

Link Dynamics and Welfare in Production Networks*

Mario Canales

UCLA & CBCh

This version: October 22, 2026.

[Link to Current Draft](#)

Abstract

How much does production network adjustment matter for aggregate welfare? Are firms' relationship choices efficient? Using the universe of firm-to-firm transactions in Chile from 2017 to 2024, I document that new relationships account for 56% of active links but only 12% of transaction value. Older relationships are larger and less likely to dissolve, and larger initial transaction size associates with longer survival. These patterns motivate a tractable endogenous production network model in which relationship formation is driven by a persistent bilateral compatibility component. In the model, cost-based Domar weights are no longer sufficient statistics because network adjustment affects welfare through changes in buyers' sourcing opportunities and the allocation of labor. However, simulation results indicate that the role of the endogenous network margin is limited. Following a 0.2 decline in log productivity among the top 5% of firms, network adjustment accounts for less than 1% of the total welfare change. Finally, the efficiency of endogenous network formation is analyzed. Theoretical results show different wedges influencing network formation decisions. Quantitatively, the gap in network formation is driven by the appropriability wedge, the difference between the private incentive to form links with the social benefit. Despite the sub-optimal allocation, the welfare gains from the second best are quantitatively modest.

*Preliminary draft, please do not circulate without permission. Refine.ink was used to check the paper for consistency and clarity.

I thank the comments and suggestions of Maria Jose Corvalan, Hugo Hopenhayn, Nico Vöigtlander, David Baqaee, Romain Wacziarg, Clemence Tricaud, Brian Wheaton. Responsibility for results, opinions, and errors is the authors' alone.

The views expressed are those of the author and do not necessarily those of the Central Bank of Chile (CBC) or its Board members. The CBC has access to anonymized information from various public and private entities, by virtue of collaboration agreements signed with these institutions. The information contained in the databases of the Chilean IRS is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of the Service. Officials of the Central Bank of Chile processed the disaggregated data from the Internal Revenue Service. Email: mario.canales.phd@anderson.ucla.edu

I Introduction

Production networks are not static. The network of buyer-seller relationships reshapes continuously as firms establish, expand and terminate customer relationships. In fact, business relationship turnover is large. Roughly half of buyer-seller relationships observed in a given year will be terminated by the end of the year.¹ These rates far exceed typical job destruction (10-12%) (Davis and Haltiwanger, 1990) and firm exit rate (about 20%) (Bartelsman et al., 2009). The substantial link turnover raises the question of how changes in trading relationships affect the economy's response to shocks. The aggregate importance of this margin of adjustment, however, depends on which relationships are created and destroyed. Changes in the network of firms affect buyers' marginal costs, and thus pricing decisions and the cost faced by downstream customers.

Additionally, firms choosing these relationships may not fully internalize their downstream effects, raising efficiency considerations. Firms choose their trading relationships based on private profitability, but the benefits of a relationship can extend beyond the firms directly involved to downstream producers and final consumers, what is known as network externality. To the extent that these effects are not reflected in sellers' private returns, incentives to create and maintain relationships may differ from those that maximize welfare. This raises the question of how large the resulting inefficiency is and how much welfare could improve through a different allocation of business relationships.

In this paper, I study how endogenous network formation influences the response of the economy to economic shocks and affects the efficiency of the production network. To answer both questions, we need to understand how firms choose which relationships to create and maintain, and how their private profitability relates to their social value. I combine empirical evidence on business relationships with a quantitative model of endogenous production networks. The empirical analysis disciplines the mechanism governing which relationships firms create and maintain. The model then is used to quantify how network adjustment changes the welfare response of economic shocks and evaluate the efficiency of the production network.

Using data covering the universe of firm-to-firm transactions in Chile between 2017 and 2024, one striking fact about relationship selection emerges. Newly formed

¹In fact, using firm-to-firm data from Chile, I document that 51% of active relationships in a given year will be terminated by the end of the year. Similar magnitudes are reported for Belgium (43%) (Dhyne et al., 2021) and Japan (50%) (Koike-Mori and Okumura, 2024)

relationships account for 56% of active links, but only for 12% of transaction value. In contrast, relationships that has been active for 5 or more years account for just 15% of active links, but 65% of transaction value. This contrast between relationship counts and transaction shares motivates a closer look at the life-cycle evolution of business links.

Two facts help discipline the fundamental process governing business relationship. First, the hazard rate of the business link decreases with how long the relationship have been active, even when controlling for buyer-year and seller-year fixed effects. This suggests that the link termination decision may be driven by a match-specific or bilateral fundamental. Second, transaction value increases sharply with relationship age, even when controlling for buyer-year and seller-year fixed effects. This suggests that the bilateral fundamental must have some productive value, challenging models where relationship heterogeneity is driven by fixed costs.² Importantly, differences in size at the beginning of the links' life-cycle translate to differences in the duration of the relationship, indicating persistence in the productive bilateral fundamental. Together, these facts imply that long-lived relationships are not simply those that avoid exit by chance, they are systematically larger and more stable.

These empirical findings motivate the inclusion of a bilateral component with productive value to model endogenous network decisions. The match-specific compatibility shifter $r_{ij,t}$ captures the suitability of the good sold by supplier j for the production process of buyer i in period t . Greater bilateral compatibility increases the transaction value, making the relationship more profitable relative to its maintenance cost. Firms' productivity and bilateral compatibility follow exogenous AR(1) processes, while firms make relationship decisions based on current-period profits. Persistence in these fundamentals generates relationship dynamics through a sequence of static equilibria.³

The endogenous network formation model features a representative household with CES preferences, and a continuum of monopolistically competitive firms producing differentiated varieties using labor and intermediate inputs through a Cob-Douglas production function. In the model, firms have to pay for each buyer. Sellers' optimal decision to create links balance the private profits generated by the relationship with a common fixed maintenance cost. Therefore, labor provided by the household is rented by firms for production and for network maintenance. The optimal relationship decision takes the form of a cutoff in bilateral compatibility,

²See [Bernard et al. \(2022\)](#); [Lim \(2018\)](#); [Dhyne et al. \(2023\)](#), for example.

³In other words, the model features a static equilibrium with dynamic states.

that is sellers serve buyers whose compatibility exceeds the relevant threshold. These decisions determine the set of inputs available to each buyer and their contribution to its production efficiency.

The model is then estimated. The parameters governing firms productivity are estimated directly from the data. On the other hand, the parameters governing bilateral compatibility and firms' network formation decision are estimated using simulated method of moments targeting the relationship life-cycle moments previously documented.

To study the aggregate implications of endogenous network adjustment, I analyze the response of the economy to exogenous shocks to firms' productivity. In contrast to models with fixed networks ([Baqaee and Farhi, 2019](#)), models with endogenous network adjustment imply that the cost-based Domar weight no longer fully characterizes welfare changes. Network adjustment opens two new channels that affect welfare: changes in the use of labor (production versus network maintenance), and changes in the sourcing capacity of the firms. The first channel emerges because to sell in the production network, firms need to hire labor to maintain and operate its relationships with customers. Therefore, shocked firms adjust their networks and also change their allocation of labor. The second channel emerges because by adjusting its connections, the sourcing capacity of buyers change, making them relatively more or less productive.

Quantitatively, however, network adjustment plays a limited role in the economy's response to productivity shocks. Following a 0.2 decline in log productivity to the top 5 percent of firms by productivity, welfare is reduced by 42%. However, network adjustment only contributes with 0.06% of the total welfare decline. The limited role of network adjustment is explained by opposing forces. On one side, economic shocks affect relationship cutoffs and therefore the mass of active links. On the other side, the newly incorporated mass of active links is close to the profitability condition, making these links marginal. The second force changes the composition of the average sourcing capacity, neutralizing part of the effect of the increase in link density.

Then, the efficiency of the decentralized equilibrium is studied. The decentralized equilibrium features two distortions. First, firm prices have markups that add up along the supply chain, leading to the double marginalization problem extensively studied in the literature. Second, firms' decision to form business relationship may differ from social planner, what is known as network externality.

To quantify how efficient is decentralized network formation, I consider the

problem of a constrained planner who chooses the production network while retaining the decentralized price mechanism conditional on the network. Notably, the problem of the constrained planner is characterized by choosing the bilateral quality cutoffs that balance the social marginal benefit of the business link from its resource cost. For each selected network, firms minimize production costs and charge the same constant markup as in the decentralized economy, while prices, expenditures, and quantities adjust to clear markets. This benchmark measures the welfare gains from coordinating relationship selection when the planner cannot directly correct pricing distortions. It therefore evaluates network formation conditional on the prevailing markup structure.

Quantitative results show that the appropriation wedge is the main driver of the gap between the decentralized and the planner threshold. This wedge arises because in the decentralized equilibrium, sellers capture exactly a fraction $1/\sigma$ of the bilateral sales as variable profits. In contrast, the constrained social planner values a link based on its marginal contribution to downstream productivity, a benefit governed by the CES love-of-variety elasticity $1/(\sigma - 1)$. Despite significant differences in network density, where decentralized thresholds are 19.3% to 23.5% higher than the planner allocation, the efficiency cost is a merely 0.015% in consumption-equivalent terms.

The paper makes three contributions.

First, it provides evidence on the the drivers of link turnover. Contrary to other models of endogenous network formation (Bernard et al., 2022; Lim, 2018; Dhyne et al., 2023), the patterns of links' life-cycle suggest that a productive bilateral component is needed. In fact, a model with only cost-side bilateral heterogeneity and no productive bilateral component cannot generate the same residual expenditure-age gradient after correctly absorbing buyer-year and seller-year components.

Second, it develops a general equilibrium framework of endogenous production networks with continuous link intensities that is consistent with the documented empirical facts. The model is then used to study the relevance of network adjustments for welfare and the propagation of shocks. The results indicate that despite significant changes in the topology of the network reacts, the welfare effect of network adjustment is very limited.

Third, it formulates and solves the problem of a constrained social planner, focusing the effect of network formation decisions on efficiency. The problem is characterized by the planner choosing bilateral cutoffs that can be compared to the decentralized equilibrium thresholds. Several sources of inefficiency are iden-

tified. Results indicate that the decentralized equilibrium underprovides business links, but the welfare gains from the constrained planner are very small.

The paper builds on the production-networks literature that provides the general-equilibrium workhorse for studying how input-output linkages shape aggregate outcomes. Seminal contributions such as [Acemoglu et al. \(2012\)](#), [Carvalho \(2014\)](#), and [Baqaee and Farhi \(2019\)](#) establish a tractable CES-based framework in which equilibrium price indices, substitution patterns, and downstream linkages determine aggregate propagation. A unifying feature of this literature is that the production network itself is taken as a primitive and static object: the set of buyer-supplier relationships does not respond endogenously to shocks or incentives. Our contribution is to retain this general-equilibrium structure-CES demand and input aggregation, monopolistic competition, and price-index propagation-while endogenizing the evolution of the network in a general setting with dynamic states.

A second literature develops models of endogenous production-network formation. [Oberfield \(2018\)](#) provides a theory of equilibrium input-output architecture in which firms' sourcing decisions jointly determine network structure, emphasizing the role of complementarities and statistical network properties. More closely related to the approach of this paper, [Lim \(2018\)](#) and [Bernard et al. \(2022\)](#) develop tractable general-equilibrium environments in which buyer-seller relationships are formed subject to fixed formation costs, generating sparse endogenous networks and extensive-margin adjustment. In these models, persistence in relationships arises primarily from cost-side frictions, while conditional on survival the economic value of a relationship is largely static. These models contradict the empirical evidence documented on link's life-cycle. Therefore, a model where relationship formation depends on a persistent match-specific productivity component is developed.

Finally, our analysis relates to work on efficiency, misallocation, and welfare in disaggregated economies. [Hsieh and Klenow \(2009\)](#) and [Baqaee and Farhi \(2019\)](#) provide general frameworks for measuring allocative efficiency and aggregate productivity in economies with heterogeneous firms and input-output linkages, typically treating the network structure as given. [Lim \(2018\)](#) also discusses how relationship frictions affect aggregate outcomes through network formation. In related work, [Arkolakis et al. \(2025\)](#) study the efficiency considerations of endogenous network formation in presence of search frictions across space. The main difference is that this paper builds the cleanest model to quantify the role of network formation on welfare. In contrast, by introducing search frictions, [Arkolakis et al. \(2025\)](#) effectively mix the network externality with matching externalities.

The remainder of the paper is organized as follows. Section II describes the data and documents the empirical facts motivating the model. Section III presents the general equilibrium framework and characterizes equilibrium. Section IV discuss the estimation of the model. Section V studies the propagation of shocks and quantifies the relevance of network adjustments for welfare. Section VI formulates the constrained planner problem and quantify the inefficiency of the decentralized equilibrium. Finally, Section VII concludes.

II Empirical Evidence

This section presents the empirical evidence that motivates the inclusion of a link-level fundamental shaping the life-cycle patterns. The main goal is to understand how selection operates at the link level.

II. A Data

The main dataset consists of firm-to-firm administrative records from Chile, covering all transactions between firms from 2017 to 2024.⁴ It includes anonymous firm identifiers for both the seller and buyer, which facilitates tracking transactions between them over time. The raw dataset contains more than 3 billion observations corresponding to every transaction between the seller and the buyer. To make the dataset easily manageable, it is then collapsed at the year level, with a resulting 124 million observations. The main variable of interest is the total ammount of all transactions for each specific buyer-seller pair within a given year.

The main datasource is matched with firms' tax records from the Chilean IRS to obtain information on the number of employees, total sales, and the economic sector of each firm. Only firms with positive sales and at least one worker remain in the final sample. Additionally, the baseline sample includes only transactions between firms in the private non-agricultural sector.

For the empirical analysis, we define the following terms. I define links or

⁴This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC) in economic and financial affairs of its competence. The CBC has access to anonymized information from various public and private entities by virtue of collaboration agreements signed with these institutions. To secure the privacy of workers and firms, the CBC mandates that the development, extraction, and publication of the results should not allow the identification, directly or indirectly, of natural or legal persons. Officials of the CBC processed the disaggregated data. All the analysis was implemented by the authors and did not involve nor compromise the Chilean IRS. The information contained in the databases of the Chilean IRS is of a tax nature originating in self-declarations of taxpayers presented to the Service; therefore, the veracity of the data is not the responsibility of the Service.

relationships as unique buyer-seller pairs that trade continuously. Therefore, relationships that exhibit long periods of inactivity are reclassified as new links. The life or total duration of a link is defined as the maximum consecutive years of trade between the buyer-seller pair, and link's tenure or age as the current number of years the relationship has been active, with the first year of the relationship considered as age or tenure 1. Finally, we define link size or value as the total amount traded (in pesos) in a given year between the buyer-seller pair. In terms of notation, when discussing the link (i, j) , this refers to the link where firm i serves as the buyer and firm j as the seller, therefore link $(i, j) \neq (j, i)$.

II. B Empirical Findings

Motivating Fact: Long-term links tend to contribute disproportionately to network sales.

Figure 1 presents the distribution of the number of links in panel 1a and the distribution of link size in panel 1b. New links, defined as links with a tenure of 1 year, constitute 58% of the total number of relationships in 2021. In contrast, long-term links, defined as buyer-seller relationships that had been active for 5 or more years, account for only 15% of the total number of links in 2021, as shown in Figure 1a. The picture depicted in figure 1b shows an opposite pattern, with new links accounting for merely 12% of the total value, while long-term links represent 65% of the total network value in 2021.


