Production Networks and R&D Allocation

Yasutaka Koike
UNC Chapel Hill
TDB-CAREE

Koki Okumura
UCLA[*]
TDB-CAREE

June 16, 2026

Abstract

This paper studies how production networks shape firms' incumbent R&D and entry incentives. We develop a dynamic model in which firms accumulate persistent buyer-supplier relationships and incumbent R&D creates new product lines that use those inherited relationships. A firm's network affects the return to R&D through the suppliers used to produce new product lines and the buyers through which those lines generate demand. The model highlights two wedges between private and social returns: private innovators do not fully appropriate the surplus created by new product lines, and they do not internalize how product-line creation changes future matching opportunities throughout the network. We discipline the model using matched Japanese firm-to-firm transaction and patent data. In our baseline specification, the decentralized economy tilts R&D too far toward entry and young incumbent firms: the social planner reduces entry by 25.4% relative to the decentralized equilibrium, reallocates incumbent R&D toward older incumbents, and raises consumption-equivalent welfare by 1.71% along the transition.

[*]We thank Ariel Burstein, David Baqaee, and Hugo Hopenhayn for excellent guidance. We also thank Andy Atkeson, Oleg Itskhoki, Lee Ohanian, Jon Vogel, Pablo Fajgelbaum, Joao Guerreiro, Gary Hansen, Pierre-Olivier Weill, and seminar participants at UCLA for valuable comments. We gratefully acknowledge the Teikoku Databank (TDB) and the TDB Center for Advanced Empirical Research on Enterprise and Economy (TDB-CAREE) for providing the datasets.

1 Introduction

How does a firm's position in a production network shape its incentive to innovate? When a semiconductor firm develops a new chip, it does not innovate in isolation: it sources materials and equipment through long-standing supplier relationships and sells the chip to buyers it already serves. Innovation is embedded in accumulated trading relationships. A large literature studies how innovation diffuses through knowledge networks (Liu and Ma, 2023), but much less is known about whether the production network itself—the web of buyer-supplier relationships through which firms source inputs and sell new product lines—changes the private return to R&D, and whether the resulting allocation of R&D effort is efficient.

This paper studies that question in a dynamic model of production networks, incumbent R&D, and entry, disciplined by matched Japanese firm-to-firm transaction and patent data. The question is allocative: given a scarce R&D resource, does the decentralized economy allocate R&D effort efficiently between entrant product-line creation and incumbent product-line expansion, and across incumbents of different ages? In our baseline specification, the answer is no. The decentralized economy tilts R&D too far toward entry and young incumbent firms. The social planner reduces entry by 25.4% relative to the decentralized equilibrium, reallocates incumbent R&D toward older incumbents with accumulated buyer-supplier relationships, and raises consumption-equivalent welfare by 1.71% along the transition from the decentralized steady state.

Three empirical facts from the matched data motivate the model. First, firms accumulate suppliers and buyers throughout the life cycle, rapidly when young and more gradually when mature. Second, trading relationships are highly persistent, so firms and their partners age together: the positive age assortativity we observe in the cross-section reflects durable relationships rather than age-selective matching. Third, firms with more patent-intensive trading partners exhibit higher own patenting. Instrumental-variable estimates that exploit partners' differential exposure to foreign patenting imply that a 10 percent increase in partner patenting is associated with a 0.6 to 0.8 percent increase in the focal firm's own patenting, and this relationship is robust to excluding direct patent-citation ties. This evidence motivates a trading-relationship channel operating through production networks. The first two facts discipline the network-formation technology; the third indicates that production-network relationships are relevant for innovation incentives. Our data contribution is to join firm-to-firm transaction networks—which prior empirical work has used mainly to study short-run shock propagation (Boehm et al., 2019; Carvalho et al., 2021; Bernard et al., 2022; Fujii et al., 2017; Bai et al., 2024)—with patent records over more than two decades, allowing us to study the longer-run joint dynamics of production networks and innovation.

We build these facts into a dynamic general-equilibrium model. Firms accumulate supplier and buyer links through random meetings and relationship persistence. Incumbent R&D creates new product lines, and our central assumption—network inheritance—is that a new product line uses the parent firm's existing relationships: it is produced with established suppliers and sold to established buyers. Entrants instead create age-zero product lines with only the initial buyer-supplier links assigned at entry. Entry and incumbent R&D draw on the same fixed aggregate R&D resource, so the model isolates the allocation of R&D effort rather than its aggregate level. To keep the problem tractable, the expected portfolio of trading relationships is summarized by age-indexed supplier and buyer stocks, preserving the key empirical margins without tracking all network identities.

The theory isolates two reasons the decentralized R&D allocation departs from the planner's. First, even for a fixed network, private product-line payoffs are shaped by markups and the contractual split of bilateral operating surplus, and need not equal the social value of the line's output. Second, innovation and entry generate a dynamic network-formation externality. New product lines bring current and future buyer-supplier matches, and private agents do not fully internalize the surplus these matches generate for buyers through additional input varieties, raising the social return to product-line creation. But product-line creation also changes the aggregate product-line stock that governs other firms' matching opportunities; when the aggregate matching technology exhibits congestion, private agents create too many product lines. The sign and size of the net distortion are therefore quantitative questions.

We estimate the model on Japanese manufacturing data using four moments: incumbent R&D intensity, buyer-supplier links per firm, the life-cycle log-degree gap between young and mature firms, and the young-firm employment share. These moments discipline, respectively, the scale of incumbent R&D, the level and life-cycle slope of network formation, and the entry technology. The estimated model also reproduces untargeted life-cycle profiles of network growth and age assortativity.

Quantitatively, the baseline specification implies that congestion in the matching technology is the dominant force. Entry has two opposing network effects. On the one hand, new firms and new product lines create additional trading opportunities. On the other hand, because matching opportunities depend on the aggregate product-line stock, additional product-line creation reduces other firms' matching rates when the matching technology is congested. Private agents do not internalize this congestion cost. The planner therefore cuts entry by 25.4% and lowers the entry share of R&D labor from

25.4% to 14.2%, raising consumption-equivalent welfare by 1.71% along the transition. A decomposition shows that most of this gain comes from reducing congestion, with a smaller positive contribution from correcting static surplus-appropriation distortions. The positive match-value externality pushes in the opposite direction: taken alone, the non-internalized surplus buyers obtain from additional input varieties would imply under-entry, as in standard expanding-variety models. In the baseline specification, however, congestion more than offsets this force, so the net entry distortion is over-entry. Splitting the overall gain across allocation margins, about two-thirds comes from reducing entry and the remaining one-third from reallocating incumbent R&D from young firms toward older incumbents with accumulated relationships. A simple policy experiment within the baseline specification reinforces this interpretation: a constant entry tax equal to 83.5% of entry costs captures 94.4% of the planner's transition gain by directly counteracting excessive entry. This policy does not implement the planner's age-dependent incumbent R&D allocation, and the over-entry result should be interpreted as a statement about how a fixed aggregate R&D resource is allocated between entry and incumbent R&D, not as a claim that the aggregate level of innovation is too high. Sensitivity exercises show that the magnitude of over-entry and the associated welfare gains, and hence the case for a large entry tax, are substantially smaller under lower matching-stock elasticities.

This paper contributes to three literatures. First, it extends the macroeconomics of production networks (Acemoglu and Carvalho, 2012; Liu, 2019; Baqaee and Farhi, 2020; Bigio and La'O, 2020; Osotimehin and Popov, 2023; Baqaee et al., 2025) by adding a dynamic R&D allocation margin to existing analyses of static allocative efficiency, shock propagation, and aggregate productivity. In our setting, accumulated buyer-supplier relationships shape the private and social returns to creating new product lines and new firms. The paper therefore links production networks to models of innovation, firm dynamics, and endogenous growth (Romer, 1990; Klette and Kortum, 2004; Akcigit and Kerr, 2018).

Second, the paper relates to work on the endogenous evolution of production networks (Oberfield, 2018; Acemoglu and Azar, 2020; Huneeus, 2020; Bernard et al., 2022; Eaton et al., 2022; Arkolakis et al., 2023). Closest to our modeling strategy are Aekka and Khanna (2025) and Asai and Nirei (2026), which use firm age to summarize life-cycle network accumulation.[1] Relative to this work, we use accumulated trading relationships as state variables for R&D and entry decisions. Our focus is therefore not on how firms search for trading partners, but on how existing relationships affect the returns to incumbent R&D and entry.

[1] Aekka and Khanna (2025) is contemporaneous work using a related age-state representation, while Asai and Nirei (2026) is a subsequent study in this line of research.

Third, we contribute to the literature on innovation externalities and R&D misallocation. Existing mechanisms emphasize knowledge spillovers across sectors or technology fields (Cai and Li, 2019; Liu and Ma, 2023; Cai and Tian, 2024), or firm heterogeneity, rents, and reallocation forces (Acemoglu et al., 2018; Aghion et al., 2025; Ayerst, 2023). Our mechanism instead operates through a trading-relationship channel in production networks. Existing suppliers affect the cost of producing new product lines, existing buyers affect the demand for those lines and the surplus they generate, and R&D and entry change both the non-internalized value of buyer-supplier matches and congestion in matching.

The rest of the paper is organized as follows. Section 2 presents the data and empirical findings on the relationship between production networks, firm age, and R&D. Section 3 develops a dynamic model of production networks, incumbent R&D, and entry based on these empirical findings. Section 4 characterizes the R&D and entry wedges. Section 5 estimates the model parameters and quantifies the transition welfare decomposition, allocation counterfactuals, and policy exercises. Finally, Section 6 concludes the paper.

2 Motivating Facts on Production Networks and Innovation

Using matched firm-to-firm transaction and patent data, we document two sets of facts that motivate the model and discipline key moments. First, life-cycle patterns show that firms' buyer-supplier networks expand with age, link growth slows later in life, and positive age assortativity mainly reflects persistent relationships. Second, firm-level regressions show that firms connected to trading partners with higher patenting also patent more themselves, and instrumental-variable estimates support a network-based interpretation of this relationship.

2.1 Data and Sample Construction

Our analysis relies on three main datasets: the TDB Center for Advanced Empirical Research on Enterprise and Economy (TDB-CAREE) archive, the IIP Patent Database, and the OECD Triadic Patent Families (TPF) database.

We obtained firm-level data from TDB-CAREE at Hitotsubashi University. The data originate from credit research reports compiled by Teikoku Databank (TDB), a major private credit research company. The dataset contains firm characteristics such as firm age, industry classification (based on codes that closely match the Japan Standard Industrial Classification), location of headquarters, number of employees, and sales revenue. Crucially for our analysis, each firm also reports the identities of its major suppliers and buyers, allowing us to construct the production network of buyer-supplier links. Because TDB-CAREE curates the full TDB universe, the coverage extends beyond publicly listed firms and is broader than sources such as Compustat.

The second dataset is the IIP Patent Database, developed for patent statistical analysis using standardized data from the Patent Office. Because the IIP Patent Database does not contain the firm identifiers needed for a direct merge with TDB, we first link it to the NISTEP database, which connects patent records to company names and identification numbers. Identification numbers from NISTEP then allow us to merge the IIP Patent Database with TDB. For firms without identification numbers, we supplement the match using firm names and addresses.

The third dataset is the OECD Triadic Patent Families (TPF) database, which combines patent applications filed with the European Patent Office (EPO), the Japan Patent Office (JPO), and the United States Patent and Trademark Office (USPTO) into patent families based on common priority applications. The primary data source for the TPF is the EPO's Worldwide Patent Statistics Database, which provides harmonized and comparable information on patents from the EPO, JPO, and USPTO. We use information on the IPC (four digits), the year of application (earliest filing date at the Japan Patent Office), and the nationality of the applicant for each patent.

The datasets play different roles in the analysis. The life-cycle facts below use manufacturing focal firms in the TDB-based firm and network data over 1998–2019. Supplier and buyer counts include all observed trading partners, regardless of the partners' industries; only the focal firm is restricted to manufacturing. The age profiles of patenting and R&D in the appendix use the same focal-firm convention. The firm-level partner-patenting regressions use the matched TDB-IIP sample over 1994–2019. The OECD TPF data enter only through the construction of foreign-patent instruments in the IV specifications. Throughout the paper, suppliers are upstream trading partners and buyers are downstream trading partners. In the TDB transaction records and regression tables, seller-side variables refer to the focal firm's suppliers.

2.2 Life-Cycle Network Patterns

The life-cycle evidence has two main components. Firms accumulate both suppliers and buyers as they age, and partner ages rise with firm age mostly because links persist. Appendix A.1 reports complementary age profiles of patenting and R&D intensity.

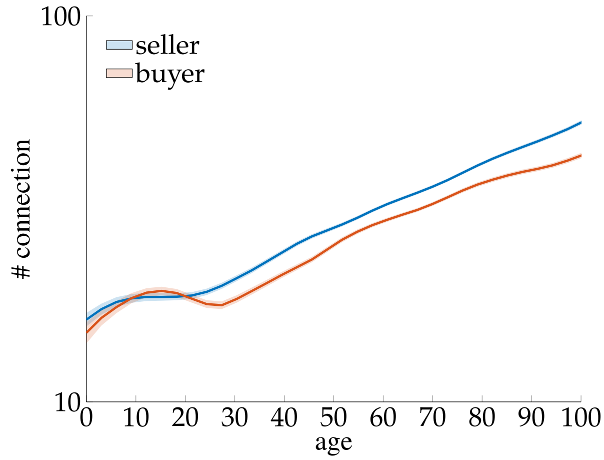

(a) Number of Links by Age

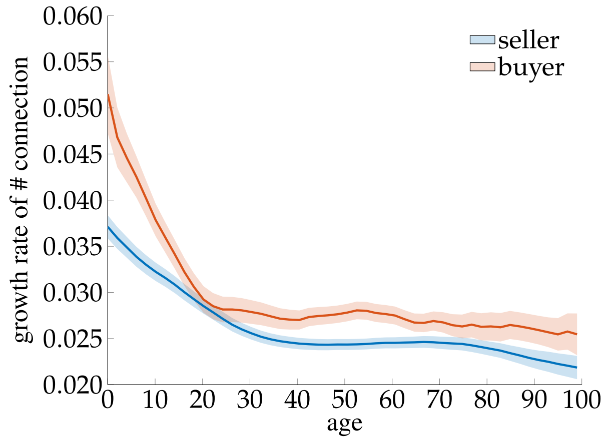

(b) Growth Rate of Links by Age

Figure 1: Life-Cycle Link Accumulation

Notes: The figure shows local linear regression estimates relating manufacturing focal firms' supplier and buyer links to firm age, with controls. The left panel plots the number of links on a log scale, and the right panel plots the corresponding growth rate. The growth rate is measured as the age derivative of log links. The sample covers 1998–2019. Trading partners are not restricted to manufacturing. The control terms include prefecture, industry, and year fixed effects. The supplier and buyer series are estimated separately. Shaded areas indicate one-standard-error bands.

2.2.1 Link Accumulation over the Firm Life Cycle

We begin by analyzing the cross-sectional relationship between firm age and the number of trading partners. Figure 1a shows local linear regressions of the numbers of suppliers and buyers on firm age. Both series increase monotonically with age. After a steep increase in the first half of the life cycle, link counts continue to rise more gradually. Figure 1b plots the corresponding growth rates, confirming that link growth is high early in the life cycle and then declines toward a low, nearly constant rate. The number of suppliers tends to slightly exceed the number of buyers, but the slopes are similar. This pattern motivates a model in which firms continue to accumulate links as they age, but at a diminishing rate.

2.2.2 Age Assortativity and Relationship Persistence

Next, we turn to the ages of trading partners conditional on firm age. Figure 2a plots trading partners' ages against firms' own ages. Older firms tend to be linked to older partners, indicating positive age assortativity in the cross-section.

To distinguish age-selective matching from relationship persistence, we compare the stock of ongoing relationships with link flows, defined as newly formed and disappearing links. Figure 2b plots partners' ages against firms' own ages among newly formed and disappearing links. In the stock, partner age rises strongly with own age; among link

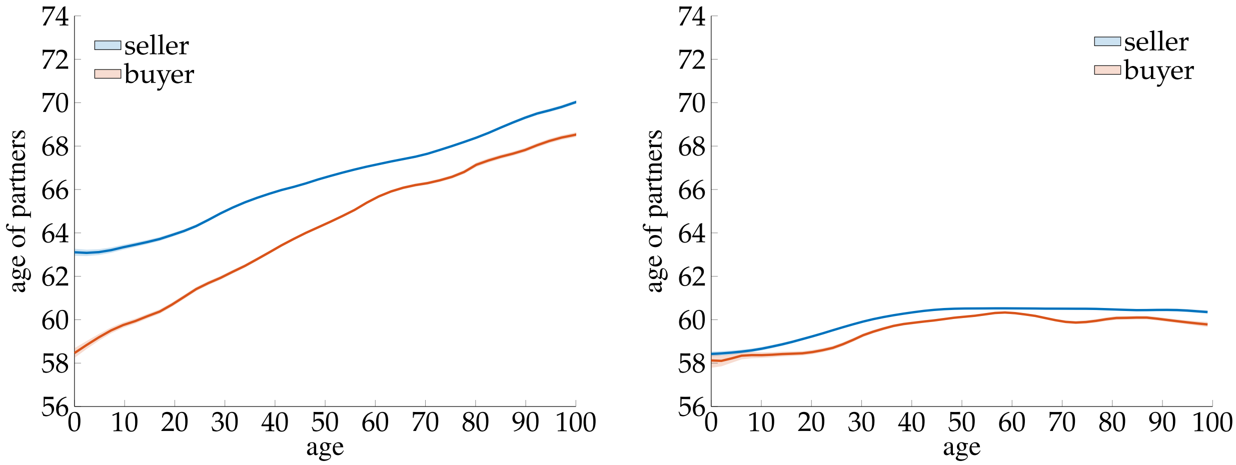

(a) Partners' Age by Firm Age

(b) Partners' Age in Flows by Firm Age

Figure 2: Positive Age Assortativity

Notes: The figure shows local linear regression estimates relating trading partners' ages to the age of manufacturing focal firms, with controls. The left panel uses the stock of observed trading relationships, while the right panel uses link flows. Link flows include newly formed and disappearing links. The sample covers 1998–2019. Trading partners are not restricted to manufacturing. The control terms include prefecture, industry, and year fixed effects. The supplier and buyer series are estimated separately. Shaded areas indicate one-standard-error bands.

flows, differences across own-age groups are only one to two years. Together with the high link survival documented in Figure A2, this pattern indicates that assortativity primarily reflects persistent relationships: firms and their partners age together.

2.3 Production-Network Exposure and Firm Patenting

We next provide firm-level evidence that production-network relationships are relevant for firms' innovation incentives. We combine firm-to-firm transaction data from TDB with patent information from the IIP Patent Database maintained by the JPO to examine whether the patenting of network-connected firms is associated with a focal firm's own patenting. The baseline matched sample spans 1994–2019 and contains 53,553 firm-year observations; the robustness sample that excludes direct patent-citation relationships contains 51,338 observations.

The matched firm-level setting has two advantages. First, instead of aggregated industry- or technology-class variables, we exploit observed transaction links between firms. Second, we can separate production-network links from direct patent-citation ties at the firm level.

We estimate regressions on firm-level data from 1994 to 2019, measuring each firm's patenting by the number of patent applications and identifying buyer-supplier links through the TDB database. The object of interest is whether a firm patents more when its suppliers (labeled seller-side partners in the transaction records) or buyers patent more, after absorbing permanent firm differences and aggregate year shocks. Let P_{it} denote patent applications by firm i, P^S_{it}/S_{it} average supplier-side partner patenting, and P^B_{it}/B_{it} average buyer-side partner patenting. Our baseline specification is

$$\log P_{it} = \beta_1 \log\left(\frac{P^S_{it}}{S_{it}}\right) + \beta_2 \log\left(\frac{P^B_{it}}{B_{it}}\right) + \alpha_i + \delta_t + \varepsilon_{it},$$

where α_i and δ_t are firm and year fixed effects. Because the specification is estimated in logs, we restrict the sample to firm-year observations with strictly positive own patent counts and strictly positive supplier-side and buyer-side partner-patenting measures.

