The other counterfactuals have much smaller negative R&D cost shares.

Projected-Positive Growth Robustness

To check whether these negative components matter for the aggregate scenario comparison, define the projected-positive policy

$$\boldsymbol{x}^+(\boldsymbol{z}) = \max\{\boldsymbol{0}, \boldsymbol{x}^u(\boldsymbol{z})\}, \tag{52}$$

where the maximum is taken component by component. The projection is used consistently in both R&D expenditure and the knowledge drift,

$$\dot{\boldsymbol{z}}^+ = (\boldsymbol{\Omega} - \delta \boldsymbol{I})\boldsymbol{z} + \mu \boldsymbol{x}^+(\boldsymbol{z}). \tag{53}$$

This adjustment holds the solved value functions and feedback rule fixed, so it is not a constrained Markov perfect equilibrium. It directly tests whether the expected economic growth rate column in the scenario table is sensitive to setting negative observed-state efforts to zero. Table 16 shows that the projected-positive expected economic growth rates are essentially unchanged. The largest difference is below 0.001 percentage points, far smaller than the scenario differences discussed in the main text.

Table 16: Projected-Positive Robustness of the 2017 Scenario Comparison

Scenario	Signed growth (%)	Projected growth (%)	Difference (pp)
CC	1.51	1.51	0.0000
CM	1.36	1.36	0.0002
CS	1.80	1.80	0.0000
MM	1.40	1.40	-0.0005
SS	1.83	1.83	-0.0008

Notes: The projected-positive policy sets $x^+(z) = \max\{0, x^u(z)\}$ and uses x^+ consistently in both R&D expenditure and the knowledge drift. This adjustment holds the solved policy fixed and does not recompute a constrained equilibrium.

C.4 BGP Positivity Diagnostics

This subsection concerns the long-run BGP interpretation of the solved closed-loop systems. It is distinct from the observed-state counterfactual comparisons in the main text, which evaluate welfare and the expected economic growth rate at the 2017 knowledge-capital distribution.

Baseline BGP Positivity Check

The long-run sign restriction concerns the dominant eigenvector of the unconstrained closed-loop system, not the observed-state expected economic growth rate column in Table 9. The dominant eigenpair solves

$$\mathbf{\Phi}\bar{z} = g\bar{z}, \qquad \mathbf{\Phi} = \mathbf{\Omega} - \delta \boldsymbol{I} + \mu^2 \widetilde{\boldsymbol{X}}. \tag{54}$$

Although $\mathbf{\Omega}$ is nonnegative under the row-receiver spillover convention, $\widetilde{\boldsymbol{X}}$ contains strategic R&D policy feedbacks and can have negative off-diagonal elements. Hence $\mathbf{\Phi}$ is not generally a Metzler matrix and need not preserve the positive orthant; a positive dominant eigenvector is therefore not guaranteed by Perron–Frobenius arguments. After orientation toward the observed knowledge-capital vector and normalization by $\sum_i |\bar{z}_i| = 1$, $\bar{z}$ is the candidate cross-sectional knowledge-capital distribution on the BGP. If it has both positive and negative entries, the candidate path would assign negative knowledge capital to some firms, so it cannot be interpreted as a positive BGP. We therefore do not use the associated firm-level BGP R&D policies or depreciation-bound calculations as economic results. This restriction is specific to the asymptotic BGP calculation; Table 9 is evaluated at the observed 2017 state. Table 17 reports this positivity check for the five counterfactual scenarios.

Table 17: Positive-State Feasibility of the Unconstrained Dominant BGP

Scenario	Neg. $\bar{z}$	Neg. $\lvert\bar{z}\rvert$ share (%)
CC	524/757	47.9
CM	442/757	49.2
CS	406/757	49.4
MM	465/757	49.1
SS	339/757	51.2

Notes: The dominant eigenvector is oriented toward the observed 2017 knowledge-capital vector and normalized by $\sum_i \lvert\bar{z}_i\rvert = 1$. Since no scenario has a positive dominant state, firm-level BGP R&D nonnegativity and depreciation lower-bound calculations are not economically interpreted.

CC Spillover-Floor Diagnostic

The baseline BGP feasibility results show that the unconstrained dominant eigenvectors are mixed-sign in all five 2017 scenarios. For the competitive benchmark, a counterfactual row-sum-preserving general spillover floor recovers a positive dominant BGP. The diagnostic replaces each off-diagonal exposure row by the mixture

$$\Omega_{ij}^{(\lambda)} = (1-\lambda)\Omega_{ij} + \lambda \frac{\sum_{k\neq i}\Omega_{ik}}{n-1}, \qquad j \neq i, \tag{55}$$

with the diagonal unchanged. Thus a fraction λ of each recipient firm's baseline spillover exposure is spread uniformly across other source firms without changing that row's total exposure.

The CC spillover-floor benchmark uses $\lambda = 0.15$. Smaller values $\lambda = 0.10$ already recover a nonnegative oriented dominant eigenvector, but their real-part eigengaps are below 10^{-3}. At $\lambda = 0.15$, the dominant eigenvector is positive with a larger eigengap, and the observed-state expected economic growth rate, R&D, and welfare remain close to the baseline competitive calculation.

Table 18 compares deterministic transitions from the same observed 2017 state under the baseline competitive system and the $\lambda = 0.15$ spillover-floor benchmark. The instantaneous drift changes by about 2.7 percent in relative ℓ_2 norm, but firm-level growth rates remain highly correlated and the normalized state-path distance remains small through 20 years. The continuation-welfare ratio evaluated at the 50-year states is about 0.978, so this exercise should be read as short- and medium-run robustness for the observed-state comparison, not as full long-run equivalence or as a re-estimated model.

Takeaway. The mixed-sign dominant eigenvectors delimit long-run BGP interpretation, not the observed-state counterfactuals. The $\lambda = 0.15$ floor is only a CC diagnostic: it is not a

Table 18: Observed-State Transition Comparison for the CC Spillover-Floor Benchmark

Years	Normalized z distance (ℓ_1)	Firm-growth corr.	Growth diff. (pp)	R&D ratio	Welfare ratio
0	0.0000	0.9983	-0.045	0.996	0.995
1	0.0002	0.9983	-0.045	0.996	0.994
5	0.0008	0.9982	-0.046	0.996	0.993
10	0.0016	0.9982	-0.045	0.995	0.990
20	0.0030	0.9980	-0.042	0.994	0.987
50	0.0062	0.9962	-0.028	0.992	0.978

Notes: The table compares deterministic transitions from the same observed 2017 state under baseline CC and the row-sum-preserving $\lambda = 0.15$ spillover-floor benchmark. The state distance is the ℓ_1 distance between states normalized by $\sum_i |z_i(t)|$. Firm-growth corr. is the cross-firm correlation of instantaneous knowledge-growth rates, $(\mathbf{\Phi} \boldsymbol{z}(t))_i / z_i(t)$, between the two paths. Growth diff. is the benchmark path minus the baseline path in percentage points. R&D is the ratio of R&D expenditure to the baseline path at the same horizon. Welfare is the ratio of continuation values, $V(\boldsymbol{z}(t)) = \boldsymbol{z}(t)^T \boldsymbol{X} \boldsymbol{z}(t)$, evaluated at the same horizon.

re-estimation or an input into the main counterfactual tables.