Therefore, most links are new, while most trade occurs between established relationships. Figure 1 presents results for a specific year, but the validity of the conclusions extends beyond this particular year. Moreover, this observation remains even when excluding the retail and wholesale sectors.⁵

The fact that long-term relationships are relatively few but account for a disproportionate share of transaction value motivates examining the roles of relationship growth and selection. Therefore, I study how transaction values and exit probabilities vary over the relationship life cycle. These patterns help discipline the relationship-level heterogeneity introduced in the model.

To analyze the life-cycle patterns of business relationships, I estimate age profiles using two specifications. First, a cross-section specification, where year fixed effects are included to control for common shocks and economic growth. And second, a two-way fixed effect econometric specification, where buyer-year and seller-year fixed effects are included to additionally control for firm heterogene-

⁵For more details about robustness, see section II. C and appendix D.

Figure 1: Distribution of Links



Notes: This figure shows the distribution of active buyer-supplier relationships by tenure in 2021. Panel (a) reports each tenure category's share of the total number of active relationships. Panel (b) reports its share of total transaction value between firms. Tenure measures consecutive years of trade in the current relationship spell, with newly formed relationships assigned tenure one. The category "5+" includes relationships active for at least five consecutive years. Numbers within the bars indicate shares expressed as fractions of the respective totals.

ity. Equation (1) presents the estimated two-way fixed effect econometric model.⁶ In the equation, Y_{ijt} denotes the outcome variable of interest between buyer i and supplier j in year t and ε_{ijt} is the error term. Relationship age records the number of consecutive periods of trade in the current spell, with the first period assigned age one. The indicator $\mathbb{1}\{age_{ijt} = a\}$ equals one when the relationship has age a , and age one is the omitted category.

$$Y_{ijt} = \alpha + \sum_{a=2}^6 \beta_a \mathbb{1}\{age_{ijt} = a\} + \varphi_{it} + \vartheta_{jt} + \varepsilon_{ijt} \quad (1)$$

In addition, the two-way FE specification includes the buyer-year φ_{it} and seller-year ϑ_{jt} fixed effects, which captures common factors or shocks to each buyer and supplier. Despite the high dimensionality of the vectors of fixed effects, these parameters can be identified. Identification relies on firms maintaining multiple relationships with buyers and suppliers each year within a connected set. Moreover, because the fixed effects are time-varying, incumbent relationships also help with the identification of the parameters of the model. In this case, the largest connected set covers more than 99.8% of the observations of the yearly relationship panel.⁷

⁶The cross-section specification changes the buyer-year and seller-year fixed effects for a year fixed effect T_t .

⁷This number varies along the specification used because of sample restrictions related to the identification of age of the relationship. However, number for the smallest of the largest connected set is presented (99.8%).

Across both specifications, the parameter of interest is β_a . This measures the average outcome variable conditional on the age of the link being a relative to links of age 1. Identification of the age coefficients requires variation in relationship ages that is not absorbed by the fixed effects. Firms maintaining multiple trading relationships provide opportunities for such variation. In the cross-section specification, β_a measures the difference in the mean outcome between relationships of age a and newly formed relationships. In the fixed-effects specification, it measures the corresponding difference after accounting for buyer-year and supplier-year effects. These coefficients describe age profiles among observed relationships.

Using as baseline these econometric specifications, three empirical facts help with the understanding of selection and growth of business relationships. The dynamic exit of business link is analyzed first, followed by the life-cycle pattern of link's sales and its persistence.

Fact 1: *Negative association between the hazard rate and the age of the link.*

Figure 2 shows how the probability of relationship termination varies with link age. The dependent variable equals one if an active relationship between buyer i and supplier j in period t is inactive in period $t + 1$, and zero otherwise. Therefore, the econometric model presented in equation (1) is a linear probability model with this dependent variable. The coefficient β_a measures the difference in termination probability between links of age a and newly formed links, the omitted category, conditional on the fixed effects included in the estimation.


As mentioned before, an important issue is to identify correctly the age of the link. To measure relationship age, the estimation sample includes only relationship spells whose beginnings are observed from 2018 onward. Relationships already active in 2017 are excluded because their starting dates are unknown. Similarly, the final year of the sample is 2023 because of the definition of the dependent variable requires the identification of relationships termination.

The dashed line in Figure 2 reports an overall annual probability of link termination of 51%. Thus, relationship turnover is substantial. For context, this rate far exceeds the rates of job destruction (10-12%) and firm exit (20%), although these statistics describe different units and need not use identical samples or denominators.

The high destruction of links, however, underneath sizable heterogeneity by link's age. The results for the estimation of the probability of exit in Figure 2 clearly show a strong negative association, indicating that older links are less likely to be terminated. In fact, the probability of new links exiting is nearly 68%, while the

probability decreases monotonically to 25% for links of 5 years of activity.⁸

Figure 2: Hazard Rate



Notes: This figure shows annual relationship termination probabilities by link's tenure. Termination occurs when a buyer-supplier relationship active in year t is inactive in year $t + 1$. Tenure measures consecutive years of trade in the current relationship spell, with newly formed relationships assigned tenure one. The blue line reports estimates with year fixed effects; the red line reports estimates controlling for buyer-year and supplier-year fixed effects. The dashed horizontal line indicates the overall annual termination rate of 51%.

The declining age profile persists after including buyer-year and supplier-year fixed effects. Additive factors common to a buyer's relationships or a supplier's relationships in a given year therefore do not fully account for the pattern. This finding is consistent with selection on persistent relationship-specific characteristics, suggesting that relationships that survive to older ages may differ systematically from newly formed relationships in ways associated with subsequent survival.

To better understand the match-specific component associated with relationship survival, I next examine the value of transactions within each relationship.

Fact 2: *Positive association between relationship age and transaction value*

Figure 3 presents the life-cycle pattern of total sales of the relationship. The dependent variable is the log of total sales between buyer i and seller j in year t , and the plot shows coefficient β_a , which is the link size differential for links of various ages relative to a new link. The sample includes relationship spells beginning in


⁸This values are for the Cross-Section specification. The values for the two-way FE specification are 65% and 32%, respectively.

2018 or later, whose starting dates can be identified using the preceding transaction history.

The results in Figure 3 show a positive association between link's age and size, meaning that older relationships have significantly larger transacted value than newly formed relationships. In fact, the cross-section specification shows an estimated differential of approximately 1.8 log points at age five, traduced into roughly six times the transaction value of a new relationship. The age profile is increasing but becomes flatter at longer tenures, indicating smaller additional differences between successive age groups.

Although including buyer-year and supplier-year fixed effects reduces the estimated age-size gradient, particularly at longer tenures, remarkably the positive association remains. Thus, additive factors common to a buyer's relationships or a supplier's relationships in a given year account for part, but not all, of the differences in transaction value across relationship ages.

Figure 3: Tenure Profile



Notes: This figure shows differences in log annual transaction value by relationship tenure, relative to newly formed relationships (tenure one). The blue line reports estimates without fixed effects, the red line reports estimates controlling for buyer-year and supplier-year fixed effects. Tenure measures consecutive years of trade in the current relationship spell. Coefficients are expressed in log points, with the tenure-one coefficient normalized to zero. The profiles reflect both relationship growth and selection among surviving relationships.

Together with the declining termination probabilities documented above, this finding is consistent with selection on persistent bilateral characteristics associated

with both transaction value and survival. However, the age profile alone does not distinguish relationships that begin larger and survive longer from relationships whose transaction values increase over time. The next analysis examines these two possibilities.

Fact 3: *Long-lived relationships are larger at birth.*

The positive association between relationship age and transaction value may reflect both initial size differences and subsequent growth. To examine these sources, I estimate equation (2), which interacts relationship age with completed lifespan.⁹ The dependent variable is the log of total transacted value between buyer i and supplier j in year t ($\ln S_{ijt}$). The main difference with equation (1) is that here we have the interaction between the age of the link (age_{ijt}) and total observed duration of the link ($life_{ij}$). This interacted term aims to capture the life-cycle difference for links of different total duration, allowing a better examination of the long-term links.

The coefficients of interest (β_{al}) measure the difference in log transaction value between relationships of age a and lifespan l and relationships that terminate after their first year, the omitted category ($a = 1, l = 1$).¹⁰ Identification of the parameters relies on buyers' and sellers' maintaining multiple relationships of different total duration surviving and aging.

$$\ln S_{ijt} = \alpha + \sum_{l=1}^6 \sum_{a=1}^l \beta_{al} \mathbb{1}\{age_{ijt} = a\} \# \mathbb{1}\{life_{ij} = l\} + \varphi_{it} + \vartheta_{jt} + \varepsilon_{ijt} \quad (2)$$

Results are presented in Figure 4. Panel 4a presents the initial link size differential, comparing transaction values in the first year across relationships with different observed lifespans (β_{1l}). Relationships that subsequently survive longer are already larger in their first year. In particular, relationships lasting five years have an initial transaction value more than three times that of relationships lasting only one year. These initial differences indicate that the positive age-size gradient cannot be attributed entirely to growth after formation. They are consistent with persistent relationship-level heterogeneity associated with both initial transaction value and subsequent survival, suggesting a role for selection in shaping the age-


⁹Equation (2) presents the two-way fixed effects specification. The cross section specification changes the buyer-year φ_{it} and seller-year ϑ_{jt} fixed effects for a year FE T_t .

¹⁰By definition, β_{al} is not estimated for $a > l$, because age of the link cannot be bigger than the total duration of the link.

size profile.

Panel 4b complements this evidence by showing the profile for links of different lifespan. For each lifespan l , the plot shows the growth as links survive by comparing the coefficient relative to the first year of each lifespan.¹¹ This normalization removes initial size differences across lifespan groups and facilitates comparison of their subsequent trajectories. Results in 4b presents the result using the two-way FE specification, while ?? presents the results using year fixed effects only. The qualitative patterns are similar across specifications. Longer-lived relationships exhibit increases in transaction value over several years, without a clearly steeper growth profile. Among relationships whose termination is observed, transaction values decline in the final year of activity. Furthermore, the observation that links shrink in the last year of their lifespan extends beyond this dataset. For example, Monarch and Schmidt-Eisenlohr (2023) find similar patterns in US import data.

Figure 4: Link's Age-Life Profile



Notes: Panel (a) compares first-year transaction values across relationships with different observed lifespans, relative to relationships lasting one year. Estimates are shown for specifications with and without buyer-year and supplier-year fixed effects. Panel (b) shows transaction-value profiles by relationship lifespan, with each profile normalized to one in its first year. Tenure measures consecutive years of trade in the current relationship spell.

The positive association between link size and duration indicates that larger links tend to persist for a longer time. This implies that long-term links first-order stochastically dominate (FOSD) short-tenured links. Figure A2 presents the cumulative distribution for links of different age. As previously discussed, as links mature they tend to be larger because they are initially bigger (selection), and because they had more time to grow. Thus, the distribution of long-term links is shifted to larger values compared to the distribution of younger links.

¹¹The comparison with the initial year coefficient is $\exp(\beta_{al} - \beta_{1l})$.

The fact that initially bigger links tend to live longer suggests a relationship between link age and survival. This insight leads us to our third fact.

II. C Robustness

This section documents a series of robustness exercises designed to verify that the three main empirical facts of the paper—the disproportionate relevance of long-tenured links, the positive tenure-size relationship, and the declining hazard profile—are not artifacts of specific sample definitions, sectoral composition, or unbalanced panels. Each robustness exercise corresponds to a set of figures in the slide deck, and all conclusions below draw directly from those empirical patterns. Across all checks, the core facts remain remarkably stable, supporting the view that link dynamics reflect underlying economic fundamentals rather than mechanical features of the data.

II. C.1 Partial-Year Effects

A first concern is that some relationships may appear short-lived simply because buyers or sellers enter or exit the sample mid-year, creating artificial “partial-year” observations. To address this concern, we recompute all three main facts redefining the age of each link according to the month we observe the first transaction. Therefore, the first year of a buyer-seller link that started trade in March includes all transactions between March of that year and February of the next year.

This restriction ensures that each observed link has the same exposure window for survival, growth, and sales accumulation since birth. It is particularly important for interpreting hazard rates and tenure profiles, which could otherwise be biased if newly formed links are mechanically more likely to appear small or exit quickly because they had less time to trade in the calendar year.

Figures [A3–A5](#) show that the robustness-adjusted patterns closely mirror the baseline results. Long-tenured links continue to account for a disproportionately high share of network sales; the positive tenure-size gradient remains strong; and hazard rates continue to decline sharply with tenure. The similarity of these patterns confirms that partial-year censoring does not drive the main facts.

II. C.2 Balanced Panel of Firms

The baseline sample allows firms to enter and exit the dataset over 2017-2024. Although firm turnover is a natural feature of the economy, unbalancedness could mechanically affect the count of long-tenured links. To ensure that firm dynamics are not confounding link dynamics, we restrict the sample to a *balanced panel* of firms present in all years.

This check is relevant because long-tenured relationships necessarily involve stable firms; thus, the balanced panel isolates whether the main facts arise solely from the behavior of firms that are always active.

The results in Figures A6–A8 replicate the core patterns. Long-tenured links remain central to total sales, the age-size gradient persists, and survival hazards continue to decline monotonically. This suggests that link dynamics are not an artifact of firm entry or exit but reflect persistent patterns within ongoing relationships.

II. C.3 Exclusion of Retail and Wholesale Sectors

Retail and wholesale firms may engage in large-volume, low-duration transactions that differ structurally from production-oriented relationships. To ensure that these sectors do not disproportionately influence the main facts, we exclude all relationships involving firms classified in retail or wholesale industries.

This robustness check is crucial because these sectors exhibit high buyer-supplier churn by design, and their presence could inflate the turnover rate or alter tenure distributions.

Figures A9–A11 show that the three core empirical facts remain intact. Long-tenured links continue to dominate value shares, the size-tenure profile is largely unchanged, and hazard rates again fall sharply with tenure. This indicates that the patterns arise within the production structure itself and are not driven by retail-wholesale supply chains.

II. C.4 Exclusion of Capital-Intensive Sectors

Because capital-intensive sectors often have longer-lived physical production chains, one may worry that they mechanically generate longer link tenures or larger link sizes. To rule out this explanation, we re-estimate the main facts excluding all capital-intensive industries.

This exercise checks whether the core patterns are driven by sector-specific technological features rather than general link dynamics.

As shown in Figures [A12–A14](#), the main facts survive the exclusion. Long-tenured links remain disproportionately relevant, size rises with tenure, and hazard rates decline. Hence, sector-specific capital intensity cannot account for the empirical patterns.

II. C.5 Cross-Sectional Robustness

Finally, we evaluate whether the main patterns arise only in panel data or whether similar relationships are visible in a single cross section. We re-estimate the tenure-size and hazard-tenure patterns avoiding the use of the buyer-year and seller-year fixed effects, using cross-sectional variation to estimate the coefficients.

This check is especially important because cross-sectional variation is often used in structural estimation and because tenure proxies may be observed differently in cross sections.

Figures [A15–A17](#) confirm that the key facts emerge even in the cross-sectional data. Link size rises with tenure, and hazard patterns remain monotonic. This demonstrates that the empirical relationships are visible in both longitudinal and static data environments.

II. D Taking Stock

The empirical evidence presented so far show that long-term business links are economically relevant: there are relatively few long-term links, but they contribute disproportionately to the value of the network (Fact 1). The relevance stems from their significantly larger size; long-term links exceed the average new link by more than seven times (Fact 2). A significant portion of this size differential arises from the initial size of the links rather than differences in their growth trajectories. Moreover, long-term links demonstrate higher persistence, exhibiting significantly smaller exit probabilities (Fact 3).