To mitigate endogeneity concerns, we instrument partner patenting with exposure to foreign patents. The instruments use the pre-existing patent-class composition of a firm's trading partners to predict how strongly those partners are exposed to foreign innovation shocks:

$$Z^S_{it} = \sum_c \frac{P^S_{ic,t-1}}{P^S_{i,t-1}} \times P^{\text{for}}_{c,t-1},$$

$$Z^B_{it} = \sum_c \frac{P^B_{ic,t-1}}{P^B_{i,t-1}} \times P^{\text{for}}_{c,t-1}.$$

Here, $P^{\text{for}}_{c,t-1}$ is the number of triadic patents in class c filed at time $t-1$ at the Japan Patent Office by foreign firms, after excluding Japanese applicants. The weights are partners' lagged patent-class shares. When lagged total supplier-side or buyer-side patents are zero, the corresponding shares are set to zero so the shift-share instruments remain well defined. The identifying variation comes from differential exposure of pre-existing trading partners to global technology movements rather than to the focal firm's own contemporaneous demand or technology shocks.

Table 1 reports ordinary least squares (OLS) and instrumental variable (IV) estimates. In the baseline sample, the coefficients on both supplier-side and buyer-side partner patenting are positive and precisely estimated. The OLS coefficients are around 0.02, while the IV coefficients range from 0.062 to 0.077. Because the specification is log-log, the IV estimates imply that a 10 percent increase in trading partners' average patenting is associated with roughly a 0.6 to 0.8 percent increase in the focal firm's own patenting. First-stage F-statistics range from 373.5 to 840.4, indicating strong instruments. Taken together, these estimates suggest that firms patent more when their trading partners patent more.

To address the concern that the results reflect direct patent-citation channels, Table A1 repeats the exercise after excluding direct patent-citation links between the applicant and

Table 1: Trading-Partner Patenting and Own Patenting

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Seller patenting per seller	0.020***		0.019***	0.068***		0.062***
	(0.003)		(0.003)	(0.018)		(0.018)
Buyer patenting per buyer		0.022***	0.020***		0.077***	0.073***
		(0.003)	(0.003)		(0.021)	(0.021)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
First-stage F				840.4	759.8	373.5
Observations	53,553	53,553	53,553	53,553	53,553	53,553

Notes: The dependent variable is log patent applications. Seller-side patenting is supplier patenting per supplier, and buyer-side patenting is buyer patenting per buyer; both are measured in logs. The sample covers 1994–2019 and is restricted to firm-year observations with strictly positive own patent counts and strictly positive seller-side and buyer-side partner-patenting measures. The excluded instruments in the IV columns are shift-share exposures to foreign triadic patents by patent class. First-stage F denotes the first-stage F-statistic reported by the IV estimation routine. Parentheses report standard errors clustered at the industry × year level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

its buyers or suppliers. The coefficients remain sizable and statistically significant. This pattern suggests that the baseline relationship is not driven solely by direct patent-citation links, and is consistent with production-network mechanisms.

These empirical findings motivate the main network and incumbent R&D ingredients of the model: gradual link accumulation, substantial relationship persistence, and production-network relationships that matter for innovation incentives.

3 A Model of Production Networks, R&D, and Entry

Building on the empirical findings from Section 2, we develop a model that captures the interaction between production networks, incumbent R&D, and entry. The model combines the product-line innovation problem in Klette and Kortum (2004) with endogenous firm entry and a firm-to-firm matching process that replicates the age-dependent production-network patterns observed in the data.

The model has three blocks. The innovation block allocates the fixed R&D resource between incumbent product-line expansion and entrant product-line creation. The matching block maps firm age and relationship persistence into age-indexed buyer-supplier links. The static production block maps those links into costs, revenues, bilateral surplus, and private product-line payoffs. A key link between the innovation and matching blocks is network inheritance: successful incumbent R&D adds a product line that uses the innovating firm's existing supplier and buyer relationships.

3.1 Model Environment

A representative infinitely lived household supplies two distinct labor endowments: one unit of production labor and one unit of R&D labor. Production labor is used in goods production and earns wage $w(t)$, while R&D labor is used only for innovation and earns wage $w_H(t)$.[2]

Households have preferences over a final consumption good,

$$U_0 = \int_0^\infty \exp\left(-\rho t\right) \log Y\left(t\right) dt,$$

where $\rho > 0$ is the discount rate and $Y(t)$ is consumption. The household faces the budget constraint

$$\dot{A}(t) \leq r(t)A(t) + w(t) + w_H(t) + \Pi(t) - P^F(t)Y(t),$$

together with the standard no-Ponzi condition, where $A(t)$ is the asset position, $r(t)$ is the interest rate, $w(t)$ and $w_H(t)$ denote wages for production and R&D labor, respectively, $P^F(t)$ is the final-consumption price index, and $\Pi(t)$ denotes aggregate net payouts from firms owned by the household. Since these payouts are lump-sum from the household's perspective, they do not affect the Euler equation. We normalize nominal final expenditure to one. The Euler equation then implies $r(t) = \rho$.

We characterize a stationary equilibrium and suppress time subscripts when there is no ambiguity.

Product Lines, Firms, and Production Networks

There is a continuum of differentiated product lines, indexed by $\omega \in \Omega(t)$. A product line is a differentiated variety; the measure of product lines evolves through creation and exit. Each product line ω is produced by a monopolist and can be used both as an intermediate input and as a final-consumption variety. A given firm may own and produce multiple

[2] The unit supply of R&D labor is a normalization: if the physical aggregate endowment were different from one, measuring R&D labor in units of that endowment would absorb its scale into the R&D efficiency parameter. Fixing the supply of R&D labor abstracts from the aggregate underinvestment margin that would arise if production workers could perform R&D or if output were directly convertible into R&D. This focus on the allocation of a fixed aggregate R&D resource follows recent work on R&D allocation and misallocation (e.g., Liu and Ma, 2023).

product lines simultaneously. A firm's age is its time since entry. Consider a firm $\iota \in \mathcal{F}$ that owns product line ω. Let $n(\iota)$ denote the number of product lines owned by firm ι. We drop the firm subscript when no confusion arises. Denote by $\mathcal{S}(\omega) \subset \Omega$ the set of product lines used as inputs for product line ω. Although buyers can be inferred from the inverse mapping of $\mathcal{S}(\cdot)$, for notational simplicity we write $\mathcal{B}(\omega) \subset \Omega$ for the set of buyers of product line ω. Once we impose network inheritance below, a product line will be summarized by the age of its owner firm rather than by its own time since birth.

The empirical network is observed at the firm-pair level, but the model tracks trading opportunities at the product-line level. When comparing the model with the data, product-line link masses are aggregated and normalized by firm mass to construct model counterparts to firm-level buyer-supplier link moments.

Production and Contracting Given Networks

Firms use production workers and intermediate inputs for production. Intermediate inputs are imperfect substitutes with a constant elasticity of substitution, $\sigma > 1$. Production workers and the composite of intermediate inputs are combined in a Cobb-Douglas aggregator with labor share $\beta \in (0, 1)$:

$$x(\omega) = \frac{1}{\beta^{\beta} \left(1 - \beta\right)^{1-\beta}} l\left(\omega\right)^{\beta} \left(\int_{\omega' \in \mathcal{S}(\omega)} x(\omega', \omega)^{\frac{\sigma-1}{\sigma}} d\omega'\right)^{\frac{\sigma}{\sigma-1}(1-\beta)},$$

where $l\left(\omega\right)$ is the demand for production workers used to produce product line ω, and $x(\omega', \omega)$ is the demand for product line ω' used to produce product line ω. The integral is taken only over linked suppliers in $\mathcal{S}(\omega)$, so a firm's network directly determines the set of intermediate varieties available for production.

A representative final-good producer aggregates product lines with the same CES elasticity of substitution σ as in the intermediate-input aggregator:

$$Y = \left(\int_{\omega \in \Omega} y(\omega)^{\frac{\sigma-1}{\sigma}} d\omega\right)^{\frac{\sigma}{\sigma-1}},$$

where $y(\omega) = Y \left(\frac{p^F(\omega)}{P^F}\right)^{-\sigma}$ is the final demand for product line ω. Gross output of a product line is used both as an input by its buyers and as a final consumption variety. Thus, for all ω, $x(\omega)$ satisfies the following market clearing condition:

$$x(\omega) = \int_{\omega' \in \mathcal{B}(\omega)} x(\omega, \omega') d\omega' + y(\omega),$$

The baseline specification uses a two-part-tariff contracting environment on buyer-supplier links. The maintained contracting assumptions are that marginal prices govern static quantities and that fixed transfers allocate the resulting bilateral operating surplus. Let $c(\omega)$ denote unit production cost. Intermediate-input transactions use the marginal price

$$q^M(\omega) = \mu^M c(\omega),$$

while final-consumption sales use

$$p^F(\omega) = \mu^F c(\omega).$$

When $\mu^M \neq \mu^F$, contracts are end-use contingent: firms purchase intermediate inputs at intermediate-input prices, households purchase final varieties at final-consumption prices, and resale across uses is ruled out. The fixed-transfer rule is an assumed surplus split with supplier share $\theta \in [0, 1]$. Since these transfers are fixed with respect to contemporaneous input quantities, they affect private product-line payoffs but not the static input allocation or aggregate resource feasibility. The baseline specification sets $\mu^M = 1$, $\mu^F = \sigma/(\sigma - 1)$, and $\theta = 1/2$.

Innovation, Entry, Network Inheritance, and Exit

Firms conduct R&D to acquire new product lines.[3] Our central reduced-form assumption is that successful innovation reuses the incumbent firm's trading relationships rather than initiating a new search process.

Assumption 1. *For a new product line ω' derived from the existing firm's product line ω,*

$$\mathcal{S}\left(\omega'\right) = \mathcal{S}\left(\omega\right),$$

$$\mathcal{B}\left(\omega'\right) = \mathcal{B}\left(\omega\right).$$

Thus, a new product line inherits the parent line's suppliers and buyers. This makes the value of incumbent R&D depend on the firm's accumulated trading relationships, while keeping all product lines within an age-a firm exposed to the same network environment.

The total R&D cost of a firm with n product lines that chooses per-product-line

[3] Our model is based on Klette and Kortum (2004), but abstracts from productivity differences and growth. Accordingly, the model treats successful R&D as product-line creation rather than replacement through creative destruction. In a network economy, a technologically superior new product line need not displace the old product line if the old producer has access to better suppliers or different buyers.

innovation rate λ follows Klette and Kortum (2004):

$$\tilde{C}_R(\lambda, n) = n w_H h_H(\lambda),$$

where $h_H(\lambda) = \frac{1}{\varphi}\lambda^\gamma$, $\varphi > 0$ is the efficiency parameter, λ is the per-product-line innovation rate, and $\gamma > 1$.[4]

New firm creation uses the R&D labor resource. A flow E of new firms enters, with aggregate R&D labor requirement

$$H_E(E) = \frac{E^{\gamma_E}}{\varphi_E}, \qquad \gamma_E > 1,\ \varphi_E > 0.$$

Each entrant starts with one age-zero product line. The parameter ζ_0 controls the initial buyer-supplier link intensity at entry, scaled by the matching shifter introduced below. The baseline specification sets $\gamma_E = \gamma = 2$, so entry and incumbent R&D have the same R&D-cost curvature while their efficiencies are estimated separately.

Let ζ denote the primitive matching efficiency; actual link formation scales with the matching shifter $\Gamma(\mathcal{N})$ introduced below. Buyer-supplier relationships are severed at rate $\delta_M > 0$ and are also removed when an endpoint firm exits at rate δ_F. Firms exit exogenously at rate δ_F. Product lines exit exogenously at rate $\delta_P > 0$; product-line exit removes the line and its attached trading links but does not by itself change the mass of firms.[5] After entry, firms face the same matching, innovation, product-line exit, and firm-exit processes as incumbents.

3.2 Stationary Decentralized Equilibrium

We characterize the stationary decentralized equilibrium in four steps. Entry and incumbent R&D determine the product-line density over firm age. The matching technology maps this density into age-indexed buyer-supplier links. The static contracting block maps those links into unit costs, revenues, bilateral surplus, and private product-line payoffs. The dynamic problem then maps those payoffs back into incumbent R&D and entry choices.

[4] This product-line-level cost function can be micro-founded by a constant-returns-to-scale innovation technology that uses R&D workers and the number of product lines as inputs. Namely, a firm ι hires $l_H(\iota)$ units of R&D workers to add one more product line at the flow rate $\Lambda(\iota) = n(\iota)^{1-1/\gamma}\left(\varphi l_H(\iota)\right)^{1/\gamma}$, where Λ is the firm-level flow rate.

[5] For accounting, a firm that loses all product lines remains in the firm population as an inactive zero-product firm, with $V^F(0, a) = 0$ and no production or R&D activity until firm exit.

Product-Line Distribution and Matching Shifter

Without the inheritance assumption, the state would include each firm's full network of suppliers and buyers. Inheritance collapses that state space. Because every product line owned by an age-a firm faces the same partner-age distribution, we index product lines by the age of their owner firm, $a \geq 0$, rather than by the line's own time since birth.

Let $f(a,t)$ denote the density of product lines owned by age-a firms at time t, $F(a,t) \equiv \int_0^a f(\tilde{a},t)d\tilde{a}$ the corresponding cumulative measure, and $\mathcal{N}(t) \equiv \int_0^\infty f(a,t)da$ the aggregate product-line stock. In the stationary equilibrium we drop time arguments and write $f(a)$, $F(a)$, and $\mathcal{N}$. Throughout the paper, $dF(a)$ denotes the measure $f(a)da$ whenever it appears under an integral.

Entrants arrive at flow rate E, and each entrant starts with one age-zero product line. Let N_f denote the mass of surviving firms. In stationarity,

$$N_f = \frac{E}{\delta_F}. \tag{1}$$

Firm mass is retained as a reporting and empirical-normalization object. It maps model link masses into links per firm in the calibration, but it is not an aggregate matching state in the baseline theory. The stationary density of product lines owned by age-a firms satisfies

$$\frac{\partial f\,(a)}{\partial a} = (\lambda(a) - \delta_F - \delta_P)\, f\,(a)\,, \qquad a > 0, \tag{2}$$

together with the boundary condition

$$f(0) = E. \tag{3}$$

The term $\lambda(a)f(a)$ does not enter as an inflow at age zero because age indexes the owner firm rather than the product line itself: successful R&D by an age-a firm creates an additional product line owned by an age-a firm. For a given innovation profile,

$$f(a) = E \exp\left\{\int_0^a (\lambda(s) - \delta_F - \delta_P)\, ds\right\}.$$

Thus E controls the scale of the product-line density, while $\lambda(a)$ controls the age profile.

Relationships are observed at the firm-pair level, whereas the search opportunities behind those relationships are product-line needs and capabilities. The baseline specification writes aggregate matching as a function of the aggregate product-line stock,

$$\Gamma(\mathcal{N}) = \left(\frac{\mathcal{N}}{\bar{\mathcal{N}}}\right)^{-\eta}, \qquad \eta \geq 0, \tag{4}$$

where $\bar{\mathcal{N}}$ is the decentralized-baseline normalization. Here Γ is a normalized matching shifter, not the aggregate matching function itself. For example, if aggregate product-line meetings scale as $\mathcal{N}^{2-\eta}$, then dividing by buyer product lines and by the supplier-age distribution delivers an age-a_s supplier access term proportional to $\mathcal{N}^{-\eta} f(a_s)$, which is normalized here as $\zeta\Gamma(\mathcal{N}) f(a_s)$. The baseline specification sets $\eta = 1$, corresponding to constant returns in aggregate product-line meetings and a scale-invariant per-product-line matching rate. The shifter Γ multiplies interior link formation and entrant boundary links.

Stationary Matched-Product Density

Lemma 1 (Stationary Matched-Product Density and Reciprocity)**.** *Let $m(a_s, a_b)$ denote the product-line mass of suppliers of age a_s attached to one buyer product line owned by an age-a_b firm. Equivalently, $m(a_s, a_b)\, da_s$ is the mass of supplier product lines whose owner firms have ages in $[a_s, a_s + da_s]$ and that are connected to one buyer product line of age a_b. The first argument indexes supplier age and the second indexes buyer age. Then $m(a_s, a_b)$ satisfies*

$$\underbrace{\frac{\partial}{\partial a_s} m(a_s, a_b) + \frac{\partial}{\partial a_b} m(a_s, a_b)}_{\text{time evolution of the matched-product density}}$$
$$= \underbrace{-\left(\delta_M + \delta_F + \delta_P\right) m(a_s, a_b)}_{\text{link and supplier-line destruction}} + \underbrace{\zeta\Gamma f(a_s)}_{\text{random matching}} + \underbrace{\lambda\left(a_s\right) m(a_s, a_b)}_{\text{inherited links from supplier R\&D}}, \qquad a_s, a_b > 0, \tag{5}$$

where $\lambda\left(a_s\right)$ is the innovation rate of an age-a_s supplier firm and $\Gamma = \Gamma(\mathcal{N})$, with two entry boundaries:

$$m(a_s, 0) = \zeta_0 \Gamma f(a_s)$$
$$m(0, a_b) = \zeta_0 \Gamma f(0) = \zeta_0 \Gamma E.$$

The first boundary condition assigns suppliers to entrant buyers according to the incumbent product-line density. The second assigns entrant suppliers to incumbent buyer lines at a rate *proportional to the entry flow. Given (f, Γ), the stationary closed form is*

$$m(a_s, a_b) = \left[\frac{\zeta}{\delta_M} - \left(\frac{\zeta}{\delta_M} - \zeta_0\right) \exp\{-\delta_M \min\{a_s, a_b\}\}\right] \Gamma f(a_s). \tag{6}$$

This closed form implies

$$\frac{m(a_s, a_b) f(a_b)}{f(a_s)} = m(a_b, a_s),$$

because the conditional kernel $m(a_s, a_b)/f(a_s)$ depends only on $\min\{a_s, a_b\}$.

For a given buyer age a_b, the law of motion for the interior matched-product density has three components. First, the mass of linked supplier product lines increases when an age-a_s supplier's R&D creates a new product line at rate $\lambda(a_s)$ and the new product line inherits the supplier's existing buyer links. Second, linked supplier product lines are lost when the relationship is severed at rate δ_M, when the supplier firm exits at rate δ_F, or when the supplier product line exits at rate δ_P.[6] Third, random matching adds age-a_s supplier product-line mass at rate $\zeta \Gamma f(a_s)$. The age weighting is therefore already embedded in the product-line mass $f(a_s)$: meetings are not uniform over the unbounded age domain. Entry operates through the inflow boundaries. New buyers at $a_b = 0$ draw $\zeta_0 \Gamma f(a_s)$ supplier product-line mass from age-a_s suppliers. Similarly, entrant suppliers have one product line and are connected to existing buyer product lines at rate $\zeta_0 \Gamma E$.

This aggregation avoids tracking each firm's full neighborhood. Conditional on buyer age, a product line faces a deterministic age-indexed mass of supplier product lines. The reciprocity relation converts supplier access per buyer line into buyer access per supplier line. The normalization is per buyer product line throughout: multiplying $m(a_s, a_b)$ by the mass of buyer product lines of age a_b would recover the total mass of links from age-a_s suppliers to age-a_b buyers. Accordingly, when an age-a_s line is viewed as a supplier, the buyer mass of age a_b attached to it is written as $m(a_b, a_s)$.

Static Contracting Block

Taking the product-line density and matched-product density (f, m) as given, the decentralized static contracting block maps the age-indexed network into unit costs, revenues, bilateral surplus, and the private flow payoff $\pi(a)$ that enters the dynamic R&D problem. The unit cost $c(a)$ summarizes the marginal cost of producing a product line, while $K^{DE}(a)$ is the gross-output cost base of that product line, normalized by final expenditure.