What economic forces account for the relevance and dynamics of long-term links? In the next section, I outline a model that positions selection as the main driver of long-term link relevance. Then, I use the model to quantify the increased value of long-term links.

III Theoretical Framework

This section presents the model used to study the evolution of production networks and its aggregate implications. The model features a static equilibrium with dynamic states, allowing for computational tractability while closely following the evolution of individual business links. For ease of exposure we omit the time subindex from the baseline model.

III. A Environment

The economy is populated by a representative household and a unit continuum of monopolistically competitive firms. Firms are indexed by their productivity type $z \in \mathcal{Z}$, distributed according to the cumulative distribution function $G(z)$. Each firm produces a differentiated variety that is used both for final consumption and as an intermediate input by other firms.

Households. The representative household has CES preferences over the differentiated varieties supplied by firms:

$$W = \left(\int_{\mathcal{Z}} x(z)^{\frac{\sigma-1}{\sigma}} dG(z) \right)^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1, \quad (3)$$

where $x(z)$ denotes final consumption of the variety produced by a firm of type z , $\mathcal{Z}$ is the set of active firms in the economy, and σ is the elasticity of substitution across varieties. We assume that varieties are imperfect substitutes and therefore $\sigma > 1$.¹²

The household supplies an inelastic aggregate endowment of labor L , owns all firms, and receives their profits. Its budget constraint is

$$\int_{\mathcal{Z}} p(z)x(z) dG(z) \leq wL + \Pi, \quad (4)$$

where $p(z)$ is the price of variety z , w is the wage, and Π denotes aggregate profits.

Firms and Production. Each firm produces output using labor and a composite of intermediate inputs. The production technology of a firm with productivity z is a Cobb-Douglas production function with constant returns to scale

$$y(z) = \kappa z l(z)^{\alpha} v(z)^{1-\alpha}, \quad \alpha \in (0, 1), \quad \kappa > 0, \quad (5)$$

¹²Because preferences are homothetic, expenditure shares depend on relative prices but not on the household's income level.

where $l(z)$ denotes labor demand, $v(z)$ is the firm's composite intermediate input, κ is a normalization constant, and z is a Hicks-neutral productivity term.¹³ The parameter α is the labor cost share, while $1 - \alpha$ is the intermediate-input cost share.

To produce, firms combine labor with intermediate inputs purchased from other firms. The intermediate-input composite of a buyer of type z is a CES aggregate over the differentiated varieties supplied by its active trading partners:

$$v(z) = \left[\int_{\mathcal{Z}} \int_{r^*(z, z')}^{\infty} (r v(z, z'; r))^{\frac{\sigma-1}{\sigma}} dH(r) dG(z') \right]^{\frac{\sigma}{\sigma-1}}, \quad (6)$$

where r is the bilateral compatibility or quality of the relationship with stationary distribution H , $v(z, z'; r)$ is the quantity of the input purchased by customer z from supplier z' when the underlying link state is r , and $r^*(z, z')$ is the endogenous bilateral compatibility cutoff that determines whether a potential relationship between a buyer of type z and a supplier of type z' is active.

Variable r captures an idiosyncratic, match-specific quality or suitability state that governs the bilateral compatibility, such as technological fit, between a buyer and a potential supplier. A higher match-specific compatibility component $r(z, z')$ makes supplier z' input more effective in the production of buyer z . Thus, relationship heterogeneity affects both the intensive margin (the amount traded conditional on an active relationship) and the extensive margin (whether the relationship is formed). The threshold $r^*(z, z')$ makes the production network endogenous: firms purchase inputs only from the subset of potential suppliers for which link formation is privately profitable, as will be explained in detail below.

The double integral in (6) provides a type-level representation of aggregation across differentiated supplier varieties. The outer integral aggregates across supplier productivity types z' . Conditional on a pair of buyer and supplier types (z, z') , there is an atomless continuum of potential buyer-supplier matches, each characterized by one realized bilateral compatibility r , drawn independently from the distribution H . By the exact law of large numbers, the cross-sectional distribution of realized match qualities within each (z, z') cohort coincides with H . The inner integral aggregates across distinct supplier relationships with different realized values of r . By restricting the integration domain to $r \geq r^*(z, z')$, it includes only profitable relationships.

For tractability, the elasticity of substitution across intermediate inputs is as-

¹³The parameter κ is given by $\kappa = \alpha^{-\alpha} (1 - \alpha)^{-(1-\alpha)}$ and is used to simplify algebraic expressions in marginal costs.

sumed to equal the elasticity of substitution across final-good varieties, σ . This restriction delivers a common demand elasticity in final and intermediate markets and therefore a common monopolistic-competition markup. It can be relaxed by introducing separate elasticities for final consumption and intermediate production.

Market clearing. In this economy, labor can be used for production (L^Y) or to create and maintain the network of firms (L^O). Therefore, the labor market clearing condition is

$$L = L^Y + L^O \quad (7)$$

where $L^Y = \int l(z) dG(z)$ is total labor demanded by firms to produce each differentiated variety, and $L^O = \int \bar{f} \ell(z, z') dG(z) dG(z')$ is total labor used to maintain and create the production network. Importantly, we assume that the relationship cost is paid only by relationships that are available and active in the current period.

At the same time, a firm's output can be consumed by the representative household $x(z)$, or by other firms as intermediate input $v(z, z')$. Therefore,

$$y(z') = x(z') + \int v(z, z') dG(z) \quad \forall z' \in G(\cdot) \quad (8)$$

where $v(z, z') = \int_{r^*(z, z')}^{\infty} v(z, z'; r) dH(r)$ is the expected input purchased by customer z from supplier z' .

III. B Decentralized Equilibrium

We characterize the decentralized equilibrium in two steps. First, we solve for prices, quantities, and input choices conditional on a given production network. We then endogenize the network by allowing firms to choose which buyer-supplier relationships to activate. For simplicity, wage is assumed as the numeraire.

Equilibrium Conditional on the Network. Consider first a fixed production network. At this stage, the cutoff schedule $r^*(z, z')$, and hence the set of active relationships, is taken as given. A relationship operates only if it is selected by the firm for trade. We characterize its endogenous determination in the next subsection.

Given the wage w (numeraire) and the price charged by suppliers, the marginal cost of a firm of type z is

$$c(z) = \frac{\mathbb{P}(z)^{1-\alpha}}{z} \quad (9)$$

where the ideal price index for the firm's intermediate input composite is given by

$$\mathbb{P}(z) = \left(\int_{\mathcal{Z}} \int_{r^*(z,z')}^{\infty} p(z')^{1-\sigma} r^{\sigma-1} dH(r) dG(z') \right)^{1/(1-\sigma)} \quad (10)$$

Conditional on an active relationship between a buyer of type z and supplier of type z' , expenditure minimization yields the bilateral intermediate input demand

$$v(z, z'; r) = r^{\sigma-1} p(z')^{-\sigma} \mathbb{P}(z)^{\sigma-1} M(z) \quad (11)$$

where $M(z)$ denotes the firms total expenditure on intermediate inputs. A higher bilateral compatibility r both raises the effective quantity delivered per physical unit and, through optimal input substitution, increases the physical quantity purchased from that supplier. Importantly, an unavailable relationship has $v(z, z'; r) = 0$.

Because each firm faces a demand elasticity equal to σ , profit maximization under monopolistic competition implies the constant markup pricing rule

$$p(z) = \frac{\sigma}{\sigma-1} c(z) \quad (12)$$

The equilibrium conditional on the network can be found by computing two fixed points. First, the backward propagation fixed point allow us to solve for input prices and marginal costs, conditional on the network. Then, the forward propagation fixed point allow us to solve for sales, total expenditures, and bilateral trade, conditional on the network.

The backward propagation equation describes how input prices depends on direct and indirect suppliers in the network, and therefore how shocks propagate downstream from suppliers to buyers. Define the selected quality-adjusted sourcing capacity $q(z, z') = \int_{r^*(z,z')}^{\infty} r^{\sigma-1} dH(r)$ as the expected contribution of active states between the buyer z and the seller z' . Note that $q(z, z')$ summarizes everything about the distribution of the bilateral compatibility that is relevant for the CES price index, including information about the network of firms.¹⁴ Then, we can express the ideal price index of intermediate purchases as

$$\mathbb{P}(z) = \left(\int p(z')^{1-\sigma} q(z, z') dG(z') \right)^{1/(1-\sigma)}$$

¹⁴In addition, we assume that $q(z, z') < \infty$, $\forall (z, z')$.

Note that by replacing the pricing scheme from equation 12, and the marginal costs from equation 9 into the price index equation 10, we have that input price index for a firm depends on it's suppliers productivity, input price index and sourcing capacity.¹⁵

$$\mathbb{P}(z)^{1-\sigma} = \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} \int \mathbb{P}(z')^{(1-\alpha)(1-\sigma)} z'^{\sigma-1} q(z, z') dG(z') \quad (13)$$

The forward propagation equation describes how a firm's demand (from the final consumer and other firms as intermediate input) depends on demand from their customers, and therefore how shocks propagate upstream from buyers to suppliers. The forward equation is presented below,¹⁶ where $S(z)$ denote total sales of firm z , P is the final-consumption price index (previously defined), and E is consumer's total expenditure.

$$S(z) = \left(\frac{\sigma}{\sigma-1} \right)^{1-\sigma} z^{\sigma-1} \mathbb{P}(z)^{(1-\sigma)(1-\alpha)} \left(\frac{E}{P^{1-\sigma}} + \frac{1-\alpha}{\sigma/(\sigma-1)} \int S(z') \mathbb{P}(z')^{\sigma-1} q(z', z) dG(z') \right) \quad (14)$$

The remaining variable we need to pin down to solve the equilibrium is nominal income E , which is identified from the labor market clearing and aggregate labor resource constraint. The proof of existence and uniqueness of the equilibrium conditional on the network is in appendix ??.

Before the description of the endogenous network equilibrium, it's important to highlight a few definitions of network variables. In the model, let $\omega(z, z')$ be the CES expenditure share of firm z on the good sold by firm z' , as presented in equation 15. Let $\tilde{\Omega}$ be the cost-based input-output matrix, where each (z, z') cell measures the direct importance of supplier z' in buyer's z production costs, as presented in equation 16. Let $\tilde{\Psi} = (I - \tilde{\Omega})^{-1}$ be the cost-based Leontief inverse that captures the total direct and indirect cost exposure of firms in the network, and let $\tilde{\Lambda} = b^T \tilde{\Psi}$ be the cost-based Domar weight, where b captures the relevance of each variety for the final consumer.

$$\omega(z, z') = \frac{p(z')^{1-\sigma} q(z, z')}{\mathbb{P}(z)^{1-\sigma}} \quad (15)$$

$$\tilde{\Omega}(z, z') = \frac{\partial \log c(z)}{\partial \log p(z')} = (1-\alpha) \omega(z, z') \quad (16)$$

¹⁵Equation 13 resembles equation 6 from Bernard et al. (2022).

¹⁶The details about the derivation of the forward equation can be found in appendix ??.

Endogenous Network Equilibrium. Consider now the case of firms optimally selecting their trade partners. We assume that sellers must pay a fixed maintenance cost $\bar{f}$ to trade with each active buyer they have. Together with the CES structure of the model and assuming that goods are substitutes (i.e., $\sigma > 1$), sellers treat each customer independently, which helps with the tractability of the network.¹⁷

The seller decision to start or maintain a relationship with a buyer will depend on that specific link profitability and the fixed maintenance cost $\bar{f}$.¹⁸ Thus, sellers will only keep profitable relationships. Let $\pi(z, z', r)$ be the variable profits if firm z' sells to firm z when the match-state is r . Given the assumptions of the model, we can express the variable profits of an individual link as a σ -fraction of the bilateral sales of the link $m(z, z', r) = p(z')v(z, z', r)$. Therefore, the share of firms z' that sell their good to firms z is given by

$$\ell(z, z') = \int \mathbb{1} \left\{ \frac{m(z, z', r)}{\sigma} \geq \bar{f} \right\} dH(r) \quad (17)$$

Note that equation 17 implies that for each bilateral pair, there is a compatibility threshold $r^*(z, z')$, defined as the minimum underlying link state that makes the bilateral relationship active. Moreover, given the distribution of r , we can find the share of active trading pairs as follows

$$\ell(z, z') = 1 - H(r^*(z, z')) \quad (18)$$

We can identify the threshold $r^*(z, z')$ by replacing the expression of bilateral sales into equation 17.¹⁹ For any bilateral trade, the cutoff is given by

$$[r^*(z, z')]^{\sigma-1} = \sigma \bar{f} \frac{q(z, z')}{\omega(z, z')} \frac{\mu}{(1 - \alpha) \Lambda(z) E}. \quad (19)$$

Equation (19) expresses relationship selection as a trade-off between maintenance costs and the revenue generated by a match. Since a supplier earns variable profits equal to $1/\sigma$ of bilateral revenue, an active relationship must generate at least $\sigma \bar{f}$ in sales to cover its maintenance cost. Holding the remaining equilibrium objects fixed, a higher maintenance cost therefore raises the minimum match quality required for trade.

The term $q(z, z')/\omega(z, z')$ captures the supplier's price relative to the buyer's

¹⁷If buyers pay a fixed maintenance cost, even if firms are gross substitutes we would have complementarities in the decision to create links because links directly affect marginal costs.

¹⁸Note that because of love-of-variety effects, buyers will always prefer to have more suppliers.

¹⁹The details of the derivation of the decentralized cutoff are in Appendix A.

sourcing alternatives. Indeed, the definition of the sourcing share implies

$$\frac{q(z, z')}{\omega(z, z')} = \left(\frac{p(z')}{P(z)} \right)^{\sigma-1}.$$

A supplier whose price is high relative to the buyer's intermediate-input price index must offer higher match quality to generate sufficient sales. Conversely, a relatively inexpensive supplier can profitably serve the buyer at a lower compatibility threshold.

The remaining term is the inverse of the buyer's intermediate expenditure:

$$\frac{\mu}{(1 - \alpha)\Lambda(z)E} = \frac{1}{M(z)}.$$

The revenue-based Domar weight $\Lambda(z) = S(z)/E$ measures the buyer's sales relative to aggregate final expenditure, so $\Lambda(z)E$ recovers its total revenue. Dividing revenue by the markup μ gives variable production cost, of which the fraction $1 - \alpha$ is spent on intermediate inputs. Buyers with greater intermediate expenditure provide a larger market for each supplier, allowing relationships with lower match quality to cover the fixed maintenance cost.

The cutoff therefore combines a revenue requirement, relative supplier competitiveness, and buyer expenditure. These interpretations hold conditional on the remaining equilibrium objects. In general equilibrium, prices, sourcing shares, and buyer revenues adjust jointly with relationship selection; the expression consequently characterizes the equilibrium threshold rather than isolated comparative statics with respect to the structural parameters.

The equilibrium network is characterized as a fixed point in the cutoff schedule. Given a conjectured quality-adjusted network $q(z, z')$, the backward and forward equilibrium systems determine prices, sales, expenditure shares, cost-based network weights, and aggregate production scale. These objects imply an updated cutoff schedule through equation (19). The updated cutoffs, in turn, determine both the share of active relationships,

$$\ell(z, z') = 1 - H(r^*(z, z'))$$

and the quality-adjusted sourcing capacity,

$$q(z, z') = \int_{r^*(z, z')}^{\infty} r^{\sigma-1} dH(r)$$

An endogenous-network equilibrium is a fixed point at which the cutoff schedule, and hence the induced network objects, remain unchanged.

IV Model Estimation

We estimate the model to quantify the contribution of bilateral compatibility to relationship dynamics and to evaluate the aggregate consequences of endogenous network selection. The parameter vector is

$$\Theta = \left\{ \sigma, \alpha, L, \bar{f}, \rho_z, \sigma_z, \rho_r, \sigma_r \right\}, \quad (20)$$

where σ is the elasticity of substitution, α is the labor share in variable production costs, L is aggregate labor supply, and $\bar{f}$ is the labor cost of maintaining a buyer–supplier relationship. The remaining parameters characterize the stochastic processes for firm productivity, z , and bilateral compatibility, r .

We combine external calibration with simulated method of moments. Three parameters are calibrated outside the model: the elasticity of substitution σ , the labor cost share α , and aggregate labor supply L .

In the baseline specification, we set $\sigma = 3$ and examine robustness to alternative values. This value lies within the range commonly used in quantitative production-network models. Under CES demand and monopolistic competition, σ determines both the elasticity of firm demand and the equilibrium markup.²⁰

We set $\alpha = 0.24$ to match the aggregate labor share in variable production costs, with the remaining share, $1 - \alpha$, corresponding to intermediate-input expenditures. A lower value of α increases the importance of intermediate inputs in production and therefore strengthens the aggregate effects of supplier quality and production-network adjustment. Finally, we normalize aggregate labor supply to $L = 100$. This normalization fixes the scale of the economy without affecting relative prices or expenditure shares.

The estimation of the remaining parameters is explained below.

²⁰Thus, $\sigma = 3$ implies a markup of 1.5 over marginal cost and an operating-profit share of revenue equal to $1/\sigma = 0.33$, before accounting for relationship costs.

IV. A Structural Estimation

We estimate the parameters governing firm productivity using a regression, and the parameters governing the bilateral compatibility and relationship costs by simulated method of moments. The targeted moments summarize the cross-sectional and dynamic properties of buyer-supplier relationships.

We assume that firms' productivity follows a stationary AR(1) process:

$$\log z_{i,t} = \rho_z \log z_{i,t-1} + \varepsilon_{i,t}^z, \quad \varepsilon_{i,t}^z \sim \mathcal{N}(0, \sigma_z^2) \quad (21)$$

where ρ_z determines the persistence of firm productivity and σ_z determines the dispersion of its innovations. The stationary distribution induced by equation (21) corresponds to $G(z)$ in the theoretical model.

We estimate the firm-productivity process outside the simulated method of moments. Specifically, we estimate a regression of labor productivity on its lag, where we set ρ_z equal to the estimated coefficient, and σ_z equal to the standard deviation of the regression residual. Standard errors are clustered at the firm level to allow for serial correlation and heteroskedasticity within firms. This approach provides a transparent empirical discipline for the productivity process and reduces the dimensionality of the structural estimation. Its main limitation, however, is that labor productivity is not a direct measure of Hicks-neutral productivity: it also reflects input intensity, prices, demand, and endogenous supplier quality. Moreover, measurement error may attenuate persistence, while unobserved permanent firm heterogeneity may increase it.

We assume that the bilateral compatibility evolves according to

$$\log r_{ij,t} = \rho_r \log r_{ij,t-1} + \varepsilon_{ij,t}^r, \quad \varepsilon_{ij,t}^r \sim \mathcal{N}(0, \sigma_r^2) \quad (22)$$

where ρ_r governs the persistence of relationship fundamentals and σ_r governs the dispersion of the match-specific innovations. Its stationary distribution corresponds to $H(r)$. Innovations to firm productivity and bilateral compatibility are assumed to be independent across firms and relationships and mutually independent. This assumption isolates relationship-specific variation from persistent changes in buyer or supplier productivity.

We identify the parameters of the match-specific compatibility shifters using moments from bilateral expenditures. Conditional on an active relationship, firm-to-firm sales satisfy

$$m_{ij,t} = r_{ij,t}^{\sigma-1} p_{j,t}^{1-\sigma} \mathbb{P}_{i,t}^{\sigma-1} M_{i,t}$$

Taking logarithms gives

$$\log m_{ij,t} = \underbrace{\log M_{i,t} + (\sigma - 1) \log \mathbb{P}_{i,t}}_{\text{buyer-year component}} + \underbrace{(1 - \sigma) \log p_{j,t}}_{\text{supplier-year component}} + (\sigma - 1) \log r_{ij,t}.$$

This expression motivates a buyer-year and supplier-year fixed-effects decomposition

$$\log m_{ij,t} = \gamma_{i,t}^B + \gamma_{j,t}^S + u_{ij,t}, \quad (23)$$

where, according to the model,

$$u_{ij,t} = (\sigma - 1) \log r_{ij,t}. \quad (24)$$

Buyer-year fixed effects absorb changes in the buyer's intermediate expenditure and input price index, while supplier-year fixed effects absorb changes in the supplier's price.

Equation (23) is observed only for active relationships. Consequently, the empirical distribution of $u_{ij,t}$ is selected by the endogenous cutoff $r_{ij,t} \geq r_{ij,t}^*$. Rather than treating the fixed-effects residuals as an unselected sample from H , the estimation procedure reproduces this selection within the simulated model. This is important because relationships with low bilateral compatibility are systematically less likely to be observed and have shorter durations.

It is important to note the similarities between equation (23) and equation (1), the regression we use to describe the empirical patterns. The structural decomposition provides the formal economic mechanism behind the empirical patterns documented in equation (1). By projecting the bilateral expenditure residuals from equation (23) onto link age, we construct age-size gradients that isolate the evolution of the latent match state. According to the structural model, the positive association between link size and age observed in the transaction data is driven exclusively by the dynamic selection of the bilateral compatibility term, $r_{ij,t}$. Consequently, the empirical age-size gradients generated by equation (1) serve as ideal targeted moments in the SMM procedure, tightly identifying the persistence ρ_r and innovation dispersion σ_r of the underlying process.

Similarly, to construct the empirical hazard rate moments, we estimate the probability of relationship destruction conditional on link tenure. Through the lens of the structural model, a link is severed when the latent bilateral compatibility $r_{ij,t}$ falls below the endogenous survival threshold $r^*(z_{i,t}, z_{j,t})$. Because this compatibility fundamental follows a persistent AR(1) process, relationships that

survive for multiple periods are positively selected. This selection implies that older links face a systematically lower probability of receiving a sequence of negative shocks large enough to trigger relationship destruction. Therefore, matching the empirically observed decreasing age-hazard profile in the SMM procedure provides another source of variation for the identification of the persistence ρ_r and volatility σ_r parameters of the bilateral compatibility process.

Table 1 shows a summary of each parameter and the moment that helps with its identification. The moments play complementary identification roles. Average creation rate primarily disciplines $\bar{f}$ relative to the location of the bilateral compatibility distribution. The bilateral expenditure-age profile disciplines ρ_r , while the frequency with which relationships cross the endogenous cutoff discipline σ_r .

Table 1: Parameter Identification

Parameter	Identifying moment	Identification logic
ρ_r	Link size-age profile	Persistence of bilateral compatibility
σ_r	Hazard-age link profile	Selection
$\bar{f}$	Creation rate	Location of the profit condition

Notes: The table reports the primary source of identification for each parameter. Identification is joint because firm sales, bilateral expenditures, and network formation are simultaneously determined in equilibrium. Bilateral expenditure residuals are obtained after absorbing buyer-year and supplier-year fixed effects.

Let $\mathcal{M}^D$ denote the vector of empirical moments and $\mathcal{M}^S(\Theta)$ the corresponding moments generated by the model. The estimator solves

$$\hat{\Theta} = \arg \min_{\Theta} \left[\mathcal{M}^D - \mathcal{M}^S(\Theta) \right]' \mathcal{W} \left[\mathcal{M}^D - \mathcal{M}^S(\Theta) \right], \quad (25)$$

where $\mathcal{W}$ is a positive-definite weighting matrix. In the baseline specification, we use a diagonal matrix based on the inverse of the square of each moment, so that they are scaled.

For a candidate parameter vector, the model is simulated as follows. First, we discretize the stationary firm-productivity distribution G and approximate the bilateral compatibility distribution H using numerical quadrature. Second, given a conjectured cutoff schedule, we solve the backward fixed point for prices and marginal costs and the forward fixed point for sales and intermediate expenditures. Third, the resulting equilibrium objects determine an updated cutoff schedule. We iterate over these steps until the cutoff schedule and quality-adjusted sourcing capacities converge.

To construct the link-age moments, we simulate the processes in equations (21) and (22) for a large panel of potential buyer-supplier matches. A relationship is active whenever

$$r_{ij,t} \geq r^*(z_{i,t}, z_{j,t}). \quad (26)$$

The simulated panel is initialized from the stationary distributions and run for a burn-in period before calculating the model moments. This ensures that creation, destruction, and tenure moments are computed from the stationary distribution of firms and relationships rather than from arbitrary initial conditions.

IV. B Model Fit and Validation

We evaluate model fit using both targeted and untargeted moments. Targeted moments establish whether the model reproduces turnover and bilateral expenditure dynamics. Untargeted moments assess whether the estimated bilateral compatibility process reproduces richer features of the data, including the assortative matching by firm size, and the concentration of purchases across suppliers.

As an additional mechanism test, we compare the baseline model with an alternative specification in which relationship dynamics are generated by shocks to fixed relationship costs rather than bilateral compatibility. The comparison evaluates whether the compatibility model more accurately reproduces the observed monotonic relationship between bilateral expenditure, tenure, and subsequent link destruction.

Estimated Parameters

Table 2 reports the point estimates for the parameters governing the dynamic structural processes of the model. The estimated process for firm productivity has a persistence parameter of $\rho_z = 0.87$ and an innovation standard deviation of $\sigma_z = 0.57$. These values are tightly aligned with the standard firm dynamics literature, which typically finds annual persistence estimates between 0.70 and 0.85.²¹

Crucially, the estimated process for bilateral compatibility reveals that relationships are highly persistent rather than transient. We estimate a persistence of $\rho_r = 0.94$ and an innovation dispersion of $\sigma_r = 0.80$. The relatively high volatility of the match-specific innovations helps the model to generate the massive gross

²¹For example, De Loecker and Syverson (2021) note that typical persistence estimates for firm productivity fall within this range using annual data.

churn and high initial hazard rates observed in the transaction data. Finally, the fixed relationship maintenance cost is estimated to be $\bar{f} = 152.43$.

Table 2: Estimated Parameters for Productivity Processes and Network Costs

Parameter	Description	Estimate
<i>Firm Productivity Process</i>		
ρ_z	Persistence of firm productivity	0.87
σ_z	Dispersion of firm productivity innovations	0.57
<i>Bilateral Compatibility Process</i>		
ρ_r	Persistence of bilateral compatibility	0.94
σ_r	Dispersion of bilateral compatibility innovations	0.81
<i>Network Maintenance</i>		
$\bar{f}$	Fixed relationship maintenance cost	152.43


Notes: This table reports the point estimates for the structural parameters governing firm-level productivity and relationship dynamics. The firm productivity (z) and bilateral compatibility (r) states follow independent AR(1) processes in logs, governed by their respective persistence (ρ) and innovation dispersion (σ) parameters. The location parameter for bilateral compatibility (μ_r) is normalized to zero to identify the fixed relationship maintenance cost $\bar{f}$, which is expressed in units of labor.

Figure 5 presents the targeted empirical moments (red squares) alongside the simulated moments from the estimated model (blue circles). Overall, the model successfully captures the core features of buyer-supplier relationships, tightly replicating both the intensive-margin growth and extensive-margin selection patterns observed in the data.

The right panel demonstrates that the model tracks the empirical residual size gradient almost perfectly. Conditional on survival, both the simulated and empirical relationships grow monotonically, expanding by approximately 1.5 log points over their first five years relative to their initial size. This precise alignment confirms that the estimated AR(1) process for bilateral compatibility accurately translates into the observed expenditure growth of maturing links.

The left panel shows that the model successfully generates the steeply declining hazard rate by link age, which is the central signature of our dynamic selection mechanism. As relationships mature, the surviving links are positively selected, they represent draws situated securely above the continuation threshold, making them significantly less likely to be severed by subsequent negative shocks. Although the model slightly underpredicts the initial destruction rate at age 1, and overpredicts the survival trajectory for established relationships.

Figure 5: Model Fit



Notes: This figure evaluates the fit of the Simulated Method of Moments (SMM) estimation by comparing targeted empirical moments (red squares) against the corresponding moments generated by the simulated model (blue circles). The left panel plots the hazard rate as a function of link age. The right panel plots the residual size gradient, defined as the average log-point difference in bilateral transaction volume relative to a relationship at age one once we control for buyer and seller unobserved characteristics. The alignment between the model and the data demonstrates that the estimated bilateral compatibility process successfully replicates both the declining exit probability and the monotonic size growth characteristic of surviving relationships.

V Welfare and the Propagation of Shocks

Having characterized the decentralized equilibrium, we now examine how firm-level shocks propagate through the production network and affect aggregate welfare. Because household preferences are homothetic and labor supply is inelastic, welfare is equal to real aggregate expenditure:²²

$$W = \frac{E}{P} \quad (27)$$

where E denotes nominal household expenditure and P is the final-consumption price index.

We will examine the economy's response to an exogenous shock to a subset of firms productivity z . Endogenous network adjustment introduces a resource tradeoff. Activating additional relationships expands firms' access to productive inputs, but also requires labor for relationship maintenance, reducing the labor available for variable production. Proposition 1 characterizes the first-order welfare effect of firm-level productivity shocks.

²²Since final consumption is the only final use in this closed economy, welfare also coincides with the (model-consistent) real GDP.

Proposition 1: *With wages as the numeraire, changes in welfare are given by*

$$d \log W = d \log L^Y + \tilde{\Lambda}^T \left[d \log z + \frac{1 - \alpha}{\sigma - 1} d \log Q \right]$$

where $d \log Q$ is the change in quality-adjusted sourcing capacity.

This result summarizes the key channels through which endogenous network formation affect welfare in the economy. When the network is fixed, only the second channel is present. Then, the effect on welfare can be summarized by the cost-based Domar weight, and there are no changes in the macro-allocation of labor, as in Baqaee and Farhi (2019).

However, when the network adjusts endogenously, productivity shocks also change relationship cutoffs. The resulting creation and destruction of links affect welfare through both the resources used in network maintenance, $d \log L^Y$, and firms' quality-adjusted access to suppliers, $d \log Q$.

The effect of the quality-adjusted network change is amplified by technology assumptions and the network. First, the quality-adjusted change is propagated directly and indirectly across the supply chain until it reaches the final consumer, which is captured by the cost-based Domar weight $\tilde{\Lambda}$. Second, the change in sourcing capacity depends on the share that firms spend in intermediate inputs $(1 - \alpha)$, and the marginal increase it generates in marginal costs $1/(\sigma - 1)$.²³

Finally, the change in the quality-adjusted sourcing capacity $d \log Q(z)$ is presented in equation 28.²⁴ The firm-level change in the sourcing capacity depends on how relevant each supplier z' is for the production process of buyer z , which is captured by the CES shares $\omega(z, z')$. The term $d \log q(z, z')$ measures the change in the bilateral sourcing capacity, and is decomposed in Lemma 1.

$$d \log Q(z) = \int \omega(z, z') d \log q(z, z') dG(z') \quad (28)$$

Lemma 1: *Define the share of active relationships between buyers of type z and suppliers of type z' as equation 18 and the average productivity contribution of an active relationship as*

$$\bar{r}(z, z') \equiv \mathbb{E} \left[r^{\sigma-1} \mid r \geq r^*(z, z') \right]. \quad (29)$$

²³Note that the increase in marginal cost takes the form of love-of-variety effect on sourcing capacity.