[6] Buyer-side firm or product-line exit removes the buyer line itself and is handled by the product-line density $f(a_b)$; it is not an additional loss term in $m(a_s, a_b)$, which is measured per surviving buyer product line.

Define

$$q^M(a) \equiv \mu^M c(a), \qquad p^F(a) \equiv \mu^F c(a),$$

where $q^M(a)$ is the intermediate-input marginal price and $p^F(a)$ is the final-consumption price. Unit costs satisfy

$$c(a_b) = w^\beta \left(\int q^M(a_s)^{1-\sigma} m(a_s, a_b)\, da_s \right)^{\frac{1-\beta}{1-\sigma}}. \tag{7}$$

The final-consumption price index and final-sales revenue, normalized by final expenditure, are

$$P^F = \left(\int p^F(a)^{1-\sigma} f(a)\, da \right)^{\frac{1}{1-\sigma}},$$
$$R^{F,DE}(a) = \left(\frac{p^F(a)}{P^F} \right)^{1-\sigma}.$$

The intermediate input share of an age-a_s supplier in an age-a_b buyer's input bundle is

$$s^M(a_s, a_b) = \frac{q^M(a_s)^{1-\sigma}}{\int q^M(\tilde{a}_s)^{1-\sigma} m(\tilde{a}_s, a_b)\, d\tilde{a}_s}.$$

By the reciprocity relation above, the mass of age-a_b buyers attached to an age-a_s supplier line is $m(a_b, a_s)$. Hence intermediate-sales revenue per age-a_s product line is

$$R^{M,DE}(a_s) = (1-\beta) \int K^{DE}(a_b) s^M(a_s, a_b) m(a_b, a_s)\, da_b.$$

Because revenues include markups over variable costs, dividing each revenue component by its corresponding markup recovers the gross-output cost base. Market clearing therefore gives the cost-base recursion and the production-labor clearing condition:

$$K^{DE}(a_s) = \frac{R^{F,DE}(a_s)}{\mu^F} + \frac{R^{M,DE}(a_s)}{\mu^M},$$
$$w = \beta \int K^{DE}(a) f(a)\, da. \tag{8}$$

The marginal prices above determine quantities, costs, and revenues. The remaining part of the assumed two-part tariff is the fixed transfer. Given the static allocation, the bilateral operating surplus generated by an age-a_s supplier linked to an age-a_b buyer is

$$S^{DE}(a_s, a_b) = \frac{1-\beta}{\sigma-1} K^{DE}(a_b) s^M(a_s, a_b),$$

The contracting assumption is that the flow transfer paid by buyer a_b to supplier a_s equals $T(a_s, a_b) = \theta S^{DE}(a_s, a_b)$. Thus, fixed transfers reallocate current bilateral operating surplus between linked buyers and suppliers without changing the static input choices characterized above.

The baseline specification sets $\mu^M = 1$, so intermediate-input quantities are priced efficiently at the margin, while final-consumption sales still carry the monopolistic markup $\mu^F = \sigma/(\sigma-1)$.

The private flow payoff generated by a product line owned by an age-a firm is

$$\begin{aligned}\pi(a) = &\left(1 - \frac{1}{\mu^F}\right) R^{F,DE}(a) + \left(1 - \frac{1}{\mu^M}\right) R^{M,DE}(a) \\ &+ \theta \int S^{DE}(a, a_b) m(a_b, a)\, da_b - \theta \int S^{DE}(a_s, a) m(a_s, a)\, da_s. \end{aligned} \tag{9}$$

The first two terms are final-sales and intermediate-sales linear markup profits. The first transfer term is received when the line sells to buyers, and the second transfer term is paid when it purchases inputs from suppliers.

Firm Value, Innovation, and Entry

Since the product-line flow payoff depends on owner-firm age through the inherited network, the firm's optimization problem can also be summarized by age and the number of product lines. A firm with $n \geq 1$ product lines and age a maximizes the value $V^F(n, a)$, with $V^F(0, a) = 0$:

$$\begin{aligned}\left(\rho + \delta_F\right) V^F(n, a) = &\underbrace{n\pi(a)}_{\text{flow payoff}} + \underbrace{V_a^F(n, a)}_{\text{age effect}} \\ &+ \underbrace{n\delta_P \left\{V^F(n-1, a) - V^F(n, a)\right\}}_{\text{product-line exit}} \\ &+ \max_{\lambda \geq 0} \left[\underbrace{n\lambda \left\{V^F(n+1, a) - V^F(n, a)\right\}}_{\text{product-line expansion}} - \underbrace{n w_H h_H(\lambda)}_{\text{R\&D costs}} \right]. \end{aligned} \tag{10}$$

The term $\delta_F V^F(n,a)$ reflects exogenous firm exit, which destroys all product lines owned by the firm. On the right-hand side, the first term is total private flow payoff, the second term is the change in firm value due to aging, the third term is the value loss from product-line exit at total rate $n\delta_P$, and the final bracket is the net gain from adding new product lines through R&D.

Linearity in the number of product lines implies that firm value can be expressed as the sum of product-line values. Flow payoffs, product-line exit risk, and R&D costs all scale with the number of product lines, and network inheritance makes all product lines within an age-a firm face the same network environment. We therefore conjecture

$$V^F(n,a) = nV(a)$$

where $V(a)$ is the value of a product line owned by an age-a firm. Substituting into (10) gives the HJB equation for $V(a)$:

$$\left(\rho + \delta_F + \delta_P\right) V(a) = \pi(a) + V_a(a) + \max_{\lambda \geq 0} \left[\lambda V(a) - w_H h_H(\lambda)\right] \tag{11}$$

The left-hand side captures discounting, owner-firm exit, and product-line exit. On the right-hand side, the first term is the static private payoff from a product line, the second term is the effect of aging on product value, and the maximization term is the net value of innovation, where $\lambda V(a)$ is the value of product-line expansion and $w_H h_H(\lambda)$ is the associated R&D cost. The first-order condition yields an optimal innovation rate:

$$\lambda(a) = \left\{\frac{\varphi}{\gamma w_H} V(a)\right\}^{\frac{1}{\gamma-1}}, \tag{12}$$

Thus the innovation rate inherits age dependence from $V(a)$.

The production labor market clears inside the static contracting block through (8). Since $h_H(\lambda(a))$ is incumbent R&D labor used per product line, the aggregate R&D labor market clearing condition is

$$1 = H_E(E) + \int h_H(\lambda(a)) f(a) da = \frac{E^{\gamma_E}}{\varphi_E} + \int \frac{\lambda(a)^\gamma}{\varphi} f(a) da. \tag{13}$$

Free entry equates the private value of an entrant's initial product line to the marginal R&D labor cost of entry:

$$V(0) = w_H H_E'(E) = w_H \frac{\gamma_E}{\varphi_E} E^{\gamma_E - 1}. \tag{14}$$

The corresponding Kuhn–Tucker form is $E \geq 0$, $w_H H_E'(E) \geq V(0)$, and $E[w_H H_E'(E) - V(0)] = 0$. Equivalently, when the entry margin is interior,

$$E = \left\{ \frac{\varphi_E}{\gamma_E w_H} V(0) \right\}^{\frac{1}{\gamma_E - 1}}.$$

Equilibrium Definition

A stationary decentralized equilibrium is a collection of functions and scalars

$$\{f, m, c, K^{DE}, V, \lambda, E, w, w_H\}$$

such that:

1. Given the private flow payoff in (9), the value function V and innovation policy λ satisfy (11) and (12).

2. Entry satisfies (14).

3. Given (λ, E), the product-line density f, matching shifter $\Gamma(\mathcal{N})$, and matched-product density m satisfy (2), (3), (4), and the matching law of motion (5) with entrant boundary conditions, equivalently the closed form (6).

4. Given (f, m), the static contracting block determines the derived objects q^M, p^F, P^F, $R^{F,DE}$, $R^{M,DE}$, s^M, S^{DE}, and π through (7)–(9).

5. R&D labor clears according to (13), with production labor clearing embedded in (8).

These equilibrium objects form the baseline for the planner-private wedge decomposition in the next section.

4 R&D Allocation and Welfare Wedges

Having established the decentralized equilibrium in Section 3, we now characterize the planner comparison and the wedges that account for the welfare gap. With endogenous entry, the allocation problem has two R&D margins. The first is the aggregate entry flow, which determines how much of the fixed R&D resource is used to create new firms and product lines. The second is the incumbent R&D profile across firm ages. We decompose the planner-decentralized gap into two broad layers: a static surplus-appropriation component and a matching-network wedge. The matching-network wedge has two components, the net value of link formation captured by match costates and the matching-stock congestion externality.

To simplify notation, we suppress environment superscripts on endogenous state variables such as (f, m, c, K) when no confusion arises; in each allocation, these objects are those induced by that allocation's R&D policy and static allocation.

4.1 From Values to R&D Allocation

This subsection translates product-line and entrant values into the two R&D allocation margins used below: aggregate entry and the age profile of incumbent R&D. The decentralized equilibrium and the planner allocations share the same allocation structure once the relevant values are specified. Let j index these allocations, set $V^{DE} \equiv V$, $w_H^{DE} \equiv w_H$, and $V_E^{DE} \equiv V(0)$. In planner allocations, V_E^j denotes the entrant value implied by the boundary conditions.

The common allocation rule follows directly from the R&D cost functions and the fixed R&D resource constraint. In the decentralized equilibrium and in each planner allocation satisfying the corresponding entry and incumbent R&D optimality conditions, entry and incumbent R&D compete for the fixed R&D resource. Entry satisfies the Kuhn–Tucker condition

$$E^j \geq 0, \qquad w_H^j H_E'(E^j) \geq V_E^j, \qquad E^j \left[w_H^j H_E'(E^j) - V_E^j \right] = 0. \tag{15}$$

The residual amount of R&D labor available to incumbents is

$$H_H^j \equiv 1 - H_E(E^j).$$

Conditional on this residual labor, if incumbent R&D is interior and $h_H(\lambda) = \lambda^\gamma / \varphi$, per-product-line incumbent R&D labor for product lines owned by age-a firms satisfies

$$h_H\left(\lambda^j(a)\right) = H_H^j \frac{\left(V^j(a)\right)^{\frac{\gamma}{\gamma-1}}}{\int \left(V^j(\tilde{a})\right)^{\frac{\gamma}{\gamma-1}} f^j(\tilde{a}) d\tilde{a}}. \tag{16}$$

This characterization separates the aggregate and cross-sectional R&D margins. Entry first determines how much of the fixed R&D resource is devoted to creating new firms; the residual H_H^j is then allocated across incumbent firm ages in proportion to the relevant product-line value. In the decentralized equilibrium, $V^j = V$ is the private product-line value from the HJB in (11), and the entry condition is the free-entry condition in (14). The social planner replaces these private values with planner costates that include the static surplus-appropriation component and the matching-network wedge.

4.2 Social Planner Values

The first-best social planner maximizes welfare by internalizing all production, entry, and matching-stock effects in the economy. Unlike the decentralized equilibrium, the social planner directly controls goods production, incumbent R&D, and entry.

The social planner maximizes the discounted utility flow

$$U_0 = \int_0^\infty \exp\left(-\rho t\right) \log Y\left(t\right) dt$$

where $Y = \left(\int y(a)^{\frac{\sigma-1}{\sigma}} dF(a)\right)^{\frac{\sigma}{\sigma-1}}$. The planner chooses goods production, incumbent R&D, and entry subject to the age-notation production technology (17), goods-market clearing (18), the production labor resource constraint (19), the laws of motion (5) and (2), the entry boundary condition (3), the entrant boundary conditions for m, and the R&D resource constraint (13). Because the intratemporal allocation is static for given (f, m), the problem separates into a static goods-allocation block and a dynamic R&D-and-entry block. The full derivation is in Section F.

For a fixed network state (f, m), the planner's static goods allocation is characterized by the following constraints. Unlike the decentralized contracting block, the planner's static allocation has no markups or fixed transfers: shadow prices reflect social marginal costs and only summarize the planner's intratemporal allocation. These constraints pin down the social cost and demand shifters, and hence the planner current-return term $R^{SP}(a)$ that enters the dynamic R&D problem.

The planner's intratemporal constraints use a_s for supplier age and a_b for buyer age:

$$x(a_b) = \frac{1}{\beta^\beta \left(1-\beta\right)^{1-\beta}} l\left(a_b\right)^\beta \left(\int x(a_s, a_b)^{\frac{\sigma-1}{\sigma}} m(a_s, a_b) da_s\right)^{\frac{\sigma}{\sigma-1}(1-\beta)}, \tag{17}$$

$$x(a_s) f(a_s) = \int x(a_s, a_b) m(a_s, a_b) f(a_b) da_b + y(a_s) f(a_s), \tag{18}$$

$$\int l(a) dF(a) = 1 \tag{19}$$

Here $x(a)$ is gross output of a product line owned by an age-a firm, $y(a)$ is final-good use, $l(a)$ is production labor, and $x(a_s, a_b)$ is the quantity supplied by an age-a_s supplier product line to an age-a_b buyer product line.

The associated planner static objects can be summarized by social cost and demand shifters:

$$c^{SP}(a_b) = \left(\int c^{SP}(a_s)^{1-\sigma}\, m(a_s, a_b)\, da_s \right)^{\frac{1-\beta}{1-\sigma}},$$
$$D^{SP}(a_s) = (1-\beta) \int [c^{SP}(a_b)]^{\frac{\beta}{1-\beta}(\sigma-1)} D^{SP}(a_b)\, \frac{m(a_s, a_b) f(a_b)}{f(a_s)}\, da_b + 1.$$

The planner price index, current-return term, and input share are

$$P^{SP} = \left(\int c^{SP}(a)^{1-\sigma} f(a)\, da \right)^{\frac{1}{1-\sigma}}, \qquad R^{SP}(a) \equiv \left(\frac{c^{SP}(a)}{P^{SP}} \right)^{1-\sigma} D^{SP}(a),$$

and

$$s(a_s, a_b) \equiv \frac{c^{SP}(a_s)^{1-\sigma}}{\int c^{SP}(\tilde{a})^{1-\sigma}\, m(\tilde{a}, a_b)\, d\tilde{a}}.$$

Proposition 1 (Planner Product-Line and Match Values)**.** *Given the network state (f, m) and the planner static allocation, the social product-line value $V^{SP}(a_s)$ solves*

$$\begin{aligned} \left(\rho + \delta_F + \delta_P - \lambda^{SP}(a_s)\right) V^{SP}(a_s) = R^{SP}(a_s) &+ V^{SP}_{a_s}(a_s) - w^{SP}_H h_H(\lambda^{SP}(a_s)) \\ &+ \bar{\mathcal{R}}_{\Gamma,f}(a_s) \\ &+ \int V^M(a_s, a_b) \left\{ \zeta \Gamma f(a_b) - \delta_M m(a_b, a_s) \right\} da_b, \end{aligned} \tag{20}$$

The planner entrant value is $V^{SP}_E = V^{SP}(0)$. Boundary link values associated with entry are included in this product-line costate through the entry boundary conditions. Equivalently, after the common costate rescaling, V^M is the costate of the matched-product density normalized by the buyer product-line density. With this normalization, $V^M(a_s, a_b)$ satisfies

$$\begin{aligned} \left(\rho + 2(\delta_F + \delta_P) + \delta_M - \lambda^{SP}(a_s) - \lambda^{SP}(a_b)\right) V^M(a_s, a_b) = V^M_{a_s}(a_s, a_b) &+ V^M_{a_b}(a_s, a_b) \\ &+ (1-\beta) R^{SP}(a_b)\, s(a_s, a_b), \end{aligned} \tag{21}$$

where $V^M(a_s, a_b)$ is the social value, per age-a_b buyer product line, of expanding that buyer's input menu by one supplier age-a_s variety. The term w^{SP}_H is the shadow wage on R&D labor, and $\bar{\mathcal{R}}_{\Gamma,f}(a_s)$ is the planner's aggregate matching-stock term. See Section F for the full derivation.

4.3 Wedge Components

The proposition provides the planner value objects that can be compared with the decentralized product-line HJB in (11). Together with the R&D allocation conditions in (15) and (16), these values determine planner entry and incumbent R&D allocation. The comparison separates the planner-private value gap into a static surplus-appropriation component, defined by the gap between private product-line payoffs and the planner static return, and a matching-network wedge. The matching-network wedge is the sum of the dynamic match-costate value of net link formation and the aggregate matching-stock component.

Comparing the decentralized static contracting system in (7)–(9) with its planner counterpart above shows that the static surplus-appropriation wedge depends on current surplus appropriation, including any marginal-pricing distortions and transfer rules. In the baseline specification, intermediate-input marginal prices are efficient because $\mu^M = 1$, but final-consumption markups and fixed-transfer shares still shape private product-line payoffs. The decomposition has three components.

Static Surplus Appropriation: The first channel is the gap between private product-line payoffs and the planner static return. In the baseline specification, this wedge includes final-consumption markups and the way bilateral operating surplus is transferred across buyer-supplier relationships.

Net Link-Flow Valuation: The first matching-network component is the match-costate value of net link formation. A new product line inherits the parent line's buyer relationships and later forms additional matches, expanding buyers' input menus in both cases. Private contracts do not price the planner continuation value of these match portfolios, so private R&D and entry choices omit part of the social value of future network surplus.

Matching-Stock Congestion: The second component of the matching-network wedge operates through the aggregate product-line stock that determines the matching shifter. Entrants and incumbents take this stock as given, but the planner internalizes that additional product lines change the matching environment faced by all other product lines. Quantitatively this is the dominant force behind lower planner entry.

For the first-best social planner, define the net link-flow component

$$W_I^{SP}(a_s) = \int V^M(a_s, a_b) \left\{ \zeta \Gamma^{SP} f^{SP}(a_b) - \delta_M m^{SP}(a_b, a_s) \right\} da_b.$$

This is a net-flow value: random matching creates links to age-a_b buyer product lines at density $\zeta \Gamma^{SP} f^{SP}(a_b)$, while link destruction removes age-a_b buyer links at density $\delta_M m^{SP}(a_b, a_s)$. Product-line exit is already included on the left-hand side through δ_P, so it does not enter as a separate link-flow term. The matching-stock component is the separate term $\bar{\mathcal{R}}_{\Gamma,f}(a_s)$ in (20). With $\Gamma(\mathcal{N}) = (\mathcal{N}/\bar{\mathcal{N}})^{-\eta}$, define

$$\bar{\mathcal{R}}_{\Gamma,f}(a) = \Gamma^{\log}_{\mathcal{N}}\Omega_\Gamma, \qquad\qquad \Gamma^{\log}_{\mathcal{N}} \equiv \frac{\partial \log \Gamma}{\partial \mathcal{N}} = -\frac{\eta}{\mathcal{N}}.$$

Here the superscript "log" indicates that the derivative applies to the log matching shifter, while the derivative is taken with respect to the aggregate product-line stock $\mathcal{N}$. When match values are positive, the exposure term below is positive; therefore $\eta > 0$ makes $\bar{\mathcal{R}}_{\Gamma,f}$ negative, because an additional product line lowers the matching shifter faced by all product lines. The exposure term is

$$\begin{aligned}\Omega_\Gamma \equiv \zeta\Gamma \iint V^M(\tilde{a},\hat{a}) f(\tilde{a}) f(\hat{a})\, d\tilde{a}\, d\hat{a} \\ + \zeta_0\Gamma E \int V^M(\tilde{a},0) f(\tilde{a})\, d\tilde{a} + \zeta_0\Gamma E \int V^M(0,\hat{a}) f(\hat{a})\, d\hat{a}.\end{aligned}$$

5 Estimation and Welfare Analysis

This section quantifies the production-network R&D allocation model using Japanese firm-level data and uses the estimated decentralized equilibrium to decompose welfare losses from endogenous entry and incumbent R&D. The baseline counterfactual economy is the contracting endogenous-entry economy, in which entry and incumbent R&D draw on the same scarce R&D resource. We proceed in three steps. First, we estimate the baseline specification and summarize targeted and untargeted model validation. Second, we compare the decentralized equilibrium with the planner allocation and decompose the welfare gap using the wedges characterized in Section 4. Finally, we evaluate simple policy and allocation exercises. The estimation disciplines incumbent R&D intensity, buyer-supplier link density, life-cycle network accumulation, and the young-firm employment share.