²⁴Note that Proposition 1 presents $d \log Q$, while here we are presenting $d \log Q(z)$. The difference is that $d \log Q$ is a vector containing all the changes in firm's sourcing capacity $d \log Q(z)$.

Then, quality-adjusted sourcing capacity satisfies

$$q(z, z') = \ell(z, z') \bar{r}(z, z'), \quad (30)$$

and its change admits the exact decomposition

$$d \log q(z, z') = \underbrace{d \log \ell(z, z')}_{\text{link expansion}} + \underbrace{d \log \bar{r}(z, z')}_{\text{match-quality composition}}. \quad (31)$$

Lemma 1 decompose changes in sourcing capacity into an extensive margin and a composition margin. The first term captures changes in the mass of active relationships, whether is more or less likely for any bilateral pair to trade given the new thresholds. The second term captures changes in the average productivity of those relationships.

Importantly, these two forces opposes each other. For example, when the cut-off falls, more relationships become active, so $d \log \ell(z, z') > 0$. However, the newly activated relationships have lower match productivity than existing relationships, reducing average active-link quality and implying $d \log \bar{r}(z, z') < 0$. The overall effect on $q(z, z')$ is nevertheless positive because each additional relationship makes a positive contribution to sourcing capacity.

The overall effect of endogenous network formation is, therefore, attenuated by the composition effect. The newly formed relationships are only marginal, decreasing the average quality of existing links. This new mechanism is key to understand the relevance of endogenous network formation for welfare.

V. A Quantifying Endogenous Network Adjustment

We quantify the aggregate relevance of endogenous network adjustment using a productivity counterfactual around the calibrated stationary equilibrium. Holding the population of firms and the distribution of relationship fundamentals fixed, we simulate a large, targeted negative shock by reducing the fundamental productivity z of the top 5% of firms by 0.2 log points. The exercise analyzes the welfare impact in steady state, where the baseline and resulting counterfactual are weighted by the baseline measure $dG(z)$.

We evaluate the role of endogenous network formation by comparing the economy before and after a negative firm-productivity shock. We consider two post-shock equilibria. In the fixed-network counterfactual, the set of active relationships

remains at its pre-shock configuration, while prices, sales, and aggregate expenditure adjust to the new productivity levels. In the endogenous-network counterfactual, firms can additionally reselect their trading relationships. Both counterfactuals therefore incorporate general-equilibrium adjustments; their difference isolates the additional welfare effect of network reselection. The exercise compares stationary allocations rather than tracing the transition between them.

Let W_0 denote baseline welfare and let W_1^F and W_1^E denote counterfactual welfare under fixed and endogenous cutoffs, respectively. In the fixed-network counterfactual, bilateral link probabilities and quality-adjusted sourcing capacities remain at their baseline values, while prices, expenditures, and production adjust to the productivity shock. In the endogenous-network counterfactual, relationship cutoffs also adjust, changing both the mass and the composition of active relationships. The aggregate contribution of endogenous network adjustment is

$$\mathcal{N} \equiv \log \left(\frac{W_1^E}{W_0} \right) - \log \left(\frac{W_1^F}{W_0} \right) \quad (32)$$

A positive value indicates that adjustment mitigates the welfare loss, while a negative value indicates amplification of the shock. Welfare in each equilibrium is computed directly as $W = E/P$. We express welfare changes in consumption-equivalent percentage terms as $100[\exp(\Delta \log U) - 1]$.

Results are presented in table 3. Relative to the pre-shock baseline, Panel A shows that welfare falls by 42.025% when the network is fixed and by 41.999% when relationships adjust endogenously. Network reselection therefore reduces the consumption-equivalent welfare loss by approximately 0.025 percentage points, measured relative to baseline welfare.

To understand the limited contribution of network adjustment, Panel B decomposes the total welfare change in the endogenous network counterfactual using Proposition 1 and Lemma 1. Proposition 1 separates the productivity, production-labor, and sourcing-capacity channels, while Lemma 1 divides the sourcing-capacity channel into link expansion and match-quality composition. We implement this decomposition by constructing a sequence of equilibria in which the magnitude of the productivity shock increases gradually from zero to its full value, always affecting the same set of firms. At each intermediate equilibrium, we recompute the endogenous objects and aggregation weights. We then numerically integrate the channel contributions between consecutive equilibria. This sequence is a computational device for evaluating the finite welfare change, rather than a transition

path.

Table 3: Welfare Comparison: Top 5% Firm Productivity Shock

	CE contribution (p.p.)
A. Welfare comparison	
Fixed network	−42.025
Endogenous network	−41.999
Gain from network adjustment	+0.025
Panel B. Endogenous-network decomposition	
Productivity	−42.018
Link expansion	+2.326
Match-quality composition	−2.271
Production labor	−0.037
Total	−41.999

Notes: This table compares the baseline equilibrium with counterfactual equilibria following a productivity shock of $\Delta \log z = -0.2$, equivalent to an 18.13% decline, applied to firms with the highest baseline productivity. The requested top 5% represents 5.0529% of firm mass on the numerical grid. In the fixed-network counterfactual, bilateral link probabilities and quality-adjusted sourcing capacities remain at baseline values, while prices and quantities adjust. The endogenous-network counterfactual also allows relationship cutoffs to adjust. All welfare entries use baseline welfare as the denominator. Panel B proportionally allocates the total consumption-equivalent welfare change across the integrated log-welfare channels; these contributions are not separate counterfactual welfare effects. Displayed sums may differ because of rounding and numerical integration.

The integrated contributions are initially expressed in log-welfare units. To report an additive decomposition in consumption-equivalent percentage points, we multiply each contribution by the common factor $100[\exp(\Delta \log W) - 1]/\Delta \log W$. The resulting contributions sum to the endogenous-network welfare change, up to the reported numerical integration residual.²⁵


Results presented in panel B helps understand the magnitude of the forces of endogenous network formation. First, using the integration procedure yields a productivity channel very close in magnitude to the fixed network channel of Panel A. This result is expected because the productivity contribution closely approximates the fixed-network welfare loss, as previously discussed.

Second, Panel B shows a significant contribution of link expansion of 2.326 percentages points of consumption equivalent welfare. This is accompanied by an increase in the aggregate probability of an active relationship, that rises from approximately 0.1704% to 0.2015%, indicating that the network becomes denser despite the adverse productivity shock. This is consistent with substitution toward suppliers whose relative attractiveness improves when the targeted firms become

²⁵The integral residual of Panel B is not presented in the table. The residual is negligible, it is 3.16×10^{-7} .

less productive. This can be clearly appreciated in Figure 6, which presents the change in the bilateral probability of trade. In blue are pairs who reduce their probability of trade with each other, while in red are the bilateral pairs that increase their probability of trade. Given the lognormal weights in the grid of firm productivity, the sum of the red areas more than compensate the decrease of link probability associated to the shock, which explains the increase in the aggregate probability of trade and supports the positive welfare contribution of this mechanism.

Figure 6: Heatmap Change in Link Probability



Notes: Each cell reports the change in the probability of an active relationship between a seller type (vertical axis) and a buyer type (horizontal axis), comparing the endogenous-network equilibrium after the productivity shock with the baseline. Productivity increases from left to right for buyers and from bottom to top for sellers. The shock reduces log productivity by 0.2 among the highest-productivity firms, representing 5.0529% of baseline firm mass. Type positions retain their baseline ordering. Red indicates an increase in link probability, blue a decrease, and white approximately no change. The color scale reports probability changes. Cells have equal visual size but represent different population masses. The change in aggregate network density therefore weights each cell by the population shares of its buyer and seller types.

Third, the match-quality composition channel reduces welfare by -2.271 CE percentage points, offsetting approximately 97.6% of the link-expansion contribution leaving a small net sourcing contribution. The reason is because marginal links that enter the sourcing capacity are of lower quality. Holding the match-quality distribution fixed, a lower relationship cutoff admits matches with lower productive intensity than the average incumbent relationship. Consequently, expanding the mass of active relationships reduces their average productive intensity. Quality-adjusted sourcing capacity still increases when a cutoff falls, but by less than the increase in relationship mass. After accounting for composition, the

net sourcing-capacity contribution is approximately 0.055 percentage points.

Finally, production labor provides an additional offset. Maintenance labor rises by approximately 18.29% relative to the baseline, reducing the labor available for production because total labor supply is fixed. This contributes -0.037 consumption-equivalent percentage points to welfare.

Overall, the results summarized in table 3 indicate that firms can smooth the negative productivity shock by adjusting their network. However, the total contribution of the new connections is almost negligible because of composition effects.

VI Efficiency of Decentralized Equilibrium

Having analyzed how endogenous network formation shapes the network and the propagation of shocks, we now turn to understand efficiency. Consider the case of a social planner deciding the production network and the allocation of resources in the economy. The planner has two distortions to consider. First, firms charge a constant markup over marginal costs that distort pricing and quantity decisions. This distortion adds up along the production network, which further distort firms' decision from social optimum given the double marginalization problem. Second, firms private decision to form connections may not consider all downstream effects, what is known as *network externality*.

To understand why the network externality emerges and quantify how efficient the decentralized equilibrium is, we turn to analyze the problem of a constrained planner. The idea is to isolate the effect of markups and double marginalization, from the social optimal decisions in terms of network formation.

VI. A Constrained Planner Problem

To characterize the efficiency properties of the decentralized equilibrium, we benchmark the market allocation against the solution of a constrained social planner. As in the decentralized market, the planner faces the fundamental technological trade-off between the quality-adjusted sourcing capacity of the production network and the fixed labor resources required to maintain it.

We can find an elegant and simple formulation of the constrained planner problem. Because every active link consumes identical maintenance labor $\bar{f}$, and because the match quality r enters the objective function monotonically, the planner's optimal policy takes the form of a threshold rule. Specifically, the plan-

ner chooses a pair-specific compatibility cutoff $r^{sp}(z, z')$ for every potential buyer-supplier pair. These cutoffs completely determine the aggregate state of the network, including the bilateral link probabilities $\ell(r^{sp}(z, z'))$, and the bilateral quality-adjusted sourcing capacities $q(r^{sp}(z, z'))$.

More precisely, the constrained planner's objective is to choose the matrix of cutoffs $R = [r^{sp}(z, z')]$ to maximize aggregate welfare, formally defined as the log difference between aggregate expenditure $E(R)$ and the aggregate price index $P(R)$. The planner controls relationship selection but takes the within-network market allocation mechanism as given. In particular, firms continue to minimize production costs and charge the constant markup $\mu = \sigma/(\sigma - 1)$. Prices and expenditures therefore adjust endogenously to the selected network; it is the pricing rule, rather than the equilibrium prices, that the planner takes as given.

Definition 1: The Constrained Planner's Problem. *Consider a benevolent social planner who observes the fundamental productivity draws z and match-specific compatibility draws r for all potential firm pairs. The constrained planner cannot overcome the technological fixed costs of relationship maintenance nor alter the CES aggregation technology, but can directly dictate which bilateral relationships are formed. The planner chooses a matrix of pair-specific compatibility cutoffs $R = [r^{sp}(z, z')]$ to maximize aggregate welfare, defined as real aggregate consumption, subject to the economy's aggregate labor resource constraint. Formally, the planner solves:*

$$\max_R \log E(R) - \log P(R) \quad (33)$$

subject to the aggregate labor market clearing condition:

$$L = L^Y(R) + L^O(R), \quad (34)$$

where $L^O(R) = \bar{f} \iint \ell(r^{sp}(z, z')) dG(z) dG(z')$ is the total labor allocated to network maintenance, $\ell(r^{sp}(z, z')) = 1 - H(r^{sp}(z, z'))$ is the probability of an active link, and aggregate labor is fixed. The aggregate price index $P(R)$ depends on the network topology through the pair-specific quality-adjusted sourcing capacity $q(r^{sp}(z, z')) = \int_{r^{sp}}^{\infty} r^{\sigma-1} dH(r)$.

This formalization precisely isolates the fundamental macroeconomic trade-off governing the extensive margin of the production network. Because the aggregate labor endowment L is fixed, any expansion of the network—achieved by lowering the pair-specific threshold $r^{sp}(z, z')$ to increase the probability of a match—mechanically drains labor resources away from physical production L^Y and diverts them into relationship maintenance L^O . This reduction in production la-

bor directly depresses aggregate expenditure $E(R)$. However, lowering the cutoff simultaneously expands the available quality-adjusted sourcing capacity q , which lowers marginal costs and ultimately reduces the aggregate price index $P(R)$. The planner's optimal cutoff matrix R^* must perfectly balance this marginal aggregate resource cost against the marginal macroeconomic price reduction, successfully internalizing the full general equilibrium propagation of every localized network adjustment.

Crucially, this formulation represents a *constrained* (or second-best) social planner because the planner's policy instrument is restricted strictly to the extensive margin of network formation, dictated by the cutoff matrix R . The planner is not permitted to choose individual transaction quantities directly or alter the intermediate input prices. Instead, the planner must evaluate aggregate welfare taking the decentralized market's pricing mechanism as given. This constraint is mathematically embedded within the functional form of the aggregate price index $P(R)$, which aggregates intermediate prices that retain the standard CES monopolistic markup, $\sigma/(\sigma - 1)$.

In contrast, a fully unconstrained, first-best planner would jointly optimize both the extensive margin of link formation and the intensive margin of input quantities, effectively eliminating the monopolistic markups to ensure that all intermediate inputs are priced exactly at their true social marginal cost. By forcing the planner to respect the decentralized pricing rule, the constrained problem cleanly isolates the specific inefficiency originating from the endogenous network formation process. It reveals how the planner optimally adjusts the overall network topology to partially compensate for the uncorrectable intensive-margin distortions propagating downstream.

To derive the social value of a link and formally establish the planner's optimality condition, we evaluate the marginal effect of expanding the measure of active relationships between a specific buyer type z and supplier type z' . Given the monotonic relationship between the cutoffs $r^{sp}(z, z')$ and the link probability $\ell(z, z')$, we interchangeably treat the planner as if choosing the threshold or the link probability.

It is clear from equation 33 that the derivative of aggregate welfare with respect to the link probability $\ell(z, z')$ separates analytically into a marginal resource cost and a marginal productivity benefit, as can be seen in equation 35.

$$\frac{\partial \log W}{\partial \ell(z, z')} = \underbrace{\frac{\partial \log E}{\partial \ell(z, z')}}_{\text{Maintenance resources}} - \underbrace{\frac{\partial \log P}{\partial \ell(z, z')}}_{\text{Network productivity benefit}} = 0 \quad (35)$$

On the expenditure side, expanding the network requires increasing the aggregate labor allocated to maintenance, L^O . Because aggregate labor is fixed, this mechanically crowds out production labor, yielding $E \propto L^Y = \bar{L} - L^O$. Taking the derivative of log expenditure with respect to the link probability $\ell(z, z')$ implies a marginal aggregate resource cost of $-\bar{f}/L^Y$. For the network allocation to be optimal, this resource drain must be perfectly offset by the marginal network productivity benefit, which operates entirely through the reduction of the aggregate price index, defined as $-\partial \log P / \partial \ell(z, z')$.

To quantify this network productivity benefit, we must trace the impact of the marginal link through the economy's CES input architecture. We can easily obtain this derivative using the results from Proposition 1.²⁶ The benefit of a marginal increase in links probability is given by

$$\frac{\partial \log P}{\partial \ell(z, z')} = \tilde{\Lambda}(z) \tilde{\Omega}(z, z') \frac{1}{1 - \sigma} \frac{r^{sp}(z, z')^{\sigma-1}}{q(z, z')} \quad (36)$$

The intuition is as follows. When the planner marginally lowers the compatibility threshold $r^{sp}(z, z')$ to increase the probability of an active link, the newly activated relationships inherently possess exactly the marginal match quality, $r^{sp}(z, z')$. Based on the properties of the CES aggregator, the absolute contribution of this marginal link to the bilateral sourcing capacity is $(r^{sp}(z, z'))^{\sigma-1}$. Expressed in proportional terms, the change in the log bilateral capacity with respect to the extensive margin expansion is $(r^{sp}(z, z'))^{\sigma-1} / q^{sp}(z, z')$. This ratio cleanly captures the pure extensive-margin quality addition relative to the buyer's existing pool of active suppliers.

Finally, we map this localized capacity expansion into the aggregate price index to close the optimality condition. The elasticity of the aggregate price index with respect to a buyer's fundamental cost structure is governed by its cost-based Domar weight, $\tilde{\Lambda}^{sp}(z)$, which summarizes the general equilibrium propagation of cost reductions through all direct and indirect downstream linkages. The buyer's total capacity scales with the specific supplier's expenditure share $\omega^{sp}(z, z')$, while

²⁶In particular, Proposition 1 has results of the total differential of the aggregate price index to the economic fundamentals of the model. Hence, we can use these results to obtain the partial derivative considering all direct and indirect effects of the network.

the resulting reduction in the buyer's marginal cost is governed by the intermediate input intensity $(1 - \alpha)$.²⁷ Importantly, each new supplier, or the marginal increase in the probability of trade, is treated as a new variety in the CES input aggregator. Therefore, the marginal gain for the buyer is valued at the CES love-of-variety effect $1/(\sigma - 1)$.

Multiplying these macroeconomic pass-through terms by the marginal link's local capacity contribution yields the exact network productivity benefit. Equating this benefit to the marginal resource cost $\bar{f}/L^Y$ produces the planner's first-order condition, proving that the optimal network expands exactly to the point where the labor cost of maintaining the marginal link perfectly equals its fully propagated, quality-adjusted contribution to aggregate cost reduction. Therefore, the optimality condition is

$$\frac{\bar{f}}{L^{Y,sp}} = \underbrace{\frac{1}{\sigma - 1}}_{\text{Elasticity of Varieties}} \underbrace{(1 - \alpha) \omega^{sp}(z, z') \frac{(r^{sp}(z, z'))^{\sigma-1}}{q^{sp}(z, z')}}_{\text{Marginal contribution of the link}} \underbrace{\tilde{\Lambda}^{sp}(z)}_{\text{Domar weight}} \quad (37)$$

To examine the determinants of the network externality, I compare the threshold from the constrained planner relative to the decentralized equilibrium. Rearranging the terms from equation 37, we get that the following expression for the threshold of the planner $r^{sp}(z, z')$

$$r^{sp}(z, z') = \left(\frac{q^{sp}(z, z')}{\omega^{sp}(z, z')} \right)^{\frac{1}{\sigma-1}} \left(\frac{\sigma - 1}{1 - \alpha} \frac{\bar{f}}{\tilde{\Lambda}^{sp}(z) L^{Y,sp}} \right)^{\frac{1}{\sigma-1}} \quad (38)$$

Proposition 2. *Let $r^*(z, z')$ denote the bilateral compatibility cutoff in the decentralized equilibrium, and let $r^{sp}(z, z')$ denote the corresponding optimal cutoff chosen by the constrained social planner. The relative inefficiency in decentralized network formation is strictly characterized by the ratio of these thresholds, which admits the following exact decomposition into three distinct wedges:*

$$\frac{r^*(z, z')}{r^{sp}(z, z')} = \underbrace{\left(\frac{\sigma}{\sigma - 1} \right)^{\frac{1}{\sigma-1}}}_{\text{Appropriation Wedge}} \underbrace{\left(\frac{(1 - \alpha) \tilde{\Lambda}^{sp}(z) L^{Y,sp}}{M^*(z)} \right)^{\frac{1}{\sigma-1}}}_{\text{Social Scale and Private Expenditure}} \underbrace{\left(\frac{q^*(z, z')/\omega^*(z, z')}{q^{sp}(z, z')/\omega^{sp}(z, z')} \right)^{\frac{1}{\sigma-1}}}_{\text{GE Relative Price Adjustment}} \quad (39)$$

Proposition 2 analytically separates the sources of extensive margin inefficiency into three channels. The first term, the Appropriation Wedge, is strictly

²⁷Remember that we defined $\tilde{\Omega}(z, z') = (1 - \alpha)\omega(z, z')$.

greater than one and originates directly from the divergence between the private and social valuation of relationship formation. In the decentralized market, a supplier captures exactly a fraction $1/\sigma$ of the bilateral sales as variable profit, which dictates their private willingness to form a link. In contrast, the constrained social planner evaluates the link based on its marginal contribution to downstream productivity, a benefit governed strictly by the CES love-of-variety elasticity $1/(\sigma - 1)$.

Crucially, while the ratio of the planner's valuation to the private valuation mathematically yields $\frac{1/(\sigma-1)}{1/\sigma} = \frac{\sigma}{\sigma-1}$, which is identical to the standard CES markup, the underlying economic friction here is distinct from intensive-margin pricing distortions. The wedge arises because the private supplier's appropriated share of the match surplus ($1/\sigma$) is systematically lower than the link's true macroeconomic value as a marginal cost reducer ($1/(\sigma - 1)$). Because the supplier cannot privately appropriate the full social value of the connection, the decentralized market requires a fundamentally different compatibility draw to justify the fixed maintenance cost $\bar{f}$, originating an inefficiency that propagates through the network.

The second term compare the scale of the firm between the decentralized and constrained planner equilibria. When link formation is decentralized, its decision depends on the buyer's intermediate expenditure $M^*(z)$, which determines the revenue available to its suppliers. The planner instead values the resulting cost reduction using the buyer's cost-based Domar weight $\tilde{\Lambda}^{sp}(z)$, its intermediate-input intensity $1 - \alpha$, and aggregate production labor $L^{Y,sp}$. The Domar weight captures how a reduction in the buyer's production cost propagates to final consumption, while production labor determines the welfare cost of diverting resources into relationship maintenance. A larger social scale relative to private expenditure therefore increases the relative cutoff, holding the remaining terms fixed. This factor incorporates both the distinction between private spending and social propagation and the changes in centrality and production labor across allocations.

The third factor captures the adjustment in relative prices induced by the change in the network. By the definition of the sourcing share, $q(z, z')/\omega(z, z') = [p(z')/P(z)]^{\sigma-1}$. Consequently, after taking the root, this factor equals the supplier's price relative to the buyer's input price index in the decentralized allocation, divided by the same relative price in the planner's allocation. A supplier that is relatively more expensive in the decentralized allocation requires a higher match quality to sustain a relationship, contributing to a higher decentralized cutoff. Network selection and match-quality composition affect this factor through equilibrium prices; it is not a separate measure of average active match quality.

The decomposition is an exact accounting identity between the two alloca-

tions. Its factors are jointly determined in equilibrium, so their individual interpretations hold conditional on the others. In particular, the appropriation wedge being greater than one does not alone establish that the decentralized network underprovides relationships for every type pair.

To have a better understanding of the sources of inefficiencies in the scale of the firm, Proposition 3 further decompose the Social Scale and Private Expenditure component.

Proposition 3. *Let $*$ denote the decentralized equilibrium and sp the constrained-planner allocation. Suppose wages are normalized to one, the markup μ and production parameter α are common across firms and allocations, and the relevant expenditures, production labor, and Domar weights are positive. Define the revenue-based Domar weight as $\Lambda^*(z) = S^*(z)/E^*$, and let $\tilde{\Lambda}(z)$ denote the cost-based Domar weight from Proposition 1. The social-scale component of Proposition 2 admits the exact decomposition*

$$\begin{aligned} \frac{(1-\alpha)\tilde{\Lambda}^{sp}(z)L^{Y,sp}}{M^*(z)} &= \underbrace{\frac{\mu\alpha}{\mu-1+\alpha} \frac{\tilde{\Lambda}^*(z)}{\Lambda^*(z)}}_{\text{Markup-induced centrality wedge}} \\ &\times \underbrace{\frac{\tilde{\Lambda}^{sp}(z)}{\tilde{\Lambda}^*(z)}}_{\text{Centrality adjustment}} \underbrace{\frac{L^{Y,sp}}{L^{Y*}}}_{\text{Production-labor adjustment}}. \end{aligned} \quad (40)$$

The first component captures the markup-induced divergence between revenue-based and cost-based centrality, evaluated at the decentralized equilibrium. Private relationship profitability depends on buyer expenditure and therefore on revenue-based firm size. The welfare effect of a sourcing improvement instead depends on how production-cost reductions propagate through the network, as summarized by the cost-based Domar weight. Markups separate these two measures. The factor $\mu\alpha/(\mu-1+\alpha)$ adjusts for the different aggregate normalizations of the two weights.²⁸

Therefore, this component measures whether a buyer's cost-propagation importance is large or small relative to the expenditure scale supporting private relationship formation. In the unit-markup benchmark, the revenue- and cost-based propagation systems coincide and this component equals one. With positive markups, its magnitude depends on the firm's position in the production network.

The second component measures the distance in firms' cost-based Domar weight

²⁸Intermediate expenditure accounts for the fraction $(1-\alpha)/\mu$ of firm revenue, whereas intermediate costs enter price propagation with coefficient $1-\alpha$.

between the decentralized allocation and the planner's allocation. A higher $\tilde{\Lambda}^{sp}(z)$ means that a given proportional reduction in the firm's production cost has a larger effect on the final-consumption price index. Holding the remaining components fixed, this increases the social value of sourcing capacity and lowers the planner's cutoff relative to the decentralized cutoff.

The third factor captures the aggregate resource feedback from relationship maintenance. If the planner's allocation has more active relationships in aggregate, maintenance absorbs additional labor and $L^{Y,sp}/L^{Y*} < 1$. The marginal welfare cost of maintaining another relationship, $\bar{f}/L^Y$, consequently rises. Holding the other factors fixed, this reduces the decentralized-to-planner cutoff ratio and limits further network expansion. Production labor therefore provides an economy-wide feedback through which the resource requirements of the selected network affect every bilateral selection decision.

Together, Propositions 2 and 3 decompose the channels of network formation inefficiency. The appropriation wedge, the markup-induced centrality wedge and the centrality adjustment are, a priori, the main reasons for inefficient network formation in presence of markups. Importantly, we studied the network externality from a constrained planner perspective, and markups also distort the incentives to form links beyond double marginalization.

Because these components are jointly determined, the decomposition does not imply universal underprovision or assign independent causal effects to its factors. It is an exact decomposition of selection differences, rather than an additive decomposition of welfare gains. The question of under or overprovision of links is of quantitative nature and is answered next.

VI. B Quantification of Decentralized Equilibrium Inefficiency

To quantify the welfare consequences of the inefficiencies arising from the endogenous network formation decisions, the problem of the constrained planner is solved. As mentioned before, the decentralized problem involves the solution of two fixed points. The first fixed point is conditional on the network and captures the economic behavior of the firms. The second fixed point concerns firms' private decision to form links given the solution of prices from the previous step. In the case of the constrained planner, the second fixed point is replaced by the optimality condition balancing the marginal social benefit and the resource cost for each relationships from equation 37. Hence, the constrained planner also involves the solution to a fixed point problem.

Table 4 presents a comparison of the results for the decentralized and the constrained planner equilibrium. The results shows that the decentralized equilibrium generates a sub-optimal allocation, though the aggregate welfare gap is quantitatively modest. Specifically, moving from the decentralized equilibrium to the constrained planner’s optimal network yields a welfare gain of 0.015% in consumption-equivalent terms. Although small, this is the estimated welfare gain the economy would have from only solving the network externality, as there are other distortions present in the model, such as markups.

The decentralized equilibrium presents a network that is thinner than the social optima. Table 4 shows that network density is 24.8% higher for the constrained planner allocation. This directly traduces into a lower aggregate price, and to less labor allocated for production. Therefore, the welfare gain from reducing the aggregate price is offset by 81.7% because of the decrease in total expenditures. Analogous to the propagation of shocks, despite a significantly denser network, the welfare gains from moving to the constrained social optimum are limited.

Table 4: Welfare Decomposition: Decentralized vs. Constrained Planner

Aggregate Variable	Decentralized	Planner	Δ (Planner - Dec)
log Welfare (W)	27.7813	27.7815	0.00015
Network Density	0.0017	0.0021	0.0004
log Price Index (P)	-22.0528	-22.0536	-0.00079
log Total Expenditures (E)	5.7286	5.7279	-0.00065
log Production Labor (L^Y)	4.6026	4.6019	-0.00065
log Maintenance Labor (L^O)	-1.3483	-1.1267	0.2216
Maintenance Labor Share ($L^O/\bar{L}$)	0.0026	0.0032	0.0006


Notes: This table decomposes the macroeconomic efficiency gains of the constrained planner relative to the decentralized equilibrium. Welfare (W) is defined as real consumption, such that $\log W = \log E - \log P$.

To understand the sources of the underprovision of business links in the decentralized equilibrium, I now examine the results for the ratio of the quality threshold using Propositions 2 and 3. Figure 7 presents the heatmap of the log of the ratio between the decentralized and planner’s threshold for every bilateral pair in the grid of seller and buyer types. The intensity of the color measures the magnitude of the gap.

Two salient features emerges. First, across all the distribution of seller and buyer types, the decentralized equilibrium underprovides links, meaning that the social planner’s threshold is lower. Second, decentralized threshold are roughly 19.3% to 23.5% higher than the planner allocation, indicating variation across type

pairs.

Figure 7: Heatmap for relative thresholds




Notes: Each cell reports $\log[r^*(z, z')/r^{sp}(z, z')]$, where r^* is the decentralized relationship cutoff and r^{sp} is the computed constrained-planner cutoff. Buyers appear on the horizontal axis and suppliers on the vertical axis, with productivity increasing from left to right and from bottom to top. All values are positive: decentralized cutoffs are higher and link probabilities are lower than in the planner allocation throughout the grid. Darker red indicates a larger log-cutoff gap. Cells have equal visual size but represent different population masses.

Propositions 2 and 3 are ideal to understand the sources of the heterogeneity in the threshold ratio. Results indicate that most of the underprovision of links emerges from the appropriability wedge. Suppliers earn variable profits equal to $1/\sigma$ of bilateral revenue, whereas the marginal sourcing benefit enters the price-index response with coefficient $1/(\sigma - 1)$. Since the model features CES demand, the appropriability wedge is constant throughout the network because it depends on the elasticity of substitution only.