5.1 Estimation and Model Validation

This subsection describes how we estimate the decentralized economy used in the counterfactual analysis and validate its life-cycle implications. We first specify the fixed and externally calibrated parameters that define the baseline specification, estimate the remaining parameters by minimum distance, report the targeted fit, and then compare untargeted life-cycle profiles with the empirical patterns.

5.1.1 Fixed and Externally Calibrated Parameters

The internal estimation is conditional on the static environment, externally calibrated dynamic hazards, R&D cost curvatures, and the baseline matching-stock elasticity. These values are reported in the lower panel of Table 2.

The CES elasticity is set to $\sigma = 3$, following Miyauchi (2024), and implies a final-consumption markup $\mu^F = 1.5$. The baseline specification sets intermediate marginal prices at cost, $\mu^M = 1$, and splits current bilateral operating surplus equally, $\theta = 1/2$. The discount rate is $\rho = 0.05$, following Peters and Walsh (2022). The production parameter $\beta = 0.33$ is calibrated to the primary-factor cost share in Japanese input-output data (Ministry of Internal Affairs and Communications, 2019); because the model has only production labor and intermediate inputs in variable production, it should not be interpreted as labor's share of value added.

The link destruction rate is set to $\delta_M = 0.08$, based on the annual link-stock survival calibration discussed in Appendix B.5. The firm exit rate is set to $\delta_F = 0.04$ using the sample-period average exit rate in the Ministry of Health, Labour and Welfare's *Employment Insurance Annual Report* (Ministry of Health, Labour and Welfare, various years). Product-line exit is set to $\delta_P = 0.06$, anchored by firm-product turnover evidence in Bernard et al. (2010) and the broader range suggested by firm-product switching and patent-renewal or obsolescence evidence (Deng, 2011; Pakes and Schankerman, 1984). The incumbent and entry R&D cost curvatures are fixed at $\gamma = \gamma_E = 2$, following Acemoglu et al. (2016) and related directed-technical-change calibrations. The baseline specification fixes the matching-stock elasticity at $\eta = 1$, the scale-invariant benchmark for aggregate product-line meetings. Under this value, per-product-line matching rates do not mechanically trend with the aggregate product-line stock. Equivalently, $\eta = 1$ is the value that would keep network formation stationary in a balanced-growth extension with a growing product-line stock. In the stationary quantitative model, however, the targeted moments discipline the level and life-cycle slope of matching, but not the aggregate matching-stock elasticity itself. Appendix H.4 therefore reports sensitivity to lower values, including $\eta = 0.1$, motivated by the one-stock interpretation of Miyauchi (2024)'s supplier-density estimate.

5.1.2 Target Moments and Minimum-Distance Estimation

The four target moments are incumbent R&D intensity, buyer-supplier links per firm, a controlled life-cycle log-degree gap, and the young-firm employment share. R&D intensity is measured as the R&D-to-sales ratio among R&D-performing manufacturing firms and is set to 0.042, close to the 4.2 percent manufacturing ratio reported for firms conducting research in the 2025 Survey of Research and Development (Statistics Bureau, Ministry of Internal Affairs and Communications, 2025). The life-cycle gap uses the controlled age-dummy estimates of Asai and Nirei (2026): supplier and buyer log-degree gaps for firms ages 1–5 relative to firms ages 50 and above, averaged across the two sides of the market. The young-firm employment-share target is the OECD DynEmp3 Japan manufacturing employment share of establishments younger than six years. It disciplines entry and the age distribution. Together, the targets discipline incumbent R&D effort, network scale, life-cycle network accumulation, and product-line entry; Section B provides the model-counterpart construction.

Given the fixed parameters and empirical targets described above, we estimate four structural parameters with a minimum-distance procedure. The estimated parameter vector is $\boldsymbol{\vartheta} = (\zeta, \zeta_0, \varphi, \varphi_E)$, where ζ governs network formation, ζ_0 captures entrant-boundary connectivity, φ measures incumbent R&D productivity, and φ_E measures entry R&D productivity. Economically, incumbent R&D intensity primarily disciplines φ, links per firm discipline ζ, the life-cycle log-degree gap disciplines entrant-boundary connectivity relative to subsequent matching, and the young-firm employment share disciplines entry productivity φ_E, although the parameters are estimated jointly. Let g^{data} denote the vector of empirical moments and $g^{\text{model}}(\boldsymbol{\vartheta})$ the corresponding model-implied moments. We choose $\hat{\boldsymbol{\vartheta}}$ to minimize

$$\sum_k r_k(\boldsymbol{\vartheta})^2,$$

where the residuals are log deviations for positive level and share moments and a level deviation for the life-cycle log-degree gap:

$$r_1 = \log\left(\frac{\text{RD}^{\text{model}}}{\text{RD}^{\text{data}}}\right), \qquad r_2 = \log\left(\frac{L^{\text{model}}}{L^{\text{data}}}\right), \qquad r_3 = \Delta^{\text{model}} - \Delta^{\text{data}}, \qquad r_4 = \log\left(\frac{s_Y^{\text{model}}}{s_Y^{\text{data}}}\right).$$

Here RD is incumbent R&D intensity, L is buyer-supplier links per firm, Δ is the controlled life-cycle log-degree gap for young firms, and s_Y is the young-firm employment share. For each candidate parameter vector, we solve the stationary endogenous-entry equilibrium and compute the implied moments deterministically.

5.1.3 Estimated Parameters and Model Fit

Table 2 reports the baseline parameter values. The upper panel contains the internally estimated parameters $(\zeta, \zeta_0, \varphi, \varphi_E)$, estimated jointly using the four moments in Table 3. The lower panel reports the fixed and externally calibrated parameters discussed above.

Table 2: Estimated and Calibrated Parameters

Parameter	Symbol	Value
Internally Estimated Parameters		
Network formation rate	ζ	86.5
Entrant-boundary connectivity	ζ_0	407
Incumbent R&D efficiency	φ	1.74×10^{-4}
Entry R&D efficiency	φ_E	3.27×10^{-6}
Fixed and Externally Calibrated Parameters		
CES elasticity	σ	3
Primary-factor share	β	0.33
Discount rate	ρ	0.05
Link destruction rate	δ_M	0.08
Firm exit rate	δ_F	0.04
Product-line exit rate	δ_P	0.06
R&D cost curvature	γ	2
Entry cost curvature	γ_E	2
Matching-stock elasticity	η	1
Intermediate marginal markup	μ^M	1
Final-consumption markup	μ^F	1.5
Transfer share	θ	0.5

Notes: Internal parameters are estimated by minimum distance using the four target moments in Table 3. The parameter β is the primary-factor cost share. All rates are annual continuous-time hazards. The contracting protocol sets intermediate marginal prices at cost, applies the monopolistic markup to final-consumption sales, and splits current bilateral operating surplus equally.

Table 3 compares the targeted moments with their model-implied counterparts. The estimated model closely matches the overall scale of buyer-supplier links and the young-firm employment share, while leaving small residuals in incumbent R&D intensity and the life-cycle log-degree gap.

5.1.4 Untargeted Life-Cycle Validation

The stock of accumulated relationships is the channel through which life-cycle network dynamics strengthen private innovation incentives. The estimated model generates front-loaded network accumulation, a high link stock among mature firms, and an innovation profile that rises with firm age, consistent with the empirical life-cycle patterns in Section 2. Appendix C reports the corresponding model-implied network and innovation profiles, and Appendix C.2 reports the model-implied age-assortativity profile.

These life-cycle diagnostics support the quantitative mechanism used below: accu-

Table 3: Targeted Moments and Model Fit

Moment	Target	Model
Incumbent R&D intensity	0.042	0.045
Buyer-supplier links per firm	27.3	27.3
Life-cycle log-degree gap	-0.457	-0.490
Young-firm employment share	0.096	0.099

Notes: Values are rounded for readability. The R&D-intensity target is incumbent R&D intensity and excludes entry R&D labor. The total R&D intensity including entry is 0.060, and the realized entry R&D share is 0.254. R&D intensity, buyer-supplier links per firm, and the young-firm employment share enter the objective as log deviations. The life-cycle target is the average of supplier- and buyer-side log-degree gaps for firms ages 1–5 relative to firms ages 50 and above, and enters the objective in levels.

mulated trading relationships make incumbent R&D age-dependent, while entry-driven product-line creation changes the aggregate matching environment faced by other firms. These forces determine whether scarce R&D is tilted toward entrants and young incumbents or toward older incumbents with accumulated production networks.

5.2 Planner Allocation and Welfare Decomposition

5.2.1 Planner Allocation and Transition Welfare

Since aggregate R&D labor is fixed, the planner's entry choice also determines how much R&D labor remains for incumbent innovation. Relative to the planner, the decentralized allocation puts too much of the fixed R&D resource into entry-driven product-line creation and, within incumbent R&D, too much toward young incumbents. In the baseline specification, decentralized entry is $E^{DE} = 9.11 \times 10^{-4}$, while the social planner chooses $E^{SP} = 6.80 \times 10^{-4}$. Planner entry is therefore 25.4% below decentralized entry.

Transition welfare gains are reported in CE units, where CE denotes consumption equivalent. For a computed transition path, let $\mathcal{U}_{\text{path}}$ denote discounted utility along the reform path and $\mathcal{U}_{DE}$ discounted utility from remaining in the decentralized steady state. We define

$$\log(1 + g^{\text{trans}}) = \rho \left(\mathcal{U}_{\text{path}} - \mathcal{U}_{DE} \right) . \tag{22}$$

Thus a reported transition CE gain is the permanent proportional increase in decentralized steady-state consumption that would make the household indifferent to the planner transition.

The planner also reallocates the scarce R&D resource away from entry: the entry R&D share falls from 25.4% in the decentralized steady state to 14.2% in the planner steady

Table 4: Entry and Transition-Welfare Decomposition

Component	Entry change (%)	CE gain (%)
Static surplus appropriation	-8.41	0.78
Net link-flow valuation	33.44	-0.51
Matching-stock congestion	-38.93	1.45
Total DE to full SP	-25.36	1.71

Notes: The decomposition starts from the baseline decentralized steady state. Static surplus appropriation replaces private product-line payoffs with planner static sources. Net link-flow valuation adds the planner match-costate term. Matching-stock congestion adds the matching-stock envelope source. The latter two are the components of the matching-network wedge.

state. Within incumbent R&D, the planner shifts R&D effort away from young incumbents and toward older incumbents whose accumulated buyer-supplier relationships raise the social value of new product lines. The transition from the decentralized steady state to the planner path raises consumption-equivalent welfare by 1.71%. The associated entry transition falls immediately at the start of the planner reform and then converges to its planner steady-state level; Figure A6 reports the transition path.

Appendix G.7 describes the transition solver, and Appendix G.8 describes the welfare computation. Decomposition switches, allocation counterfactuals, and policy exercises are described in Appendix H.

5.2.2 Wedge Decomposition

Table 4 decomposes the planner gap into the static correction and the two matching-network corrections characterized in Section 4. These components should be read as counterfactual planner corrections, not as mutually orthogonal structural shocks. In economic terms, the static correction changes how current product-line surplus is valued in innovation and entry decisions. Net link-flow valuation adds the future surplus created by the new relationships attached to product lines. Matching-stock congestion internalizes how additional product lines change matching opportunities for the rest of the network.

The first column reports the effect of each component on stationary entry, as a percentage change relative to decentralized entry. The second column reports Shapley contributions to transition welfare in CE percentage terms.

The congestion component is the largest positive contributor to transition welfare. It contributes 1.45% in CE terms and reduces entry by 38.9%. The net link-flow valuation component pushes in the opposite direction for entry: the match-costate term raises entry by 33.4% because additional product lines generate future network-surplus value.

These entry responses do not translate one-for-one into transition-welfare contributions. The entry response isolates the direction of the entry incentive, while the Shapley contribution evaluates the transition-welfare effect conditional on the other planner corrections and the fixed R&D resource. Conditional on the other corrections, the Shapley contribution of net link-flow valuation to transition welfare is negative, equal to −0.51% in CE terms. The two matching-network components are not independent: the interaction between net link-flow valuation and matching-stock congestion is positive, so the matching-stock correction is more valuable when network-surplus values are also internalized. The social planner combines these forces, reducing entry by 25.4% and raising transition welfare by 1.71% in CE terms.

5.2.3 Sensitivity to the Matching-Stock Elasticity

This sensitivity is important because the targeted moments discipline the level and life-cycle slope of matching, but not the aggregate matching-stock elasticity itself. The size of the over-entry result therefore depends on the baseline matching-stock elasticity. Appendix H.4 re-estimates the targeted parameters at $\eta = 0.1$ and recomputes the planner comparison as a lower-elasticity case. With $\eta = 0.1$, a value motivated by the one-stock interpretation of Miyauchi (2024), planner entry is only 0.8% below decentralized entry and the transition CE gain is 0.02%. The baseline specification therefore reports the quantitative implications of the scale-invariant matching benchmark. Lower matching-stock elasticities attenuate the entry wedge and the associated welfare gains, while the decomposition clarifies the competing match-value and matching-stock forces.

5.3 Allocation Counterfactuals and Policy Exercises

We use the planner comparison for two exercises. First, we decompose the transition gain across the aggregate entry margin and the age allocation of incumbent R&D. Second, we ask how much of the planner gain can be captured by a simple uniform entry tax.

5.3.1 Entry versus Incumbent-Age Allocation Margins

To split the 1.71% planner transition gain across R&D margins, we compare allocation counterfactuals that combine the decentralized and planner paths for two objects: the aggregate entry path and the incumbent R&D age-share path. This accounting isolates the entry and incumbent-age allocation margins; it is not intended as an implementable policy experiment. The entry margin contributes about 1.15% in CE terms, or 68% of the combined allocation-margin contribution, while the incumbent R&D age-allocation margin contributes about 0.55% in CE terms, or 32%. The entry margin is therefore the larger source of the allocation gain, but roughly one-third of the gain still comes from correcting the decentralized tilt of incumbent R&D toward young incumbents.

5.3.2 Uniform Entry Tax

The policy exercise asks how much of the planner gain can be captured by an actual constant entry tax. We solve decentralized perfect-foresight transitions under constant uniform entry taxes. In the baseline policy exercise, the welfare-maximizing tax is $\tau_E = 1.84$, an 83.5% entry tax. It raises consumption-equivalent welfare by 1.62%, capturing 94.4% of the planner CE gain. The residual gap is almost entirely a transition-flow allocation gap rather than a terminal steady-state gap.

For comparison, the decentralized-to-planner entry change corresponds to a 34% entry-cost-equivalent wedge. This number is an entry-margin accounting measure, not the policy tax. It translates the planner entry gap into entry-cost units holding other margins fixed, and should not be read as a full decentralization of the planner allocation: it leaves the planner's age-specific incumbent-R&D reallocation and static surplus-allocation wedges outside the calculation. The welfare-maximizing tax above is chosen from decentralized transition paths and includes general-equilibrium and dynamic responses.

The policy implication is conditional on the matching-stock elasticity. In the scale-invariant baseline, a constant entry tax captures most of the welfare gain because it directly leans against over-entry, the dominant distortion in that economy. Under lower matching-stock elasticities, Appendix H.4 shows that the planner's entry reduction is much smaller, and a separate $\eta = 0.1$ tax grid selects a much smaller entry tax, about 7%. Because the planner transition gain in that case is only 0.02%, the implied policy-gain shares are numerically unstable. Thus, the large entry-tax prescription should be interpreted as a feature of the scale-invariant baseline, not as a policy magnitude that is invariant across matching technologies. Even in the baseline policy exercise, the tax remains a one-dimensional compromise: it cannot reproduce the planner's state-dependent entry values or the age-dependent incumbent R&D values that move R&D away from young incumbents and toward older incumbents along the transition.

6 Conclusion

This paper studies how production networks shape firms' incumbent R&D and entry incentives. The empirical evidence shows that firms accumulate buyer-supplier relationships over the life cycle, that persistent links generate strong age patterns in production networks, and that firms connected to more patent-intensive trading partners patent more themselves. Guided by these facts, we develop a dynamic model in which incumbent R&D creates new product lines that inherit existing trading relationships, while entry competes with incumbent R&D for the same scarce R&D resource.

In the baseline specification, the quantitative results point to an over-allocation of R&D toward entrants and young incumbents, with over-entry as the central margin. The social planner reduces entry by 25.4% relative to the decentralized economy, reallocates incumbent R&D away from young incumbents and toward older incumbents, and raises consumption-equivalent transition welfare by 1.71%. The dominant force is the matching-stock congestion externality: private agents take the matching environment as given, while the planner internalizes how product-line creation changes matching opportunities throughout the network. The match-costate network channel works in the opposite direction for entry, raising the social value of creating new product lines, but it does not overturn the congestion force in the baseline specification.

The policy exercises are best interpreted as allocation diagnostics rather than as a stand-alone policy recommendation. Within the baseline specification, a constant uniform entry tax captures 94.4% of the planner consumption-equivalent transition welfare gain because it leans directly against the over-entry force. This policy conclusion is quantitatively tied to the baseline matching-stock elasticity: under lower elasticities, the planner's entry reduction and the welfare gains from entry taxation are much smaller. Even in the baseline specification, the entry tax is only a one-dimensional instrument: it does not reproduce the planner's age-dependent incumbent R&D allocation away from young incumbents and toward older incumbents, and the remaining gap is a transition-flow allocation gap rather than a terminal steady-state entry gap. A full decentralization would require richer, state- and time-dependent instruments for entry values, incumbent R&D values, and static surplus allocation. The broader implication is that production networks affect not only the propagation of shocks or knowledge, but also the allocation of a fixed R&D resource between entry and incumbent product-line expansion, and across firms over the life cycle.

References

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018): "Innovation, Reallocation, and Growth," *The American economic review*, 108, 3450–3491.

Acemoglu, D., U. Akcigit, D. Hanley, and W. R. Kerr (2016): "Transition to Clean Technology," *Journal of Political Economy*, 124, 52–104, 10.1086/684511.

Acemoglu, D., and P. D. Azar (2020): "Endogenous production networks," *Econometrica: journal of the Econometric Society*, 88, 33–82.

Acemoglu, D., and V. M. Carvalho (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

Aekka, A., and G. Khanna (2025): "Endogenous Production Networks and Firm Dynamics," *Working paper*.

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2025): "Good Rents versus Bad Rents: R&D Misallocation and Growth," *NBER Working Paper 34190*, 10.3386/w34190.

Akcigit, U., and W. R. Kerr (2018): "Growth through Heterogeneous Innovations," *The journal of political economy*, 126, 1374–1443.

Arkolakis, C., F. Huneeus, and Y. Miyauchi (2023): "Production Network Formation, Trade, and Welfare," *NBER Working Paper 30954*, 10.3386/w30954, Revised January 2026.

Asai, H., and M. Nirei (2026): "Delayed Network Formation of Young Firms: Micro Estimates and Macro Consequences," CREPE Discussion Paper, Center for Research and Education in Program Evaluation, University of Tokyo, https://www.crepe.e.u-tokyo.ac.jp/results/2025/CREPEDP167.pdf, First circulated February 2025; revised February 2026.

Ayerst, S. (2023): "Innovator Heterogeneity, R&D Misallocation and the Productivity Growth Slowdown," *Working paper*.

Bai, X., J. Fernández-Villaverde, Y. Li, and F. Zanetti (2024): "The Causal Effects of Global Supply Chain Disruptions on Macroeconomic Outcomes: Evidence and Theory," *NBER Working Paper 32098*, 10.3386/w32098.