Although the appropriability wedge is the most relevant channel to explain the size of the inefficiency, the remaining channels accounts for the heterogeneous patterns. Figure 8 presents the heatmap for the social versus private scale of the firm in Panel 8a, and the heatmap with the general equilibrium relative price effects in Panel 8b. The social-scale component is negative for all buyer types. Accordingly, the planner's cost-based social scale, $(1 - \alpha)\tilde{\Lambda}^{sp}(z)L^{Y,sp}$, is smaller than decentralized intermediate expenditure $M^*(z)$ throughout the grid. This component partially offsets the appropriation term, although it is not large enough to

Figure 8: Heatmaps from Proposition 2



Notes: The panels report two additive components of $(\sigma - 1) \log[r^*(z, z') / r^{sp}(z, z')]$ from Proposition 2. Panel (a) shows $\log[(1 - \alpha) \tilde{\Lambda}^{sp}(z) L^{Y,sp} / M^*(z)]$, comparing the planner's cost-based social scale with decentralized buyer expenditure. This component depends only on buyer type and is therefore constant across suppliers within each column. Panel (b) shows $\log[(q^* / \omega^*) / (q^{sp} / \omega^{sp})]$, capturing changes in supplier prices relative to buyers' input price indices. Positive values increase the decentralized-to-planner cutoff ratio; negative values reduce it. Buyers appear on the horizontal axis and suppliers on the vertical axis, with productivity increasing from left to right and from bottom to top. Each panel has its own color scale. Entries are cutoff-decomposition terms, not welfare contributions or changes in link probabilities.


reverse the direction of the cutoff gap. Its negative contribution is particularly pronounced for buyers at the low-productivity end of the distribution.

The relative-price component takes both signs. Because $q/\omega = [p(z')/P(z)]^{\sigma-1}$, this component captures changes in supplier prices relative to buyers' input price indices. It amplifies the cutoff gap for some pairs and reduces it for others. In the numerical solution, its contribution is negative for the pairing of the highest-productivity supplier with the lowest-productivity buyer, and positive for the opposite pairing. This variation illustrates how a common appropriation distortion produces heterogeneous selection gaps through equilibrium relative prices.

The negative contribution of the social scale component can be further analyzed using Proposition 3. Results are presented in Figure 9, where Panel 9a shows the markup-induced centrality wedge, and Panel 9b shows the centrality adjustment wedge. The remaining term in Proposition 3 is the Production labor wedge, which is by definition common to all type pairs. Results show that the production labor wedge is relatively small and negative because the denser planner network requires more maintenance labor, reducing the labor available for production.

The markup-induced centrality term is strongly negative for the least produc-

Figure 9: Heatmaps from Proposition 3



Notes: The panels report two components of the social-scale decomposition in Proposition 3. Panel (a) shows the markup-induced centrality term, $\log[\mu\alpha/(\mu-1+\alpha)] + \log[\tilde{\Lambda}^*(z)/\Lambda^*(z)]$, where $\Lambda^*(z) = S^*(z)/E^*$ is the revenue-based Domar weight and $\tilde{\Lambda}^*(z)$ is the cost-based Domar weight. Both are evaluated at the decentralized equilibrium. Panel (b) shows the centrality adjustment, $\log[\tilde{\Lambda}^{sp}(z)/\tilde{\Lambda}^*(z)]$, measuring the change in cost-based centrality between the decentralized and planner allocations. Together with the omitted, uniform production-labor term $\log(L^{Y,sp}/L^{Y*})$, these components sum to the social-scale term in Proposition 2. Positive values increase the decentralized-to-planner cutoff ratio; negative values reduce it. Buyers appear on the horizontal axis and suppliers on the vertical axis, with productivity increasing from left to right and from bottom to top. Each panel uses its own color scale. Entries are additive contributions to $(\sigma-1)\log(r^*/r^{sp})$.

tive buyers and marginally positive at the highest-productivity end of the grid. At the decentralized allocation, the normalized cost-propagation importance of low-productivity buyers is therefore smaller than their private expenditure scale would suggest. This moderates the incentive for planner expansion relative to the appropriation coefficient alone.

In contrast, the centrality-adjustment term largely works in the opposite direction. It is relatively high and positive for the least productive buyers and slightly negative for the most productive buyers. The planner's network raises the cost-based centrality of low-productivity buyers while slightly reducing that of the most productive buyers. For low-productivity buyers, this adjustment partially offsets the negative markup-induced centrality term. The two components should therefore be distinguished: one describes the discrepancy between revenue-based and cost-based importance at the decentralized allocation, whereas the other describes how cost-based importance changes when the network is reallocated.

Taken together, Propositions 2 and 3 show that widespread underprovision in the computed allocation is driven primarily by the appropriation component,

with buyer-specific valuation and relative-price adjustments modifying its magnitude. The relatively small net social-scale term conceals larger, partly offsetting changes in markup-induced centrality and equilibrium centrality. These selection differences coexist with a modest aggregate welfare gain because changing thresholds affects both productive sourcing opportunities and the resources required to maintain relationships.

VII Conclusion

This paper provides evidence and a theoretical framework to understand the aggregate consequences of the microeconomic forces shaping network formation. Using comprehensive firm-to-firm transaction data from Chile, the empirical analysis exploits business relationship life-cycle patterns to understand the link-level fundamental driving the evolution of the network. Guided by the empirical facts, a model of endogenous network formation with a match-specific quality shifter is built.

The theoretical model is used to quantify the relevance of network adjustment for the propagation of shocks, and how efficient are the network adjustment decisions. In both cases, results indicate that despite high link churning, the welfare effects are quantitatively small.

For the propagation of shocks, even for a large negative shock accompanied by a significant adjustment in the network, the welfare change is almost completely explained by the fixed network. Therefore, the role of endogenous network adjustment is limited because of opposing forces that dampens the effect of network adjustment. On one side, the creation of new links (network extension) mitigates the negative effect of the productivity shock. On the other side, because new links are close to the margin of the activation threshold, the average quality decreases significantly.

For the efficiency of network formation decisions, a constrained planner is used as benchmark. The planner's problem is characterized and solved. The results indicate that the decentralized network is entirely underprovided of business links. The main source of inefficiency is the appropriability wedge, where sellers can appropriate only a fraction of the revenues of the link, while the social benefits are larger. Despite the generalized underprovision of links, the welfare cost is low.

Taken together, the results show that substantial relationship turnover and widespread underprovision of links need not imply large aggregate gains from

network adjustment. In the calibrated economy, the productive contribution of marginal relationships and their resource requirements are key to understanding this disconnect.

References

- Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi**, “The network origins of aggregate fluctuations,” *Econometrica*, 2012, 80 (5), 1977–2016.
- Arkolakis, Costas, Federico Huneus, and Yuhei Miyauchi**, “Production network formation, trade, and welfare,” Technical Report, CESifo 2025.
- Baqae, David Rezza and Emmanuel Farhi**, “The macroeconomic impact of microeconomic shocks: Beyond Hulten’s theorem,” *Econometrica*, 2019, 87 (4), 1155–1203.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta**, “Measuring and analyzing cross-country differences in firm dynamics,” in “Producer dynamics: New evidence from micro data,” University of Chicago Press, 2009, pp. 15–76.
- Bernard, Andrew B, Emmanuel Dhyne, Glenn Magerman, Kalina Manova, and Andreas Moxnes**, “The origins of firm heterogeneity: A production network approach,” *Journal of Political Economy*, 2022, 130 (7), 1765–1804.
- Carvalho, Vasco M**, “From micro to macro via production networks,” *Journal of Economic Perspectives*, 2014, 28 (4), 23–48.
- Davis, Steven J and John Haltiwanger**, “Gross job creation and destruction: Microeconomic evidence and macroeconomic implications,” *NBER macroeconomics annual*, 1990, 5, 123–168.
- Dhyne, Emmanuel, Ayumu Ken Kikkawa, Magne Mogstad, and Felix Tintelnot**, “Trade and domestic production networks,” *The Review of Economic Studies*, 2021, 88 (2), 643–668.
- , —, **Xianglong Kong, Magne Mogstad, and Felix Tintelnot**, “Endogenous production networks with fixed costs,” *Journal of International Economics*, 2023, 145, 103841.
- Hsieh, Chang-Tai and Peter J Klenow**, “Misallocation and manufacturing TFP in China and India,” *The Quarterly journal of economics*, 2009, 124 (4), 1403–1448.
- Koike-Mori, Yasutaka and Koki Okumura**, “Production Networks and R&D Misallocation,” *unpublished, UCLA*, 2024.
- Lim, Kevin**, “Endogenous Production Networks and the Business Cycle,” in “in” 2018.

Loecker, Jan De and Chad Syverson, “An industrial organization perspective on productivity,” in “Handbook of industrial organization,” Vol. 4, Elsevier, 2021, pp. 141–223.

Monarch, Ryan and Tim Schmidt-Eisenlohr, “Longevity and the value of trade relationships,” *Journal of International Economics*, 2023, 145, 103842.

Oberfield, Ezra, “A theory of input–output architecture,” *Econometrica*, 2018, 86 (2), 559–589.

A Derivation of the Decentralized Relationship Cutoff

Let z denote the buyer's productivity and z' the supplier's productivity. Match quality is denoted by r , the supplier's price by $p(z')$, and the buyer's intermediate-input price index by $\mathbb{P}(z)$. Firms charge the constant markup $\mu = \sigma/(\sigma - 1)$, with $\sigma > 1$. Throughout, wages are normalized to one, so the maintenance cost of an active relationship is $\bar{f}$. All equilibrium objects below refer to the decentralized allocation.

Bilateral revenue and private profitability. Conditional on an active relationship, CES demand implies bilateral sales revenue

$$m(z, z', r) = \left(\frac{p(z')}{r \mathbb{P}(z)} \right)^{1-\sigma} M(z) = r^{\sigma-1} \left(\frac{p(z')}{\mathbb{P}(z)} \right)^{1-\sigma} M(z), \quad (41)$$

where $M(z)$ is the buyer's total expenditure on intermediate inputs.

Because price is μ times marginal cost, variable profit is the fraction

$$1 - \frac{1}{\mu} = \frac{1}{\sigma} \quad (42)$$

of sales revenue. Net profit from the relationship is therefore

$$\pi(z, z', r) = \frac{m(z, z', r)}{\sigma} - \bar{f}. \quad (43)$$

Since bilateral revenue increases with r , the supplier maintains the relationship if and only if $r \geq r^*(z, z')$. At an interior cutoff, the marginal relationship earns zero net profit:

$$\frac{1}{\sigma} [r^*(z, z')]^{\sigma-1} \left(\frac{p(z')}{\mathbb{P}(z)} \right)^{1-\sigma} M(z) = \bar{f}. \quad (44)$$

Rearranging gives the price-based cutoff:

$$r^*(z, z') = \frac{p(z')}{\mathbb{P}(z)} \left(\frac{\sigma \bar{f}}{M(z)} \right)^{\frac{1}{\sigma-1}}. \quad (45)$$

Representation using sourcing capacity and shares. Quality-adjusted bilateral sourcing capacity is

$$q(z, z') = \int_{r^*(z, z')}^{\infty} r^{\sigma-1} dH(r). \quad (46)$$

The buyer's intermediate-input price index satisfies

$$\mathbb{P}(z)^{1-\sigma} = \int q(z, z') p(z')^{1-\sigma} dG(z'). \quad (47)$$

Define the sourcing-share density with respect to the supplier distribution G as

$$\omega(z, z') = \frac{q(z, z') p(z')^{1-\sigma}}{\mathbb{P}(z)^{1-\sigma}}, \quad \int \omega(z, z') dG(z') = 1. \quad (48)$$

For type pairs with positive sourcing capacity,

$$\frac{\omega(z, z')}{q(z, z')} = \left(\frac{p(z')}{\mathbb{P}(z)} \right)^{1-\sigma}. \quad (49)$$

Substituting into the zero-profit condition yields

$$[r^*(z, z')]^{\sigma-1} = \sigma \bar{f} \frac{q(z, z')}{\omega(z, z')} \frac{1}{M(z)}. \quad (50)$$

Representation using revenue-based Domar weights. Let $S(z)$ denote total firm revenue, including intermediate and final sales. Total variable production cost is $S(z)/\mu$. Cost minimization allocates the fraction $1 - \alpha$ of this cost to intermediate inputs, implying

$$M(z) = \frac{1 - \alpha}{\mu} S(z). \quad (51)$$

Let E denote aggregate household expenditure and define the revenue-based Domar weight as

$$\Lambda(z) = \frac{S(z)}{E}. \quad (52)$$

Consequently,

$$M(z) = \frac{1 - \alpha}{\mu} \Lambda(z) E. \quad (53)$$

Substitution into (50) gives

$$\boxed{[r^*(z, z')]^{\sigma-1} = \sigma \bar{f} \frac{q(z, z')}{\omega(z, z')} \frac{\mu}{(1 - \alpha) \Lambda(z) E}.} \quad (54)$$

This expression uses the revenue-based weight $\Lambda(z)$, rather than the cost-based propagation weight $\tilde{\Lambda}(z)$ appearing in the welfare decomposition.

The cutoff increases with the supplier's price relative to the buyer's input price index and with maintenance costs, and decreases with the buyer's intermediate ex-

penditure, holding the remaining equilibrium objects fixed. Because prices, expenditure, and sourcing shares themselves depend on relationship selection, these expressions characterize an equilibrium cutoff rather than a solution solely in terms of primitives.

Aggregate expenditure and production labor. We can express aggregate expenditure in terms of production labor. With wages normalized to one, cost minimization implies

$$L^Y = \frac{\alpha}{\mu} \int S(z) dG(z), \quad \int S(z) dG(z) = \frac{\mu}{\alpha} L^Y. \quad (55)$$

Aggregate profits net of relationship maintenance costs are

$$\Pi = \frac{\mu-1}{\mu} \int S(z) dG(z) - L^O = \frac{\mu-1}{\alpha} L^Y - L^O, \quad (56)$$

where L^O is maintenance labor. Households receive all labor income and aggregate profits, so their expenditure satisfies $E = \bar{L} + \Pi$. Using labor-market clearing, $\bar{L} = L^Y + L^O$, yields

$$E = L^Y + L^O + \frac{\mu-1}{\alpha} L^Y - L^O = \frac{\mu-1+\alpha}{\alpha} L^Y. \quad (57)$$

Maintenance wage payments cancel against maintenance expenses in aggregate income accounting. Maintenance nevertheless reduces expenditure by absorbing labor that would otherwise be available for production.

Substituting (57) into (54) gives

$$\boxed{[r^*(z, z')]^{\sigma-1} = \frac{\sigma\mu\alpha}{(1-\alpha)(\mu-1+\alpha)} \frac{\bar{f}}{\Lambda(z)L^Y} \frac{q(z, z')}{\omega(z, z')}}. \quad (58)$$

Equivalently, the decentralized cutoff is

$$r^*(z, z') = \left[\frac{\sigma\mu\alpha}{(1-\alpha)(\mu-1+\alpha)} \frac{\bar{f}}{\Lambda(z)L^Y} \frac{q(z, z')}{\omega(z, z')} \right]^{\frac{1}{\sigma-1}}. \quad (59)$$

B Decomposition of the Social-Scale Component

This appendix derives the decomposition in Proposition 3. A superscript $*$ denotes the decentralized equilibrium, while sp denotes the constrained-planner allocation. Let $S^*(z)$ denote firm revenue, E^* aggregate household expenditure, and

$M^*(z)$ intermediate expenditure. The revenue-based Domar weight is $\Lambda^*(z) = S^*(z)/E^*$, whereas $\tilde{\Lambda}^*(z)$ denotes the cost-based Domar weight defined in Proposition 1. These weights are distinct. Wages are normalized to one, and the markup μ and production parameter α are common across firms and allocations.