Baqaee, D., A. Burstein, C. Duprez, and E. Farhi (2025): "Consumer Surplus From Suppliers: How Big Is It and Does It Matter for Growth?" *Econometrica*, 93, 2043–2081, 10.3982/ECTA22672.

Baqaee, D. R., and E. Farhi (2020): "Productivity and Misallocation in General Equilibrium," *The quarterly journal of economics*, 135, 105–163.

Bernard, A. B., E. Dhyne, G. Magerman, K. Manova, and A. Moxnes (2022): "The Origins of Firm Heterogeneity: A Production Network Approach," *The journal of political economy*, 130, 1765–1804.

Bernard, A. B., S. J. Redding, and P. K. Schott (2010): "Multiple-Product Firms and Product Switching," *American Economic Review*, 100, 70–97, 10.1257/aer.100.1.70.

Bigio, S., and J. La'O (2020): "Distortions in Production Networks*," *The quarterly journal of economics*, 135, 2187–2253.

Boehm, C. E., A. Flaaen, and N. Pandalai-Nayar (2019): "Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake," *The review of economics and statistics*, 101, 60–75.

Cai, J., and N. Li (2019): "Growth through inter-sectoral knowledge linkages," *The Review of Economic Studies*, 86, 1827–1866.

Cai, J., and C. Tian (2024): "Innovation and Endogenous Knowledge Network Dynamics," *Working paper*.

Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021): "Supply Chain Disruptions: Evidence from the Great East Japan Earthquake," *The quarterly journal of economics*, 136, 1255–1321, 10.1093/qje/qjaa044.

Deng, Y. (2011): "A Dynamic Stochastic Analysis of International Patent Application and Renewal Processes," *International Journal of Industrial Organization*, 29, 766–777.

Eaton, J., S. S. Kortum, and F. Kramarz (2022): "Firm-to-Firm Trade: Imports, Exports, and the Labor Market," *NBER Working Paper 29685*, 10.3386/w29685.

Fujii, D., Y. Saito, and T. Senga (2017): "The Dynamics of Inter-firm Networks and Firm Growth," *RIETI Discussion Paper 17-E-110*.

Huneeus, F. (2020): "Production network dynamics and the propagation of shocks," *Working paper, Duke University*.

Imani, Y., and A. Ohyama (2022): "The Roles of Structured Management in the Formation of Transactional Relationships," TDB-CAREE Discussion Paper Series, TDB Center for Advanced Empirical Research on Enterprise and Economy, Hitotsubashi University, https://hermes-ir.lib.hit-u.ac.jp/hermes/ir/re/73180/070careeDP-E-2107.pdf.

Klette, T. J., and S. Kortum (2004): "Innovating Firms and Aggregate Innovation," *The journal of political economy*, 112, 986–1018.

Liu, E. (2019): "Industrial Policies in Production Networks*," *The quarterly journal of economics*, 134, 1883–1948.

Liu, E., and S. Ma (2023): "Innovation Networks and R&D Allocation," *NBER Working Paper 29607*, 10.3386/w29607, Revised January 2023.

Ministry of Health, Labour and Welfare (various years): "Employment Insurance Annual Report," https://www.mhlw.go.jp/bunya/koyou/koyouhoken02/annual01.html, Various fiscal years, 1998–2019.

Ministry of Internal Affairs and Communications (2019): "2015 Input-Output Tables for Japan," https://www.e-stat.go.jp/en/stat-search/files?cycle=0&layout=datalist&page=1&tclass1val=0&toukei=00200603&tstat=000001130583, Input table and input coefficients.

Miyauchi, Y. (2024): "Matching and agglomeration: Theory and evidence from Japanese firm-to-firm trade," *Econometrica: journal of the Econometric Society*, 92, 1869–1905, 10.3982/ECTA19697.

Oberfield, E. (2018): "A theory of input-output architecture," *Econometrica: journal of the Econometric Society*, 86, 559–589.

OECD (2021): "DynEmp3 Database," http://oe.cd/dynemp, DynEmp project: Measuring business dynamics, start-ups and job creation.

Osotimehin, S., and L. Popov (2023): "Misallocation and Intersectoral Linkages," *Review of economic dynamics*, 51, 177–198, 10.1016/j.red.2022.12.005.

Pakes, A., and M. Schankerman (1984): "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources," in *R&D, Patents, and Productivity*, ed. by Z. Griliches, 73–88: University of Chicago Press.

Peters, M., and C. Walsh (2022): "Population growth and firm-product dynamics."

Romer, P. M. (1990): "Endogenous Technological Change," *The journal of political economy*, 98, S71–S102.

Statistics Bureau, Ministry of Internal Affairs and Communications (2025): "2025 Survey of Research and Development: Summary of Results," https://www.stat.go.jp/data/kagaku/kekka/kekkagai/pdf/2025ke_gai.pdf, Table 2-3.

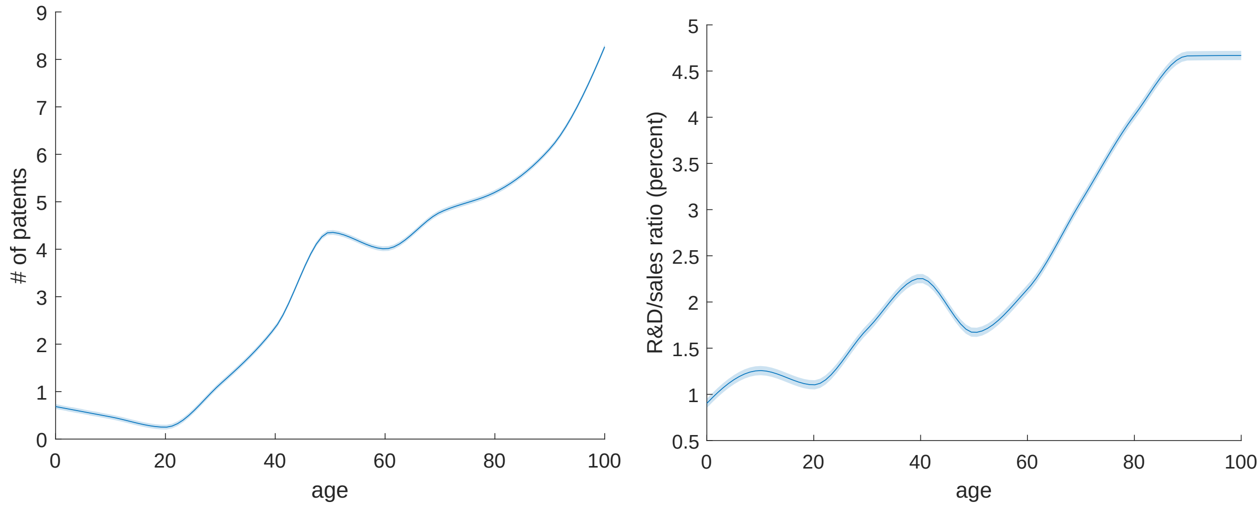

(a) Patent Applications

(b) R&D/Sales Ratio

Figure A1: R&D and Age

Notes: The figure plots age profiles of two measures of innovative and R&D activity for manufacturing focal firms. The left panel uses the number of patent applications, and the right panel uses the R&D-to-sales ratio, measured in percent. The sample covers 1998–2019.

Appendix for "Production Networks and R&D Allocation"

This appendix is organized as follows. Appendix A reports additional empirical evidence, and Appendix B describes the estimation targets. Appendix C reports model-implied life-cycle profiles, while Appendix D derives the closed-form stationary densities. Appendices E–F provide the decentralized and planner derivations, Appendix G describes the numerical implementation, and Appendix H reports the decomposition switches, allocation-margin counterfactuals, policy exercises, and sensitivity checks underlying the welfare results.

A Additional Empirical Evidence

A.1 R&D and Age

Figure A1 presents local linear regressions relating firm age to two measures of innovative and R&D activity: the number of patent applications and the R&D-to-sales ratio. Both measures increase monotonically with age. This pattern complements the network facts in Section 2 by showing that R&D activity also rises over the firm life cycle.

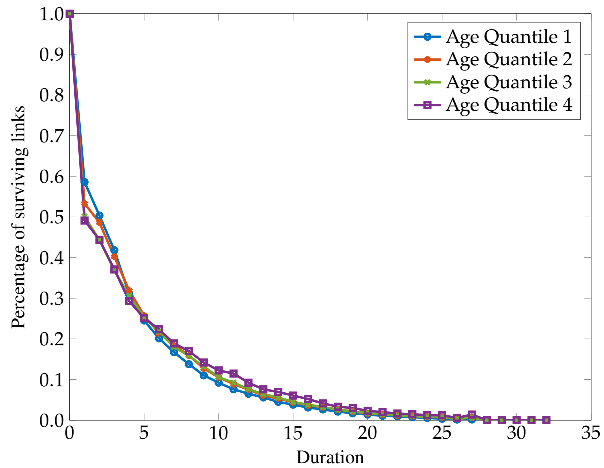

(a) Link Survival by Supplier's Age

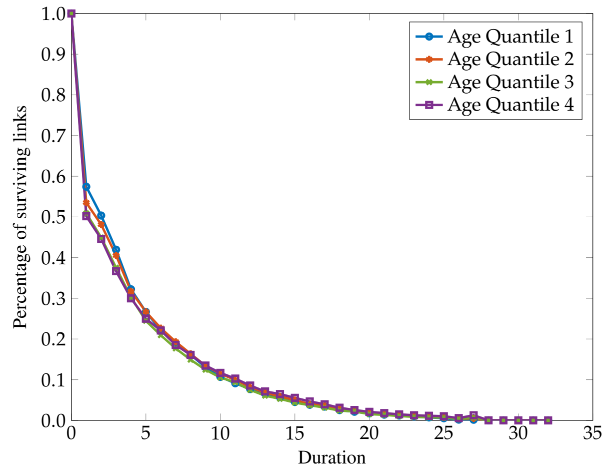

(b) Link Survival by Buyer's Age

Figure A2: Survival of Trading Links

Notes: The figure plots the share of trading links that remain observed at each relationship duration. Duration is measured in years. To reduce censoring in the duration profile, the sample is restricted to links where (1) the link is first observed more than one year after both firms first enter the sample and (2) the link is no longer observed more than one year before either endpoint firm leaves the sample. The left panel groups links with manufacturing suppliers by supplier age quartile, and the right panel groups links with manufacturing buyers by buyer age quartile, with age measured when the link is first observed. Trading partners on the other side of the link are not restricted to manufacturing. Age Quantile 1 is the youngest quartile and Age Quantile 4 is the oldest quartile.

A.2 Empirical Link Survival

Figure A2 complements the life-cycle evidence in Section 2 by documenting the persistence of observed buyer-supplier relationships. The figure is a persistence diagnostic; the annual stock-survival moment used to calibrate δ_M is described in Appendix B.5.

A.3 Excluding Direct Patent-Citation Links

Table A1 repeats the firm-level patenting regressions after excluding buyer-supplier relationships that are also connected to the focal firm through direct patent citations. The sample falls from 53,553 to 51,338 firm-year observations. The estimated coefficients remain positive and statistically significant. The OLS estimates are close to the baseline estimates, and the IV estimates remain economically meaningful, with large reported first-stage F-statistics. These results reduce the concern that the baseline relationship mainly reflects direct knowledge links between firms and their trading partners.

Table A1: Trading-Partner Patenting and Own Patenting, Excluding Direct Patent-Citation Links

	OLS			IV		
	(1)	(2)	(3)	(4)	(5)	(6)
Seller patenting per seller	0.018***		0.017***	0.049**		0.043**
	(0.003)		(0.003)	(0.019)		(0.019)
Buyer patenting per buyer		0.020***	0.019***		0.081***	0.079***
		(0.003)	(0.003)		(0.023)	(0.023)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
First-stage F				805.3	690.5	339.0
Observations	51,338	51,338	51,338	51,338	51,338	51,338

Notes: The dependent variable is log patent applications. Seller-side patenting is supplier patenting per supplier, and buyer-side patenting is buyer patenting per buyer; both are measured in logs. The sample covers 1994–2019, excludes direct patent-citation links between the focal firm and its buyers or suppliers, and is restricted to firm-year observations with strictly positive own patent counts and strictly positive seller-side and buyer-side partner-patenting measures. The excluded instruments in the IV columns are shift-share exposures to foreign triadic patents by patent class. First-stage F denotes the first-stage F-statistic reported by the IV estimation routine. Parentheses report standard errors clustered at the industry × year level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

B Estimation Targets and Calibration

This section describes the empirical targets, model counterparts, and auxiliary calibrations used in Section 5.

B.1 Incumbent R&D Intensity

The first target is incumbent R&D intensity, denoted RD^{data}. It is measured as the empirical R&D-to-sales ratio among R&D-performing manufacturing firms and equals 0.042 in the baseline estimation, close to the manufacturing ratio reported for firms conducting research in the 2025 Survey of Research and Development (Statistics Bureau, Ministry of Internal Affairs and Communications, 2025). We use R&D-performing firms because this moment is intended to discipline the active incumbent R&D margin rather than the extensive margin of whether a firm conducts any R&D. The model counterpart is incumbent R&D spending divided by aggregate sales in the decentralized stationary equilibrium:

$$\text{RD}^{\text{model}} = \frac{\int w_H h_H(\lambda(a)) f(a)\, da}{\int r(a) f(a)\, da}.$$

This moment excludes entry R&D labor and disciplines the level of incumbent R&D costs and the R&D productivity parameter φ. Entry R&D is reported separately in the diagnostics and is disciplined primarily by the young-firm employment-share target through φ_E.

B.2 Buyer-Supplier Links per Firm

The second target is the average number of buyer-supplier links per firm, denoted L^{data}. In the baseline target, $L^{\text{data}} = 27.3$. The model counterpart integrates the matched-product density over supplier and buyer ages and normalizes by the stationary firm mass $N_f = E/\delta_F$:

$$L^{\text{model}} = \frac{\int\int m(a_s, a_b) f(a_b)\, da_s\, da_b}{N_f}.$$

The target is based on the manufacturing JP MOPS sample linked to TDB's COSMOS2 business-to-business transaction data in Imani and Ohyama (2022). That study reports mean supplier and customer (buyer) counts of 28.0 and 26.6 firms, respectively, for Japanese manufacturing firms with establishments in Japan and at least 30 employees. The underlying transaction records are directed buyer-supplier relationships, but the target is not a double count of both directions. We construct an average link degree per firm by averaging the reported supplier and customer (buyer) counts, and set $L^{\text{data}} = (28.0 + 26.6)/2 = 27.3$.

We use the mean rather than the median because this moment is an aggregate link mass normalized by firm mass, and therefore corresponds to network density rather than the typical firm's degree. The factor $f(a_b)$ converts the per-buyer-product-line object $m(a_s, a_b)$ into aggregate buyer-supplier links attached to buyer product lines. Because the model aggregates over product lines and does not track partner identities within firms, this moment should be interpreted as matching aggregate link density rather than the exact distribution of unique firm-level trading partners. This moment disciplines the overall scale of network links, primarily through ζ and ζ_0.

B.3 Life-Cycle Log-Degree Gap

The third target summarizes life-cycle network accumulation using the controlled life-cycle estimates in Asai and Nirei (2026). Their Figure 1 reports age-dummy coefficients for log supplier degree and log buyer degree, normalized relative to firms ages 50 and above. Because the empirical specification adds one before taking logs, the model counterpart uses $\log(1 + C)$ rather than $\log C$ for this target.

Let β_a^S and β_a^B denote the supplier- and buyer-side age-dummy coefficients. The baseline uses ages 1–5 and excludes age 0 because Asai and Nirei (2026) report that the age-zero estimate is imprecise. The data moment is

$$\Delta^{\text{data}} = \frac{1}{2}\left(\frac{1}{5}\sum_{a=1}^{5}\beta_a^S + \frac{1}{5}\sum_{a=1}^{5}\beta_a^B\right) = -0.457.$$

The model counterpart first constructs supplier and buyer degree profiles for a firm of age a. In the notation below, the first argument of $m(a_s, a_b)$ is supplier age and the second argument is buyer age:

$$C^S(a) = \int m(a_s, a)\, da_s, \qquad C^B(a) = \frac{\int m(a, a_b) f(a_b)\, da_b}{f(a)}.$$

Define $\ell^S(a) = \log(1 + C^S(a))$ and $\ell^B(a) = \log(1 + C^B(a))$. The mature reference groups are firm-mass weighted means over ages 50 and above:

$$\bar{\ell}^S_{50+} = \frac{\int_{50}^{\infty} \ell^S(a) f(a)\, da}{\int_{50}^{\infty} f(a)\, da}, \qquad \bar{\ell}^B_{50+} = \frac{\int_{50}^{\infty} \ell^B(a) f(a)\, da}{\int_{50}^{\infty} f(a)\, da}.$$

The model moment is

$$\Delta^{\text{model}} = \frac{1}{2}\left[\left(\frac{1}{5}\sum_{j=1}^{5}\ell^S(j) - \bar{\ell}^S_{50+}\right) + \left(\frac{1}{5}\sum_{j=1}^{5}\ell^B(j) - \bar{\ell}^B_{50+}\right)\right],$$

where the young-age values are interpolated to integer ages 1 through 5. This moment disciplines the life-cycle slope of network accumulation separately from the aggregate link scale. In the numerical implementation, R&D intensity, links per firm, and the young-firm employment share enter the loss as log deviations, while the life-cycle gap enters in levels. A corresponding young-to-mature gap computed from our TDB link-age profile is used only as an additional diagnostic and is not a baseline target.

B.4 Young-Firm Employment Share

The entry-efficiency parameter is disciplined by the employment share of young firms. The target is the OECD DynEmp3 Japan manufacturing employment share of establishments younger than six years (OECD, 2021):

$$s_Y^{\text{data}} = 0.096.$$

For Japan, DynEmp3 reports the young-employment indicator for manufacturing establishments. The model has no separate establishment margin. We therefore treat this object as an age-distribution target for production employment and map it to the production-employment share of model firms below age six:

$$s_Y^{\text{model}} = \frac{\int_0^6 L^{\text{prod}}(a) f(a)\, da}{\int_0^\infty L^{\text{prod}}(a) f(a)\, da}.$$

In the contracting static block, production labor is recovered from the gross-output cost base as $L^{\text{prod}}(a) = \beta K^{DE}(a)$. Entry R&D labor is included in total-R&D diagnostics, but not in this production-employment target.

B.5 Calibration of Link Destruction

The empirical object is a firm-to-firm relationship, not a product-line-level link. The baseline calibration uses a one-year link-stock survival measure: among trading links observed in year t with at least one manufacturing endpoint and for which both endpoint firms remain observable in year $t+1$, approximately 92% remain observed in year $t+1$. We interpret this link-stock survival object as the link-specific survival component in the model, leaving firm exit and product-line obsolescence to δ_F and δ_P. This convention keeps δ_M distinct from firm exit and product-line exit. In the model, δ_M destroys an individual buyer-supplier link, while δ_F and δ_P remove endpoint product lines. The annual link-stock survival rate of approximately 0.92 implies

$$\delta_M = -\log(0.92) \simeq 0.083.$$

We therefore set the baseline link destruction rate to $\delta_M = 0.08$ per year. In discrete annual terms, this corresponds to a one-year link dissolution probability of $1 - \exp(-0.08) \simeq 0.077$.

Figure A2 reports a separate completed-spell duration diagnostic. The sample restriction reduces censoring in the plotted duration profiles, but it also selects relationships whose

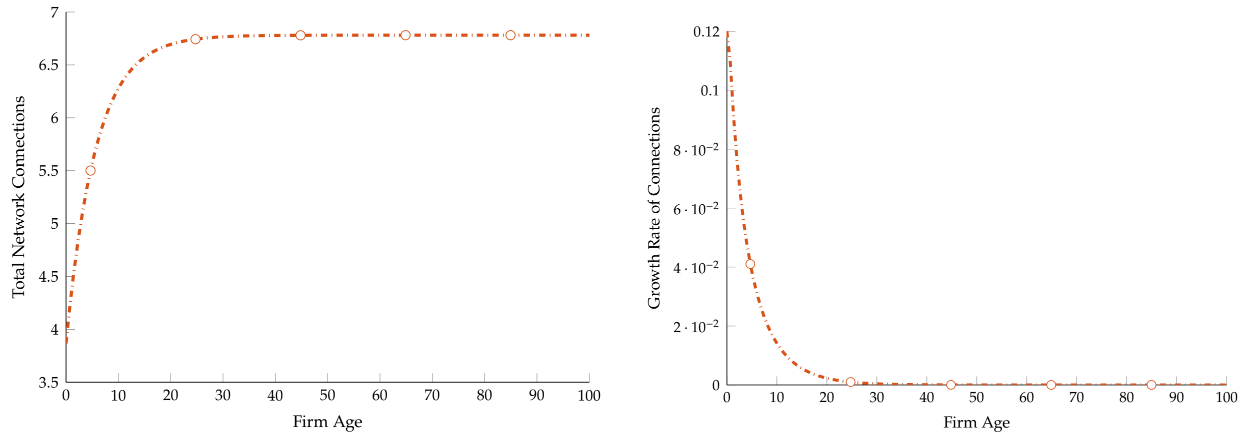

Figure A3: Model-Implied Life-Cycle Network Profiles

Notes: The left panel shows the model-implied stock of buyer-supplier links by firm age. The right panel shows the model-implied growth rate of that link stock by firm age, computed as the age derivative of log total links. The decentralized economy in the baseline specification is the contracting economy with transfer share $\theta = 1/2$.

starts and endings are both observed inside the sample window. These completed-spell profiles imply faster link turnover than the one-year stock-survival measure, so we use them as a high-turnover diagnostic rather than as the baseline calibration target.

C Model-Implied Life-Cycle Profiles

C.1 Life-Cycle Network and Innovation Profiles

This subsection reports model-implied life-cycle profiles that complement the targeted moment fit in Table 3. The estimation targets the overall link stock and a young-versus-mature life-cycle gap, but it does not target the full age profile point by point. Figure A3 shows that the estimated model generates a steep early increase in the stock of buyer-supplier links and a high, nearly flat link stock among mature firms. Link growth is front-loaded and declines rapidly with age, matching the qualitative life-cycle pattern documented in Section 2.

The stock of accumulated relationships is the channel through which life-cycle network dynamics strengthen private innovation incentives. Figure A4 reports the model-implied innovation rate by firm age. The profile rises with age and reflects the accumulation of buyer-supplier relationships. The planner profile is included for comparison with the welfare exercises in the main text: relative to the social planner, the decentralized profile in the baseline specification sets higher innovation rates for the youngest incumbents and

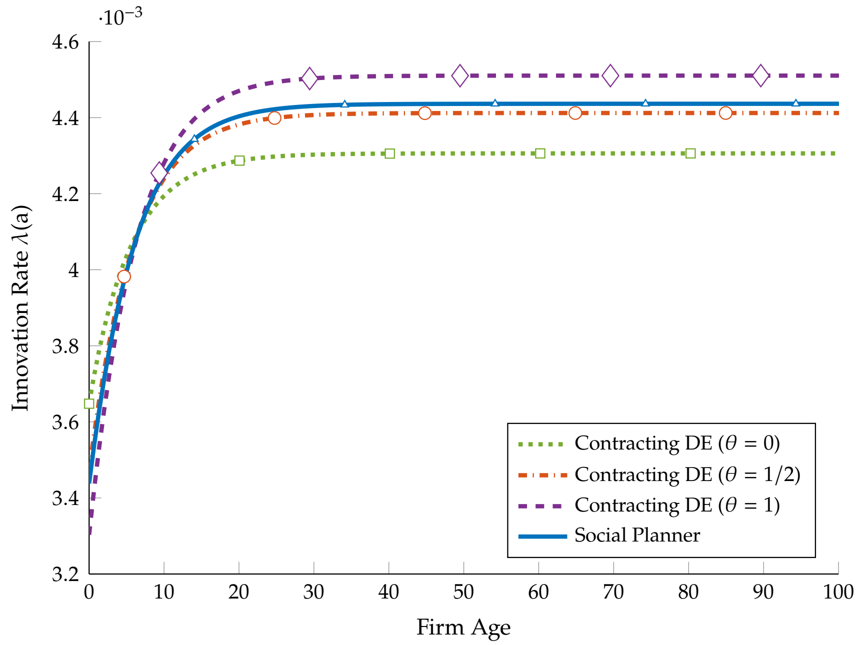

Figure A4: Model-Implied Innovation by Firm Age

Notes: The figure shows the model-implied innovation rate by firm age. The decentralized economy in the baseline specification is the contracting economy with transfer share $\theta = 1/2$; the $\theta = 0$, $\theta = 1$, and planner profiles are shown as allocation counterfactuals for comparison with the welfare analysis. In the model, accumulated trading relationships strengthen innovation incentives over the firm life cycle.

lower rates for older incumbents.

C.2 Model-Implied Age Assortativity

This subsection reports an additional untargeted life-cycle diagnostic. The same relationship-persistence mechanism that generates life-cycle link accumulation also implies positive age assortativity: as firms and their trading partners remain linked, they age together. Figure A5 plots the model-implied average partner age against own age. The upward-sloping profile shows that the estimated model generates positive age assortativity, consistent with the empirical pattern documented in Section 2.

D Closed-Form Stationary Densities

Fix an entry flow E, an innovation profile $\lambda(a)$, and a matching shifter Γ. These objects may be those induced by the decentralized equilibrium or by a planner allocation. The stationary product-line density $f(a)$ and matched-product density $m(a_s, a_b)$ then admit closed-form solutions. Here a_s and a_b denote the ages of the supplier and buyer firms

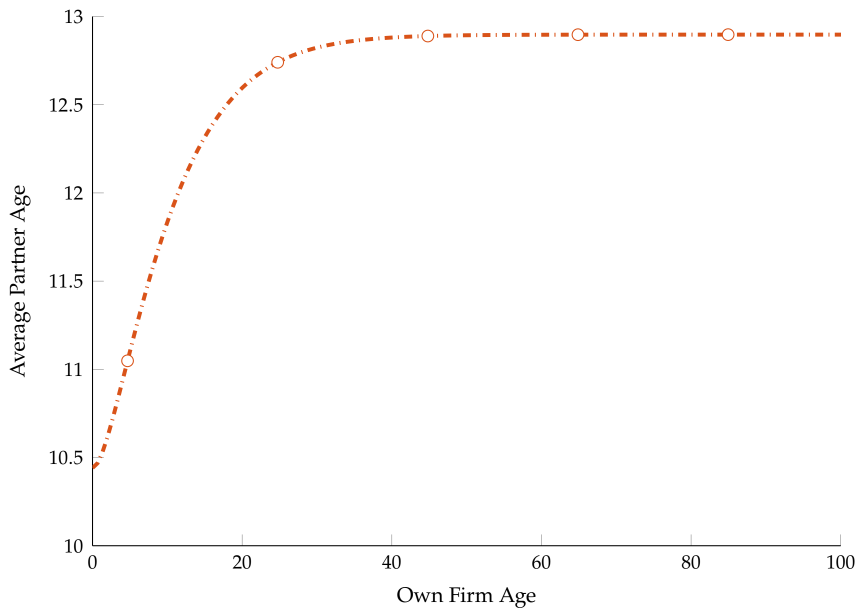

Figure A5: Model-Implied Age Assortativity

Notes: The figure plots the model-implied average age of a firm's trading partners against the firm's own age. Partner ages are averaged over model-implied buyer-supplier link masses conditional on the firm's own age. The decentralized economy in the baseline specification is the contracting economy with transfer share $\theta = 1/2$.

associated with a matched product-line pair. Throughout this section, $\delta_M > 0$.

D.1 Product-Line Density

For $a > 0$, the product-line density in (2) satisfies

$$\frac{\partial}{\partial a} f(a) = (\lambda(a) - \delta_F - \delta_P)\, f(a)$$

with boundary condition $f(0) = E$. Integrating gives

$$f(a) = E \exp\left(\int_0^a \{\lambda(\tilde{a}) - \delta_F - \delta_P\}\, d\tilde{a} \right). \tag{23}$$

A stationary allocation additionally requires the implied product-line mass to be integrable.

D.2 Matched-Product Density

For the matched-product density, on the interior $a_s > 0$ and $a_b > 0$, (5) is

$$\frac{\partial}{\partial a_s} m(a_s, a_b) + \frac{\partial}{\partial a_b} m(a_s, a_b) = (\lambda(a_s) - \delta_F - \delta_P - \delta_M)\, m(a_s, a_b) + \zeta\Gamma f(a_s),$$

with boundary conditions

$$m(a_s, 0) = \zeta_0 \Gamma f(a_s), \qquad m(0, a_b) = \zeta_0 \Gamma f(0) = \zeta_0 \Gamma E.$$

Entry appears as boundary inflow at age zero; the interior transport equation tracks the evolution of matched product lines after entry. Define the normalized object

$$\bar{m}(a_s, a_b) \equiv \frac{m(a_s, a_b)}{f(a_s)}.$$

This normalization isolates the conditional matching kernel from the supplier-side product-line mass. Using (23),

$$\begin{aligned}
\frac{\partial}{\partial a_b} m(a_s, a_b) &= f(a_s) \frac{\partial}{\partial a_b} \bar{m}(a_s, a_b), \\
\frac{\partial}{\partial a_s} m(a_s, a_b) &= f(a_s) \left[\frac{\partial}{\partial a_s} \bar{m}(a_s, a_b) + (\lambda(a_s) - \delta_F - \delta_P)\, \bar{m}(a_s, a_b) \right].
\end{aligned}$$

Substituting these expressions into the law of motion for $m(a_s, a_b)$ and dividing by $f(a_s)$ yields

$$\frac{\partial}{\partial a_s} \bar{m}(a_s, a_b) + \frac{\partial}{\partial a_b} \bar{m}(a_s, a_b) = -\delta_M \bar{m}(a_s, a_b) + \zeta\Gamma. \tag{24}$$

Equation (24) characterizes the normalized conditional object $\bar{m}(a_s, a_b)$, not the matched-product density itself. Hence a non-decaying tail in $\bar{m}(a_s, a_b)$ does not imply an infinite mass of matches: the economically relevant density is $m(a_s, a_b) = \bar{m}(a_s, a_b) f(a_s)$, which inherits the tail behavior of the supplier-side product-line density $f(a_s)$. The boundary conditions imply

$$\bar{m}(a_s, 0) = \bar{m}(0, a_b) = \zeta_0 \Gamma.$$

Let $\ell \equiv \min\{a_s, a_b\}$. Along the characteristic through (a_s, a_b), define

$$h(s) \equiv \bar{m}(a_s - \ell + s, a_b - \ell + s), \qquad s \in [0, \ell].$$

Then (24) implies

$$h'(s) = -\delta_M h(s) + \zeta\Gamma, \qquad h(0) = \zeta_0\Gamma.$$

Solving this ordinary differential equation gives

$$\bar{m}(a_s, a_b) = \Gamma\left[\frac{\zeta}{\delta_M} + \left(\zeta_0 - \frac{\zeta}{\delta_M}\right)\exp\left(-\delta_M \min\{a_s, a_b\}\right)\right].$$

Therefore,

$$m(a_s, a_b) = \left[\frac{\zeta}{\delta_M} - \left(\frac{\zeta}{\delta_M} - \zeta_0\right)\exp\left(-\delta_M \min\{a_s, a_b\}\right)\right]\Gamma f(a_s). \tag{25}$$

The closed form shows that, conditional on $(f(a_s), \Gamma)$, the innovation profile, supplier-side firm exit, and supplier product-line exit affect $m(a_s, a_b)$ only through the supplier-side mass $f(a_s)$ and the matching shifter Γ. After normalizing by $f(a_s)$, the innovation term and the δ_F and δ_P exit terms in the matched-product law cancel with the corresponding terms in the KFE for $f(a_s)$. The normalized matching kernel depends on ages only through $\min\{a_s, a_b\}$. Its age decay is governed by link destruction δ_M, while its level is pinned down by ζ, ζ_0, and Γ. This closed-form expression also implies the reciprocity relation

$$\frac{m(a_s, a_b) f(a_b)}{f(a_s)} = m(a_b, a_s),$$

which ensures consistency between supplier- and buyer-side normalizations of the same bilateral match mass.

E Decentralized Equilibrium: Private Product-Line Value

This section derives the private product-line value used in the decentralized equilibrium. The key simplification is that incumbent R&D creates a new product line inside the same age-a firm, and the new line inherits the firm's existing supplier and buyer relationships. Hence the value of a successful incumbent innovation at age a is $V(a)$. Entry is different: a new entrant starts at age zero, and its boundary links are embedded in the age-zero state through the boundary conditions for m. The free-entry condition therefore values entry through $V(0)$.

Because product lines enter the firm's payoff and transition law additively, firm value is homogeneous in the number of product lines:

$$V^F(n,a) = nV(a).$$

Substituting this representation into (10), dividing by n, and imposing stationarity with the household interest rate equal to ρ yields

$$\left(\rho + \delta_F + \delta_P\right) V(a) = \pi(a) + V_a(a) + \max_{\lambda \geq 0} \left[\lambda V(a) - w_H h_H(\lambda)\right].$$

Here $\pi(a)$ is the private product-line flow payoff from the static contracting block in (9). This is the product-line HJB in (11). Maximizing the last term gives the decentralized incumbent R&D policy in (12).

F Social Planner Derivations

We derive the planner's optimality conditions from the primitive objective $\log Y$ using an intratemporal Lagrangian for the static allocation and a Hamiltonian for the laws of motion. The derivation keeps endogenous entry, the entry R&D cost $H_E(E)$, and the matching shifter $\Gamma(\mathcal{N})$ explicit. It proceeds in three steps: static first-order conditions, dynamic costate and policy conditions, and the static cost and demand shifters used in the numerical implementation. Throughout this appendix, the first argument of $x(\cdot,\cdot)$, $m(\cdot,\cdot)$, and $V^M(\cdot,\cdot)$ is the supplier age and the second argument is the buyer age; when both sides appear, we write them as a_s and a_b. Dividing the primitive costates by a common age-independent factor and normalizing the match costate by buyer density yields the value functions V^{SP} and V^M stated in Proposition 1.

F.1 Objective and Constraints

The social planner maximizes $\int_0^\infty e^{-\rho t} \log Y(t)\, dt$ subject to the production technology (17), the goods market clearing condition

$$x(a_s)\, f(a_s) = y(a_s)\, f(a_s) + \int x(a_s, a_b)\, m(a_s, a_b)\, f(a_b)\, da_b, \tag{26}$$

the laws of motion (2), (5), the entry boundary condition (3), the entrant boundary conditions for $m(a_s, a_b)$, the matching shifter (4), and the factor constraints (19), (13). In (26), a_s is in the supplier position and the integral runs over all buyer ages a_b.

F.2 Intratemporal Lagrangian

With states (f, m) fixed, the intratemporal Lagrangian is

$$\begin{aligned}\mathcal{L}_0^{SP} =\ & \log Y \\ & + \int \chi(a_b)\left[\mathcal{A}_Y\, l(a_b)^{\beta}\, J(a_b)^{\frac{\sigma(1-\beta)}{\sigma-1}} - x(a_b)\right] f(a_b)\, da_b \\ & + \int p^{SP}(a_s)\left[x(a_s)\, f(a_s) - y(a_s)\, f(a_s) - \int x(a_s, a_b)\, m(a_s, a_b)\, f(a_b)\, da_b\right] da_s \\ & + \nu_L\left[1 - \int l(a)\, f(a)\, da\right],\end{aligned}$$

where $\mathcal{A}_Y \equiv 1/[\beta^{\beta}(1-\beta)^{1-\beta}]$, $J(a_b) \equiv \int x(a_s, a_b)^{(\sigma-1)/\sigma}\, m(a_s, a_b)\, da_s$ is the CES intermediate bundle of buyer a_b, $\chi(a)$ is the shadow value of relaxing buyer a's production constraint, $p^{SP}(a)$ is the planner's shadow-price multiplier for good a rather than a decentralized transaction price, and ν_L is the shadow wage of production labor.

Imposing the entry boundary condition for f and the entrant boundary conditions for $m(a_s, a_b)$ separately on $a_s = 0$ and $a_b = 0$, the Hamiltonian over the interior states is

$$\begin{aligned}\mathcal{H}^{SP} = \mathcal{L}_0^{SP} & + \int \xi_f(a)\left[-\frac{\partial f}{\partial a} + (\lambda(a) - \delta_F - \delta_P)\, f(a)\right] da \\ & + \iint \xi_m(a_s, a_b)\Big[-\frac{\partial m}{\partial a_s} - \frac{\partial m}{\partial a_b} \\ & \qquad + (\lambda(a_s) - \delta_F - \delta_P - \delta_M)\, m(a_s, a_b) + \zeta\Gamma f(a_s)\Big] da_s\, da_b \\ & + \xi_\lambda\left[1 - H_E(E) - \int h_H(\lambda(a))\, f(a)\, da\right]. \end{aligned} \tag{27}$$

The boundary conditions are $f(0) = E$, $m(a_s, 0) = \zeta_0\Gamma f(a_s)$, and $m(0, a_b) = \zeta_0\Gamma E$. The shifter $\Gamma = \Gamma(\mathcal{N})$ is held fixed in the coordinate derivatives below; the separate envelope effect from $\Gamma(\mathcal{N})$ is added in Section F.6.

F.3 Static First-Order Conditions

FOC w.r.t. $y(a)$. Differentiating $\log Y$ gives

$$p^{SP}(a) = \frac{y(a)^{-1/\sigma}}{\int y(\tilde{a})^{(\sigma-1)/\sigma}\, f(\tilde{a})\, d\tilde{a}}.$$

FOC w.r.t. $x(a)$. From the production and resource constraints:

$$\chi(a) = p^{SP}(a). \tag{28}$$

FOC w.r.t. $l(a)$.

$$p^{SP}(a)\,\beta\,x(a) = \nu_L\,l(a). \tag{29}$$

FOC w.r.t. $x(a_s, a_b)$: CES optimality. Define the expenditure share

$$s(a_s, a_b) \equiv \frac{x(a_s, a_b)^{(\sigma-1)/\sigma}}{J(a_b)}.$$

The per-match quantity $x(a_s, a_b)$ enters buyer a_b's production through $J(a_b)$ and supplier a_s's resource constraint. The FOC yields

$$p^{SP}(a_b)\,(1-\beta)\,x(a_b)\,\frac{x(a_s, a_b)^{-1/\sigma}}{J(a_b)} = p^{SP}(a_s). \tag{30}$$

Multiplying both sides by $x(a_s, a_b)$, we obtain the useful static identity:

$$p^{SP}(a_s)\,x(a_s, a_b) = p^{SP}(a_b)\,(1-\beta)\,x(a_b)\,s(a_s, a_b). \tag{31}$$

F.4 Matched-Product Density Envelope and Match-Value Flow

This subsection derives the current-payoff term that enters the costate equation for the matched-product density $m(a_s, a_b)$. The love-of-variety decomposition is an intermediate step: it separates the gross CES variety benefit from the resource cost and then reduces the envelope derivative to the net match-value flow.

Define the reduced current return

$$\Psi^{SP}(f, m) \equiv \max_{y,\,x,\,l,\,\{x(a_s,a_b)\}} \mathcal{L}_0^{SP}.$$

By the envelope theorem, holding the optimal static allocations fixed, $m(a_s, a_b)$ enters $\mathcal{L}_0^{SP}$ in exactly two places: (i) buyer a_b's intermediate bundle $J(a_b)$, and (ii) supplier a_s's resource constraint.

Gross variety benefit. A marginal increase in the matched-product density $m(a_s, a_b)$ expands buyer a_b's input variety:

$$\frac{\partial x(a_b)}{\partial m(a_s, a_b)} = \frac{\sigma(1-\beta)}{\sigma-1}\, x(a_b)\, s(a_s, a_b).$$

Valued at $p^{SP}(a_b)\, f(a_b)$ using (28), the gross variety benefit is

$$\text{Variety benefit} = p^{SP}(a_b)\, f(a_b)\, \frac{\sigma(1-\beta)}{\sigma-1}\, x(a_b)\, s(a_s, a_b).$$

Resource cost. Supplier a_s must deliver $x(a_s, a_b)$ units through this match:

$$\text{Resource cost} = p^{SP}(a_s)\, x(a_s, a_b)\, f(a_b).$$

Net match-flow decomposition. Combining the gross variety benefit and the resource cost:

$$\frac{\partial \Psi^{SP}}{\partial m(a_s, a_b)} = f(a_b)\left[p^{SP}(a_b)\, \frac{\sigma(1-\beta)}{\sigma-1}\, x(a_b)\, s(a_s, a_b) \;-\; p^{SP}(a_s)\, x(a_s, a_b)\right].$$

The first term is the gross variety benefit for buyer a_b, and the second is the resource cost for supplier a_s.

Simplification via CES identity. Using the static identity (31) to substitute the resource cost:

$$\frac{\partial \Psi^{SP}}{\partial m(a_s, a_b)} = \frac{1-\beta}{\sigma-1}\, p^{SP}(a_b)\, f(a_b)\, x(a_b)\, s(a_s, a_b).$$

The factor $\sigma/(\sigma-1) - 1 = 1/(\sigma-1) > 0$ is the net CES variety surplus. This net flow payoff is the primitive source term in the match-value costate equation below. It is positive because an additional supplier variety raises the buyer's CES input bundle by more than the resource cost of the shipment evaluated at the planner's shadow prices. In the decentralized allocation, the private supplier is compensated for the shipment but does not capture this net buyer-side variety surplus, which is why the corresponding match-value term appears in the planner's innovation value but not in the private HJB.

F.5 Product-Line Density Envelope

$f(a)$ enters $\mathcal{L}_0^{SP}$ through four terms. First, the objective $\log Y$ contributes $\frac{\sigma}{\sigma-1}\, p^{SP}(a)\, y(a)$. Second, when age a is the buyer, the resource constraint for intermediate purchases contributes

$$-\int p^{SP}(\tilde{a})\,x(\tilde{a},a)\,m(\tilde{a},a)\,d\tilde{a} = -p^{SP}(a)\,(1-\beta)\,x(a),$$

using the aggregate intermediate-use identity implied by (31), the CES share definition, and (26). Third, when age a is the supplier, the resource constraint contributes $p^{SP}(a)\,[x(a)-y(a)]$. Fourth, the labor constraint contributes $-\nu_L\,l(a) = -p^{SP}(a)\,\beta\,x(a)$, using (29). Summing:

$$\begin{aligned}\frac{\partial \Psi^{SP}}{\partial f(a)} &= \frac{\sigma}{\sigma-1}\,p^{SP}(a)\,y(a) + p^{SP}(a)\,[x(a)-y(a)] \\ &\quad - p^{SP}(a)\,(1-\beta)\,x(a) - p^{SP}(a)\,\beta\,x(a) \\ &= \frac{p^{SP}(a)\,y(a)}{\sigma-1}. \end{aligned} \tag{32}$$

F.6 Aggregate Matching-Stock Envelope

The coordinate envelopes above hold the matching shifter Γ fixed. In the baseline specification, however, Γ depends on the aggregate product-line stock $\mathcal{N} = \int f(a)\,da$:

$$\Gamma(\mathcal{N}) = \left(\frac{\mathcal{N}}{\bar{\mathcal{N}}}\right)^{-\eta}.$$

Private agents take this shifter as given, while the planner internalizes how an additional product line changes the matching intensity faced by all product lines. The derivative of the log shifter with respect to any age-a product line is

$$\frac{\delta \log \Gamma(\mathcal{N})}{\delta f(a)} = -\frac{\eta}{\mathcal{N}}. \tag{33}$$

Let $V^M(a_s, a_b) = \xi_m(a_s, a_b)/(K f(a_b))$ be the normalized match value used below. The normalized value exposure to a proportional change in Γ is

$$\begin{aligned}\Omega_\Gamma \equiv \Gamma\Bigg[&\zeta \iint V^M(a_s, a_b) f(a_s) f(a_b)\,da_s\,da_b \\ &+ \zeta_0 E \int V^M(a_s, 0) f(a_s)\,da_s + \zeta_0 E \int V^M(0, a_b) f(a_b)\,da_b\Bigg]. \end{aligned} \tag{34}$$

The first term is the value exposure of interior match formation; the last two terms are the corresponding entrant-as-buyer and entrant-as-supplier boundary exposures. Combining

(33) with (34) gives the product-line matching-stock source

$$\bar{\mathcal{R}}_{\Gamma,f} = -\eta \frac{\Omega_\Gamma}{\mathcal{N}}. \tag{35}$$

In the product-line-stock specification this source is common across ages, because every product line contributes one unit to $\mathcal{N}$. It is the matching-stock congestion term in (20).

F.7 Dynamic First-Order Conditions

Match value HJB. The costate equation for $\xi_m(a_s, a_b)$ in the interior ($a_s, a_b > 0$) gives:

$$\begin{aligned}\big(\rho + \delta_F + \delta_P + \delta_M - \lambda(a_s)\big)\, \xi_m(a_s, a_b) = \frac{\partial}{\partial a_s}\xi_m(a_s, a_b) + \frac{\partial}{\partial a_b}\xi_m(a_s, a_b) \\ + \frac{1-\beta}{\sigma - 1}\, p^{SP}(a_b)\, f(a_b)\, x(a_b)\, s(a_s, a_b).\end{aligned} \tag{36}$$

Here $\lambda(a_s)$ is the supplier's innovation rate, since the first argument denotes supplier age. The primitive costate $\xi_m(a_s, a_b)$ is the shadow value of increasing the matched-product density $m(a_s, a_b)$, whereas the main text uses the buyer-density-normalized match value.

Costate Normalization and Planner Values. This step maps the primitive costates into the planner value objects used in the main text. The primitive flow in (36) involves the shadow price and gross output of the buyer, $p^{SP}(a_b)x(a_b)$. Since the planner's shadow price is proportional to the full marginal cost shifter, the definition $D^{SP}(a) \equiv x(a)\,(P^{SP})^{1-\sigma}\, c^{SP}(a)^{\sigma}$ implies that $p^{SP}(a)x(a)$ is proportional to $R^{SP}(a)$ for any age a. Choose the positive age-independent normalization constant K so that

$$\frac{p^{SP}(a)\, x(a)}{(\sigma - 1)K} = R^{SP}(a), \qquad R^{SP}(a) \equiv \left(\frac{c^{SP}(a)}{P^{SP}}\right)^{1-\sigma} D^{SP}(a).$$

The exact value of K depends only on the normalization of the planner's multipliers. Because K is age-independent, it only fixes units: dividing all costates by K preserves V^{SP}/w_H^{SP}, the innovation FOC, and the entry FOC.

The following aggregation turns primitive costates into an incumbent product-line innovation value. An age-a innovator enters the supplier position, so the inherited buyer-link portfolio is weighted by $m(a, a_b) f(a_b)/f(a)$, the mass of age-a_b buyer product lines attached to one age-a supplier product line. Define the normalized value functions:

$$\begin{aligned}
V^{SP}(a) &\equiv \frac{\xi_f(a)}{K} + \int V^M(a, a_b)\,\frac{m(a, a_b) f(a_b)}{f(a)}\,da_b, \\
V^M(a_s, a_b) &\equiv \frac{\xi_m(a_s, a_b)}{K f(a_b)}, \\
w_H^{SP} &\equiv \frac{\xi_\lambda}{K}.
\end{aligned}$$

In the definition of $V^{SP}(a)$, innovator a enters the supplier position (first argument of V^M and m), and the integral runs over all buyers a_b.

Dividing the primitive costate equation by $K f(a_b)$ gives:

$$\begin{aligned}
\big(\rho + 2\,(\delta_F + \delta_P) + \delta_M - \lambda(a_s) - \lambda(a_b)\big)\, V^M(a_s, a_b) = \frac{\partial}{\partial a_s} V^M(a_s, a_b) + \frac{\partial}{\partial a_b} V^M(a_s, a_b) \\
+ (1-\beta)\, R^{SP}(a_b)\, s(a_s, a_b),
\end{aligned}$$

where

$$s(a_s, a_b) = \frac{c^{SP}(a_s)^{1-\sigma}}{\int c^{SP}(\tilde{a})^{1-\sigma}\, m(\tilde{a}, a_b)\, d\tilde{a}}.$$

The additional $\delta_F + \delta_P - \lambda(a_b)$ term on the left-hand side comes from differentiating the buyer density $f(a_b)$ in the normalization.

FOC w.r.t. $f(a)$. Differentiating Ψ^{SP} with respect to $f(a)$, holding Γ fixed (see Section F.5):

$$\begin{aligned}
\rho\, \xi_f(a) = \frac{p^{SP}(a)\, y(a)}{\sigma - 1} + \frac{\partial}{\partial a}\xi_f(a) + \xi_f(a)\,(\lambda(a) - \delta_F - \delta_P) \\
- \xi_\lambda\, h_H(\lambda(a)) + \zeta\Gamma \int \xi_m(a, a_b)\, da_b + K\bar{\mathcal{R}}_{\Gamma,f}(a).
\end{aligned}$$

In the ξ_m integral, innovator a is in the supplier position (first argument of ξ_m). The source $\bar{\mathcal{R}}_{\Gamma,f}(a)$ is the normalized derivative of the matching shifter with respect to the product-line stock. It is defined in (35); its primitive contribution to the ξ_f equation is multiplied by the common normalization K.

Derivation of $V^{SP}(a)$'s HJB. The definition of $V^{SP}(a)$ involves the buyer-link mass $\int V^M(a, a_b)\, m(a, a_b) f(a_b)/f(a)\, da_b$, where innovator a enters the supplier position. Evaluating the ξ_m HJB (36) with supplier a and buyer a_b, the primitive flow is

$$\frac{1-\beta}{\sigma-1}\, p^{SP}(a_b)\, f(a_b)\, x(a_b)\, s(a,a_b).$$

Multiply by $m(a,a_b)/f(a)$ and integrate over a_b. Using the static identity (31) with supplier a and buyer a_b:

$$p^{SP}(a_b)\, x(a_b)\, s(a,a_b) = \frac{p^{SP}(a)\, x(a,a_b)}{1-\beta}.$$

Then applying the reciprocity relation $m(a,a_b)\, f(a_b) = m(a_b,a)\, f(a)$ and the resource constraint $\int x(a,a_b)\, m(a_b,a)\, da_b = x(a) - y(a)$:

$$\frac{1}{f(a)} \int \frac{1-\beta}{\sigma-1}\, p^{SP}(a_b)\, f(a_b)\, x(a_b)\, s(a,a_b)\, m(a,a_b)\, da_b = \frac{p^{SP}(a)\,[x(a)-y(a)]}{\sigma-1}.$$

Combining with the ξ_f flow $= p^{SP}(a)\, y(a)/(\sigma-1)$ from (32):

$$V^{SP}(a) \text{ flow} = \frac{p^{SP}(a)\, y(a)}{\sigma-1} + \frac{p^{SP}(a)\,[x(a)-y(a)]}{\sigma-1} = \frac{p^{SP}(a)\, x(a)}{\sigma-1} = K\, R^{SP}(a).$$

The transition terms aggregate in the same way. If the entrant boundary $m(a,0) = \zeta_0 \Gamma f(a)$ is represented explicitly, the ξ_f equation contains the boundary contribution $\zeta_0 \Gamma \xi_m(a,0)$, while the integrated ξ_m equation contains

$$-\frac{m(a,0)}{f(a)} \xi_m(a,0) = -\zeta_0 \Gamma \xi_m(a,0),$$

so the boundary terms cancel. Random match formation contributes

$$\zeta\Gamma \int \frac{\xi_m(a,a_b)}{K}\, da_b = \int V^M(a,a_b) \zeta \Gamma f(a_b)\, da_b,$$

whereas relationship destruction contributes

$$-\frac{\delta_M}{K f(a)} \int \xi_m(a,a_b) m(a,a_b)\, da_b = -\int V^M(a,a_b) \delta_M m(a_b,a)\, da_b,$$

using $V^M(a,a_b) = \xi_m(a,a_b)/[K f(a_b)]$ and the reciprocity relation $m(a,a_b) f(a_b)/f(a) = m(a_b,a)$. The δ_F and δ_P terms multiply the whole supplier product-line value and remain on the left-hand side rather than appearing as separate net link-flow terms. After rescaling by K and applying the normalized match-value and reciprocity identities, the V^{SP} HJB is:

$$\begin{aligned}\left(\rho + \delta_F + \delta_P - \lambda(a)\right) V^{SP}(a) &= R^{SP}(a) + \frac{\partial}{\partial a} V^{SP}(a) - w_H^{SP}\, h_H(\lambda(a)) \\ &\quad + \bar{\mathcal{R}}_{\Gamma,f}(a) + \int V^M(a, a_b) \left\{ \zeta \Gamma f(a_b) - \delta_M\, m(a_b, a) \right\} da_b,\end{aligned} \tag{37}$$

where $R^{SP}(a) = (c^{SP}(a)/P^{SP})^{1-\sigma} D^{SP}(a)$. This matches the main text equation (20). The product-line exit rate δ_P is absorbed into the product-line value on the left-hand side. It therefore does not appear in the explicit net link-flow term, which tracks formation and destruction of individual buyer links conditional on the supplier product line surviving.

Here $\partial V^{SP}/\partial a$ collects the deterministic age drift of the primitive costates, while the explicit integral collects stochastic buyer-link creation and destruction.

FOC w.r.t. $\lambda(a)$**.**

$$\lambda^{SP}(a) = \left[\frac{\varphi}{\gamma\, w_H^{SP}}\, V^{SP}(a) \right]^{\frac{1}{\gamma - 1}}.$$

FOC w.r.t. entry E**.** Entry is a boundary control. A marginal entrant uses entry R&D labor and creates an age-zero product line. In the normalized planner value, $V^{SP}(0)$ is the value of this boundary object: it includes the age-zero product-line costate and the entrant's initial buyer-link portfolio. In primitive costates, the explicit E-dependent match-boundary term is the entrant-as-supplier boundary $m(0, a_b) = \zeta_0 \Gamma E$, which links the entrant supplier to existing buyers. The entrant-as-buyer boundary $m(a_s, 0) = \zeta_0 \Gamma f(a_s)$ is tied to the age-zero product-line boundary and is incorporated in $V^{SP}(0)$ after applying the boundary conditions. After integrating the age derivatives in (27) by parts, the E-dependent boundary and resource terms are

$$\xi_f(0) E + \zeta_0 \Gamma E \int \xi_m(0, a_b)\, da_b - \xi_\lambda H_E(E).$$

The primitive first-order condition is therefore

$$\xi_\lambda H_E'(E) = \xi_f(0) + \zeta_0 \Gamma \int \xi_m(0, a_b)\, da_b.$$

Using the definition of $V^{SP}(0)$, dividing by the common normalization K gives

$$w_H^{SP} H_E'(E) = V^{SP}(0). \tag{38}$$

The planner KKT condition is $E \geq 0$, $w_H^{SP} H_E'(E) \geq V^{SP}(0)$, and $E[w_H^{SP} H_E'(E) - V^{SP}(0)] = 0$. For $H_E(E) = E^{\gamma_E}/\varphi_E$, an interior entry margin satisfies

$$E = \left\{ \frac{\varphi_E}{\gamma_E w_H^{SP}} V^{SP}(0) \right\}^{\frac{1}{\gamma_E - 1}}. \tag{39}$$

F.8 Cost and Demand Shifters

We finally translate the static first-order conditions into the cost and demand shifters used to compute R^{SP} in the planner value equations. The primitive multipliers determine the social cost and demand shifters that characterize the intratemporal allocation. From constant-returns cost duality in the static block, the planner's shadow price is proportional to the full marginal cost shifter:

$$p^{SP}(a) \propto c^{SP}(a).$$

Substituting this relation into the CES input-demand FOC (30) gives

$$x(a_s, a_b) = (1-\beta)\, c^{SP}(a_s)^{-\sigma}\, c^{SP}(a_b)^{\frac{\sigma-\beta}{1-\beta}}\, x(a_b).$$

Inserting this into the production function (17) for buyer a_b and cancelling $x(a_b)$ from both sides yields the cost recursion:

$$c^{SP}(a_b) = \left(\int c^{SP}(a_s)^{1-\sigma}\, m(a_s, a_b)\, da_s \right)^{\frac{1-\beta}{1-\sigma}}. \tag{40}$$

Relabeling the left-hand-side age as a, substituting into the planner goods-market clearing condition (26), and defining

$$D^{SP}(a) \equiv x(a)\, [P^{SP}]^{1-\sigma}\, [c^{SP}(a)]^{\sigma}$$

gives

$$D^{SP}(a) = (1-\beta) \int [c^{SP}(a_b)]^{\frac{\beta}{1-\beta}(\sigma-1)} D^{SP}(a_b)\, \frac{m(a, a_b) f(a_b)}{f(a)}\, da_b + 1.$$

From the labor market clearing condition (19):

$$(P^{SP})^{1-\sigma} = \beta \int c^{SP}(a)^{1-\sigma} \, D^{SP}(a) \, f(a) \, da,$$

and the CES price index is

$$P^{SP} = \left(\int c^{SP}(a)^{1-\sigma} \, f(a) \, da \right)^{\frac{1}{1-\sigma}}. \tag{41}$$

These are the social planner analogs of the decentralized static cost and demand objects. Unlike the decentralized contracting block, the planner prices intermediate and final uses at social marginal cost and does not use current fixed transfers to allocate bilateral operating surplus.

G Numerical Implementation

The analytical formulas in Section D are used directly in the numerical implementation for the stationary densities. We evaluate the product-line density $f(a)$ from its closed-form expression using cumulative trapezoidal integration, compute Γ from the resulting aggregate product-line stock, compute the matched-product density $m(a_s, a_b)$ pointwise from its closed-form expression conditional on $(f(a_s), \Gamma)$, and use finite-difference methods for the HJB equations. Deterministic aging in the HJB equations is discretized with a one-sided derivative that uses the continuation value at the next age grid point. Under our sign convention, this derivative is implemented as a forward difference in each age dimension. Below, D_a denotes the resulting one-sided finite-difference matrix, including the boundary row at the upper age grid point; on two-dimensional grids, D_{a_s} and D_{a_b} denote copies of this matrix applied along the supplier and buyer dimensions.

G.1 Stationary Equation Blocks and Fixed-Point Variables

The stationary calculation separates common state equations from allocation-specific equilibrium conditions. The common state equations map a candidate policy pair (λ, E) into the product-line density, matched-product density, and matching shifter. The decentralized equilibrium and the planner allocation then differ in the static allocation, value equations, and policy update used to determine (λ, E, w_H).

Common state blocks. For any candidate (λ, E), the product-line density $f(a)$ is computed from the KFE and entry boundary, (2) and (3), with firm mass given by (1). Given f, we compute $\mathcal{N}$, the matching shifter Γ from (4), and the matched-product density $m(a_s, a_b)$ from the stationary kernel (6). These state laws are common to the decentralized equilibrium and the planner; they are simply evaluated at the policy objects generated by the relevant allocation problem.

Decentralized stationary equilibrium. Given (f, m, Γ), the decentralized block computes the contracting static block in (7)–(9), including costs, input shares, revenues, transfers, and the private flow payoff. The private product-line value then solves the HJB in (11). The outer update uses the private incumbent R&D policy (12), the free-entry condition (14), and R&D labor clearing (13) to update (λ, E, w_H).

Stationary planner allocation. The planner uses the same state laws, but recomputes the intratemporal allocation from the first-best static block (40)–(41). The dynamic block consists of the planner product-line value $V^{SP}(a)$ and match value $V^M(a_s, a_b)$, which satisfy (20) and (21). The planner policy update applies the incumbent R&D allocation condition in (16) together with the planner entry condition (38).

Fixed-point variables. In both stationary calculations, the outer fixed point is over the policy and R&D-wage objects (λ, E, w_H). State densities, static allocations, values, and reporting statistics are recomputed conditional on the current policy guess. Objects such as $\mathcal{N}$, N_f, P^F, P^{SP}, R^{SP}, input shares, transfers, W_I, and $\bar{\mathcal{R}}_{\Gamma,f}$ are therefore derived objects rather than independent fixed-point variables.

The primitive entry flow E is used throughout the fixed point, including in the boundary condition $f(0) = E$, the entry FOC, and the R&D resource constraint. When entry paths are reported relative to the decentralized baseline, we apply the reporting normalization $\widehat{E} = E/E^{DE}$ after solving the primitive equilibrium.

G.2 Product-Line Density f

We discretize a on an evenly spaced grid $\{a_i\}_{i=1}^{N_a}$ with grid spacing Δ_a and $a_1 = 0$. Rather than solving a separate finite-difference system for the stationary product-line density, we evaluate the closed-form solution in (23):

$$f_i = E \exp\{I_i\}, \qquad I_i \approx \int_0^{a_i} \{\lambda(\tilde{a}) - \delta_F - \delta_P\}\, d\tilde{a},$$

where I_i is computed by cumulative trapezoidal integration with $I_1 = 0$. By construction, this procedure imposes $f_1 = E$ and preserves nonnegativity of the product-line density. The same procedure applies under planner allocations after replacing $(\lambda(a), E)$ with the relevant policy objects and recomputing Γ from the induced product-line stock.

G.3 Matched-Product Density m

Given the supplier product-line density $\{f_i\}_{i=1}^{N_a}$, we compute $\mathcal{N}$, $\Gamma = (\mathcal{N}/\bar{\mathcal{N}})^{-\eta}$, and the matched-product density pointwise from (25). For each supplier-buyer age pair (a_i, a_j),

$$m_{ij} \equiv m(a_i, a_j) = \left[\frac{\zeta}{\delta_M} - \left(\frac{\zeta}{\delta_M} - \zeta_0\right) \exp\left(-\delta_M \min\{a_i, a_j\}\right)\right] \Gamma f_i.$$

By construction, this direct evaluation enforces the boundary conditions $m(a_i, 0) = \zeta_0 \Gamma f(a_i)$ and $m(0, a_j) = \zeta_0 \Gamma E$. Hence, once f and Γ are available, the full two-dimensional object $m(a_s, a_b)$ is obtained without solving a separate finite-difference system or iterating forward in buyer age. The static contracting and planner blocks then use this density to compute costs, input shares, revenues, transfers, and value sources.

G.4 Private Value Function V

Given (f, m, Γ) and the static contracting objects, the private product-line value solves the HJB in (11). On the age grid, the implementation uses the one-sided age-derivative matrix D_a and solves the implicit linear system

$$\left[(\rho + \delta_F + \delta_P)I - \Lambda - \mathsf{D}_a\right] V = \pi - w_H h_H(\lambda),$$

where $\Lambda = \text{diag}(\lambda(a_i))$, π is the vector of private flow payoffs, and $h_H(\lambda)$ is evaluated pointwise. Policy iteration updates λ and w_H from the private R&D FOC and the R&D resource constraint.

G.5 Match-Value Function V^M

The planner match value $V^M(a_s, a_b)$ solves the two-dimensional HJB in (21). The implementation discretizes supplier and buyer age on the same grid, applies one-sided finite differences in both age dimensions, and solves the resulting sparse linear system

$$\mathcal{M}^{SP} V^M = b^M,$$

where the source vector has entries $b^M(a_s, a_b) = (1-\beta)R^{SP}(a_b)s(a_s, a_b)$. The matrix $\mathcal{M}^{SP}$ contains the discount, exit, link-destruction, innovation, and age-drift terms from (21). This compact representation is used inside the planner policy iteration before computing the net link-flow source $W_I(a)$.

G.6 Planner Value and Policy Updates

Given V^M, we compute the net link-flow source

$$W_I(a_i) = \sum_j V^M(a_i, a_j)\left\{\zeta\Gamma f_j - \delta_M m(a_j, a_i)\right\}\Delta_a,$$

and compute Ω_Γ and $\bar{\mathcal{R}}_{\Gamma,f}$ from (34)–(35). The planner product-line value then solves the discretized version of (37). With the same one-sided age-derivative matrix D_a,

$$\mathcal{A}^{SP} \equiv \left(\rho + \delta_F + \delta_P\right) I - \Lambda^{SP} - \mathsf{D}_a,$$

and the source vector

$$b_i^{SP} = R_i^{SP} + \bar{\mathcal{R}}_{\Gamma,f}(a_i) + W_I(a_i) - w_H^{SP} h_H(\lambda_i^{SP}),$$

the stationary planner value update is

$$\mathcal{A}^{SP} V^{SP} = b^{SP}.$$

The outer policy iteration updates incumbent R&D, entry, and the R&D resource multiplier from

$$\lambda_i^{SP} = \left\{\frac{\varphi V_i^{SP}}{\gamma w_H^{SP}}\right\}^{\frac{1}{\gamma-1}}, \qquad E^{SP} = \left\{\frac{\varphi_E V_1^{SP}}{\gamma_E w_H^{SP}}\right\}^{\frac{1}{\gamma_E-1}},$$

and

$$1 = H_E(E^{SP}) + \sum_i h_H(\lambda_i^{SP}) f_i \Delta_a.$$

Here V_1^{SP} is the value at the age-zero grid point. The same primitive entry units are used in $f_1 = E$, the entry policy, and the R&D resource constraint. The displayed entry update is the positive interior case of (39); if the lower bound binds, the numerical update is interpreted through the KKT condition in (38) rather than as an unconstrained entry flow.

G.7 Finite-Horizon Transition Solver

The transition exercises solve finite-horizon perfect-foresight paths. The same finite-horizon structure is used for decentralized policy transitions, including uniform entry-tax paths, with the private value equations and the tax-modified entry condition replacing the planner policy update as appropriate. For the decentralized-to-planner transition, the initial state is the calibrated decentralized steady state,

$$f_0 = f^{DE}, \qquad m_0 = m^{DE}, \qquad N_{f,0} = N_f^{DE},$$

and controls may jump at the reform date. At terminal date T, the calculation uses the stationary continuation values associated with the relevant terminal environment, but does not impose the terminal state itself. For the planner transition, terminal-state distances such as $\|f_T - f^{SP}\|$, $\|m_T - m^{SP}\|$, and $|N_{f,T} - N_f^{SP}|$ are reported as horizon diagnostics; the decentralized policy transitions use the analogous distances to their relevant terminal steady states.

The transition calculation iterates on paths for entry, incumbent R&D, and the R&D wage. Given a candidate path $(E_n, \lambda_{n,i}, w_{H,n})$, it first updates the state variables forward. Let $t_n = n\Delta t$ and let a_i denote the age grid. The implementation uses a uniform grid with $\Delta t = \Delta a$, so one time step moves a product line by one age cell. The boundary conditions are

$$f_{n,1} = E_n, \qquad m_{n,i,1} = \zeta_0 \Gamma_n f_{n,i}, \qquad m_{n,1,j} = \zeta_0 \Gamma_n E_n,$$

where Γ_n is computed from the current product-line stock. The product-line density is then advanced along characteristics:

$$f_{n+1,i+1} = f_{n,i} \exp\left\{(\lambda_{n,i} - \delta_F - \delta_P)\,\Delta t\right\}.$$

The transition update for m_t uses the forward state equation rather than the stationary closed-form kernel. In supplier-normalized form, define

$$G^f_{n,i} = \frac{f_{n+1,i+1}}{\max\{f_{n,i}, \varepsilon\}}, \qquad Q^M_{n,i} = \zeta \Gamma_n f_{n+1,i+1} \frac{1 - \exp(-\delta_M \Delta t)}{\delta_M}.$$

Here ε is a numerical floor used only to define the ratio when the lagged supplier density is zero. The interior matched-product density is then updated by

$$m_{n+1,i+1,j+1} = \exp(-\delta_M \Delta t) G^f_{n,i} m_{n,i,j} + Q^M_{n,i}.$$

The source term $Q^M_{n,i}$ is interpreted by continuity as $\zeta\Gamma_n f_{n+1,i+1}\Delta t$ at $\delta_M = 0$. Firm mass is updated by

$$N_{f,n+1} = \exp(-\delta_F \Delta t) N_{f,n} + E_n \frac{1 - \exp(-\delta_F \Delta t)}{\delta_F}.$$

Given the forward state path, the calculation computes the static block at each date and solves the relevant value equations backward from the terminal continuation values. For planner transitions, the backward equations are the finite-horizon analogs of (21) and (20), with the same one-sided age-derivative matrices used in the stationary value equations and with T-date continuation values replacing the infinite-horizon continuation terms. For decentralized policy transitions, the backward step instead uses the private value equation and the policy-specific entry condition. The policy update then applies the entry and incumbent R&D FOCs date by date and chooses $w_{H,n}$ so that the R&D resource constraint holds.

To stabilize the fixed point, the iteration applies damping to the logs of E_n, $w_{H,n}$, and $\lambda_{n,i}$, and then recomputes the state path from the damped controls. We assess convergence using policy fixed-point residuals, R&D-resource residuals, weighted state changes, and terminal state gaps. Before reporting transition results, we verify that the initial state reproduces the decentralized steady state, the terminal gaps are small relative to the reported transition effects, and the constant decentralized path reproduces the analytic decentralized value $U^{DE} = \rho^{-1} \log Y^{DE}$.

G.8 Transition Welfare Computation

For any computed path, household utility is

$$U = \int_0^\infty e^{-\rho t} \log Y_t \, dt, \qquad Y_t = \frac{1}{P_t}.$$

With a terminal stationary continuation at date T, the discretized path value is computed as the finite-path integral plus a terminal tail. Let $y_n = \log Y_n$, $s_n = (y_{n+1} - y_n)/(t_{n+1} - t_n)$, and $\Delta t_n = t_{n+1} - t_n$. The piecewise-linear quadrature on each interval is

$$\int_{t_n}^{t_{n+1}} e^{-\rho t} \left[y_n + s_n (t - t_n) \right] dt = e^{-\rho t_n} \left[y_n \frac{1 - e^{-\rho \Delta t_n}}{\rho} + s_n \frac{1 - e^{-\rho \Delta t_n}(1 + \rho \Delta t_n)}{\rho^2} \right].$$

The stationary terminal tail is

$$U^{tail} = \frac{e^{-\rho T}}{\rho} y_T.$$

The decentralized steady-state path is evaluated with the same welfare routine by feeding the routine a constant decentralized path; the analytic value $U^{DE} = \rho^{-1} \log Y^{DE}$ is used as a check. The consumption-equivalent gain relative to the decentralized steady state is then

$$g^{trans} = \exp\left\{\rho\left(U^{trans} - U^{DE}\right)\right\} - 1,$$

which is the computation underlying (22).

H Counterfactual and Sensitivity Exercises

This appendix reports the exercises that support the welfare results in Section 5. It first describes the transition decomposition used to separate the planner-private wedge into static, net-link-flow, and matching-stock components. It then documents the allocation counterfactuals that isolate entry from incumbent-age R&D allocation, the uniform-entry-tax exercise, and the sensitivity analysis for the matching-stock elasticity.

H.1 Transition Decompositions and Planner Entry Path

The transition exercises in Section 5 use the welfare computation described in Section G.8. To construct the source-block decomposition, we evaluate planner transition variants in which only selected wedge components are active: the static surplus-appropriation source and the two matching-network sources, the match-costate net-link-flow source W_I and the aggregate matching-stock source $\bar{\mathcal{R}}_{\Gamma,f}$. The static component in Table 4 is the transition difference between the baseline decentralized allocation and the static-only planner variant. The network and matching-stock components are then computed as a two-source Shapley decomposition conditional on the static source being active.

The source-block switches are accounting decompositions of the planner-private wedge rather than re-estimated economies. The static-only case replaces private product-line payoffs with the planner static source while omitting the two planner-only network sources. The transfer-share case $\theta = 1$ is a related static appropriability diagnostic. Under the baseline two-part-tariff pricing convention, it matches the normalized static innovation and entry incentives after a common payoff normalization, but it does not price either $W_I(a)$ or $\bar{\mathcal{R}}_{\Gamma,f}$.

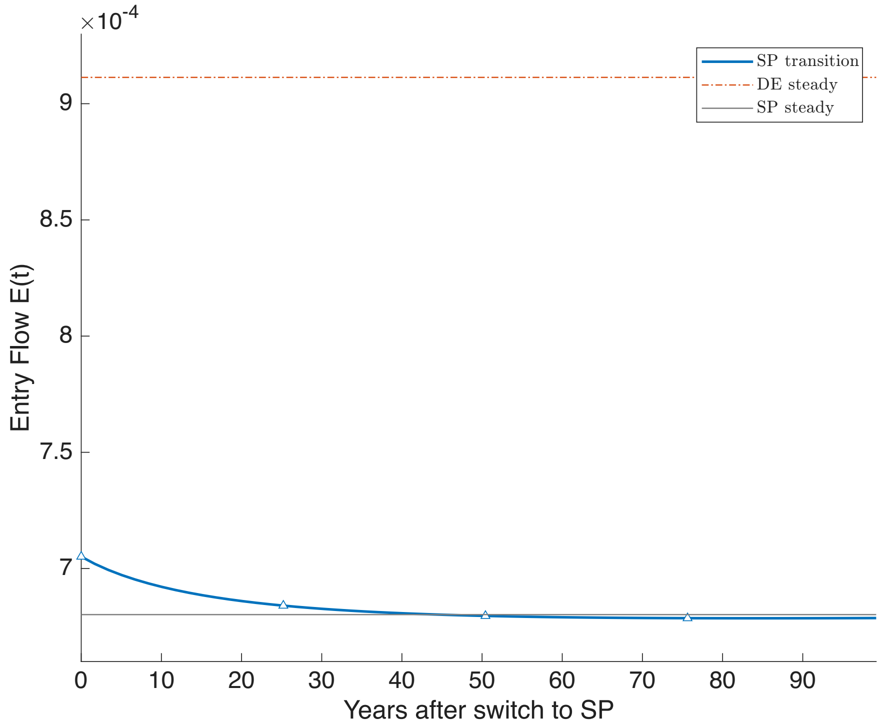

Figure A6: Entry Transition from Decentralized Entry to the Social Planner

Notes: The figure plots the entry flow over the first 100 years of the decentralized-to-planner transition. The transition starts from the baseline decentralized steady state, and the horizontal lines mark the decentralized and planner steady states. Entry jumps downward at the start of the reform, consistent with the dominant matching-stock correction documented in the decomposition.

The no-net-link-flow case sets $W_I(a) = 0$ directly, which isolates the match-costate valuation channel without changing the baseline matching technology. In the stationary primitive economy, the corresponding restriction would be the knife-edge condition $\zeta = \delta_M \zeta_0$. The no-matching-stock case sets the envelope source $\bar{\mathcal{R}}_{\Gamma,f} = 0$, whose primitive analogue is $\eta = 0$. These restrictions are diagnostics for the decomposition; the baseline specification keeps the estimated matching parameters and $\eta = 1$.

Figure A6 plots the transition path for entry under the planner reform. Entry falls immediately and then converges toward the planner steady state, consistent with the dominant matching-stock correction in Table 4.

H.2 Entry and Incumbent-Age Allocation Counterfactuals

The entry-versus-age allocation counterfactuals use allocation paths rather than private or planner FOC fixed points. For each date, define the incumbent R&D labor age share

$$q_t(a) = \frac{h_H(\lambda_t(a)) f_t(a)}{H_{H,t}}, \qquad H_{H,t} = 1 - H_E(E_t).$$

Each counterfactual path selects the entry path E_t from either the decentralized or planner transition, which pins down entrant R&D labor $H_E(E_t)$, and selects the incumbent age-share path $q_t(a)$ from either transition. The residual incumbent R&D labor $H_{H,t}$ is allocated across ages using the selected $q_t(a)$, and $\lambda_t(a)$ is reconstructed from that allocation and the current product-line density. The state laws are then updated forward and evaluated with the planner static block. Thus this counterfactual isolates entry versus incumbent-age R&D allocation from static pricing and surplus-appropriation distortions. The skilled-wage series on these paths is used only to recover the implied R&D effort; it is not an equilibrium wage that supports the counterfactual allocation through FOCs.

H.3 Entry-Cost Accounting and Uniform Entry Tax

For entry-cost-equivalent accounting, a target log-entry contribution $\Delta \log E$ maps into the proportional private entry-cost multiplier

$$\tau_E^{\text{eq}} = \exp\{-(\gamma_E - 1)\Delta \log E\},$$

so the reported entry-cost-equivalent rate is $\tau_E^{\text{eq}} - 1$. This mapping is a one-dimensional accounting device, not a decentralization result.

The planner-implied wedge calculation evaluates private entry and incumbent R&D FOCs along the planner transition path. Negative implied private cost multipliers mean that nonnegative proportional cost wedges alone cannot support the planner path; full decentralization also requires value or appropriability transfers for entry and incumbent R&D, together with an instrument for the static surplus-allocation wedge.

For the uniform-entry-tax policy exercise, the path is generated by decentralized FOCs rather than by imposing planner allocations. A constant entry-cost multiplier τ_E changes the private entry condition to

$$V_t^{DE}(0) = \tau_E w_{H,t}^{DE} H_E'(E_t),$$

Table A2: Sensitivity to the Matching-Stock Elasticity

Matching-stock elasticity η	Entry reduction (%)	Steady-state CE (%)	Transition CE (%)	Stock entry component (%)	Stock welfare component (%)
0.1	0.81	-0.03	0.02	-1.80	-0.16
1	25.36	2.56	1.71	-38.93	1.86

Notes: Each row reports the planner comparison after estimating the four internally estimated parameters at the reported η. Entry reduction is the percentage reduction in planner steady-state entry relative to the decentralized steady state. CE gains are consumption-equivalent percentage gains relative to the decentralized allocation. Matching-stock components are Shapley contributions from the matching-stock source. The abbreviation CE denotes consumption-equivalent units. The $\eta = 0.1$ row is motivated by a one-stock interpretation of Miyauchi (2024)'s supplier-density elasticity.

while incumbent R&D, entry, and the R&D wage continue to update from private FOCs and the same aggregate R&D resource constraint. Comparing these decentralized tax paths with the planner path gives the reported share of the planner transition gain captured by the tax.

H.4 Matching-Stock Elasticity Sensitivity

The baseline specification fixes the matching-stock elasticity at $\eta = 1$, the scale-invariant benchmark for aggregate product-line meetings. Because the decentralized baseline normalizes $\Gamma = 1$, the targeted steady-state moments used in the minimum-distance estimation do not by themselves identify η. To quantify the role of this benchmark elasticity, Table A2 re-estimates $(\zeta, \zeta_0, \varphi, \varphi_E)$ at $\eta = 0.1$ and recomputes the planner comparison.

The table shows that the magnitude of the over-entry result is quantitatively tied to the scale-invariant baseline. In the $\eta = 0.1$ case, motivated by Miyauchi (2024), the planner reduces entry by only 0.8%, and the transition CE gain falls to 0.02%. These results support treating $\eta = 1$ as a scale-invariant baseline rather than as a data-identified elasticity.

For the $\eta = 0.1$ case, an auxiliary uniform-entry-tax grid selects a much smaller tax, about 7% of entry costs. Because the planner transition gain is only 0.02% in this case, the denominator in policy-gain shares is small and the shares are not numerically stable. The exercise therefore shows that the entry-tax magnitude is specific to the scale-invariant baseline and is substantially attenuated under lower matching-stock elasticities.