Step 1: Private expenditure and firm revenue. Under constant-markup pricing, variable production cost equals revenue divided by μ . Cost minimization allocates the fraction $1 - \alpha$ of variable cost to intermediate inputs. Therefore,

$$M^*(z) = \frac{1 - \alpha}{\mu} S^*(z) = \frac{1 - \alpha}{\mu} \Lambda^*(z) E^*. \quad (60)$$

Step 2: Aggregate expenditure and production labor. Aggregate production labor satisfies

$$L^{Y*} = \frac{\alpha}{\mu} \int S^*(z) dG(z). \quad (61)$$

Aggregate profits net of maintenance costs are

$$\Pi^* = \frac{\mu - 1}{\mu} \int S^*(z) dG(z) - L^{F*} = \frac{\mu - 1}{\alpha} L^{Y*} - L^{F*}. \quad (62)$$

The household budget is $E^* = \bar{L} + \Pi^*$. Combining this identity with labor-market clearing, $\bar{L} = L^{Y*} + L^{F*}$, gives

$$E^* = \frac{\mu - 1 + \alpha}{\alpha} L^{Y*}. \quad (63)$$

Step 3: The baseline social-private scale gap. Using (60),

$$\frac{(1 - \alpha) \tilde{\Lambda}^*(z) L^{Y*}}{M^*(z)} = \frac{\mu \tilde{\Lambda}^*(z) L^{Y*}}{\Lambda^*(z) E^*}. \quad (64)$$

Substituting (63) then yields

$$\boxed{\frac{(1 - \alpha) \tilde{\Lambda}^*(z) L^{Y*}}{M^*(z)} = \frac{\mu \alpha}{\mu - 1 + \alpha} \frac{\tilde{\Lambda}^*(z)}{\Lambda^*(z)}}. \quad (65)$$

This identity compares social propagation and private expenditure at the same decentralized allocation. The scalar factor converts revenue-based expenditure units into production-labor units. The ratio of Domar weights captures the distinction between a firm's importance for cost propagation and its revenue relative to final expenditure. No equality between the two Domar weights is imposed.

Step 4: Changes between allocations. Multiply and divide the social-scale component by $\tilde{\Lambda}^*(z)L^{Y*}$:

$$\frac{(1-\alpha)\tilde{\Lambda}^{sp}(z)L^{Y,sp}}{M^*(z)} = \frac{(1-\alpha)\tilde{\Lambda}^*(z)L^{Y*}}{M^*(z)} \frac{\tilde{\Lambda}^{sp}(z)}{\tilde{\Lambda}^*(z)} \frac{L^{Y,sp}}{L^{Y*}}. \quad (66)$$

Applying (65) establishes Proposition 3:

$$\boxed{\begin{aligned} \frac{(1-\alpha)\tilde{\Lambda}^{sp}(z)L^{Y,sp}}{M^*(z)} &= \underbrace{\frac{\mu\alpha}{\mu-1+\alpha} \frac{\tilde{\Lambda}^*(z)}{\Lambda^*(z)}}_{\text{Baseline social-private scale gap}} \\ &\times \underbrace{\frac{\tilde{\Lambda}^{sp}(z)}{\tilde{\Lambda}^*(z)}}_{\text{Centrality adjustment}} \underbrace{\frac{L^{Y,sp}}{L^{Y*}}}_{\text{Production-labor adjustment}}. \end{aligned}} \quad (67)$$

The first factor is evaluated entirely at the decentralized equilibrium. The second records the change in the buyer's cost-based centrality, and the third records the change in aggregate labor available for production. The decomposition is an exact multiplicative identity, not a set of independent comparative statics. Neither the baseline scale gap nor the centrality adjustment is necessarily greater than one. If the planner's allocation uses more maintenance labor, the production-labor adjustment is below one.

C Numerical Decomposition of Welfare Changes

This appendix describes how we decompose finite welfare changes using Proposition 1, Lemma 1, and Corollary 1. The procedure distinguishes the contributions of productivity, production labor, link expansion, and match-quality composition. Because Proposition 1 is a differential identity, we evaluate its components along a sequence of equilibria connecting the baseline and shocked economies. This allows propagation weights and sourcing shares to adjust with the shock.

Counterfactual equilibria

Let δ denote the full log-productivity shock and let $v(z)$ indicate the firms targeted by the experiment. We index firms by their baseline types and define

$$\log z_t(z) = \log z + t\delta v(z), \quad t \in [0, 1]. \quad (68)$$

The targeted set is selected in the baseline and remains fixed. The measure G continues to assign the same mass to each group of firms; it is not recomputed by evaluating the baseline density at shocked productivities. All other primitives, including the distribution of match quality, are held fixed. A proportional productivity increase of g corresponds to $\delta = \log(1 + g)$.

At each t , we solve the decentralized endogenous-network equilibrium. Prices, sales, relationship cutoffs, and labor allocation therefore adjust jointly. The parameter t indexes the magnitude of the counterfactual shock, not calendar time. The resulting sequence is a numerical construction for decomposing a comparative-static welfare change.

We assume that the equilibria belong to a differentiable branch with positive production labor and positive bilateral sourcing capacities. This is a local regularity requirement and does not establish global uniqueness of equilibrium. Throughout, z indexes buyers and z' indexes suppliers in bilateral objects.

The welfare channels in Proposition 1

With wages as the numeraire, Proposition 1 gives

$$d \log W = d \log L^Y + \tilde{\Lambda}^T \left[d \log z + \frac{1 - \alpha}{\sigma - 1} d \log Q \right], \quad (69)$$

where $\tilde{\Lambda}$ denotes the cost-based propagation weights and the sourcing-capacity differential satisfies

$$d \log Q_t(z) = \int \omega_t(z, z') d \log q_t(z, z') dG(z'). \quad (70)$$

The sourcing shares $\omega_t(z, z')$ are evaluated at the equilibrium associated with t and satisfy $\int \omega_t(z, z') dG(z') = 1$.

Integrating equation (69) gives

$$\log W_1 - \log W_0 = C_L + C_z + C_Q, \quad (71)$$

where

$$C_L = \log L_1^Y - \log L_0^Y, \quad (72)$$

$$C_Z = \int_0^1 \int \tilde{\Lambda}_t(z) \delta v(z) dG(z) dt, \quad (73)$$

$$C_Q = \frac{1-\alpha}{\sigma-1} \int_0^1 \int \int \tilde{\Lambda}_t(z) \omega_t(z, z') \frac{\partial \log q_t(z, z')}{\partial t} dG(z') dG(z) dt. \quad (74)$$

The labor contribution measures the welfare effect of changes in labor available for production as relationship maintenance adjusts. The productivity contribution values the exogenous productivity changes using the propagation weights prevailing at each equilibrium. The sourcing contribution measures the effect of changes in bilateral productive capacity, accounting for both sourcing shares and propagation through the network.

Link expansion and match-quality composition

Lemma 1 decomposes bilateral sourcing capacity as

$$q_t(z, z') = \ell_t(z, z') \bar{r}_t(z, z'), \quad (75)$$

where

$$\ell_t(z, z') = \Pr(r \geq r_t^*(z, z')), \quad (76)$$

$$\bar{r}_t(z, z') = \mathbb{E} \left[r^{\sigma-1} \mid r \geq r_t^*(z, z') \right]. \quad (77)$$

Here $\bar{r}$ is the average productive intensity of an active relationship, rather than its average quality level.

Taking logarithms yields the exact identity

$$d \log q_t(z, z') = d \log \ell_t(z, z') + d \log \bar{r}_t(z, z'). \quad (78)$$

Substituting into equation (74) gives

$$C_Q = C_\ell + C_{\bar{r}}, \quad (79)$$

with

$$C_\ell = \frac{1-\alpha}{\sigma-1} \int_0^1 \int \int \tilde{\Lambda}_t(z) \omega_t(z, z') \frac{\partial \log \ell_t(z, z')}{\partial t} dG(z') dG(z) dt, \quad (80)$$

$$C_{\bar{r}} = \frac{1-\alpha}{\sigma-1} \int_0^1 \int \int \tilde{\Lambda}_t(z) \omega_t(z, z') \frac{\partial \log \bar{r}_t(z, z')}{\partial t} dG(z') dG(z) dt. \quad (81)$$

Consequently, the welfare change admits the four-channel decomposition

$$\boxed{\log W_1 - \log W_0 = C_L + C_z + C_\ell + C_{\bar{r}}.} \quad (82)$$

The link-expansion contribution records changes in the probability of an active relationship. The composition contribution records changes in average productive intensity among active relationships. These components describe jointly determined selection responses; they are not separate counterfactual experiments.

The role of the marginal relationship

Corollary 1 sharpens the interpretation of this decomposition. Holding the quality distribution H fixed, define

$$\theta_t(z, z') = \frac{r_t^*(z, z')^{\sigma-1}}{\bar{r}_t(z, z')}. \quad (83)$$

Assume that H has a continuous density h , that the cutoff is positive, and that the relevant probabilities and moments are positive and finite. Differentiating the cutoff integrals gives

$$d\ell = -h(r^*) dr^*, \quad dq = -(r^*)^{\sigma-1} h(r^*) dr^*. \quad (84)$$

Using $q = \ell \bar{r}$ therefore yields

$$d \log q = \theta d \log \ell, \quad (85)$$

$$d \log \bar{r} = -(1 - \theta) d \log \ell. \quad (86)$$

For $\sigma > 1$, active relationships have productive intensity at least as large as that of the marginal relationship, so $0 < \theta \leq 1$.

A reduction in the cutoff expands the active set but admits relationships with lower productive intensity than the existing average. Link expansion therefore increases sourcing capacity, while composition offsets part of that increase. The ra-

tio θ measures the fraction of the proportional link expansion that translates into a proportional capacity increase. Conversely, raising the cutoff reduces link probability while increasing average productive intensity among surviving relationships.

The corollary provides equivalent expressions for the integrated contributions:

$$C_Q = \frac{1-\alpha}{\sigma-1} \int_0^1 \int \int \tilde{\Lambda}_t(z) \omega_t(z, z') \theta_t(z, z') \frac{\partial \log \ell_t(z, z')}{\partial t} dG(z') dG(z) dt, \quad (87)$$

$$C_{\bar{r}} = -\frac{1-\alpha}{\sigma-1} \int_0^1 \int \int \tilde{\Lambda}_t(z) \omega_t(z, z') [1 - \theta_t(z, z')] \frac{\partial \log \ell_t(z, z')}{\partial t} dG(z') dG(z) dt. \quad (88)$$

These expressions also provide a numerical check on the decomposition obtained directly from Lemma 1. Aggregate contributions depend on the pattern of cutoff changes and their welfare weights; their signs need not follow the response of any particular type pair.

Numerical implementation

We divide the full shock into K equal increments and solve equilibria at $t_k = k/K$, for $k = 0, \dots, K$. Consecutive productivity vectors differ by

$$\Delta_k \log z(z) = \frac{\delta}{K} v(z). \quad (89)$$

Each equilibrium is initialized using the preceding solution. All solutions must satisfy the equilibrium convergence criteria before entering the decomposition.

Define the combined bilateral welfare weight

$$A_k(z, z') = \frac{1-\alpha}{\sigma-1} \tilde{\Lambda}_{t_k}(z) \omega_{t_k}(z, z'). \quad (90)$$

For any positive equilibrium object x , write $\Delta_k \log x = \log x_{t_{k+1}} - \log x_{t_k}$. We ap-

proximate the integrated contributions using endpoint averages:

$$\widehat{C}_z = \sum_{k=0}^{K-1} \int \frac{\widetilde{\Lambda}_{t_k}(z) + \widetilde{\Lambda}_{t_{k+1}}(z)}{2} \Delta_k \log z(z) dG(z), \quad (91)$$

$$\widehat{C}_\ell = \sum_{k=0}^{K-1} \iint \frac{A_k(z, z') + A_{k+1}(z, z')}{2} \Delta_k \log \ell(z, z') dG(z') dG(z), \quad (92)$$

$$\widehat{C}_{\bar{r}} = \sum_{k=0}^{K-1} \iint \frac{A_k(z, z') + A_{k+1}(z, z')}{2} \Delta_k \log \bar{r}(z, z') dG(z') dG(z). \quad (93)$$

The labor contribution is computed directly from the endpoints, $C_L = \log L_1^Y - \log L_0^Y$.

Because the factorization $q = \ell \bar{r}$ also implies

$$\Delta_k \log q = \Delta_k \log \ell + \Delta_k \log \bar{r}, \quad (94)$$

the two numerical sourcing components sum to the sourcing contribution computed from $\Delta_k \log q$, up to numerical precision, when identical weights are used.

For the corollary-based check, we replace the integrand weights by endpoint averages of the products $A\theta$ and $A(1 - \theta)$, respectively. These alternative calculations need not coincide exactly with the Lemma 1 calculations at finite K , because the corollary is a differential identity. They should converge to the same contributions as the integration grid is refined.

In the lognormal specification, link probabilities and productive capacities are evaluated using analytical cutoff integrals. This avoids the artificial jumps in selection generated by a finite quality grid. Baseline, intermediate, and final equilibria use the same integration method.


We assess accuracy by increasing K and reporting the accounting residual

$$R_K = \log W_1 - \log W_0 - \left[C_L + \widehat{C}_z + \widehat{C}_\ell + \widehat{C}_{\bar{r}} \right]. \quad (95)$$

Under the stated regularity conditions, this residual should approach zero as numerical integration and equilibrium solutions become sufficiently accurate.

Interpretation. The decomposition is additive in log welfare units. Its components are evaluated along the proportional log-productivity experiment in equation (68); individual contributions can depend on the specified counterfactual se-

Figure A1: Distribution of Links with observed birth



Note:

quence, although their exact sum equals the endpoint welfare change.

In particular, C_z uses propagation weights that adjust under endogenous network selection. It need not equal the finite welfare change from a separately solved fixed-network counterfactual. The decomposition and the fixed-network comparison therefore answer complementary questions: the former accounts for the channels operating within the endogenous equilibrium response, while the latter measures how the total response changes when network adjustment is restricted.

D Robustness






Figure A5: Robustness (Partial-Year): Hazard and Turnover by Link Tenure

Figure A2: Link Tenure and FOSD



Note:

Figure A3: Robustness (Partial-Year): Count Share by Link Tenure



(a) Count Share, 2021

(b) Sales Share, 2021

Note:






Figure A8: Robustness (Balanced Panel): Hazard and Turnover by Link Tenure

Figure A4: Robustness (Partial-Year): Link Value by Link Tenure



Note:

Figure A6: Robustness (Balanced Panel): Count Share by Link Tenure



Note:






Figure A11: Robustness (Excluding Retail/Wholesale): Hazard and Turnover by Link Tenure

Figure A7: Robustness (Balanced Panel): Link Value by Link Tenure



Note:

Figure A9: Robustness (Excluding Retail/Wholesale): Count Share by Link Tenure



Note:






Figure A14: Robustness (Excluding Capital-Intensive Sectors): Hazard and Turnover by Link Tenure

Figure A10: Robustness (Excluding Retail/Wholesale): Link Value by Link Tenure



Note:

Figure A12: Robustness (Excluding Capital-Intensive Sectors): Count Share by Link Tenure



Note:






Figure A17: Robustness (Cross Section): Hazard and Turnover by Link Tenure

Figure A13: Robustness (Excluding Capital-Intensive Sectors): Link Value by Link Tenure




Note:

Figure A15: Robustness (Cross Section): Count Share by Link Tenure



Note:

Figure A16: Robustness (Cross Section): Link Value by Link Tenure



(a) Count Share, 2021

(b) Sales Share, 2021

Note: